A POLICY FAILURE ANALYSIS OF SALMON ENHANCEMENT PROGRAMS

Ray Hilborn

February 1975

WP-75-11

A POLICY FAILURE ANALYSIS OF SALMON ENHANCEMENT PROGRAMS

Ray Hilborn

Introduction

The Canadian government has established a policy of enhancing natural salmon runs on the west coast. The basic concept of enhancement for commercial species is to provide additional artifical spawning grounds. In effect this creates new salmon stocks. The Fulton River spawning channels are the best example currently in operation; more such developments are being considered. There are several potential problems with such stock enhancement facilities. In this paper I wish to consider long range problems associated with achieving an optimal exploitation of both enhanced and natural stocks. I have discussed this problem earlier (Hilborn, 1974) and used a deterministic model to find what would happen to a natural salmon stock being harvested simultaneously with an enhanced stock with a higher productivity. Briefly, the problem is that in order to optimally harvest the combined stocks, the natural stock (with a lower productivity) would be kept at lower stock levels, thus subjecting it to a higher probability of random extinction. This concept is summarized in figure 1, which shows the equilibrium stock level of the natural stock when a combination of natural and enhanced stocks are harvested at maximum sustained yield. The larger and more productive the enhanced stock is made, the lower is the equilibrium size of the natural stock.

This model was deterministic; in nature there is a very high variance in productivities. Walters (1975) has looked at optimal exploitation rates for stochastic models of a single stock and derived several alternative policies for maximization of yield or minimization of variance of yield. My approach was to use the same stochastic dynamic programming optimization technique, but I applied it to a combination of natural and enhanced stocks. The optimal policies thus derived were analyzed by a new technique for policy failure analysis. The technique described in detail later consists of taking a single management policy and asking what happens in the event of a disaster. The two types of disaster I consider in this paper are 1) complete failure of the enhanced stock, and 2) two consecutive generations with very poor productivity.

Policies Analyzed

I have considered five possible management strategies. In all cases I assume a single natural stock with a Ricker equilibrium density of two million and a productivity of 1.3, and an enhanced stock with a Ricker equilibrium density of 2 million and a productivity of 1.8. The five management policies considered were:

1) Long term maximized yield using dynamic programming optimization.

2) Maximization of the following objective function:

Objective = the harvest + 2 * the natural stock size. (This objective function should prevent the natural stock from ever reaching very low levels).

3) A harvest curve (derived by dynamic programming) designed to minimize the variance of the harvest around 1.9 million fish per year.

4) A constant harvest rate of .594, which is the optimum long term harvest rate for a deterministic population. See Hilborn (1974) for equations.

5) A maximum yield policy (from dynamic programming) for the natural stock, with no enhancement at all.

For all of the policies except 4, stochastic dynamic programming was used to determine the actual harvest policies. This is the best method currently available for complex non-linear dynamic models. All programs and conceptual development were done independently from those of Walters (1975), and our results were identical for the single stock case under policies 1, 2, and 5. This gives us greater confidence than usual with our own programming.

The next section presents the technique of policy failure analysis used and then applies it to a very simple case, our five salmon policies. This is primarily an exercise in methodology. Now that we are satisfied that it works, we will later apply the methodology to a more realistic salmon model which keeps track of the age classes, has adults returning at four and five years, etc.

Policy Failure Analysis

Policy failure consists of an unexpected occurrence in the managed system which disrupts maximization of the objective function. Such failure may be due to natural events such as poor weather, disasters, etc., or man made changes or restrictions outside our control as system managers. For instance, the decision to build a hydro development on an important salmon stream made by another agency would be a policy failure to a salmon manager. Some kinds of policy failure are explicitly taken into account in stochastic dynamic programming situations. For instance, several years of poor productivity are a possible stochastic outcome recognized in the optimization. In general, the kinds of policy failure we wish to consider will be external to the model and we will have to artifically cause the failure to happen in the model. We then see how the system, as represented by the model, would respond to this form of failure.

In this salmon analysis, the two years of bad productivity, or weather, are implicitly optimized using stochastic dynamic programming. We consider this a policy failure only to explicitly look at the time stream of payoffs if we do get these two bad years. The total enhancement failure is completely external to the model and is more typical of the types of policy failure usually comsidered with this type of analysis.

There are three steps in the analysis of policy failure. First, we must decide which types of policy failure we wish to consider; second, we must assess the subjective probability of each of these failures occurring; and third, we must find a set of techniques for assessing the consequences of the failure. The end product of policy failure analysis should be a table listing for every policy, the possible forms of policy failure, the probability of failure, and the cost of failure (table 1).

Defining the objective functions and the types of policy failure is a task best suited for system managers in concert with systems analysts. There are no formal rules for this step in the analysis and I will not consider it further. Calculating the probabilities of the failures occurring is also a difficult task. If the policy failure is a natural event, some form of historical time series analysis may prove the best technique. If the failure is a man made one, deciding the probability of failure is a subjective judgment and is probably best left up to the management agency.

Having ignored the first two steps in policy failure analysis, we believe we can offer some good techniques for assessing the cost of policy failure. To measure this cost, we must first define what the payoffs are so that we know what we lose by a policy failure. This again touches on the question of objective functions, and for salmon we used the total annual catch as the measure of payoffs. We have a much more sophisticated method of measuring payoffs for complex systems such as the budworm, and this method is described elsewhere. Given our payoffs (total catch), we ask what happens when a policy failure occurs.

We now must introduce the concept of manager's time scale (MTS). MTS is a measure of over what period the manager responsible is interested in what happens to the system. If the system itself is rapidly changing and policy failures will happen over a short period, for instance a strike in a municipal sewage treatment plant, then the MTS is very short. If the system is a much slower one and problems arise slowly and have long effects, then the MTS will be much longer. An example of this might be an erosion prevention program, or forest management, both of which have long time periods associated with management. The MTS is also a function of the institutional framework of the management agency. If the persons responsible for responding to policy failure change rapidly, then the MTS will tend to be much less than if the same person tends to be in charge for long periods of time. Given these considerations, the persons performing the policy failure analysis must select what they believe the appropriate MTS, but the policy failure analysis can be done for several possible MTS's and the results compared. For the salmon analysis we have chosen five generations (20-25 years) as the appropriate time scale.

The purpose of choosing a MTS is that when we ask: "What happens to our payoffs if this type of policy failure occurs?", we must have a time scale in which to assess the consequences of the failure. Our technique is to run the model for the MTS under each type of policy failure and measure the payoffs under that failure. This is a bit more complicated than meets the eye. The cost of policy failure greatly depends on the state of the system when policy failure occurs, and the state of the system at the time of policy failure. This in turn depends on the management tactics being used. Our technique involves running the model for many intervals (5000 years) under each management option to assess the long term payoffs over the MTS. This must be repeated many times so that the state of the system at the point of policy failure will assume a frequency distribution similar to the long term frequency distribution. For complex cases like the budworm, discrete states are defined and the long term probability of being in that state is multiplied times the cost of failure if the system was in that state (this whole procedure for the budworm is described elsewhere).

We can now construct the first table of cost of policy failure (table 2). For a simple objective function such as annual catch it is fairly easy to see what happens under policy failure from this table. However, there is a further step in the analysis: We shall attempt to directly measure the "resilience" of various management tactics. Without going into an in-depth review of resilience, let me define a resilient strategy as one whose payoffs are not reduced by a policy failure. Let us scale everything from zero to one so that a strategy that loses no payoff by policy failure has a "resilience" of one and a policy that loses the maximum amount of payoff has a resilience of zero. Thus resilience is defined as

1.0 - (payoffs before policy failure - payoffs after policy failure).

The payoffs must also have been scaled between zero and one. What I have used as the maximum was the highest payoff found under any management strategy, which for this study is the long term payoffs under the maximum yield strategy (A). Thus we can present a new payoff table (table 3) with all payoffs scaled between zero and one, and from this table calculate a resilience table (table 4). A slight problem with this analysis is that any strategy which does not have a long term payoff of 1.0, cannot have a resilience of zero, even if the stocks are completely wiped out. We might alternatively define the resilience as the proportion payoffs lost under policy failure. The basic question is whether we are interested in the absolute magnitude of payoff loss, or the relative one.

In more complex ecological systems it is possible to produce irreversible effects due to some management practices and policy failures. The only irreversible effect possible for this salmon model is the total elimination of a stock, which does not happen under any of our proposed management tactics. For systems where irreversible changes do occur, we want to assess the long term cost of the policy failure as well as the cost during the MTS. To do this we must run the model for a very long period after policy failure, again repeating it many times to approximate the natural distribution of states at the point of policy failure. This would produce an additional column at the bottom of each table, listing long term benefits after a policy failure.

Discussion

Despite the simplifying assumptions used in this model, we can draw some useful conclusions from the results in tables 2, 3, and 4. It is clear that policy 1, the long term yield optimization, produces the highest yield under all policy failure. This is not surprising, considering the technique of dynamic programming used: the rules for optimal yield have been worked out for situations when the enhanced stock is at low levels, or when there are two consecutive generations of poor productivity. The second policy, maintenance of old stocks, does not look particularly good. The size and productivity of the natural and enhanced stock used here never brought the natural stock near extinction, so the yield after policy failure was not better for this policy than the maximum yield. The minimized variance policy looks very good. Although the long term yield is considerably lower than the maximum yield, there are many benefits to maintaining a somewhat constant harvest. The fleet may not have the capacity to harvest at the highest possible rates and the canneries may not be able to process the really big runs. Both the fishermen and the canners may well be willing to sacrifice a little in long term yield for a much more reliable income. Walters (1975) has discussed this also. Under the two types of policy failure considered here, the minimized variance policy is particularly good. It is very resilient to both these failures (see table 4), and the actual harvests are not substantially lower than the maximized yield policy. The fifth management policy was included mostly for comparison.

The fixed harvest rate policy is clearly inferior to the dynamic programming optimization of policy 1. This is natural and really not worth any more discussion. Since there was no enhanced stock to fail, it has a resilience of 1.0 to enhancement failure. The resilience to bad weather was high because the changes were small relative to the value used as the maximum. If the ratio method of calculating resilience (mentioned earlier) had been used, then the resilience of the no-enhancement policy would have been comprable to that of the maximum yield policy for two stocks.

It is clear that the best policy is either the maximum yield or minimized variance. The choice is up to the decision makers. This analysis makes it clear what is sacrificed in total yield for a more steady income. A distribution of incomes similar to that presented by Walters (1975) might prove a useful addition when presenting these options to a policy maker. We are now examining the possibilities of an automatic insurance system which would allow the fishermen to be paid back in bad years for money accumulated in good years. However, this does not resolve the problem of cannery capacity. We shall test these conclusions against the more complex model, but from our current understanding of the system it is difficult to see how our conclusions will differ.

FIGURE 1. EQUILIBRIUM DENSITY OF OLD STOCK (× 1000)
SI = 600000 AI = 1.5

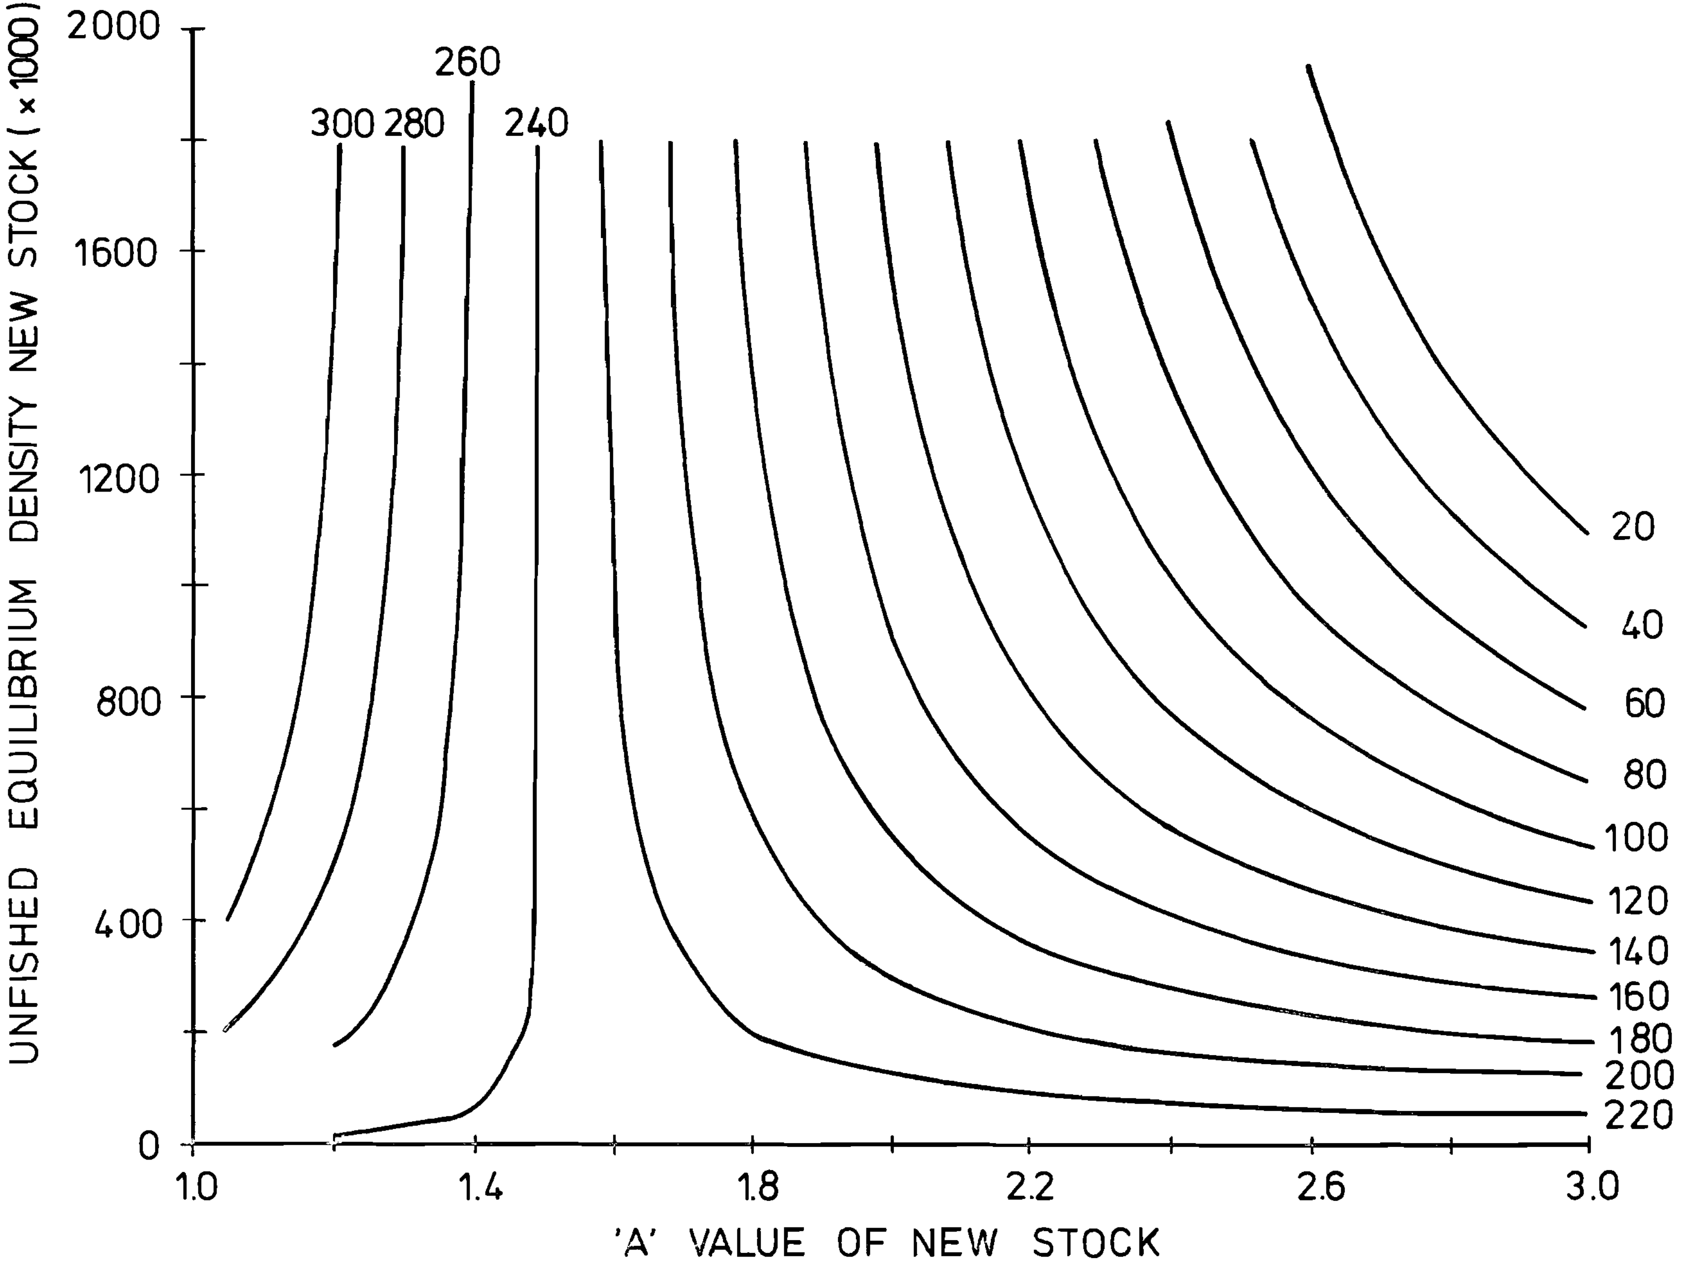

Table 1

IDEAL POLICY FAILURE ANALYSIS TABLE

		POLICY 1	POLICY 2	POLICY 3
POLICY FAILURE 1	PROBABILITY			
	COST			
POLICY FAILURE 2	PROBABILITY			
	COST			
POLICY FAILURE 3	PROBABILITY			
	COST			

Table 2

BENEFITS
(AVERAGE ANNUAL CATCH IN MILLIONS)

MANAGEMENT POLICY

	A MAXIMIZE YIELD	B MAINTAIN OLD STOCK	C MINIMIZE VARIANCE	D FIXED HARVEST RATE	E NO ENHANCEMENT ONLY OLD STOCK
LONG RUN AVERAGE	2.50	2.15	1.82	2.36	1.01
5 YEARS FOLLOWING ENHANCEMENT FAILURE	1.03	.87	.99	.92	1.03
5 YEARS FOLLOWING 2 VERY-BAD-WEATHER YEARS	1.77	1.56	1.56	1.62	.71

Table 3

BENEFITS SCALED TO A MAXIMUM OF 1.0

A	B	C	D	E
1.0	.86	.73	.94	.40
.41	.35	.40	.37	.41
.71	.62	.62	.65	.28

Table 4

RESILIENCE INDICATORS

	A	B	C	D	E
RESILIENCE OF LONG TERM BENEFITS	1.0	.86	.73	.94	.40
RESILIENCE TO ENHANCEMENT FAILURE	.41	.49	.67	.43	1.0
RESILIENCE TO BAD WEATHER	.71	.76	.89	.71	.88

References

Hilborn, R. 1974. Stock Enhancement in Salmon and Maintenance of Historic Runs. IIASA WP-74-63.

Walters, C.J. 1975 Optimal Harvest Strategies for Salmon in Relation to Environmental Variability and Uncertainty about Production Parameters. IIASA WP-75-4