

Optimal Harvest Strategies for Salmon in Relation to Environmental Variability and Uncertainty About Production Parameters

H

HH

RH I

Walters, C.J.

IIASA Working Paper

WP-75-004

1975

Walters CJ (1975). Optimal Harvest Strategies for Salmon in Relation to Environmental Variability and Uncertainty About Production Parameters. IIASA Working Paper. IIASA, Laxenburg, Austria: WP-75-004 Copyright © 1975 by the author(s). http://pure.iiasa.ac.at/id/eprint/424/

Working Papers on work of the International Institute for Applied Systems Analysis receive only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute, its National Member Organizations, or other organizations supporting the work. All rights reserved. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage. All copies must bear this notice and the full citation on the first page. For other purposes, to republish, to post on servers or to redistribute to lists, permission must be sought by contacting repository@iiasa.ac.at

OPTIMAL HARVEST STRATEGIES FOR SALMON IN RELATION TO ENVIRONMENTAL VARIABILITY AND UNCERTAINTY ABOUT PRODUCTION PARAMETERS*

Carl J. Walters**

January 1975

WP-75-4

* Research supported by Environment Canada and by the International Institute for Applied Systems Analysis.

** Institute of Animal Resource Ecology, The University of British Columbia.

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.

.

Abstract

A method is developed for incorporating the effects of environmental variability and judgmental uncertainty about future production parameters into the design of optimal harvest strategies, expressed as curves relating stock size and exploitation rate. For the Skeena River Sockeye, the method suggests that optimal strategies are insensitive to judgmental uncertainty about the Ricker Stock production parameter, but are very sensitive to management objectives related to the mean and variance of catches. Best possible tradeoffs between mean and variance of catches for the Skeena River are developed and a simplified strategy is suggested for improving mean catch while reducing year to year variation.

1. Introduction

Pacific Salmon management in recent years has been based on the concept that maximum sustained yield can be obtained by holding escapements at some constant level determined by analysis of the stock-recruitment relationship. Larkin and Ricker (1964), and Tautz, Larkin, and Ricker (1969) showed that such fixed escapement strategies should result in higher mean yields than fixed exploitation rate strategies in the face of high stochastic variation in productivity. However, Allen (1973) has stressed the need to look at other possible management strategies expressed as relationships between harvest and stock size; he shows for the Skeena River that fixed escapement strategies should result in unnecessarily high variance in catches from year to year, and he develops alternative relationships that should cut the variance of catches nearly in half with only about a 15% reduction in mean catch.

The intent of this paper is to present a set of optimal harvest strategies for salmon, based on tradeoffs between the mean and variance of catches. The Skeena River is used as an example, and the optimal strategies are developed by using stochastic dynamic programming. This formidable sounding optimization technique is actually a relatively simple method for testing the multitude of possible future stock changes that harvest and environmental variability may produce, weighting each future change by its probability of occurrence.

Since the technique has seen little application in biology, section II gives an intuitive introduction to stochastic dynamic programming. Section III presents a variety of harvest strategies for the Skeena River, under different assumptions about environmental variability and using different management objectives, and examines possible management strategies in relation to current management practice on the Skeena River. Section IV analyses potential tradeoffs between mean and variance of catches, and suggests an overall optimal strategy for the Skeena River. Hopefully it is demonstrated that optimal management policies should bear no clear relationship either to the current (fixed escapement) practice or to the strategy alternatives suggested by Allen (1973).

II Stochastic Dynamic Programming

The basic concept of dynamic programming was introduced by Richard Bellman in the 1940's (See Bellman, 1961; Bellman

-2-

and Dreyfus, 1962; Bellman and Kalaba, 1965). It is an optimization technique for systems in which a series of decisions must be made in sequence, where each decision affects the subsequent system state and thus each future decision. Two key ingredients are necessary to apply the method: a dynamic model to predict the next state of the system given any starting state and any decision, and an objective function to specify the value of the return obtained in one time step for any state-decision combination. In stochastic problems, the dynamic model must specify not a single future state but instead must specify probabilities for each new state that might arise after one time step from any starting state-decision combination.

The dynamic model

Following most authors on salmon management theory, the simple Ricker model is used in this study as the necessary dynamic model:

$$N_{t+1} = S_t e^{\alpha (1-S_t)}$$
[1]

where

N_{t+1} = Stock (recruitment) after one generation, in standard stock units (approximately 2,000,000 for Skeena Sockeye)

 S_{+} = Escapement or spawing population, in stock units

a = stock production parameter, assumed to be a random
variable.

If S_t is held fixed, e^{α} represents the net stock productivity or recruitment excess. This factor arises in nature as a product of several survival factors that vary randomly but may be considered more or less independent of one another. Thus, α , the logarithm of e^{α} is a sum of random variables and should be normally distributed by the Central Limit Theorem of basic statistics. Allen (1973) provides good empirical justification for this assumption using data from the Skeena River. If S_t is written as

$$S_{+} = N_{+}(1 - u_{+}) \qquad 0 \le u_{+} \le 0 \ 1.0 \qquad [2]$$

where u_t is the exploitation rate, or <u>decision variable</u>, then we have the first basic ingredient for dynamic programming. The objective is to find an optimal <u>relationship</u> between u_t and N_t , by examining sequences of decisions where the next state arising from any $N_t - u_t$ combination is predicted with the Ricker model using an appropriate probability distribution for α .

As an alternative to the Ricker model, we could simply specify a separate empirical or judgmental probability distribution of recruitment for each conceivable spawing stock (in other words, treat the stock-recruitment relationship as a Markov process). However, even for the Skeena River Sockeye there is insufficient data to meaningfully interpolate recruitment probabilities for high and low spawningstocks (Figure 1). The Ricker model appears to be as good a way as any for extrapolation to extreme stock sizes.

The objective function

The other basic ingredient, the objective function, may take a variety of forms. For maximizing mean harvest, we can take it to be simply $u_t \cdot N_t$. If variance is important, we can instead try to minimize the variance around some desired catch level; for each time step the relative contribution to variance is then

$$(u_t \cdot N_t - \mu)^2$$

where μ is the desired catch level. Note that if μ is arbitrarily increased to high values that cannot be achieved in nature, the variance contribution at each step becomes essentially linear in $u_t \cdot N_t$. This means mathematically that minimizing the sum over time of squared deviations from high μ values tends toward being equivalent to maximizing $u_t N_t$, as μ is increased. Thus by changing μ we can generate a series of objective functions that range from variance-minimizing to harvest maximizing as μ is increased (this point will be clarified in Section IV).

The computational procedure

Given the basic ingredients above, the next step required for dynamic programming is to approximate the continuous variables $u_{+}^{}$, $N_{+}^{}$ and α by a series of discrete, representative

-5-

levels or states. The concept here is the same as is used in solving differential equations by taking short discrete time steps. By trial and error, it was found necessary for this study to use 30 discrete population levels, each representing an increment of .05 stock units ($N_t = 0.0, 0.05, 0.1, ..., 1.45$), 30 discrete exploitation rates at intervals of 0.03 ($U_t = 0.0, 0.03, 0.06, ..., 0.82$), and 10 discrete α values (α discretization will be presented in Section III).

The reader is referred to figure 2 for the following explanation. Suppose we look at any discrete stock size at some time step, and think about applying many possible harvest rates to it (left hand "decision branches" in Figure 2). For each harvest rate a return (harvest or contribution to variance) can be computed, but the recruitment subsequently resulting from this escapement will be uncertain (right hand "probability branches" in Figure 2). Suppose that we specify probabilities for each possible new stock size that might be produced, and suppose that we already know (somehow) what future returns can be expected for each of these new stock sizes. Then for each harvest rate, we can find an expected overall value: it is simply the return this year, plus the sum of products of probabilities of getting new stock sizes times the expected future returns for these new sizes. In other words, we take each possible future and weight it by its probability of occurrence to give an expected value for future returns; this expected future value is added to this year's return to give

-6-

the overall value for the harvest rate-present stock combination, for the particular time step under consideration. The process can be repeated for each possible harvest rate, and afterwardsit is a simple matter to choose which rate gives the best overall return.

We can next choose another stock size, and try many possible harvest rates on it. Again providing that we already know what future returns can be expected for each new stock size that might result, and that we can associate a probability with each possibility, it is a simple matter to choose the best harvest rate for this second stock size.

The whole process is repeated for a third stock size, a fourth, and so on until the optimal harvest rate for every reasonable stock size has been computed. The result is a set of stock-harvest combinations that can be plotted against one another as a smooth curve; this curve is called the <u>optimal</u> control law for the time step under consideration.

The real trick in dynamic programming is to get the expected future returns for each new stock size that can result for each starting $u_t - N_t$ combination. This trick, the key discovery of Richard Bellman, is remarkably simple: we work <u>backwards</u> in time from an arbitrary end point (t = K). Values are assigned to different stock sizes at this endpoint, and these values are used to look ahead at the endpoint from one time step backward (t = K - 1). After getting overall values for each stock size one step back from the endpoint, we can

-7-

then move back another step (t = K - 2), and look ahead to the values just computed for t = K - 1. This <u>backward recursion</u> process is repeated over and over (t = k - 3, K - 4, etc.)

After several backward recursion steps, a phenomenon emerges that forms the central basis for this paper: the endpoint values cease to have any effect, and the optimal exploitation rate for each stock size becomes independent of the time step. The optimal control law or harvest strategy curve is then said to have stabilized; this usually occurs within 10 - 20 steps for the Ricker model. Certain computational tricks are necessary to insure that the stable control law is valid, since the new stocks produced at each forward look may not correspond exactly to any that have already been examined for the next time step forward. This interpolation problem is solved by being careful to examine enough discretized stock sizes and exploitation rates.

The key feature of stochastic dynamic programming is that it explicitly takes account of all the possible futures that are considered likely enough to be assigned probabilities of occurrence. Furthermore, it makes no difference whether these probabilities are chosen to represent judgmental uncertainty (Raiffa, 1968) about deterministic parameters, or true stochastic variation in parameter values, or some combination of these sources of uncertainty.

III Optimal Strategy Examples

This section develops a set of judgmental probability

-8-

distributions for the α parameter of equation 1, using the Skeena River Sockeye as an example. These probability distributions are then used to demonstrate the form of optimal harvest curves obtained by the procedures outlined above, for different objective functions. Simulation results are presented to show the likely consequences of applying the harvest curves, in terms of probability distributions of catches and stock sizes. Finally, alternative harvest curves are compared to actual management practice on the Skeena River.

α distributions for the Skeena River

Using the data in Figure 1, a set of empirical α values can be computed as

$$\alpha_{i} = \ln \left(\frac{R_{i}}{S_{i}} \right) / \left(1 - \frac{S_{i}}{S_{e}} \right)$$

where i is the data point

R_i, S_i are the recruitment and spawner values
S_e is the replacement number of spawners in the
 absence of harvest.

 S_e was taken to be 2,000,000 spawners, and the results for α are presented in Figure 3, top panel. As Ricker (1973) points out, there has been a decrease in the mean value of α in recent years. With some imagination, one might conclude that the frequencieshad been drawn from a normal distribution; luckily, no such assumption is necessary in order to apply stochastic dynamic programming.

The bottom panel of Figure 3 shows three judgmental probability distributions that a decision maker might draw after examining the top panel. These test distributions are all truncated at 0 and 2.3, for computational convenience (test runs showed that extreme values have little effect for the present problem). The distribution marked "pessimistic" (for obvious reasons) assumes an even distribution of α values in the future. The distribution marked "natural" is the author's artistic (?) rendition of the data, weighting recent years more heavily. The "optimistic" distribution might be drawn by a decision maker who believes that the good production rates of recent years (Figure 1) will continue in the future due to better management practices of some sort. An important concept behind these distributions is that the stochastic dynamic programming solution can be made to take a variety of intuitive judgments into account, beyond the hard facts of past observations.

Form of the optimal solution

The judgmental probability distributions in Figure 3, combined with equations (1) and (2) and with several objective functions, were used to obtain a variety of optimal solutions. For the computer freaks, I used a PDP 11/45; each solution required about 100 sec of computer time (30 N_t levels × 30 u_{t} levels × 10 probability levels × 20 time steps). The discrete N_t - u_{t} optimal solutions were connected as smooth curves for presentation here.

-10-

Let us first examine the dome shaped band of optimal harvest curves indicated by horizontal shading in Figure 4. All three curves were generated by trying to minimize the objective function $(H - .6)^2$, that is by trying to minimize the variance of catches around a mean value of 0.6 million fish. The top curve represents the strategy that should be followed if the optimistic probability curve for α (Figure 3) is considered best; the lower two curves represent optimal strategies for the natural and pessimistic α probabilities of Figure 3, respectively. The most important conclusion to be drawn from these curves is that the optimal strategy (for minimizing $(H - .6)^2$) is quite insensitive to the judgmental probability distribution for α , except when stock size is between 0.4 and 1.0 million fish. In hindsight, it is easy to give intuitive reasons for the shapes of the curves: very low stocks should not be fished since recovery will be showed, and high stocks should be fished lightly so as to avoid high, variance-generating catches. An assumption of the Ricker curve becomes important for high stock sizes, namely that large numbers of spawners will not result in very low recruitment in later years.

Similar results are obtained for the objective of trying to minimize the variance of catches around a mean value of 1.0 million fish (vertical shaded curves in Figure 4). Again the prediction is that low stocks should not be fished at all, while high stocks should receive moderate exploitation. The most interesting curves in Figure 4 are for the maximum harvest objective function. These curves essentially call for a constant escapement of around 0.8 - 1.0 million spawners, as suggested by earlier authors. Also, the optimal strategy is almost independent of the judgmental probability distribution for α . In other words, current management polices on the Skeena River should result, if they can be followed, in maximum average catches even if the future distribution of α values is quite different from what it has been. Predicted catch and stock size distribution

Since the stochastic optimal solutions are based on the assumption that there is no certain future population trend, the anticipated returns by applying them are best presented as probability distributions. The simplest way to approximate these distributions is by making very long simulation runs, using equations (1) and (2), with an appropriate random number generation procedure for α values.

Figure 5 presents catch distributions from 5000 year simulation trials, for the optimal harvest curves from Figure 4 that should be used if the "natural" α distribution is considered most credible. Results are also presented for a harvest curve shown in Figure 7, that was obtained by trying to minimize the variance of catches around a mean value (not achievable) of 2.0 million fish. The results in the top panel of Figure 5 were generated by actually using the "natural" distribution to choose different α values for each simulated year; the results in the bottom panel were generated by choosing α values from a normal distribution with mean 1.3 and standard deviation 0.5 (after Allen, 1973). The results are quite similar, again suggesting that the optimal strategies should be insensitive to the realized future distribution of α values. The roughness of the curves for the "natural" α distribution is due to the numerical approximation procedure used in the simulation program.

Thereshould be an additional benefit from the varianceminimizing strategies, as shown in Figure 6. The variance of recruitment stock sizes increases progressively, and the mean stock size decreases for strategies that place more emphasis on maximizing mean catch. This is a surprising result, since the catch maximizing strategies tend to produce stabilized escapements.

Comparison to actual Management Practice

Catch and escapment statistics kindly provided by F. E. A. Wood, Environment Canada, were used to compute actual harvest rates for the Skeena River Sockeye (Figure 7). It is apparent that management practice in recent years has been able to follow the best fixed escapement policy quite closely. The optimal harvest curves in Figure 7 (all for "natural" α assumption) represent a spectrum of possible objectives based on trying to minimize the variance of catches around a series of increasing values.

For the 15 year period before 1970, Figure 7 suggests that

management practice more closely followed a strategy of trying to minimize the variance of catches. The correlation could be purely spurious, but it is tempting to speculate. Management decisions are open to pressure from the industry to allow higher catches in low stock years, and the industry may be unwilling to accept excessively high catches in the good years. If fishing decisions have been affected in these ways in recent years, one wonders about the Wisdom of pursuing fixed escapement policies. This question is the central topic of the following section.

IV Tradeoffs between Mean and Variance of Catches

The results in Allen (1973) and Figures 5 and 6 clearly suggest that management strategies can be devised to significantly reduce the variance of catches without intolerable losses in average yield. The aim of this section is to quantify the best possible tradeoff relationship between mean and variance of catches, so that the question of what is "intolerable" can be subjected to open negotiation. This analysis leads to a simplified optimal harvest law that can be practically implemented as an alternative to fixed escapement policies.

Definition: The Pareto Frontier

It is necessary to introduce a concept at this point that may be unfamiliar. Suppose one picks a value for the variance of catches, and then asks for the maximum mean catch that can be obtained at this level of variance. Presumably there is some answer to this question, and some optimal harvest strategy

-14-

that will do the job. One can then pick another variance value and ask the same question about mean catch. If one demands 0.0 variance in catches from the Skeena River, then the maximum mean catch is not likely to exceed about 0.4 million. On the other hand, if one says that any variance is tolerable, then he can be presented with the maximum harvest strategy from Figure 7 with its associated mean value. The set of variancemean combinations that can be generated in this way is known as a <u>Pareto Frontier</u>. In any decision problem where there are tradeoffs between different kinds of benefits, the highest achievable combinations are said to define the Pareto Frontier. Presumably the only management strategies worth considering are those which generate points along the frontier.

The variance minimizing objective functions used to obtain the harvest curves of Figure 4 and 7 are asking essentially the same questions, but in reverse; for any desired mean value, they ask for a minimum variance harvest curve. Unfortunately, stochastic dynamic programming does not permit us to ask the questions the other way around without doing excessive additional computation. As we ask for higher and higher mean values with the variance-minimizing objective functions, the optimal solutions place more and more weight on getting higher catches, and correspondingly less on reducing variation (which is always large if the desired mean value is impossibly high).

-15-

Application to the Skeena River Sockeye

Thus the harvest strategies in Figure 7 should generate (approximately) values along the mean-variance Pareto Frontier. Figure 8 presents this frontier for two possible α distributions. Points along the upper frontier were obtained by 5000 year simulations with "natural" α probabilities and associated optimal harvest curves, while points along the lower frontier were obtained by simulating with the pessimistic α probabilities and their associated harvest curves. Observed catchvariance combinations for the past two decades have been well below the potential suggested by the "natural" α distribution. Since the catch-variance combination since 1960 has been well above the pessimistic frontier, and stocks have increased steadily over this period, the pessimistic frontier is clearly too conservative. The main suggestion of Figure 8 is that the average catch of the past decade could be either:

- (1) maintained with an extreme reduction in variance (using an $(H .8)^2$ strategy curve)
- (2) increased by 25% (O.2 million fish) while maintaining the same variance (using an $(H 2)^2$ strategy curve)
- (3) or increased by (perhaps) 39% (0.3 million fish)

while increasing the variance by about 50%. The average catch over the 1970-1974 period has actually been around 0.9 million fish, as it should be according to figure 8, but a variance estimate for this short period would hardly be meaningful.

A simplified strategy for practical implementation

The optimal strategy curves based on variance minimization would be difficult to implement in practice, since they call for very good control of annual exploitation rates. Figure 7 suggests that such control is not yet available, even if it were possible to negotiate a best point along the Pareto Frontier of Figure 8. Thus a simplified strategy is suggested in Figure 7. This strategy recommends to:

- (1) take no harvest from stocks less than 0.5 millionfish
- (2) use exploitation rates between 0 and 50% for stocks between 0.5 and 1.0 million fish
- (3) use a 50% exploitation rate for all stock size above 1.0 million.

This strategy should result in a mean-variance combination (Figure 8 and 9) nearly on the frontier of best possible combinations, with a mean catch (0.94 million fish) near the 1970-74 observed average and a 20% reduction in variance from the 1955-1974 average. By calling for a fixed exploitation rate (and thus fixed effective fishing effort) most of the time, the simplified strategy should be less costly to implement since it should not require close monitoring of escapements during each fishing season.

-17-

V. Conclusions

While I have concentrated on the Skeena River as an example, the methods outlined in this paper should be applicable in many fisheries situations. The stochastic programming solutions can be performed with any stock model that has relatively few state variables (<7 for modern computers), and it is certainly possible to design more complex objective functions to take a variety of cost and benefit factors into account.

To summarize the previous sections:

- Stochastic dynamic programming provides a mechanism for incorporating judgmental uncertainty about production parameters into the design of optimal managemental strategies.
- (2) Optimal strategy curves (exploitation rate versus stock size) are relatively insensitive to the judgemental probability distribution for the Ricker stock production parameter.
- (3) Optimal strategy curves are very sensitive to changing management objectives related to mean and variance of catches.
- (4) Strategies for reducing the variance of harvests should also lead to higher and more predictable stock sizes.
- (5) Potential tradeoffs between mean and variance of catches can be quantified along a Pareto Frontier for decision negotiations.

(6) Simplified strategy curves can be developed that give nearly optimal results.

Acknowledgements

The ideas in this paper are the result of discussions with scientists at the International Institute for Applied Systems Analysis, Vienna, especially Jim Bigelow, John Casti and Sandra Buckingham. Special thanks to Peter Larkin and F. E. A. Wood for suggesting the direction to look.

References

- [1] Allen, K. R. 1973. The influence of random fluctuations in the stock-recruitment relationship on the economic return from salmon fisheries. Conseil Internat. Pour L'Exploration De La Mer, Rapport#164, pp 351-359.
- [2] Bellman, R. 1961. Adaptive Control Processes: A guided tour. Princeton Univ. Press, Princeton, N. J.
- [3] Bellman, R., and Dreyfus, S. 1962. Applied dynamic programming. Princeton Univ. Press, Princeton, N. J.
- [4] Bellman, R., and Kalaba, R. 1965. Dynamic Programming and modern control theory. Academic Press, New York.
- [5] Raiffa, H. 1968. Decision analysis: introductory lectures on choices under uncertainty. Addison-Wesley, Reading, Massachusetts.
- [6] Ricker, W. E. 1964. Stock and recruitment. J. Fish. Res. Bd. Canada, 11:559-623.
- [7] Ricker, W. E. 1973. Two mechanisms that make it impossible to maintain peak-period yields from stocks of pacific salmon and other fishes. J. Fish. Res. Bd. Canada, 30:1275-1286.
- [8] Shepard, M. P., Withler, F. C., McDonald, J., and Aro, K. V. 1964. Further information on spawing stock size and resultant production for Skeena Sockeye. J. Fish. Res. Bd. Canada, 21:1329-31.
- [9] Tautz, A., Larkin, P. A., and Ricker, W. E. 1969. Some effects of simulated long-term environmental fluctuations on maximum sustained yield. J. Fish. Res. Bd. Canada, 26:2715;26.



Figure 1. Stock-recruitment relationship for the Skeena Sockeye. From Shepard, et al. (1964), with recent points from unpublished data provided by F. E. A. Wood, Environment Canada.

,







Figure 3. Observed distribution of the Ricker production parameter α (equation 1) for Skeena River Sockeye, and three judgmental probability distributions for possible future values.





Figure 5. Predicted probability distributions of catches using the "natural" optimal strategies of Figure 4.



Figure 6. Predicted probability distribution of stock sizes associated with the catches of Figure 5.



Figure 7. Optimal strategies for different objectives compared to actual management practice on the Skeena River, to the optimal fixed ascapement policy, and to a simplified alternative.



mean and variance of catches, for the Skeena River (explanation in text). Pareto Frontier for best possible combinations of Figure 8.



Figure 9. Predicted probability distributions of catches using the simplified strategy curve in Figure 7 as opposed to the best fixed escapement strategy. Recent actual catches are shown for comparison.