

ADAPTIVE CONTROL OF FISHING SYSTEMS

Carl J. Walters

Ray Hilborn

September 1975

WP-75-114

Working Papers are not intended for distribution outside of IIASA, and are solely for discussion and information purposes. The views expressed are those of the author, and do not necessarily reflect those of IIASA.

Adaptive Control of Fishing Systems

Carl J. Walters*
Ray Hilborn*

Abstract

This paper discusses some formal techniques for deciding how harvesting policies should be modified in the face of uncertainty. Parameter estimation and dynamic optimization methods are combined for the Ricker stock recruitment model to show how exploitation rates should be manipulated to give more information about the model parameters. In general, harvesting rates should be lower than would be predicted by the best fitting recruitment curve unless this curve predicts that the stock is very productive. A decision procedure is developed for comparing alternative stock recruitment models; when applied to the Fraser River sockeye salmon, the procedure indicates that an experimental increase in escapements would be worthwhile. It appears that there is considerable promise for extending these methods and procedures to cases where the stock size is unknown and where fishing effort is poorly controlled.

*Institute of Animal Resource Ecology, University of British Columbia, Vancouver, Canada, and International Institute for Applied Systems Analysis, Laxenburg, Austria.

I. Introduction

A variety of dynamic models have been used in recent years to establish fishery regulations and to catch quotas; increasingly, elaborate monitoring systems have been developed to provide sound statistical estimates of model parameters. Most models are used only to predict optimum equilibrium harvest rates, though there have been a few recent attempts to develop harvest strategy curves or "control laws" that specify optimum harvest rates for non-equilibrium situations (Allen, [1]; Walters, [5]). Given a time series from which model parameter estimates have been derived, it has often been assumed that the best management strategy is to act as though these estimates were actually correct: that is, insert the estimates into the model (or into several alternative models), generate a yield curve or an isopleth diagram that reveals an apparently optimum harvest policy, and recommend that this policy be followed.

Little attention has been paid to the problem that, by following the apparently optimum policy, the fishery might be brought to an equilibrium that is neither truly optimal nor productive of the type of data necessary to determine the true optimum. Luckily, most fisheries have gone through a period of more or less uncontrolled development of the point of obvious overexploitation, allowing us to interpolate an optimum regime from data on a broad range of stock sizes. Another problem subsequently arises: when a fishery has been held near some equilibrium for a long period of time, how much confidence can we have in older

data from the non-equilibrium period? Environmental carrying capacities may have changed, selection by the fishery may have produced new genetic types capable of responses, or the old data may be simply unreliable. Implicit in many research programs is the assumption that detailed biological studies on populations near equilibrium will allow a priori determination of optimum harvesting policies, thus making it unnecessary to introduce trial and error changes or large scale experiments in harvest rates; this assumption appears to be naive and unjustified at present.

This paper addresses the question of how harvesting decisions should be modified to take account of statistical uncertainty. In seeking a formal framework for dealing with this question, we have been drawn to the literature on control system theory, where the problem is addressed under the heading of "adaptive" or "dual" control (Larson, [8]). To simplify the discussion, we will concentrate primarily on situations where the stock-recruitment relationship (as opposed to growth and natural mortality) is the critical determinant of potential yield. The analysis is divided into two parts: a) we look at the case where a simple model, the Ricker stock-recruitment curve, is assumed to be the correct functional form and only the model parameters are uncertain; and b) we examine more general cases where the form or shape of the stock-recruitment function is uncertain, stock sizes are not directly measurable, and fishing effort is poorly controlled.

II. Adaptive Control with the Ricker Model

As indicated by several authors with many examples in Parrish [11], the simple model developed by Ricker [14] has been used often in the analysis of stock-recruitment relationships:

$$R_t = S_{t-1} \exp \left\{ \alpha - \beta S_{t-1} + v_t \right\} \quad (1)$$

where

R_t = recruits (adults) at end of generation t ;

S_{t-1} = spawners at the start of generation t ;

α = a stock production parameter ;

β = equilibrium stock parameter (equilibrium stock in absence of fishing is equal to $1/\beta$);

v_t = a random environmental factor normally distributed with mean 0.0 and variance σ^2 .

For the discussion that follows, it is critical that v_t , the noise factor, be normally distributed. There is good empirical evidence for this assumption for a sockeye salmon population on the Skeena River (Allen, [1]). A theoretical justification can be constructed by noting that e^{v_t} can be viewed as a random survival factor resulting from several independent and multiplicative environmental factors operating in series. (Thus v_t represents a sum of several random factors and should be normally distributed by the central limit theorem.)

Let us assume that the management objective is to maximize the sum of discounted catches over time:

$$\max \sum_{t=0}^{\infty} C_t e^{-\delta t} \quad (2)$$

where

$$C_t = \text{catch} = R_t - S_t ;$$

δ = a discount rate .

The discount rate is critical in adaptive control problems, since without $\delta = 0$ we would put all management emphasis on obtaining improved information for the distant future, regardless of the cost in terms of lost yields in the near future. It is known (Allen [1]; Walters [5]) that, for the objective in equation (2) when α and β are not uncertain, the optimum management policy is to allow a fixed escapement each year: choose

$$C_t = \begin{cases} R_t - \tilde{S} & \left(\text{if } R_t > \tilde{S} \right) \\ 0 & \left(\text{if } R_t \leq \tilde{S} \right), \end{cases} \quad (3)$$

where \tilde{S} is the optimum escapement, computed from α and β (Ricker [15]).

Ordinarily, we would recommend that management actions be based on estimates of \tilde{S} computed from regression estimates $\hat{\alpha}_t$ and $\hat{\beta}_t$. The Ricker model can be rewritten (after Dahlberg, [3]) in the form:

$$\ln\left(\frac{R_t}{S_{t-1}}\right) = \alpha - \beta S_{t-1} + v_t ; \quad (4)$$

this is a linear regression for ($y = \alpha + \beta x$) with:
 $y = \ln\left(\frac{R_t}{S_{t-1}}\right)$ and $x = -S_{t-1}$. We would probably ignore some useful information that comes from regression analysis,

namely, the parameter covariance matrix:

$$\hat{P}_t = \begin{bmatrix} \hat{\sigma}_\alpha^2 & \hat{\sigma}_{\alpha\beta} \\ \hat{\sigma}_{\alpha\beta} & \hat{\sigma}_\beta^2 \end{bmatrix}, \quad (5)$$

that measures our uncertainty about the parameter estimates given data to time t . Further, under the assumption that the v_t are normally distributed and independent of one another, it can be shown that $\hat{\alpha}_t$, $\hat{\beta}_t$, and the elements of \hat{P}_t , constitute a set of "sufficient statistics". That is, there is no other function or manipulation of data that can give additional information about the underlying true Ricker parameters.

The objective of adaptive control analysis in this case is to show how the choice of escapement S_t should be related to \hat{P} as well as to $\hat{\alpha}$ and $\hat{\beta}$. The analysis can be formulated as a problem in stochastic dynamic optimization (Walters [5]): given the system state at any time as measured by $\{R_t, \hat{\alpha}_t, \hat{\beta}_t, \hat{P}_t\}$, what choice of C_t will give the best expected combination of present return and future returns, recognizing that a variety of possible future states may occur because of random events?

To solve problems of this type, we must be able to formulate a model that specifies how each of the state variables (R , $\hat{\alpha}$, etc.) will change in relation to the variety of stochastic outcomes that may occur between times t and $t + 1$. We must also be able to assign probabilities to each of these stochastic outcomes. Future recruitment states (R_{t+1}) can be predicted with the Ricker model, but analogous predictive

formulae are required for the statistical parameters. The following section shows how these formulae can be derived from a special form of regression analysis.

A. Recursive or Adaptive Parameter Estimation

Suppose we begin at time $t = 0$ no data but with prior estimates $\hat{\alpha}_0$ and $\hat{\beta}_0$. We might wish to assign no confidence to these estimates, which is equivalent to saying that we recognize σ_α^2 and σ_β^2 to be very large, or:

$$\hat{P}_0 = \begin{bmatrix} L & 0 \\ 0 & L \end{bmatrix}, \quad (6)$$

where L is some large number (e.g. 10^6). In Bayesian statistical terms, we are in effect assigning a "diffuse prior" distribution for α and β (Raiffa and Schlaifer, [13]). With starting conditions such as these, it can be shown that ordinary regression analysis can be written in a special "recursive" format (Young [22]). The general format is presented here, since it may be of interest outside the adaptive control context.

In general, linear regression equations are written in the form:

$$y_i = \sum_{j=1}^m a_j x_{ij} + e_i$$

where

y_i = dependent observations;

x_{ij} = independent variables; and

e_i = error terms.

This form can be written more compactly in vector notation. That is:

$$y_i = \underline{a}^T \underline{x}_i + e_i, \quad (7)$$

where \underline{a} and \underline{x}_i represent the vectors (a_1, a_2, \dots, a_m) and $(x_{i1}, x_{i2}, \dots, x_{im})$, respectively. Using this notation, we could write the common regression formulae in recursive form as:

$$\hat{\underline{a}}_n = \hat{\underline{a}}_{n-1} - \frac{\hat{P}_{n-1} \underline{x}_n}{\sigma^2 + \underline{x}_n^T \hat{P}_{n-1} \underline{x}_n} (\hat{\underline{a}}_{n-1}^T \underline{x}_n - y_n) \quad , \quad (8a)$$

and

$$\hat{P}_n = \hat{P}_{n-1} - \frac{\hat{P}_{n-1} \underline{x}_n \underline{x}_n^T \hat{P}_{n-1}}{\sigma^2 + \underline{x}_n^T \hat{P}_{n-1} \underline{x}_n} \quad , \quad (8b)$$

where $\hat{\underline{a}}_n$ and \hat{P}_n refer to the parameter and parameter error covariance estimators, respectively, after the n^{th} data point is acquired, and σ^2 is the regression error variance. These formulae allow new data points to be added to a regression analysis without tedious computations involving matrix inversion. Estimation for the Ricker model can be written in the recursive form with:

$$y_n = \ln\left(\frac{R_t}{S_{t-1}}\right)$$

$$\underline{x}_n = \begin{bmatrix} 1 \\ -S_{t-1} \end{bmatrix} \quad (9)$$

$$\hat{\underline{a}}_n = \begin{bmatrix} \hat{\alpha}_t \\ \hat{\beta}_t \end{bmatrix} ;$$

similar transformations can be developed for a variety of other fisheries models.

Equations (8a) and (8b) are critical for the adaptive control formulation developed in Section III. Note that the change in parameter uncertainty from any observation or time step to the next, as measured by $\hat{P}_n - \hat{P}_{n-1}$, depends on σ^2 , \hat{P}_{n-1} , and on the choice of \underline{x}_n , i.e. on the choice of S_{t-1} for the Ricker model. Changes in the parameter estimates as measured by $\hat{a}_n - \hat{a}_{n-1}$ depend on:

- a) the level of uncertainty as measured by \hat{P}_{n-1} and σ^2 ;
- b) the choice of \underline{x}_n ; and
- c) the a priori prediction error, $D_n = (\underline{a}_{n-1}^T \underline{x}_n - y_n)$.

The a priori prediction error D_n is the difference between the observed y_n and its predicted value using the latest \underline{x}_n data and the older or prior parameter estimates, \hat{a}_{n-1} . This prediction error (that is the only uncontrolled or stochastic input into the \hat{a} and \hat{P} changes for any time step), can be rewritten as two error components:

$$D_n = (\underline{a}_{n-1}^T \underline{x}_n - y_n) + (\hat{a}_{n-1}^T - \underline{a}_{n-1}^T) \underline{x}_n \quad . \quad (10)$$

The first component is the deviation of y_n from the true model, while the second component represents deviation of the parameter estimates from the true value. If the regression errors v_t are normally distributed, both of these error components are normally distributed; thus D_n should have a normal distribution with mean zero and variance. That is:

$$\sigma_{D_n}^2 = \sigma^2 + \underline{x}_n^T \hat{P}_{n-1} \underline{x}_n \quad . \quad (11)$$

With data available up to any time step, we can compute probabilities for different values of D_n , and thus have different parameter estimates at the next time step. This is known as preposterior analysis in statistical decision theory (Raiffa [12]). For the Ricker model, D_n is interpreted as:

$$D_n = \hat{\alpha}_{t-1} - \hat{\beta}_{t-1} S_{t-1} - \ln\left(\frac{R_t}{S_{t-1}}\right) . \quad (12)$$

Having chosen a value for D_n , with its associated probability, we can predict R_t by solving equation (12) as:

$$\hat{R}_t = S_{t-1} \exp \left\{ \hat{\alpha}_{t-1} - \hat{\beta}_{t-1} S_{t-1} - D_n \right\} . \quad (13)$$

(given D_n)

This is the original Ricker model, with an error component that reflects not only the noise v_t but also the uncertainty about α and β .

The adaptive regression equations (8) can be modified to "forget" older data. There are two simple types of modifications: a) exponential past weighting of data, based on the assumption that all data become progressively less reliable as they become older; and b) parameter variance incrementation, based on the more specific assumption that the parameters do vary in some random or unspecifiable systematic way over time. For exponential past weighting, we define a discount factor V_d that represents the value of any observation relative to the next one that is obtained; for example, if we want to assume that an observation at time $t - 1$ is

worth 90% as much as an observation at time t , then $V_d = 0.9$. Using this discount factor, equations (8) are simply modified by changing the denominator terms $\sigma^2 + \mathbf{x}_n^T \hat{\mathbf{P}}_{n-1} \mathbf{x}_n$ to:

$$V_d \sigma^2 + \mathbf{x}_n^T \hat{\mathbf{P}}_{n-1} \mathbf{x}_n \quad (14)$$

When there is reason to believe that one or more parameters are changing over time, the estimation is modified by introducing a parameter variation matrix Q , where the elements of the matrix are chosen to reflect the expected rate of change in the parameters. For example, if we believe that the Ricker β parameter may change about 10% per year from an average value of about 10^{-6} (while the α parameter is stable), we could set:

$$Q = \begin{bmatrix} 0 & 0 \\ 0 & 10^{-14} \end{bmatrix} \quad (15)$$

In statistical terms, the elements of Q are interpreted as variances on a random walk process; thus a 10% change from a base of 10^{-6} represents a standard deviation of 10^{-7} , or a variance of $(10^{-7})^2$. The Q matrix is introduced into equations (8) simply by replacing every $\hat{\mathbf{P}}_{t-1}$ with $\hat{\mathbf{P}}_{t-1}^*$ where

$$\hat{\mathbf{P}}_{t-1}^* = \hat{\mathbf{P}}_{t-1} + Q \quad (16)$$

The choice of V_d or Q is not particularly critical; the major effect in both cases is to prevent $\hat{\mathbf{P}}_t$ from going to zero over time, so that new observations can continue to effect changes in $\hat{\mathbf{a}}$.

B. Adaptive Decision Structure and Optimization

The problem of adaptive control and optimization for the Ricker model can be visualized in terms of the decision tree shown in Figure 1. At any point in time, the manager is faced with a recruitment R_t , a summary of past data in terms of $\hat{\alpha}_t$ and $\hat{\beta}_t$, and with uncertainties in terms of σ^2 and the elements of \hat{P}_t . He must choose a harvest C_t ; there are many possible choices, but the optimization problem can be approximated by looking at a reduced, discrete set of possibilities (e.g. $C_t = 0$, $C_t = 0.1R_t$, $C_t = 0.2R_t$).

Given any choice of C_t , there are many possible random outcomes; these can be summarized in terms of discrete deviations D_t from the regression predictors of $y = \ln(R_{t+1}/S_t)$. The reasoning is as follows:

- a) Given C_t , S_t is calculated as $R_t - C_t$;
- b) $\hat{\alpha}_t$, $\hat{\beta}_t$, and S_t are used to make a regression prediction y_{t+1} [equations (7) and (9)];
- c) Probabilities for different outcomes $Y_{t+1} = \hat{y}_{t+1} + D_t$ are computed from the probability distribution for D_t , that is normal with mean zero and variance given by equation (11);
- d) Each outcome y_{t+1} is inserted, along with C_t , into the recursive regression equations (8) to obtain new estimates $\hat{\alpha}_{t+1}$, $\hat{\beta}_{t+1}$, \hat{P}_{t+1} ; and
- e) Since $y_{t+1} = \ln(R_{t+1}/S_t)$, R_{t+1} for each outcome y_{t+1} is given as $S_t e^{y_{t+1}}$ [equivalent to equation (13)].

If we know the total future value for being in any next state $\{R_{t+1}, \hat{\alpha}_{t+1}, \hat{\beta}_{t+1}, \hat{P}_{t+1}\}$, we can calculate the expected value for making a C_t decision. This expected value is the sum of products of probabilities of next states, times the values associated with these next states, plus the value of C_t itself. The difficulty is that we cannot immediately assign a value for each of the next states, since that state is itself a starting point for another decision tree similar to Figure 1. If we look ahead a few time steps, the number of branching possibilities becomes essentially infinite. There is a partial way out of this problem using the "backward recursion" procedure of dynamic programming. A simplified discussion of this procedure is given in Walters [5]; the basic idea is that we begin the optimization calculations at some time point far enough in the future that the discounted values after that point can be neglected. We then move backwards towards the present, evaluating decisions at each of the time steps in terms of future values that have just been computed for the next time steps forward.

Unfortunately, even dynamic programming involves formidable computation problems. If at each of the time steps we examine only ten discrete values for each of the six state variables $(R_t, \hat{\alpha}_t, \hat{\beta}_t, \sigma_{\alpha}^2, \sigma_{\beta}^2, \sigma_{\alpha\beta})$, 10 catch levels and 10 values of D_t , we must compute about 10^8 solutions for equations (8) and (13). The problem can be reduced somewhat by using special computation procedures (Larson [8]), but there is a clear need for different ways of looking at the problem. The optimization need only be carried out for a few representative

values of the environmental variance σ^2 and the discount rate δ in order to get a complete adaptive picture for the Ricker model. Given σ^2 and δ , the optimization automatically arrives at best harvest rates for all stock size-parameter value-parameter uncertainty combinations, in the form of a multidimensional "control law." (Allen [1] and Walters [5] have referred to one dimensional versions of this control law as "strategy curves.")

A further point worth noting is the size of adaptive optimization problems. Let us suppose that instead of the Ricker model we wish to analyze some model with three parameters (say, α_1 , α_2 , and α_3). Even if we can put this model into the linear regression form with normally distributed errors, the number of state variables for the dynamic optimization is 10 ($R_t, \hat{\alpha}_1, \hat{\alpha}_2, \hat{\alpha}_3, \sigma_{\alpha_1}^2, \sigma_{\alpha_2}^2, \sigma_{\alpha_3}^2, \sigma_{\alpha_1\alpha_2}, \sigma_{\alpha_1\alpha_3}, \sigma_{\alpha_2\alpha_3}$). This is too large a problem for even the best modern computers to handle.

C. Solutions for Special Cases

Instead of carrying out the tedious and expensive computations for the full adaptive optimization, we elected to examine two special cases that appear to be of management interest and that should reveal the general flavor of the full solution. These cases are shown in Figure 2, and reflect two extreme situations:

Case 1: The fishery is just beginning, and the stock is near natural equilibrium; β can be treated as known and the main uncertainty is about α .

Case 2: The fishery has been holding spawning stocks at low levels for many years; α is well known and the main uncertainty is about β . Many Pacific salmon fisheries seem to fit this case; environmental carrying capacities may have changed considerably in recent years.

In either of the cases the size of dynamic optimization problem is reduced considerably by treating one parameter as known. In Case 1, the stock and recruitment data can be expressed in stock units relative to the natural equilibrium: the Ricker model can be written in the form:

$$R_{t+1} = S_t e^{\alpha(1-S_t) + v_t} ,$$

where the system state vector for optimization becomes $\{R_t, \hat{\alpha}_t, \sigma_\alpha^2\}$, and the variables in the adaptive regression equations become:

$$y_t = \ln\left(\frac{R_t}{S_{t-1}}\right)$$

$$x_t = (1 - S_{t-1}) \tag{17}$$

$$\hat{p}_t = \hat{\sigma}_\alpha^2 .$$

In Case 2, the Ricker model is assumed to maintain its usual form, the optimization state vector becomes $\{R_t, \hat{\beta}_t, \sigma_\beta^2\}$, and the adaptive regression variables become:

$$y_t = \ln \left(\frac{R_t}{S_{t-1}} \right) - \tilde{\alpha}$$

$$x_t = -S_{t-1} \tag{18}$$

$$\hat{P}_t = \hat{\sigma}_\beta^2 ,$$

where $\tilde{\alpha}$ is the reasonably certain estimate of α .

Several dynamic programming solutions for the simplified cases were carried out on a PDP 11/45 computer system. Each of the solutions required about five hours of computer time, which is not excessive considering the broad range of stock-parameter-uncertainty combinations that must be evaluated. By trial and error, we discovered that it was necessary to use 10 discrete levels for each of the variables (R_t , $\hat{\alpha}_t$ or $\hat{\beta}_t$, $\hat{\sigma}_\alpha^2$ or $\hat{\sigma}_\beta^2$), and to move backwards in time around 20 steps (generations); finer state intervals and more time steps did not change the solutions.

Representative results for Case 1 (α uncertain) are shown in Figure 3. Each of the isopleth diagrams show optimal harvest rates for a cross section through the $R_t - \hat{\alpha}_t$ plane at one uncertainty (σ_α^2) level. The most striking feature of these results is that optimal harvest rates are nearly independent of $\hat{\alpha}_t$ for large σ_α^2 . What we expected to see was some indication that spawning populations should be reduced (high exploitation) when α is uncertain; by equations (8) and (17), we would expect the greatest reduction in uncertainty by conducting such an "experiment." As the cross section for high uncertainty ($\sigma_\alpha^2 = 4.0$) in Figure 3 shows, experiments involving high exploitation rates are optimal only if $\hat{\alpha}_t$ is also large; indeed it appears that the best strategy is to avoid high harvest experiments

when α is low and the stock size is large. The optimization also takes into account the possibility that low spawning stocks will reveal α to be small; thus a period of recovery without harvest will be necessary. Examining the low uncertainty ($\sigma_\alpha^2 < 0.1$) diagram in Figure 3, it is apparent that the optimal harvest rate for any stock size is insensitive to α , no matter what the environmental noise σ^2 (Walters [5] obtained a similar result). We should not expect the optimal harvest strategy to depend greatly on σ_α^2 if this strategy is nearly independent of α in the first place.

The results for Case 2 (β uncertain) indicate a similar pattern; the optimum harvest strategy is quite insensitive to $\hat{\beta}$ when σ_β^2 is high (Figure 4). Examining equations (8) and (18), we would expect high spawning stocks to produce the greatest reduction in uncertainty about β ; yet the optimization balances the value of low exploitation (high S_t) experiments against the loss in immediate yields that such experiments would entail. Low harvest experiments are called for only when there is intermediate uncertainty about β .

III. Selection Among Alternative Models.

The analysis in Section II took two sources of uncertainty into account: random environmental variation, and uncertainty about production parameters. Section III explores a third type of problem: uncertainty about the basic functional form of the stock recruitment relationship. As an example, consider the data in Figure 5 on "off-cycle" runs of sockeye salmon (Oncorhynchus nerka) in the Fraser River. Several subpopulations of sockeye in this river system exhibit cyclic dominance (Ward and Larkin, [20]), with very large "cycle" runs every four years

(1962, 1966,.....) that apparently follow a different stock-recruitment relationship from the off-cycle runs. Escapement levels in the off-cycle years have apparently been chosen under the assumption that high spawning populations may result in lowered recruitment because of overutilization (space, oxygen, etc.) of spawning areas; Figure 5 does not support this assumption, at least when the whole river system is treated as a single population unit. Also, the overspawning phenomenon should have resulted in damping or destruction of higher cycle years during the early development of the fishery. However, there is no evidence of this (Ward and Larkin [20]); the off-cycle years sustained annual catches of around 4 million fish until the destructive Hell's Gate Slides of 1911. The best fitting Ricker curve for the data (curve η_1 in Figure 5; $\alpha = 1.9$, $\beta = 0.44$) does predict that production would decline for spawning stock above 2 million, but it seems equally reasonable to assume that the correct relationship is a saturating curve of the "Beverton Holt" type (Ricker [15]). That is:

$$R_t = \frac{e^{v_t}}{\beta + \alpha/S_{t-1}}, \quad (19)$$

where

$\alpha = 1/(\text{maximum recruits per spawner});$

$\beta = 1/(\text{maximum recruits ever possible});$

$e^{v_t} = \text{random environmental survival factor as in equation (1)}.$

A visual fit to this relationship is shown in Figure 5 as curve η_2 ; the parameters ($\alpha = 0.1237$, $\beta = 0.1025$) were chosen so as to:

- (a) closely match the Ricker curve through the available data; and
- (b) predict an equilibrium stock (8.5 million) that seems reasonable considering early catch records.

Whatever the fitting procedure and even allowing for decrease in production for high spawning stocks (dotted lines off curve η_2 in

Figure 5), significant improvements in yield could be obtained if the η_2 curve is correct. The question is: should an experiment (reduced harvests for one or more years) be conducted to test this possibility?

In principle this question could be addressed with the optimization approach introduced in Section II. The stock-recruitment relationship can be written in the form:

$$R_t = \theta_1 f_1(S_{t-1}) + \theta_2 f_2(S_{t-1}) + \dots \quad (20)$$

where θ_i represents model selection parameters that take only the values 0 or 1 and are constrained as $\sum \theta_i = 1$ (so that all but one of the θ_i must equal 0), and the f_i are alternative models such as the Ricker [equation (1)] and the Beverton-Holt [equation (19)]. Wood [21], Smallwood [17] and others have shown that it is possible to calculate the probability that each $\theta_i = 1$ (model i is correct), given that the true model is among the alternatives represented. These probabilities along with parameter estimates and measures of uncertainty for each of the alternative models can be formed into an extended vector of state variables. Unfortunately, the number of variables involved makes dynamic programming optimization impractical. Thus some drastic simplifications and approximations are necessary in order to trace the most likely statistical outcomes and the most promising decision possibilities.

Since full adaptive control analysis is not feasible, the remainder of this section attempts to develop a simplified procedure for designing and evaluating experimental harvesting regimes intended to discriminate between alternative production models. The procedure is modified from a general approach suggested by Bard [2], and involves the following five basic steps.

- (a) Identification of a series of models, or of alternative

"states of nature" $\eta_1, \eta_2, \dots, \eta_m$ that are to be compared. Alternative η_1 might be the Ricker model, η_2 might be the Beverton-Holt model, η_3 might be a simple free hand curve extrapolating from existing data, and so forth.

(b) Assignment of prior or judgemental probabilities

$p^*(\eta_1), p^*(\eta_2) \dots$ to each of the alternative states of nature. These probabilities might be derived through some statistical procedure, or they may represent simple intuition. Reflecting on the Fraser River data, we might for example assign p^* (Ricker model) = 0.7, and p^* (Beverton-Holt model) = 0.3. Though both models fit the data about equally well, this probability assignment would give some weight to the common arguments about overspawning.

(c) Identification of a series of alternative harvesting

experiments $\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_n$, each of which would be reasonably certain to discriminate between the alternative models but would require different lengths of time to complete. For the Frazer River example, some reasonable alternatives are

- \mathcal{E}_1 - Continue the present escapement policy (S_t around 1.0 million/yr.) indefinitely. That is, do not experiment and hope that luck will eventually provide the necessary data.
- \mathcal{E}_2 - Allow escapements intermediate between the optima for the alternative models; considering all environmental variability, this experiment would probably not give definite results for at least 20 years.
- \mathcal{E}_3 - Allow the optimum escapement ($S_t = 2.0$ million) for the Beverton-Holt model. At this escapement level, any tendency for overspawning should be apparent within five

years while it is unlikely that environmental circumstances would combine for that long to give consistently low recruitments if the saturating model was true.

In general, the experiments should reflect tradeoff between small harvest manipulations that require a long time to yield definitive results and large harvest manipulations that yield results quickly. The length of experiment required at any escapement or stock level can be assessed by examining expected variability around the alternative stock-recruitment models at that escapement level.

(d) Calculations of expected long-term returns for each of the combinations of experiment and state of nature. The elements of Table 1 below must be evaluated:

Table 1

		EXPERIMENT			
		\mathcal{E}_1	\mathcal{E}_2	\mathcal{E}_n
TRUE STATE	η_1	V_{11}	V_{12}	V_{1n}
	η_2	V_{21}	V_{22}	
	η_m	V_{m1}		V_{mn}

Here V_{ij} represents the expected total value of all future harvests, given that experimental strategy \mathcal{E}_j is applied and the true state of nature is η_i . Let us suppose that strategy \mathcal{E}_j involves allowing a certain escapement \tilde{E}_j for τ_j years. If τ_j is chosen carefully, we should be reasonably certain of detecting that η_i is the true state of nature after the τ_j years; we should be unwilling to accept η_i as the true state until the τ_j years have elapsed. (Any experiment not

meeting these qualifications should not be included). Thus, V_{ij} can be calculated as the expected value of two discounted sums: (a) the sum of catches during the τ_j experiment years, given that escapement \tilde{E}_j is applied and model η_i is correct; and (b) the sum of catches after the τ_j 'th year, given that the optimum escapement for model η_i is followed thereafter. The first component reflects the short-term impacts of the experiment, while the second component reflects long term benefits. The simplest way to calculate V_{ij} is to do a whole series of simulation trials, each using model η_i and a different sequence of random environmental inputs. For each of the trials we calculate:

$$V'_{ij} = \sum_{t=1}^{\tau_j} (C_t | \tilde{E}_j) e^{-\delta t} + \sum_{t=\tau_j}^T (C_t | \hat{E}_i) e^{-\delta t} ,$$

where $(C_t | \tilde{E}_j)$ is defined as the catch in year t using experimental escapement \tilde{E}_j , and $(C_t | \hat{E}_i)$ as the catch in year t given the optimum escapement \hat{E}_i for model η_i (δ is the discount rate). It should not be necessary to perform more than about 20 trials of length $T \approx 50$ generations for reasonable discount rates. V_{ij} is found as the average of V'_{ij} across these trials. To develop the entire strategy-experiment table, it is necessary to do about $m \times n \times 20$ simulation trials. This is a trivial computing exercise.

(e) Selection of the experiment with maximum expected benefits. Each of the columns of the strategy-experiment Table 1 gives the returns to be expected from one experiment for each of the possible states of nature. The overall value for the experi-

ment is simply the sum of these returns weighted by the prior probabilities for the η_i . That is:

$$\text{expected value of experiment } j = \sum_{i=1}^m p^*(\eta_i) V_{ij} \quad (21)$$

The best experiment is that which has the maximum expected value. The key point about this selection procedure is that it takes into account all possible states of nature in evaluating each of the proposed experiments.

The five steps outlined above lend themselves well to a gaming situation in which the resource manager is asked to devise alternative recruitment models, to assess their probabilities, and to evaluate alternative experimental schemes. The most critical point in the analysis is the identification of appropriate durations for alternative experiments. If the manager is overly pessimistic (e.g. if he assumes that some experiment will require τ_j years to be certain which model is correct when in fact fewer years are required), perfectly good experiments may appear poor in relation to those states of nature for which the experimental escapement \tilde{E}_j is far from optimal. On the other hand, the manager may be overconfident, and may suggest a short experiment that in reality would simply result in loss of yield with no improvement in understanding about the system. In a gaming situation, the best τ_j for any proposed \tilde{E}_j can be evaluated quickly by facing the manager with several stochastic simulations from each of the possible true models while noting how long it takes to be sure which model is being used in each

of the simulations. More precisely, the analyst sets up a series of trials. For each of the trials the manager chooses an escapement \tilde{E}_j , and the analyst secretly chooses a model η_i . A stochastic simulation with η_i is then initiated and carried forward in time until the manager positively identifies the results as coming from η_i . The number of simulation steps required for the various trials can be plotted in relation to the choices of \tilde{E}_j . This plot will reflect the manager's subjective definition of "positive" results, analogous to his choice of confidence limit probabilities (e.g. 90% vs. 95%) in ordinary statistical problems. The degree of random variation introduced in each of the trials should reflect uncertainty about the model parameters as well as expected environmental variation by using the variance relationship in equation (11) or its subjective equivalent.

To test the procedure, we carried out a gaming analysis on the Fraser River problem with one of the authors acting as manager and the other as analyst. Two alternative models were considered:

η_1 = Ricker curve from Figure 5;

η_2 = Beverton-Holt curve from Figure 5.

By examining the data and following the trial procedure for τ_j outlined in the previous paragraph, we arrived at the following set of experiments:

\mathcal{E}_1 - allow an escapement of 1.0 million indefinitely. (Thus do not experiment);

\mathcal{E}_2 - allow an escapement of 1.5 million for 15 years;

\mathcal{E}_3 - allow an escapement of 2.0 million for 5 years;

\mathcal{E}_4 - allow an escapement of 3.0 million for 3 years.

Simulation trials to evaluate the V_{ij} were performed,

assuming a discount rate of 4% per generation (i.e. 1% per year for Fraser sockeye). The results are given below in Table 2.

Table 2

Expected Future Values for Alternative Harvesting Experiments on Off-Cycle Fraser River Sockeye.*

		<u>Experiment</u>			
		\mathcal{E}_1	\mathcal{E}_2	\mathcal{E}_3	\mathcal{E}_4
TRUE STATE	η_1	77.2	77.8	75.4	71.7
	η_2	92.4	108.2	110.9	110.6

*Elements of the table are discounted sums of catches, in millions of fish.

When only two or three alternative states of nature are to be compared (as in the Fraser example), the analysis can be presented in an elegant form that simplifies the problem of assigning subjective probabilities to the alternative models. Suppose we make a graph where the abscissa is V_{1j} (expected value of experiment j given that η_1 is true) and the ordinate is V_{2j} (expected value of experiment j given that η_2 is true). Each of the experiments can be plotted as a point on this graph (Figure 6). Points that are close to the ordinate represent experiments or policies that are good if η_2 is true, and poor if η_1 is true; points near the abscissa represent policies that are good if η_1 is true, and poor if η_2 is true. If we designate those experiments that can distinguish between the η_i as "effective experiments," then the graphical representation allows us to sort out a smaller subset of "efficient experiments" that are best for at least some values of $p^*(\eta_i)$. In Figure 6, experiments \mathcal{E}_2 and \mathcal{E}_3 are efficient, while experiments \mathcal{E}_1 and \mathcal{E}_4 are inefficient since \mathcal{E}_2 and \mathcal{E}_3 have better expected values for

all assignments of $p^*(\eta_i)$. We are trying to find that experiment \mathcal{E}_j which maximizes the "objective function":

$$y = p^*(\eta_i)V_{1j} + (1 - p^*(\eta_1))V_{2j} .$$

This objective function can be represented as a series of straight lines on Figure 6, with higher lines representing larger y values and greater slopes representing larger values of $p^*(\eta_1)$. To find the optimum experiment, we move the lines downward (choose lower y values) until it first touches a "feasible point" representing some \mathcal{E}_j . The key point is that we can identify ranges of $p^*(\eta_1)$ for which any efficient \mathcal{E}_j is optimal, simply by changing $p^*(\eta_1)$ so as to steepen or flatten the objective function line. The results of the process on the Fraser River example are given in Table 3 below.

		Table 3
		Range of $p^*(\text{Ricker})$ For Which
		\mathcal{E}_j is Best
Experiment	\mathcal{E}_j	_____
	\mathcal{E}_1	none
	\mathcal{E}_2	0.4 - 1.0
	\mathcal{E}_3	0.0 - 0.4
	\mathcal{E}_4	none

Thus, the manager does not need to precisely specify his judgement about $p^*(\eta_1)$ as a single number.

The Fraser River test results suggest the following conclusions:

- (a) Present management policy for off-cycle year is not optimal for either of the models shown in Figure 5; some alternative (and perhaps unclearly specified)

model is in use, or the real management objectives are unrelated to maximization of discounted long-term catches.

- (b) A modest experiment involving increased escapements (1.5 million) should cause no serious problems if the Ricker model is correct; it may, in fact, result in considerably higher yields over the long run.
- (c) A more drastic experimental policy involving escapements of 2.0 million for several years would be more advisable if there is considerable confidence that the true stock-recruit relationship is similar to the curve η_2 in Figure 5.

To test the effect of discounting rate on these conclusions, we reevaluated Table 2 for $\delta = 1\%$, $\delta = 10\%$, $\delta = 20\%$, and $\delta = 30\%$ per generation. For δ greater than 20%, the tests suggested that the modest experiment is the best unless $p^*(\text{Ricker})$ is less than 0.2. For $\delta = 1\%$, the drastic experiment becomes the best alternative unless $p^*(\text{Ricker})$ is greater than 0.6.

IV. Extensions and Generalizations.

Previous sections have dealt mostly with uncertainty about stock recruitment relationships. Two major assumptions have been the discussion: a) stock size is directly measurable without error; and b) fishing effort is fully controllable to conform with biological recommendations. Since these assumptions are often not justified, Section IV attempts to show how the concepts and methods introduced previously could be extended to include these additional sources of uncertainty.

A. Schaeffer Production Model

The idea of using logistic population growth assumptions as a basis for production modelling was first made popular by Schaeffer [16] in his studies of Pacific tunas. In its simplest form, the "Schaeffer Model" can be written as:

$$N_{t+1} = N_t + \alpha N_t (1 - \beta N_t) - C_t \quad , \quad (22)$$

where

N_t = stock size, usually in biomass units;

α, β = production parameters with similar definitions as
in the Ricker Model;

C_t = total catch.

Noting that equation (22) is not directly usable (since N_t is not observable for most populations), Schaeffer and others have assumed a simple "observation model" to accompany the dynamic model:

$$N_t \doteq \frac{q_t}{c} = \frac{C_t/E_t}{c} \quad , \quad (23)$$

where

q_t = catch per unit effort;

E_t = some effort measure having units (boats) x (time
fishing per
boat);

c = catchability coefficient.

Substituting the observation model (23) for (22), we obtain an expression containing only observable quantities and parameters:

$$\frac{q_{t+1}}{c} = \frac{q_t}{c} + \alpha \frac{q_t}{c} \left(1 - \beta \frac{q_t}{c} \right) - c_t \quad ;$$

this can be simplified to give:

$$q_{t+1} - q_t = \alpha q_t - \frac{\alpha\beta}{c} q_t^2 - c c_t \quad . \quad (24)$$

For parameter estimation and adaptive control analysis, this version of the Schaeffer Model can be cast into the recursive regression format with:

$$y_t = q_t - q_{t-1} \quad (\text{or } y_t = q_t)$$

$$\underline{x}_t = \begin{bmatrix} q_{t-1} \\ q_{t-1}^2 \\ c_{t-1} \end{bmatrix} \quad (25)$$

$$\underline{\hat{a}}_t = \begin{bmatrix} \hat{\alpha}_t \\ \widehat{\alpha\beta/c}_t \\ \hat{c}_t \end{bmatrix} \quad \left(\text{or } \underline{\hat{a}}_t = \begin{bmatrix} \widehat{(1+\alpha)}_t \\ \widehat{\alpha\beta/c}_t \\ \hat{c}_t \end{bmatrix} \right) .$$

Though there is no reason here to expect such statistical properties as normally distributed errors, the regression format at least provides a unified framework for evaluating parameter uncertainties. Also, the format automatically provides a

simple state estimator, $\hat{N}_t = q_t/\hat{c}_t$. More complicated, statistically nonlinear versions of equation (24) can be devised using more realistic observation models than equation (23). Recursive nonlinear estimation techniques are beginning to appear in the literature under the general heading "extended Kalman filters" (Young, [22]).

The Schaeffer Model gives a remarkably good fit to historical data for many large fisheries, as shown in Figure 7. In fitting these data, we used the linear regression scheme in equations (24) to obtain estimates of α, β , and c . In all cases the data had already been corrected for changing vessel efficiency (changing c); thus it was not necessary to introduce discounting of old data or a parameter variation matrix Q (see Section II) into the regression equations. It is clear that further adaptive control work for the Schaeffer Model is justified, and we intend to develop more complete analyses in a future paper.

B. Incomplete Control of Fishing Effort

The fisheries literature abounds with biological models and equilibrium yield analyses; almost no attention has been paid to the dynamics of the predator-prey system that results from incomplete control of economic investment. Fishery fleets have basic "reproductive" (investment) and "mortality" (disinvestment) relationships that in principle make them similar to any predator population (Smith [18]; Gatto, et al. [5]). In the absence of investment control, many fishing fleets have developed to the point where pressure for short-term economic and social welfare benefits has made it virtually im-

possible to implement biologically sound long-term policies; the current state of the International Whaling Commission is a good example.

To pursue the predator-prey analogy, we may find it useful to think of management controls directed at fishing effort as generating a "reachable region" of stock size and investment combinations (Figure 8) around the "nominal trajectory" of development that would occur without management. Investment control may occur as subsidies to increase the rate of investment, or as taxes and direct regulations to reduce it. If the fishery fleet can operate economically at very low stock sizes, and if only small control decreases in effort are possible each year, it may be impossible to move the fishery to a state where maximum sustained yield is possible. This problem may become serious if the first incremental controls are not applied until the fishery is well developed. One is reminded of the adage about ounces of protection and pounds of cure.

Limitations on control changes from one time step to the next may be represented in dynamic optimization by including the control level (effort, harvest rate, etc.) as an additional state variable. For example, if the system state without effort limitation was represented as $(R_t, \hat{\alpha}_t, \sigma_\alpha^2)$, the state with limitation would be $(U_{t-1}, R_t, \hat{\alpha}_t, \sigma_\alpha^2)$ where U_{t-1} is the exploitation rate from the previous time step. Instead of looking at all possible harvest rates at each of the time steps for each of the $(R_t, \hat{\alpha}_t, \sigma_\alpha^2)$ combinations, the optimization would only

examine harvest rates U_t over an interval:

$$U_{t-1} - k_1 \leq U_t \leq U_{t-1} + k_2$$

where

k_1 = maximum permissible annual decrease in exploitation rate;

k_2 = maximum permissible annual increase in exploitation rate.

Similar constraints can be applied in generating experimental harvest regimes for the analyses mentioned in Section III.

A series of interesting issues arise concerning the selection of appropriate values for the control limits k_1 and k_2 . The maximum rate of increase in harvest, k_2 , will depend on private and public willingness to invest in the fishery and on the availability of fishing gear to be transferred from alternative fisheries. The maximum rate of decrease, k_1 , will depend on the regulatory power vested in the management agency, on the willingness of the agency to accept responsibility for immediate economic and social hardships, and on the expected profitability of the fishery. Fisheries agencies are beginning to face these political and economic issues, and optimization formulations may provide a useful focus for formal debate even if no quantitative solutions are attempted.

V. Summary.

This paper moved from formal optimization analysis of trivial models to broader approaches for experimental management. The formal analysis was conducted in order

to discover simple principles that might be applied in more complex and realistic fisheries situations where uncertainty is a key factor. The principles that emerged are intuitively reasonable: (a) when production parameters are uncertain, the harvest rates used should be lower than would be supposed if only the available parameter estimates were considered; and (b) when the equilibrium stock size is uncertain but production rates are well understood, the harvest rates should be lowered. When the general form of the production relationship is uncertain and when management control is limited, formal optimization becomes impossible: however, a gaming procedure may help to devise and to evaluate alternative management strategies. The gaming procedure involves defining a series of possible models, selecting a series of effective experiments, and calculating the optimum experiment under subjective prior probabilities for all models. This technique is an immediately useable solution for complex fisheries problems. Future work is continuing to overcome the computational obstacles encountered in formal optimization of complex models.





Figure 1. Decision possibilities and random outcomes for each of the time steps in adaptive control analysis for Ricker model.

CASE 1: DEVELOPING FISHERY, β KNOWN AND α UNCERTAIN.



CASE 2: OLDER FISHERY WITH UNRELIABLE DATA ON NATURAL STOCKS, α KNOWN AND β UNKNOWN.



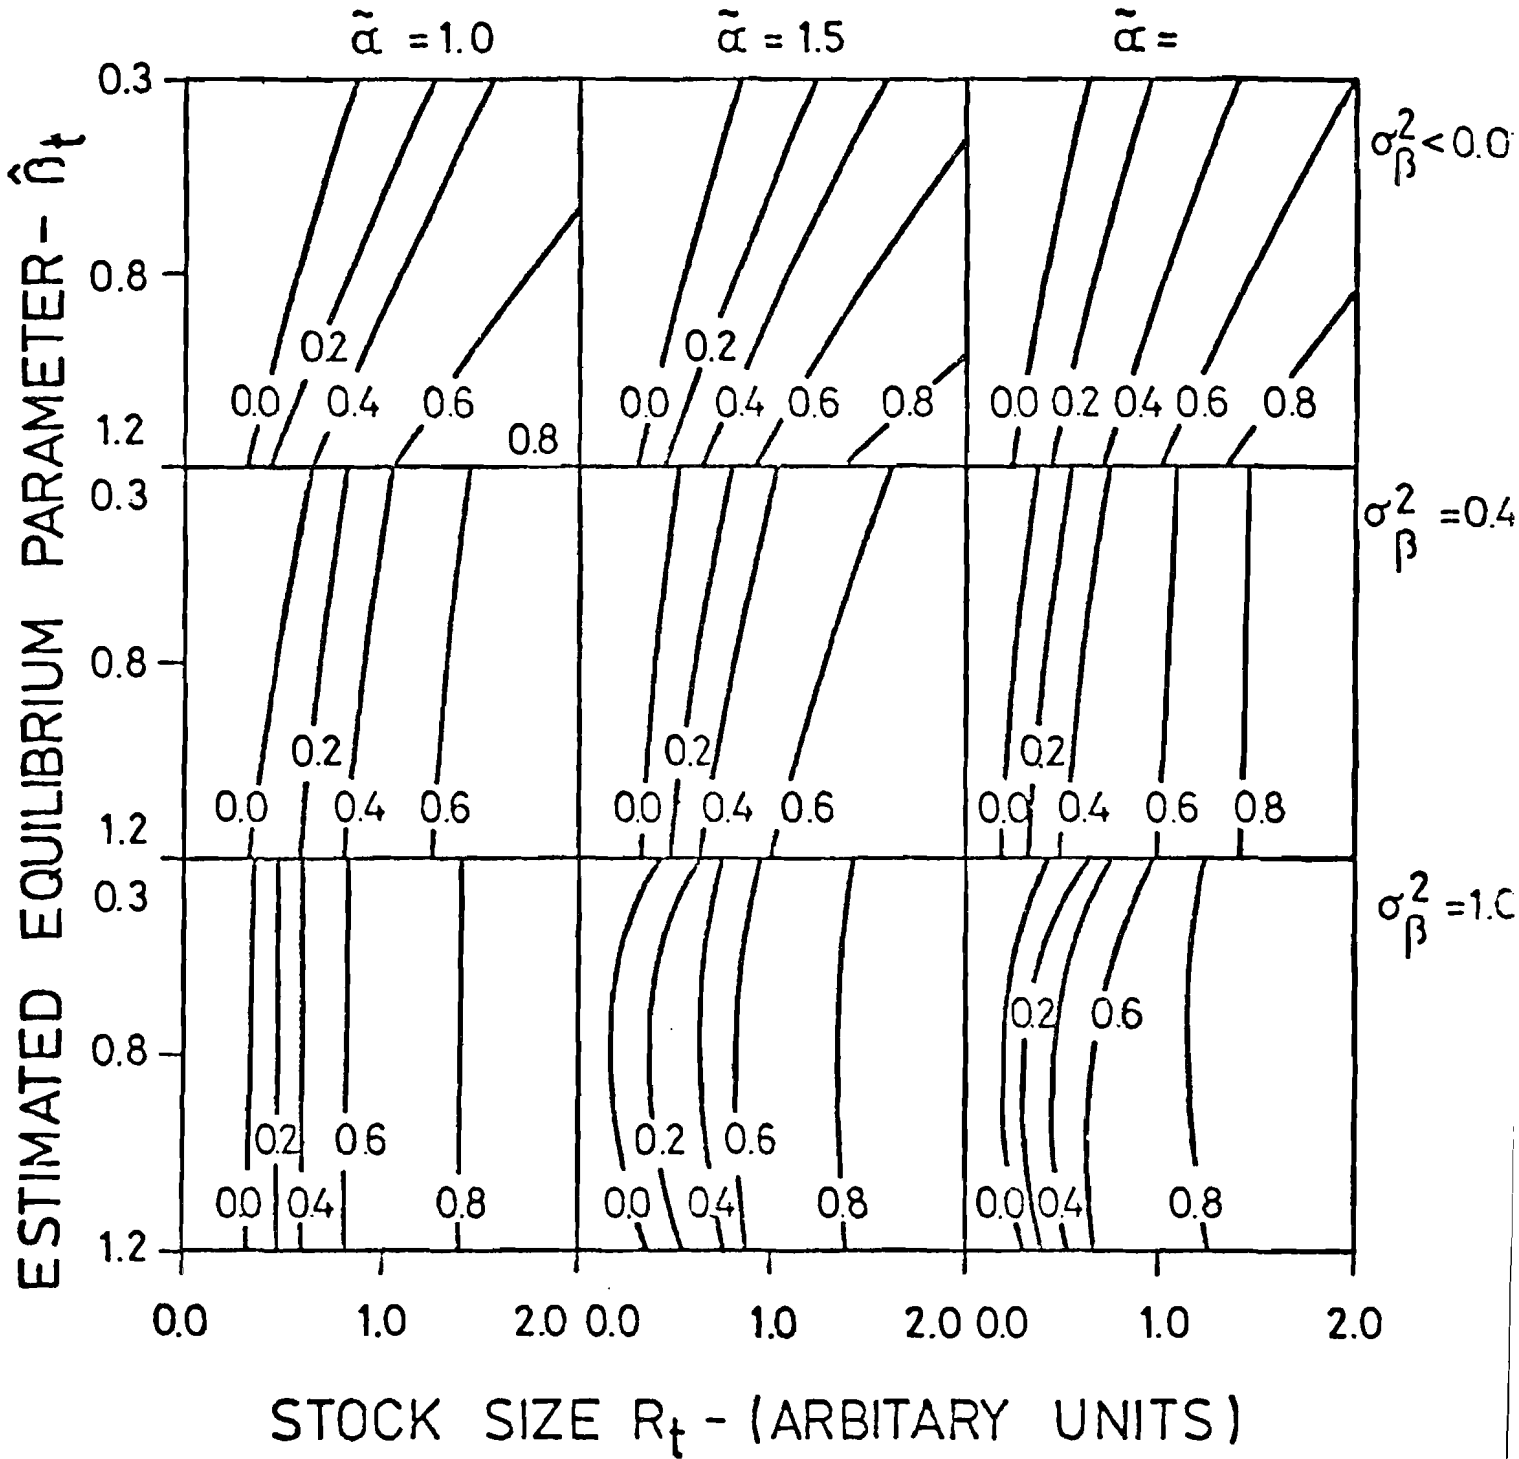


Figure 2. Management situations that permit simplified adaptive optimization.




Note: Results obtained with environmental variance $\sigma^2 = 0.5$ and discount rate $\delta = 4\%$ per generation.

Figure 3. Optimal exploitation rates for various stock sizes, production rate estimates ($\hat{\alpha}_t$), and uncertainties about α (σ_α^2) assuming the Ricker model form is correct and equilibrium stock is known (see case 1, Figure 2).




Note: Results obtained with environmental variance $\sigma^2 = 0.5$ and discount rate $\delta = 4\%$ per generation.

Figure 4. Optimal exploitation rates for various stock sizes, equilibrium population sizes ($1/\hat{\beta}_t$), and uncertainties about $\hat{\beta}$ ($\sigma_{\hat{\beta}}^2$), assuming Ricker form is correct and $\hat{\alpha}$ is known (see case 2, Figure 2).




Note: Data shown are for 1939-1973, omitting every fourth or cycle year beginning in 1942. η_1 --least squares fit to Ricker model; η_2 --visual fit to Beverton-Holt model. Graph axis in millions of fish.

Figure 5. Alternative stock-recruitment models for Fraser River Sockeye Salmon, off-cycle years.



Note: Graph units are discounted (4%/generation) sums of catches, in millions of fish.
 $S_1 - S_4$ denote escapement experiments defined in text. Total value line is given by value = $p^*(\eta_1)$ [value given η_1] + $p^*(\eta_2)$ [value given η_2].


Figure 6. Values of alternative Fraser River harvesting experiments if the alternative models η_1 and η_2 in Figure 5 are correct.



Note: The entire data set for each case was used to obtain parameter estimates; simulated value for the previous year (rather than from previous year's data). Data were obtained from following sources: fin whales - International Whaling Commission [7]; halibut - Southward [19]; cod - Garrod [4]; haddock - Grosslein and Hennemuth [6]; yellowtail - Lux [9]; sardine - Marr [10];

Figure 7. Observed changes in some fisheries compared to simulated trends predicted by the Schaeffer model (equations 13, 23) using parameter estimator from regression analysis.

GOOD SITUATION



PATHOLOGICAL SITUATION




Figure 8. Imperfect control of fishing effort may result in a restricted "reachable region" of stock size and effort combinations as a fishery develops. The point "0" denotes optimum equilibrium state.

References

- [1] Allen, K.R. The Influence of Random Fluctuations in the Stock-recruitment Relationship on the Economic Return from Salmon Fisheries. Rapp. No. 164. Conseil International Pour L'Exploration De La Mer, 1973, pp. 351-359.
- [2] Bard, Yonathan. Nonlinear Parameter Estimation. New York, Academic Press, p. 341.
- [3] Dahlberg, M.L. Stock-and-Recruitment Relationships and Optimum Escapements of Sockeye Salmon Stocks of the Chignik Lakes, Alaska. Rapp. No. 164. Conseil International Pour L'Exploration De La Mer, 1973, pp. 98-105.
- [4] Garrod, D.J. "Empirical Assessments of Catch/Effort Relationships in North Atlantic Cod Stocks." ICNAF Res. Bull. No. 6 (1969), pp. 26-34.
- [5] Gatto, M., Rinaldi, S., and Walters, C. A Predator-Prey Model for Discrete-Time Commercial Fisheries. IIASA-RR-75-5. Laxenburg, Austria, International Institute for Applied Systems Analysis, 1975, p. 38.
- [6] Grosslein, M.D., and Hennemuth, R.C. Spawning Stock and Other Factors Related to Recruitment of Haddock on Georges Bank. Rapp. No. 164, Conseil International Pour L'Exploration De La Mer, 1973, pp. 77-88.
- [7] International Whaling Commission. Annual Report for 1967-68, 1969.
- [8] Larson, R.E. State Increment Dynamic Programming. New York, American Elsevier Publishing Co., 1968, p. 300.
- [9] Lux, F.E. "Landings Per Unit of Effort, Age Composition, and Total Mortality of Yellowtail Flounder, Limanda ferruginea (Storer), off New England." ICNAF Res. Bull. No. 6 (1969) pp. 35-43.
- [10] Marr, J.C. "The Causes of Major Variations in the Catch of the Pacific Sardine, Sardinops caerulea (Girard)," Proceedings of the World Scientific Meeting on the Biology of Sardines and Related Species, Rome, Food and Agricultural Organization of the United Nations (1969), pp. 667-791.

- [11] Parrish, B.B. (Ed). Fish Stocks and Recruitment. Rapp. No. 164, Conseil International Pour L'Exploration De La Mer, 1973.
- [12] Raiffa, H. "Decision Analysis: Introductory Lectures on Choices Under Uncertainty, Reading (Massachusetts), Addison-Wesley, 1968.
- [13] Raiffa, H., and Schlaifer, R. Applied Statistical Decision Theory, Cambridge, Massachusetts, Massachusetts Institute of Technology Press, 1961, p. 356.
- [14] Ricker, W.E. "Stock and Recruitment," Journal Fish. Res. Bd., Canada 11: (1954), pp. 559-623.
- [15] Ricker, W.E. Critical Statistics from Two Reproduction Curves, Rapp. No. 164, Conseil International Pour L'Exploration De La Mer, 1973, pp. 333-340.
- [16] Schaeffer, M.B. "Some Aspects of the Dynamics of Populations Important to the Management of the Commercial Marine Fisheries," Bulletin of the International Trop. Tuna Commission, No. 1 (1954), pp. 27-56.
- [17] Smallwood, R. "A Decision Analysis of Model Selection," IEEE Trans. on Systems Science and Cybernetics, Vol. SSC-4, No. 4, 1968.
- [18] Smith, V.L. "Economics of Production from Natural Resources," American Economic Review, No. 58 (1968), pp. 409-431.
- [19] Southward, G.M. A Simulation of Management Strategies in the Pacific Halibut Fishery, Report No. 47, International Pacific Halibut Commission, 1968.
- [20] Ward, F.J., and Larkin, P.A. Cyclic Dominance in Adams River Sockeye Salmon, Progress Report 11, Pacific Salmon Fisheries Commission, 1964, p. 116.
- [21] Wood, E.F. A Bayesian Approach to Analyzing Uncertainty Among Stochastic Models, IIASA-RR-74-16. Laxenburg, Austria, International Institute for Applied Systems Analysis, 1974, p. 19.
- [22] Young, Peter. "Recursive Approaches to Time Series Analysis." Institute of Mathematics and its Applications, 10 (1974), pp. 209-224.