ADAPTIVE CONTROL PROBLEMS IN RENEWABLE RESOURCE DEVELOPMENT

Carl J. Walters

January 1975 WP-75-7

September 1974

ADAPTIVE CONTROL PROBLEMS IN RENEWABLE RESOURCE DEVELOPMENT

Carl J. Walters

A major problem in ecological management is to devise strategies for dealing with the development of resources whose response characteristics are completely unknown when the development begins. To visualize the problem, imagine a large and featureless vista of open ocean with a fishing boat in the middle distance. Imagine also that you, as the fishery manager, have only a small row boat without even a periscope to look beneath the waves. The fishing boat you see is the first of a potential fleet, and your task is to decide how large that fleet should eventually be. Essentially your only source of information is the fleet itself: you can examine the catch, but your facilities will permit only a glimpse of the biological system from which it comes. Other fishery managers have dealt with the same problem on other fishes, but their experiences can give you only a qualitative idea of what to expect.

My objective in this paper is to state the problem as a question of adaptive control, in hopes that techniques and concepts of control systems theory may at least give us a better idea of how to go about the trials and errors of learning how resource systems work. I emphasize that my concern is with the development phase of resource management, not with the identification of optimum long-range equilibria. There is an extensive literature on the latter subject ("theory of fishing", etc.), but this literature largely presupposes a substantial data base acquired through some unspecified development period.

Historically the development process has been haphazard at best, as I will show with examples below. Ecological managers have usually done little more than demand caution (low development rates) until sufficient data has accumulated for the equilibrium models to be applied. Luckily, this attitude has usually been ignored in favor of economic interests, and development has proceeded far enough for many exploited systems to show their full range of responses. If heeded, the cautious viewpoint of traditional ecologists might never lead to effective management for most ecological systems.

THE ADAPTIVE CONTROL PROBLEM

a. Dynamical form

The development control problem may be stated as follows. Supose a resource system whose state can be described by a vector x_t, and suppose that dynamic change in this system is given by

$$x_{t+1} = S x_t + g(x_{t-\tau}) - cE_t x_t$$

where

- $g(u)$ is a stochastic growth function (recruitment function) depending on state of system τ years in the past
- S is a matrix of survival rates
- c is a "catchability coefficient" that would be constant over time in the absence of technological innovation
- E_t is the total harvesting effort or some other measure of exploitation intensity. Usually this measure has units: (number of harvesters) X (harvesting time per year).

When the development begins, g will be o, or fluctuating aroung an average value of o, and E_t will be small. For most one dimensional cases, the mean value of g will have the general shape shown in figure 1. Also for many population management problems, X need take vector form only to describe the age distribution, and the first age class (x_1 say) has the g form shown in figure 1 when related to a linear combination a'X of the other age classes.

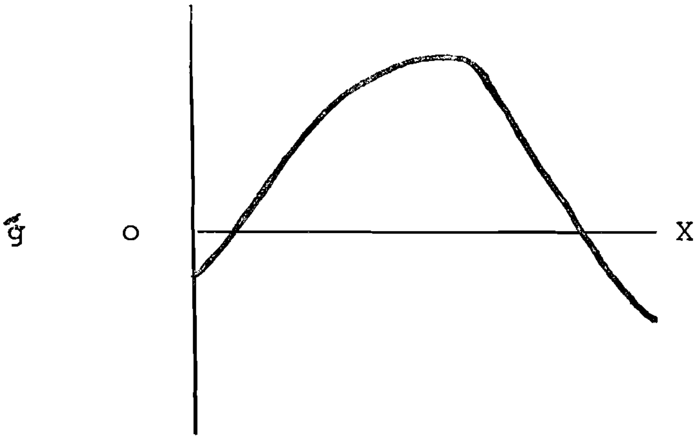

Figure 1. General form of the growth function

A variety of functional forms for g have been used with some empirical success in fisheries. For example, in salmon management where S=o (animals die after reproducing), the "Ricker-curve" is usually used:

$$g(x) = xe^{a(1-x/b)}$$

where a and b are empirical constants relating to maximum population growth rate and equilibrium population in the absence of harvest. For many high seas fisheries and land animal populations, the even simpler "Schaeffer model" has appeared satisfactory:

$$g(x) = ax(1-x/b), \quad S = 1$$

This model has led to a number of simplified guidelines for fisheries management in developing countries; it predicts for example that maximum sustained biological yields should be obtained by holding populations at 1/2 b; where b is the unexploited or natural density.

There is a key point about most observations and models for g. This point is shown in figure 2: harvest rates or effort levels that maximize g tend to be very close to levels that will cause sudden and dramatic ecological collapse.

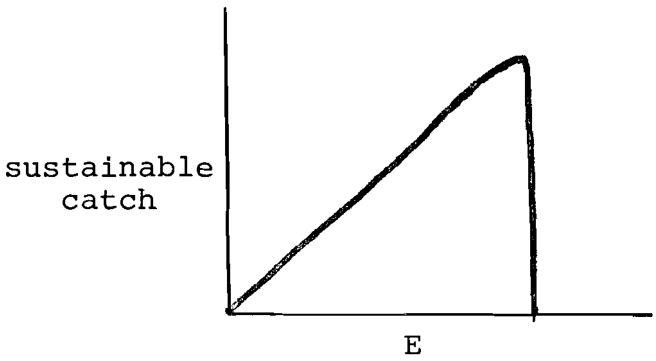

Figure 2. Equilibrium catch levels as a function of effort.

The dynamic term $cE_t x_t$ represents annual catch or harvest when x is one dimensional. For multiplè or age distributed populations, c becomes a diagonal matrix, and the total annual catch is given by $E_t \sum_i c_{ii} x_i$. Additional terms can be included in this sum to represent differential sizes or values of the various catch components. It should be mentioned that the product form $cE_t x_t$ is a poor approximation to catch-effort relationships when cE_t is large. A better form in such cases is (one dimensional case):

$$C_t = \frac{cE_t x_t}{cE_t - \ln S}\left(1 - e^{-(cE_t - \ln S)}\right)$$

Hopefully this is a bothersome and trivial complication that can be ignored in developing most arguments about control strategies.

b. Observability

In most cases x_t is not directly observable; salmons are a notable exception since they can be counted in their spawning rivers. It is usually possible to measure only the total catch ($cE_t x_t$) and the effort, E_t. Estimates of S and c can be gradually developed over time (at great expense) by tagging or marking animals to form groups of known size whose fate can be followed in the catches; such methods are not reliable for obvious statistical reasons unless catches are large relative to x. The "catch per effort" (cx_t) is often taken as a surrogate measure for population size when c is unknown; this is usually considered a dangerous policy since c is subject to technological change (often hard to observe), stochastic variation due to effects like weather on the harvesters, and nonlinearity in the catch-effort relationship(see above).

c. Control Variables

The most obvious control variable for exploitation systems is E_t, which can be manipulated either through the number of harvesters or the length of the harvesting season. In the absence of any controls, E_t can be expected to have some natural dynamics of economic development:

$$E_{t+1} = i\left(\frac{C_t}{E_t}, \frac{C_{t-1}}{E_{t-1}}, \ldots\right) - d\left(\frac{C_t}{E_t}, \frac{C_{t-1}}{E_{t-1}}, \ldots\right)$$

where i and d represent investment and disinvestment rates as a function of past success rates. The $\{E, x\}$ system can be viewed as a predator-prey interaction, and historical experience ("tragedy of the commons") suggests that this interaction is likely to either be dynamically unstable or have an undesirable stable equilibrium involving low biological productivity and economic overinvestment.

E_t is likely to be most easily controlled during the initial development period for a new resource. Later on, social and economic infrastructure induced by the development is likely to force management decisions more and more into the political arena. In any case there are limits on the control action that can be taken in any year:

$$E_{t-1} - \alpha_1 \leq E_t \quad E_{t-1} + \alpha_2$$

where

- α_1 is the maximum economically or socially tolerable decrease that could be made in any year
- α_2 is the maximum effort increase that could be achieved based on economic investment or construction limitations. (α_2 is in some way related to the i function mentioned above, or at least to limits on governmental willingness to subsidize development.)

It should be clear that there exists a hierarchy of control tactics for achieving desired changes in E_t, but the first step should be to make clear the desirable outcomes of such tactics in terms of overall changes in E.

Rather than direct regulation of E_t, it is tempting to consider more subtle control measures that operate on the i and d functions (for example, taxation and subsidies). It should be possible to design such measures so as to change the basic character of the predator-prey interaction between E and x, so as to produce desirable stability properties even in the absence of explicit control actions.

d. Optimization Criteria

The most obvious objective function to use for control would be

$$\max \sum_{t=0}^{\infty} cE_t x_t (1-d)^t$$

where d = discount rate.
This is simply the long run biological catch from the system. A trivial extension would be to use measures of economic

return, eg $cp'E_t x_t$ where p is a price vector.

However, most management agencies are charged with maintenance of yields in perpetuity. Most management now proceeds by trying to identify an optimum equilibrium (or limit cycle) for X_t and E, (say x_e, E_e), based on information gained during the development phase. x_e, E_e may be chosen based on a variety of objectives, as shown for one dimension in figure 3. Possible situations and

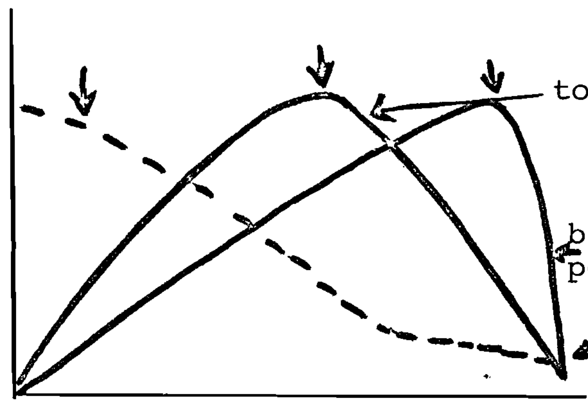

Figure 3. Alternative equilibrium objectives

tradeoffs have been thoroughly analysed from economic and ecological viewpoints, in terms of optimization, satisficing, risk, and so on. The key point is that decision procedures already exist for choosing x_e, E_e given that g and c are known; the key problem is to control development so as to most rapidly and (or) safely build up the information base necessary for these procedures to work.

Let me try to state this alternative formulation more precisely (the following is due largely to Sergio Rinaldi). Suppose that at any time t during the development it is possible to estimate the parameter set

$$\hat{\theta}_t = \left\{ \hat{c},\ \hat{S},\ \hat{a},\ \hat{b},\ \text{etc.} \right\}_t$$

using all information collected up to time t, and past experience with similar systems. Presumably some variance estimates or confidence limits can be placed on the set at any time:

$$\sigma^2_{\hat{\Theta}} = \left\{ \sigma^2_{\hat{c}_t} , \sigma^2_{\hat{s}_t} , \text{etc.} \right\}$$

Given the set $\hat{\Theta}_t$ and a set of criteria for defining the optimum equilibrium situation, it should be possible to calculate

$$\hat{x}_e = f_1 \ (\hat{\theta}_t)$$

$$\hat{E}_e = f_2 \ (\hat{\theta}_t)$$

as conditional targets for management action during time step t (and perhaps also t=1, t=2, etc.). Presumably E_t could then be chosen so as to move the system most rapidly (subject to the constraints mentioned above) towards this conditional target. For example, some $\hat{E}_t$ could be computed by dynamic programming under the assumption that $\hat{\theta}_t$ are the true values. Essentially this approach has been taken with the spruce budworm case study at IIASA. The question, however, is: <u>should E_t be deviated from $\hat{E}_t$ in order to obtain better estimates of θ?</u> In other words, to what extent should management actions be explicitly directed toward the acquisition of better information?

It might be fruitful to look at this question in terms of the phase space dynamics of E and x. As shown in Figure 4, suppose there is a natural E,x trajectory that would be followed in the absence of any management actions (curve A). Are there feasible alternative trajectories (eg,B) that are in some sense better than the unmanaged trajectory? From a strategic viewpoint, can such alternatives be classified for different dynamical situations (eg, degrees of stochastic variation) so as to provide guidelines for managers who have at least some prior knowledge about the structure of the phase space?

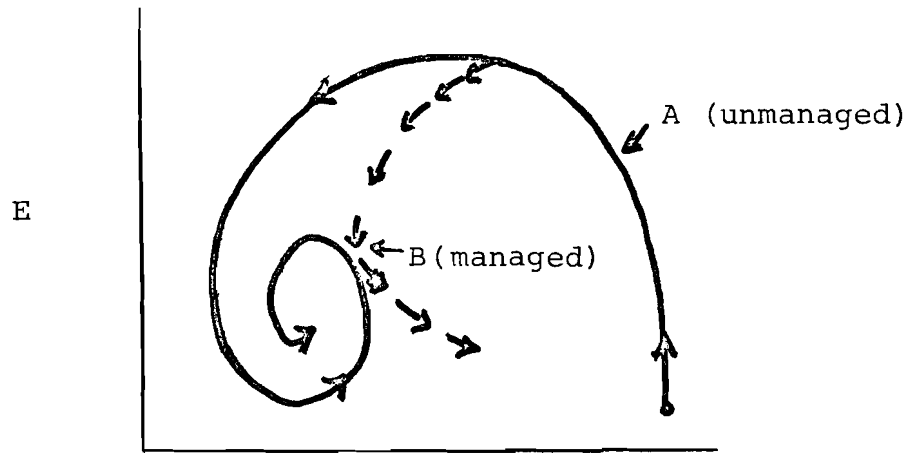

Figure 4. Managed versus natural development trajectories

A NAIVE ANALYSIS

Obviously the problem described above is not a simple one; hidden in almost every assertion that I have made are a variety of cases, and I have certainly not defined all possible options for measurement and control. No mention has been made of the management costs for various control actions, nor have I touched on problems of risk evaluation in relation to alternative development schemes. Before discussing these problems I would like to naively suggest a few directions that analysis of the problem might take. In the following discussion, I will assume that x is one dimensional.

Consider the E, x phase plane, and note that I have implicitly been thinking about controls that operate on E such that

$$E_t = f(E_{t-1}, \text{C/E in past years}) + U_t$$

where U_t is a control variable such that

$$-\alpha_1 \leq U_t \leq \alpha_2$$

(see "constraints" section above). The f function is investment rate minus disinvestment rate. Constraints on changes in E imply that only a limited variety of development patterns can occur, relative to the "nominal" development path that would occur with no management (figure 5). The band of possible states, or sampling combinations, increases with stochastic variation in g. Holling has argued that systems with more natural stochastic variation are more "resilient" to management actions; this may be true, but in such systems we have also had a better chance to sample the E,x phase space and thus to make more

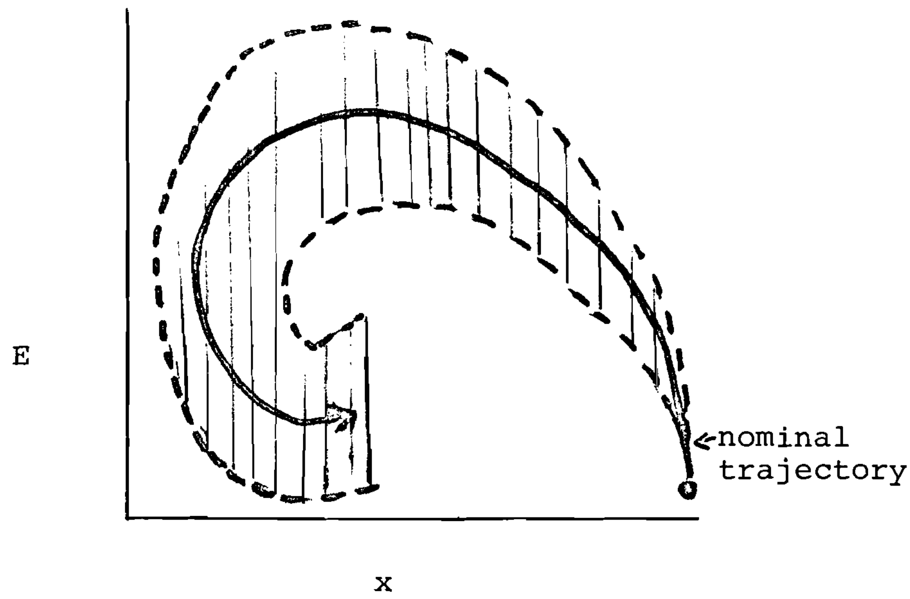

Figure 5. "Reachability" region for resource development

perceptive decisions about the consequences of control actions. In any case, recognition of the "reachable" region in figure 5 helps to narrow the problem of finding a best path, and points out that some sampling of the system will occur just due to normal economic development. If there is some a priori reason to expect boundaries in the phase space (such that E,x combinations outside these boundaries might lead to extinction of x), the reachable region might be narrowed to reflect such beliefs.

Also, we can define some measures of how valuable it would be to reach any point in the E,x plane for at least one time step. For example, we might define an information gain measure for each point in the space:

$$I\ (E_t,\ x_t) = \sigma^2_{\hat{\Theta}_t} - \sigma^2_{\hat{\Theta}_{t-1}}$$

From experience, we would expect isopleths of I to look as shown in figure 6. This figure simply asserts that high E levels are more informative than low ones, and changes in

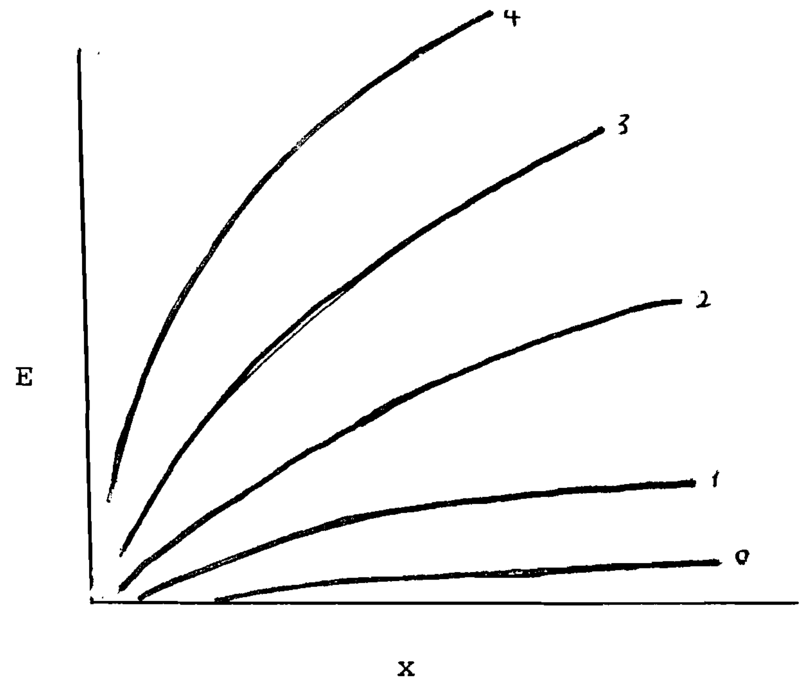

Figure 6. Information gain as a function of system state

low populations tell us more than changes in high ones (for example, about resilience properties). I (E,x) cannot be defined for any real problem without detailed analysis of the particular sampling systems and estimation procedures involved; however, it is fair to ask what can be said about optimal growth trajectories given only the qualitative form shown in figure 6.

Analogous to the information measure, we can define a one-step benefit function R over E,x, assuming no management control. If primary concern in the development is with biological yield,

$$R\ (E_t,\ X_t \mid U_t = 0) \approx c\ E_t x_t .$$

For most resource systems, we can reasonably expect R to have the qualitative form shown in figure 7. Again it is fair to ask whether knowledge only of this qualitative pattern gives any insight about optimal choice of a trajectory in the feasible band defined in figure 5.

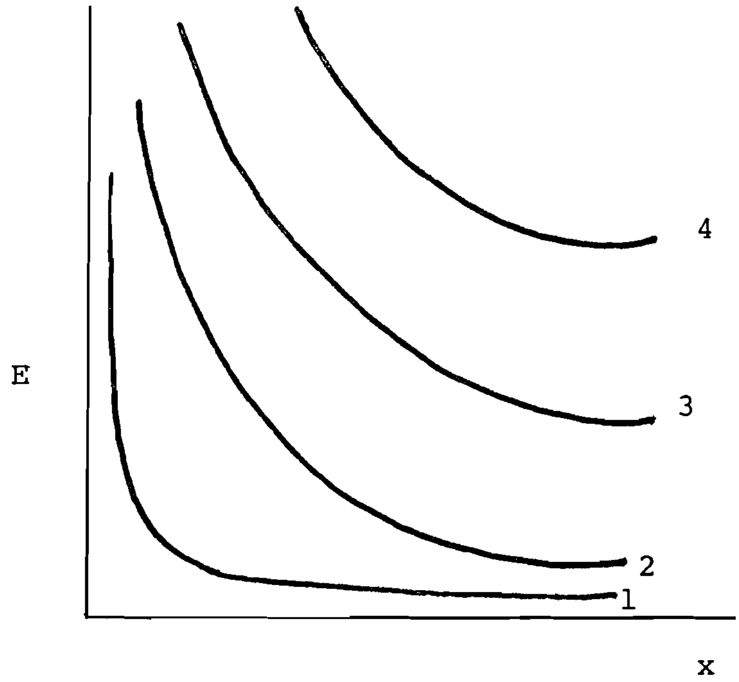

Figure 7. One time step return function

It should be possible to find a feasible trajectory that maximizes

$$I_T = \sum_{t=0}^{\infty} I\ (E_t, x_t)$$

given only the qualitative form for I in figure 6. Likewise it should be possible to find a feasible trajectory that maximizes

$$R_T = \sum_{t=0}^{\infty} R\ (E_t, X_t)$$

given only the qualitative form for R in figure 7. It may well turn out that these alternative trajectories do not depend on particular assumptions about g. It might then turn out that comparison of the I_T and R_T trajectories gives some insight about how to trade off between the need for information and the desire to maintain short-run returns.

Some useful insights might also be obtained by looking at a few extreme cases. For example, it is easy (I think) to define the optimal development strategy for g completely deterministic asymptotically stable but of unknown form (figure 8). In this case, the time lag τ would appear to be a critical management variable, since it affects the overshoot before E_e is detected. Moderate stochastic

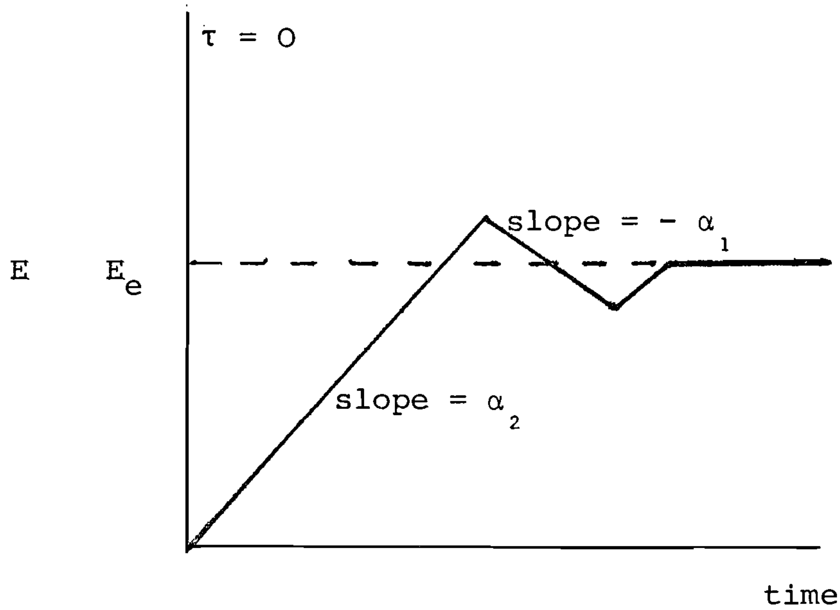

Figure 8. Optimal E trajectory in a trivial determinative case

variation in this case might even lead to more rapid convergence to E_e, since a wider range of points along the g curve are likely to be sampled simply by accident.

EXAMPLES OF THE DEVELOPMENT PROCESS

Figures 9 and 10 show two extreme examples of fisheries development. The depletion of the Antarctic Fin Whale stocks shown in Figure 9 is considered to be one of the saddest examples of resource mismanagement in the 20th century. Biologists had a good idea of the fin whale g function and recognized the likelihood of stock collapse as early as 1955, when the period of peak catches began. However, the "α_1" for effort control by the International Whaling Commission was relatively small; the final decline in catches was due mostly to economic collapse of a large proportion of the whaling industry. A complicating factor with the whale example is that the industry was able (in part) to maintain itself far into the stock collapse by fishing another species (sei whale); also the industrial development process was largely stimulated by another species (blue whale). In hindsight, we can now estimate the maximum sustainable yield for the fin whales; a key point about figure 9 is that the dramatic decline was produced by catches only 20% above this maximum sustainable level.

The history of Skeena River (B.C.) salmon management shown in figure 10 illustrates extreme stochastic variation in g. This variation is apparently due at least in part to variation in water flows in the rivers where the salmon are born. Also there are persistent effects of the time lag between birth and maturity; large adult populations ("runs") in one year tend to produce large populations 4-5 years later (Sockeye) or 2 years later (pink). Salmon are unique fisheries in that the total population (x) can be observed in any year (after harvest) as the catch plus the number of adults escaping catch and reaching the spawning grounds. It is not clear from figure 10 that overexploitation of the Skeena populations ever occurred; the reductions in exploitation rate after 1930 were explicitly due to management control rather than economic factors. The control was initially exerted by limitation of fishing time rather than the number of boats, with sad economic consequences: the 1950's and 1960's saw a very large fishing fleet with each boat barely covering its costs over the short fishing season. In the last few years, license limitation has severely reduced the fleet size so as to roughly maximize profit per boat. Thus in contrast with the fin whales, where economic interests predominated, it appears that early salmon management placed too much weight on biological considerations.

Figure 9. Historical drops in the Antarctic Fin Whales

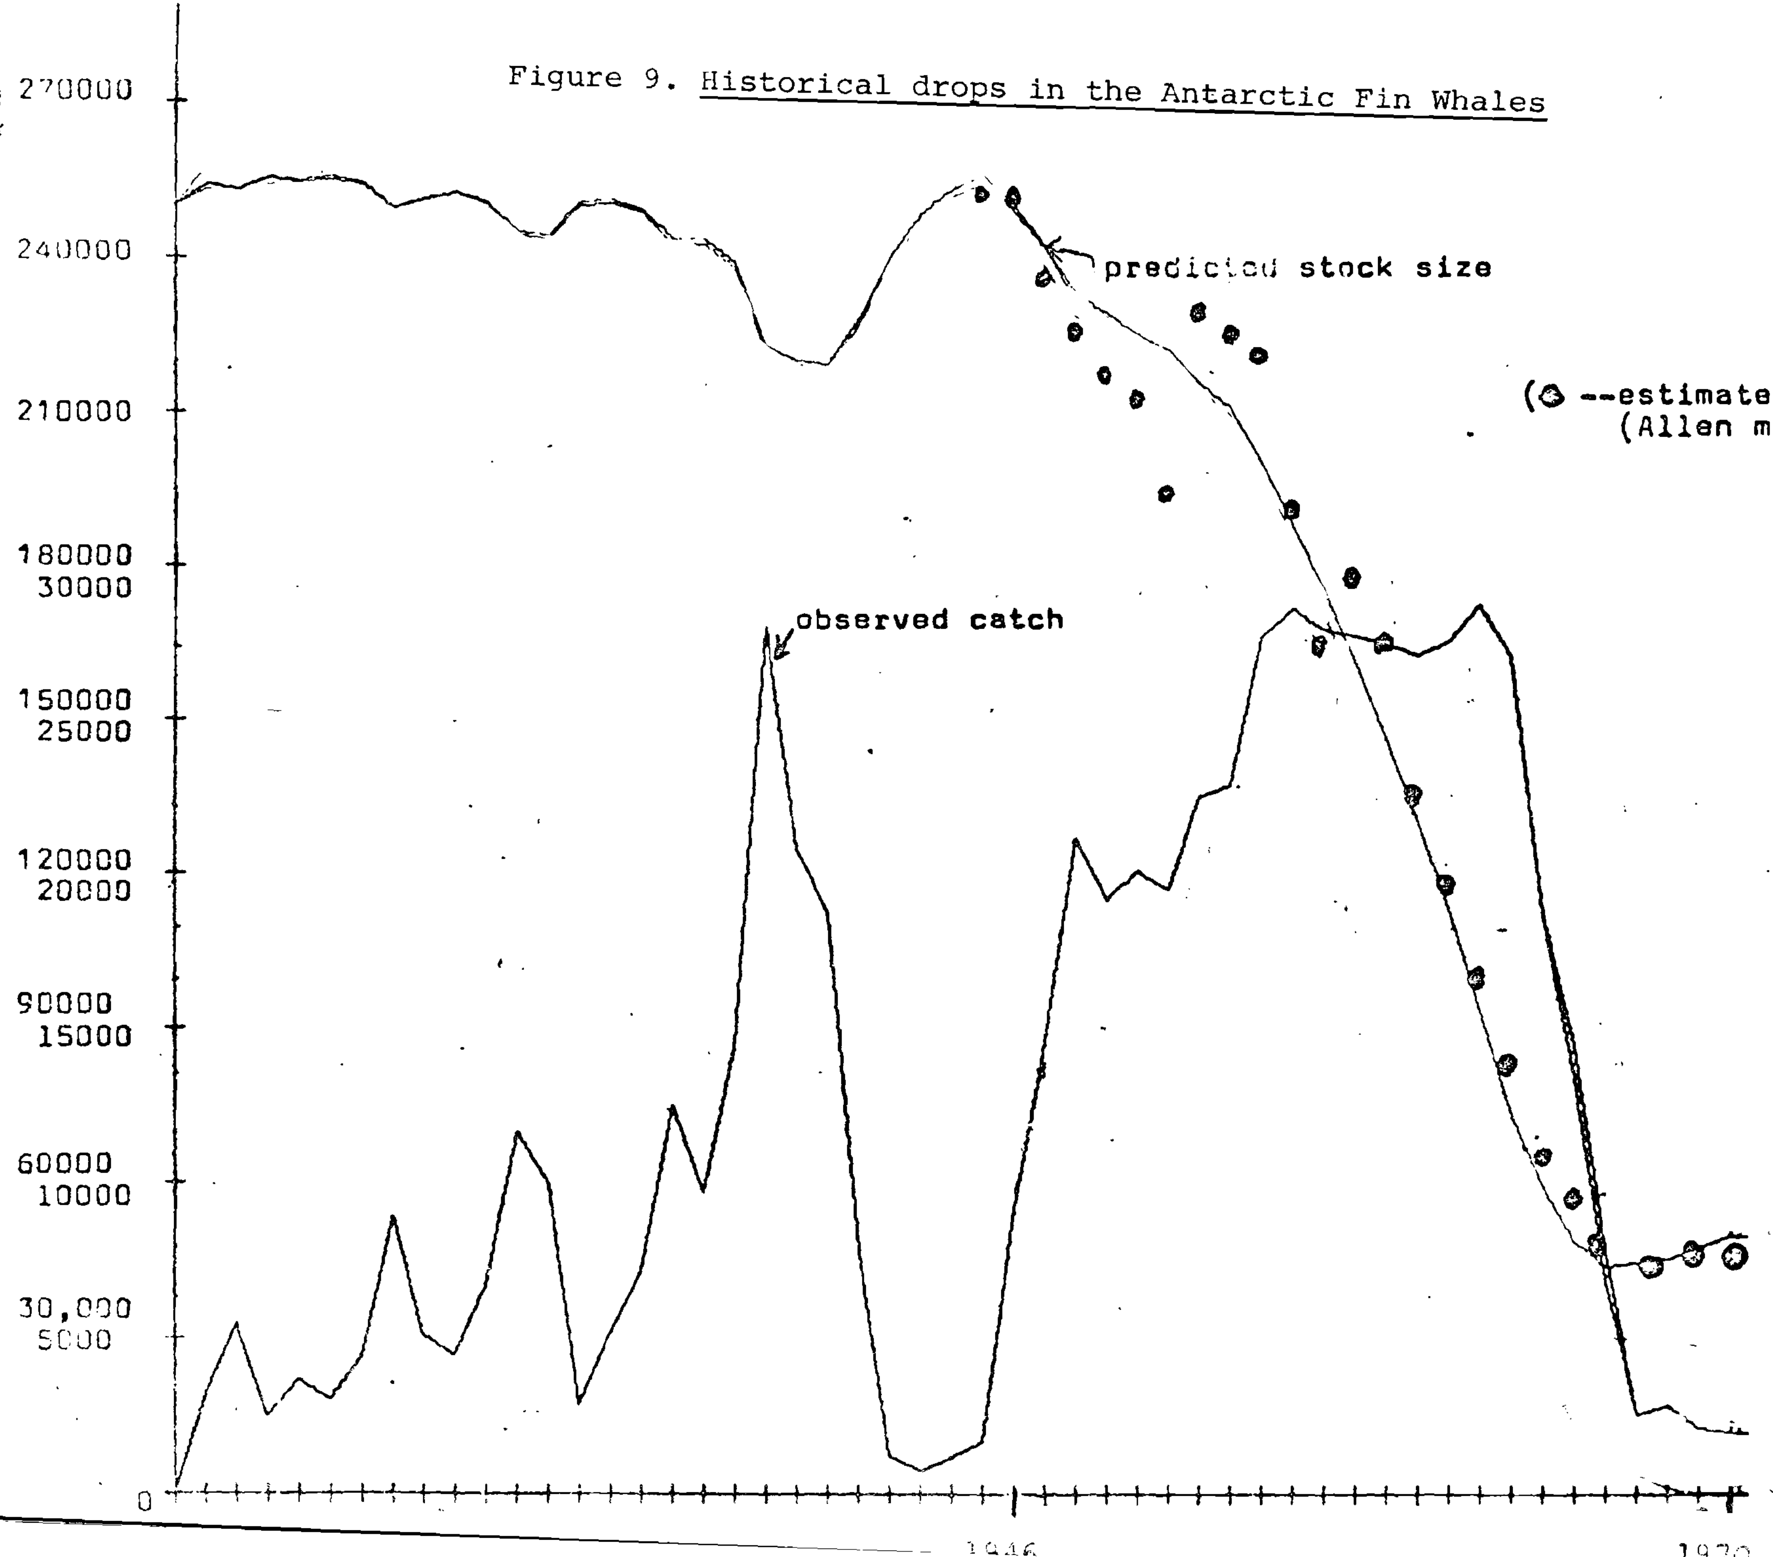

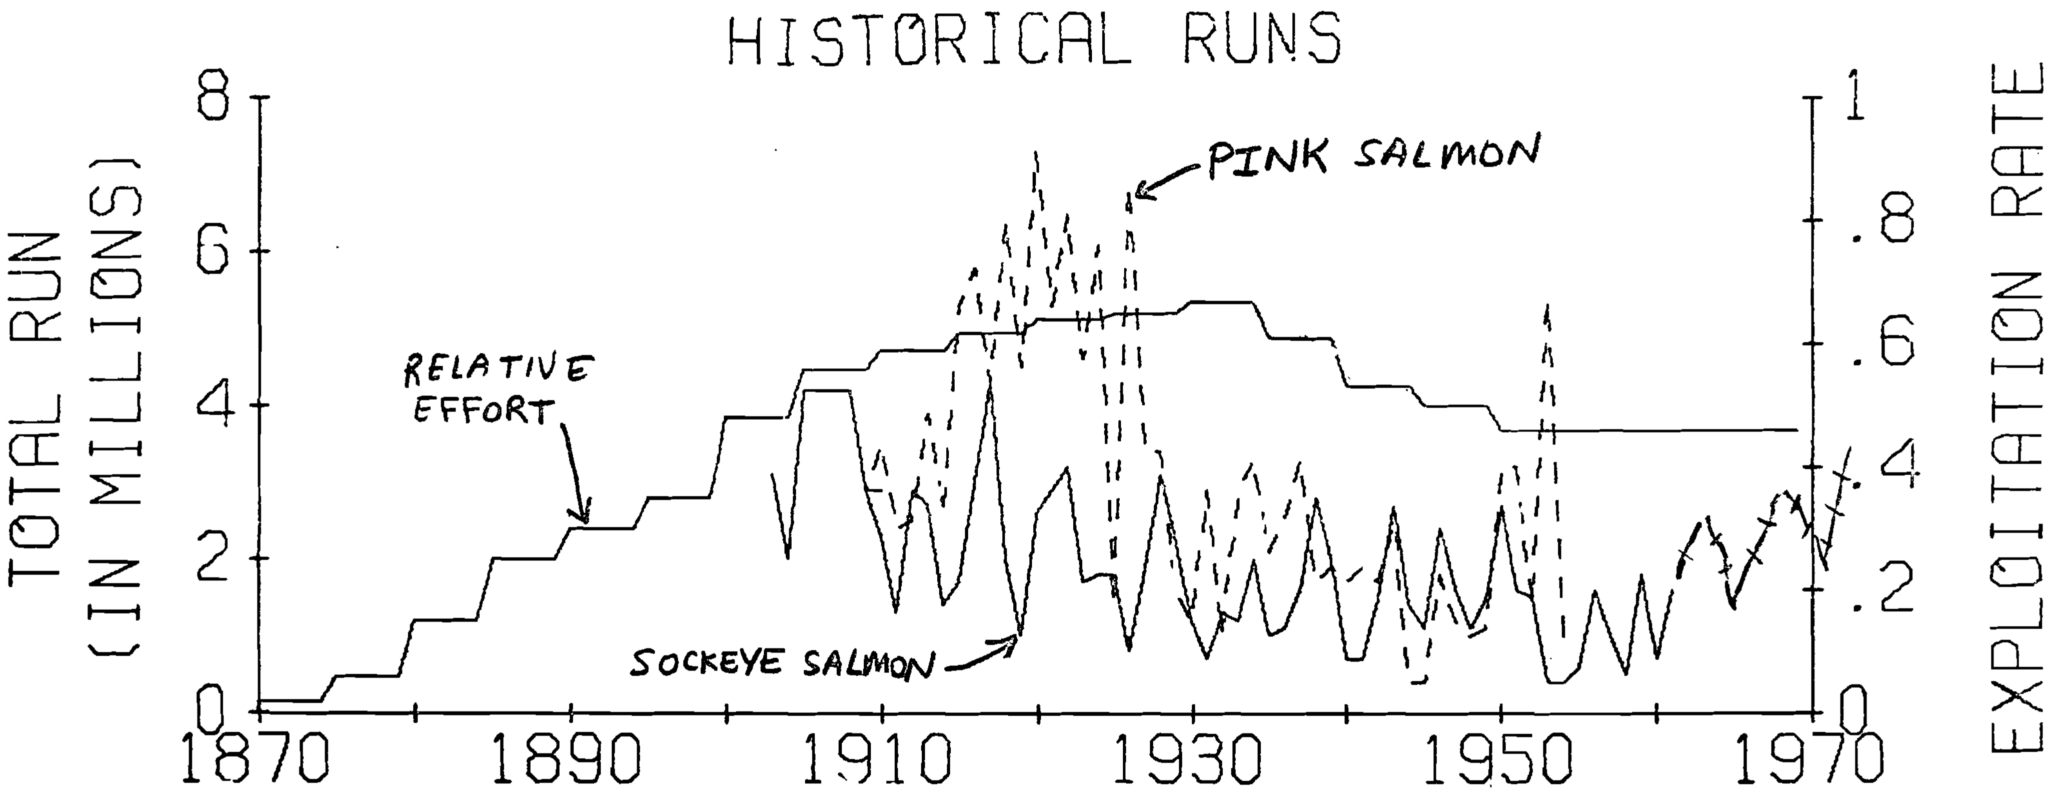

Figure 10. Historical changes in Skeena River Salmon

Table 1. E,x phase behaviour indicators for Whales and Salmon

BIOLOGICAL	FIN WHALES	SKEENA SALMON
(1) x_e/x_o	0.8	0.5
(2) g_{max}/x_e	0.05-0.1	0.5
(3) σ_{gmax}/g_{max}	0.0	~0.3
(4) τ	<3 years	2-4 years
(5) T_e	30-40 years	10-20 years
ECONOMIC		
(6) x_∞/x_o	0.2-0.4	0.1-0.3
(7) $\max\Delta E/E$	1.0	0.2
CONTROL		
(8) α_1/E	0.0	0.02-0.1
(9) α_2/E	not known (1.0?)	not known (0.2?)

DEFINITION:

(1) most productive population as a proportion of natural size

(2) harvest proportion giving maximum sustained yield

(3) coefficient of variation in annual production at optimum stock size

(4) time lag before growth rate change can be sensed

(5) T_e - time to recover to optimum level from worst likely exploitation, assuming E = 0 over recovery period

(6) Equilibrium population as a fraction of natural population, assuming asymptotic stability of E,x equilibrium and no management controls or targets

(7) maximum relative increase in effort per year

(8)-(9) maximum feasible yearly changes in effort due to management actions (α_1 = decrease)

Table 1 compares some indices of E,x phase plane behaviour for the whale and salmon examples. These are reasonable guesses at best, most fisheries cases that have been documented in the literature appear to be somewhere between these extremes. The point of table 1 is to indicate that quite general adaptive control strategies for new resources might emerge from examination of a relatively small set of parameter combinations or examples.

There has been much talk recently about "resilience limits", combinations of E,x which are likely to lead to extinction of x. There is no really convincing evidence that such limits exist in fisheries systems. For the Skeena Salmon, there is reason to believe that extinction might occur if the stocks were depleted to below about 5% of their natural levels. Some whale stocks appear to be recovering from depletion to 10% of natural size. This is not to say that sudden and dramatic population changes do not occur (witness figure 9); many examples like the fin whales could be drawn from the fisheries literature.

PARAMETER ESTIMATION IN THE DEVELPING SYSTEM

In previous sections I have tried to give an overview of the development control problem and to suggest intuitive ways of approaching it through strategic indications like the "reachability set" of figure 5. A key assumption underlying the discussion so far has been the existence of procedures for parameter and variance estimation ($\hat{\theta}$ and σ_{θ}^{2} of the previous section). In this section I explore that assumption further, using some results developed by Sergio Rinaldi.

Resource managers use three basic kinds of information in developing models and parameter estimates:

(1) Overall catch and effort statistics

(2) Catch, effort, and other sampling statistics on closed subpopulations of x created by tagging and marking

(3) Indicator statistics developed in special sampling programs (eg. egg counts for herring)

In the following discussion I will concentrate on the first category. Overall catch-effort statistics are about always available and are the cheapest to collect. There is an extensive literature on the use of these statistics for parameter estimation in particular cases, especially in relation to the Schaeffer and Ricker g models. The "estimated stock" sizes for fin whales in Figure 9 come from one of the most general procedures now available, developed by K.R. Allen (J.F.R.B., 1966) for the case when x is an age-structure vector. What follows is essentially a generalization of his results.

Consider the general dynamic model on page 2, for x one-dimensional, and suppose that it is only possible to observe

E_t (total effort)

and $y_t \equiv cx_t$ (catch per effort)

suppose further that g can be written in the form

$$g(x) = ax \cdot g_1(x,b)$$

where $g_1(x,b)$ is a monotonically decreasing function of x and a parameter b that controls the rate of decrease. Most existing fisheries production models can be cast in this form:

Ricker: $g_1 = e^{-bx}$

Schaetter: $g_1 = 1-bx$

Beverton-Holt: $g_1 = 1/(1+bx)$

Biologically, the idea is that ax represents fecundity or maximum relative growth rate, while g_1 represents the regulatory processes that prevent unbounded population increase.

Using these suppositions, the general dynamic model can be written in terms of the (observable) catch per effort, y_t:

$$y_{t+1} = S_{yt} + ay_{t-\tau}g_1(y_{t-\tau}, \frac{b}{c}) - cE_t y_t$$

For any arbitrary choice of b/c, this equation constitutes a linear regression model with dependent variable y_{t+1} and independent variables

$$\begin{bmatrix} y_t \\ y_{t-\tau}g_1(y_{t-\tau}, b/c) \\ E_t y_t \end{bmatrix}$$

The parameters directly estimable from the regression are S, a, and c. The regression can be solved interatively with different estimates of b/c to eventually arrive also at a least squares estimator for b. A key point is that a single, well known estimation procedure can be used with a wide variety of assumptions about g; the regression procedure also gives variance estimates for each parameter. Recursive procedures are available to update the estimators at each t without going through the full regression analysis with its messy details like matrix inversion. In cases where only weight yields are measurable, the estimator for S becomes instead an estimator of growth plus survivial rate. In cases where cE_t is large so that $cE_t x_t$ is a poor model for catch, the exponential "catch curve" (pg 4) can be used and c estimated iteratively along with b.

The regression formulation suggests some exciting possibilities for measuring the information gained along any path through the E,x phase plane. One obvious conclusion is that y and E must change considerably along the path- otherwise the regression variables are linearly dependent. This is equivalent to saying that "you don't learn nothing if you don't do nothing." Stochastic variation in g should have the effect of apparently decoupling the regression variables: it will be interesting to explore the tradeoff between increased independence of the regression variables versus increased error variance as a function of the degree of stochastic variation. Under what conditions does the variation help more than it hurts?

Given variance estimates at any time step for $\hat{S}$, $\hat{a}$, $\hat{b}$ and $\hat{c}$, it should be possible to compute (if necessary, by Monte Carlo methods) a variance estimate for the target equilibrium point

$$\left\{\hat{E}_e, \hat{x}_e\right\} = f\,(\hat{S}, \hat{a}, \hat{b}, \hat{c}).$$

Even better, it should be possible to compute a region of uncertainty (at some policy-determined confidence level) for the location of E_e, x_e (figure 11). The size of this

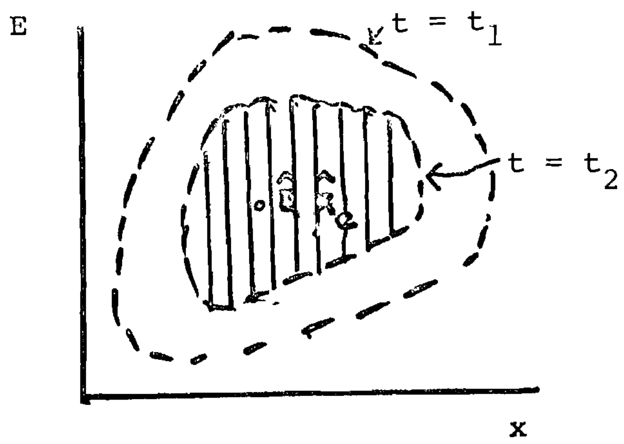

Figure 11. Region of uncertainty for estimates of optimum State

region should help provide guidelines for when to begin taking management actions to stabilize the system. If the $f(\hat{S}, \hat{a}, \hat{b}, \hat{c})$ function is structured from the dynamic model (eg. so as to maximize equilibrium catch), computation of its uncertainty region should automatically take into account the sensitivity of E_e, x_e to the various parameters. Note that the size of the uncertainty region at any time will depend on the choice of f. If the function is chosen to reflect strong aversion to risks (eg, low x values), the uncertainty region ceases to be meaningful.

The studies of Allen (mentioned above) and recent results by Sergio Rinaldi for a complex, multi-age salmon model give me hope that the methods outlined above can be extended to many cases where x is multi-dimensional.

Also summary measures like the uncertainty region should continue to be meaningful. The next step is to examine how statistical procedures can be combined with optimal control concepts to produce integrated methods for dealing with the unknown.

CONCLUSION

Rather than delve further into the integrated problem of adaptive control, it is reasonable to stop and ask whether the definitions and formalisms presented so far have anything to offer to the policy maker. Two concepts in particular have emerged that I believe can be of considerable value.

The first concept is the "reachability region" defined graphically in Figure 5. This region summarizes the interrelations of several factors: feasible limits of control acts, effects of stochastic variation, rates of economic development, and resilience of the biological system to increasing exploitation. To estimate it, the policy maker must clearly state his beliefs not only about the biological system, but also about the political and economic constraints on his actions. In short, estimation of the reachability region forces the policy maker to be realistic about the management strategies available to him.

The second concept is the "uncertainty region" (figure 11) for estimates of optimal system state. To estimate this region the policy maker must clarify his objectives (as expressed in the function for computing E_e, x_e), and he must also decide acceptable confidence levels. If computable, the uncertainty region provides the policy maker with an objective measure of how much information and understanding has been accumulated at any point in the development process. Even if we can devise no formal rules for trading off between uncertainty and the desire for temporal stability, the uncertainty region should at least provide a better reference frame for making intuitive decisions.