CHARACTERIZATION OF THE OPTIMUM DATA ACQUISITION AND MANAGEMENT OF A SALMON FISHERY AS A STOCHASTIC DYNAMIC PROGRAM[1]

GARY E. LORD[2]

ABSTRACT

The optimum data acquisition and management of a typical Bristol Bay sockeye salmon fishery have been expressed as a problem in statistical decision theory. Optimality has been defined as that set of sequential decision rules that minimizes the Bayes risk over the duration of the run. Economic losses or costs are ascribed to acquisition of catch and escapement data in such a manner that an optimal data acquisition scheme can be defined in addition to defining the set of optimal management strategies.

The management and inshore harvesting of salmon stocks characteristically consist of several interrelated phases. Stock assessment and run profiling are of importance both to the management personnel and to segments of the industry in planning their respective operations. Prior knowledge of the run size and time profile is useful in the preliminary planning of a management strategy that will in some way permit the escapement of the desired number of spawners. Similarly, such information also serves the industry in planning the level of its anticipated activities (Mathews, 1967). The theoretical investigation described here was motivated by specific consideration of the data gathering and management schemes currently applied in the Bristol Bay sockeye salmon fisheries. However, the formulation is relatively abstract and of sufficient generality so that, properly interpreted, it may apply to a variety of fishery situations which are evolutionary or time varying in nature. Indeed, it is this dynamic aspect of the problem that is at once the crucial feature of the analysis and also the principal source of analytical and computational difficulty.

Rothschild and Balsiger (1971) treated the optimum management of the Kvichak fishery of Bristol Bay as a problem in linear programming in which the various entities comprising the run could be optimally allocated, subject to various constraints on escapement, sex ratios, etc., over the duration of the run. Optimality was chosen as that set of allocation rules which maximized the economic return, expressed as a linear objective function, subject to the satisfaction of a set of linear inequality constraints. This formulation and solution as a linear program are particularly powerful since the solution, easily obtained by standard techniques, is particularly rich in interpretive detail.

The formulation and solution as a linear program suffer from several disadvantages, as Rothschild and Balsiger noted in their paper. First, the solution is deterministic in that it assumes precise knowledge of the run size and its time profile. In actual practice, although there is considerable investment in stock assessment and run forecasting, the resulting estimates are subject to considerable variation. Second, not all of the constraints are "firm," i.e., inviolable. This applies especially to the escapement. Presumably a unique optimum escapement for each actual realization of the run exists, but it is not necessarily imperative that this escapement be attained. Instead it may be assumed that an escapement below the optimum is accompanied by some economic loss, suitably discounted, for the diminished future returns. Similarly, an excessive escapement will result in a loss due to the decreased catch and, if in the right-hand tail of a dome-shaped (Ricker

[1] Contribution No. 385, College of Fisheries, University of Washington, Seattle, Wash.

[2] Fisheries Research Institute, University of Washington, Seattle, WA 98195.

Manuscript accepted April 1973.

type) spawner-return curve, there will be an additional loss due to decreased returns. Finally, the linear programming formulation is essentially static although Hillier and Lieberman (1967) discuss certain techniques that are applicable to a limited class of stochastic and dynamic allocation problems. However, the only such technique that would be readily applicable to the fishery management case is chance constrained programming in which the admissibility of "soft" constraints, i.e., constraints which may be violated with certain allowable probabilities, is permitted.

In the subsequent analysis an attempt is made to simulate directly the inherently stochastic and dynamic nature of the management of a typical salmon fishery. A brief discussion of the assumptions made as well as a comparison with the linear programming formulation of Rothschild and Balsiger will also be given. This comparison of methods should be considered to be somewhat subjective and reflects, to a certain extent, the author's own opinions and predilections. The interested reader may be able to arrive at more meaningful comparisons and conclusions after studying the respective analyses in more detail.

We will assume here gear of fixed selectivity with regard to sexes and year classes. This corresponds to the status quo with respect to Bristol Bay although Rothschild and Balsiger showed that an optimum allocation among the various entities comprising the run was economically advantageous, particularly in the case of altered sex ratios. The fixed selectivity assumption, which results in an allocation based only on total numbers of fish, is made principally in the interests of tractability although it is possible to generalize the loss functions and probability densities to include the various individual entities.

Hillier and Lieberman list and discuss the basic features which serve to characterize dynamic programming problems. The principal characteristics will be repeated here, paraphrased slightly, and it will be shown here and in the subsequent analysis that the salmon fishery management problem conforms quite naturally to the class of problems for which dynamic programming is applicable.

1) The problem can be divided into stages with a policy decision required at each stage. This is obviously the case in Bristol Bay where the stages consist of discrete fishing periods, for each of which a management decision must be made. Discreteness is not an essential feature, however, since continuous time allocation problems may also be treated by dynamic programming techniques.
2) Each stage has associated with it a (possibly infinite) number of states. The state of the system is somewhat difficult to characterize precisely. It will be sufficient to treat the state of the system, in this case the salmon fishery, at the start of any stage as reflecting the true state of nature, e.g., run size, time profile, migration patterns, etc., as well as the effects of all previous policy decisions through the preceding time period. Closely related is:
3) The effect of the policy decision at each stage is to transform the current state into a state associated with the next stage. We will generalize this slightly to include sequentially acquired data as an additional quantity serving to characterize the state of the system and the transformation from one state to the next. The remainder of the characterizing features enumerated by Hillier and Lieberman are related to the very fundamental "Principle of Optimality," the statement and discussion of which will will be deferred until the section on Discussion.

The close correspondence of the concepts of dynamic programming, and also the closely related sequential statistical decision theory, to the problem of salmon fishery management suggests that together they provide potentially powerful tools for the description and simulation of such processes. A caveat is appropriate here, however. In general there will be the loss of a considerable portion of the economic "fine structure" of the problem, particularly in comparison with the solution as a linear program. The solution of the linear programming primal and its dual allows one to make inferences beyond simply the attainment of the optimum. This was emphasized adequately by Rothschild and Balsiger in their identification and discussion of the various shadow prices, etc. A more serious reservation concerns the very serious analytical and computational difficulties to be anticipated. It is a truism of dynamic programming that many more problems may be formulated than may be solved, and it is not at all certain at present whether the salmon fisheries problem falls into the soluble category. Thus, the present discussion will be confined to the presentation of the theory, which is self-contained. The very difficult problems of formulation of the loss functions and the selection of optimum decision rules will be the subject of subsequent investigations.

THEORY

It is helpful to think of the salmon run, its assessment, and its management as evolutionary processes in time. Prior to the start of fishing the management biologist has at his disposal certain prior information, such as preseason forecasts, on which to base his early management strategy. As the run proceeds, additional data are gathered so that, as his knowledge of the true state of nature increases, he may modify his strategy to conform more closely to the optimum strategy. This will now be developed more formally.

Assume that the entire run occurs over a total of m discrete nonoverlapping time intervals. If n_i is the number of fish entering the fishery on the ith day then the total run size, N_{tot}, is given by

$$N_{\text{tot}} = \sum_{i=1}^{m} n_i. \tag{1}$$

Define a parameter vector $\underset{\sim}{\theta}$, of arbitrary dimension, that is assumed to characterize all relevant details of the run. As a specific example, we could define $\underset{\sim}{\theta}$ of the m-dimensional vector $(n_1, n_2, \ldots n_m)$, i.e., the ith component of $\underset{\sim}{\theta}$ is the number of fish entering the fishery on the ith day. More generally, we can leave $\underset{\sim}{\theta}$ arbitrary and write $n_i = n_i(\underset{\sim}{\theta})$. For each known $\underset{\sim}{\theta}$ there exists some known set of optimum allocation rules $\eta_i(\underset{\sim}{\theta})(i = 1, \ldots m)$ where η_i is the optimum fraction of the fish to be allocated to the catch on the ith day. For example, the linear programming formulation of Rothschild and Balsiger provided a set of optimum allocation rules based on a fixed total run size and a two parameter time profile proposed by Royce (1965).

Let D be a finite set of decision rules and let δ_i, a member of D, be the decision adopted on the ith day. Typically the set D consists of such management decisions as fishery opening or closing, fishing area limitations, etc. For each δ_i there will be an actual allocation $\hat{\eta}_i$ where, in general, both η_i and $\hat{\eta}_i$ will be random variables. The former will depend on the true (unknown) state of nature, $\underset{\sim}{\theta}$, while the latter will also be a random function of $\underset{\sim}{\theta}$ as well as of the decision taken, δ_i. As the actual and optimum allocations differ, various economic losses will be assumed to accrue, and the average or expected loss will be these losses averaged over all possible outcomes. This will be developed more formally after considering the various loss functions.

We may postulate the existence of an overall loss function that reflects economic losses from all sources. For our purposes we will consider the loss as arising only from 1) the cost of data acquisition, 2) the catch, and 3) the escapement. Since the catch and the escapement are complementary quantities, their sum comprising the total run, we could consider either one individually as is done in a subsequent example using the Ricker (1958) spawner-return relation. However, an individual treatment of each permits the separate discussion of loss functions that are linear, additive, and separable, as for the catch, and those which are nonlinear and not additive, as will be postulated for the escapement.

First consider the cost of data acquisition. Generally, the sampling schemes to be used and the level of effort for each are selected prior to the run. There may be in-season variations, such as occasional stream surveys, etc. but the cost associated with these is much less than that allocated prior to the run. We denote the experimental design symbolically by ζ and, in line with the above argument, there is associated with ζ some fixed cost $C(\zeta)$. The quanitity ζ, which is an abstract designation of the experimental design, will also appear as a conditioning quantity when we consider the various probability density functions associated with the sampled quantities, i.e., the distribution of the sample estimates will depend on, among other things, the manner in which the data are acquired.

Next is the loss associated with the catch. Consider the start of the kth time period where $1 \leq k \leq m$ but otherwise arbitrary. If the economic value of the catch is assumed to be linear and additive, we may write

$$L_X(k) = \sum_{i=1}^{k-1} \mathrm{v}_i \, [\eta_i(\underset{\sim}{\theta}) - \hat{\eta}_i] \, n_i(\underset{\sim}{\theta}) \qquad (2)$$

as the loss through the (k - 1)st time period for the catch.[3] Here v_i is the unit value of the fish for the ith interval. Note that if $\hat{\eta}_i$ exceeds η_i, i.e., the actual catch exceeds the optimum catch, the loss function is negative and becomes a gain function. However, this apparent benefit must be offset by the loss associated with the corresponding decreased escapement. If this were not the case, then this would not be the correct optimum since any departure from the optimum must result in a nonnegative incremental loss.

Finally consider the loss function for the escapement. This function cannot be considered to be linear or additive since the average number of returns per unit of spawners escaping on any particular day will be a function primarily of the final value of the total escapement. This is a consequence of the fact that late spawners may interfere with the redds of the earlier arrivals and thus diminish the returns for this group. Also, excessive escapement may lead to increased competition for food among the fry to the general detriment of the population as a whole. Thus, to a first approximation, the loss function for the escapement depends only on the final values of the actual and optimum escapements. Symbolically this may be written as

$$L_E(m) = \phi\,(E, \hat{E}) \qquad (3)$$

where

$$E = \sum_{i=1}^{m} (1 - \eta_i)\, n_i(\underset{\sim}{\theta}) \qquad (4a)$$

is the optimum total escapement while

$$\hat{E} = \sum_{i=1}^{m} (1 - \hat{\eta}_i)\, n_i(\underset{\sim}{\theta}) \qquad (4b)$$

is the actual escapement and $\phi\,(\cdot,\cdot)$ denotes some suitable functional form. An even more general formulation is possible if the possibility is admitted that the magnitude and timing of the arrivals on the spawning beds are also significant. The loss function must still be expressed in terms of the entire run but the functional form would be of the type

$$L_E(m) = \phi'\,(E_1, E_2, \ldots E_m; \hat{E}_1, \hat{E}_2, \ldots \hat{E}_m) \qquad (3')$$

where E_i and $\hat{E}_i$ are, respectively, the optimum and actual escapements for the ith time period. However, the determination of the optimum total level of escapement, which is necessary to characterize $\phi\,(E, \hat{E})$, is a subject of current research and is by no means resolved at present. Thus, the characterization of a function of the generality of $\phi'(E_1, E_2, \ldots E_m;\ \hat{E}_1, E_2, \ldots \hat{E}_m)$ must await further biological data.

The Bayes risk is defined as the average or expected loss where the averaging is over all possible outcomes and an optimum strategy will be defined as that strategy that minimizes the Bayes risk.[4] An expression for the Bayes risk will now be constructed that is appropriate for the salmon management problem just out-

[3] If the capacity of the cannery becomes limiting, as may happen in cycle years, the loss function for the catch may be written in slightly more general form as

$$L_X(K) = \sum_{i=1}^{k-1} \mathrm{v}_i \left\{ \eta_i n_i(\theta) - \min\,|\,\hat{\eta}_i n_i(\theta), \mathrm{cap}(i)\,| \right\} \qquad (2')$$

where cap(i) denotes the cannery capacity for the ith time period. This relation implies that any actual catch that exceeds the cannery capacity will not decrease the corresponding loss. This function is no longer linear but it is still additive. Note also that no cost has been ascribed to the additional economically nonproductive fishing effort. Cannery capacity was one of the constraints imposed by Rothschild and Balsiger in their paper.

[4] The specification of the minimum Bayes risk as the criterion of optimality, while a reasonable one, is somewhat arbitrary. Other criteria are in common use, most notably the "Minimax," in which the optimum strategy is that which minimizes the maximum risk.

lined. Consider the start of the kth time period where $1 \leq k \leq m$ but otherwise arbitrary. The management biologist has at his disposal data observed through the $(k - 1)$st period which will be designated by $(\underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1})$. This is an abstract designation for data which may typically be in the form of catch reports, catch per unit effort data, tower counts, etc. By convention $\underset{\sim}{Y}_0$ represents the preseason information, such as the high-seas forecast, that is obtained prior to the start of the run. However, all future outcomes must be considered since the loss function for the escapement is formulated in terms of the final state of the system. The vector $\underset{\sim}{\theta}$, which characterizes both the run and the corresponding set of optimum allocation rules, is, from the biologists' point of view, an unknown parameter whose value he is attempting to infer. Generally $\underset{\sim}{\theta}$ may be considered to have some underlying prior distribution which may be inferred from historical data, etc. As additional data are gathered, the probability density of $\underset{\sim}{\theta}$ may be successively updated to reflect this additional information. Thus the probability density for $\underset{\sim}{\theta}$ at the beginning of the kth period may be written as

$$f_k\,(\underset{\sim}{\theta} \mid \underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1}; \zeta)$$

where the prior data, $(\underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1})$, and the manner in which it is obtained, ζ, appear as conditioning quantities.

Now consider the distribution of the actual allocation $\hat{\eta}_i$. Generally $\hat{\eta}_i$ will be a random variable whose distribution will depend on the action taken, δ_i, and on the true state of nature, $\underset{\sim}{\theta}$. Thus the probability density of $\hat{\eta}_i$ may be written in conditional form as $g_i(\hat{\eta}_i|\delta_i, \underset{\sim}{\theta})$. This tacitly assumes that the allocation resulting from a decision taken during any particular time period is independent of the outcomes during any other time period which in turn implies that the individual fish is vulnerable during only a single time period. This condition is generally fairly well satisfied in most of Bristol Bay where the fishing districts are relatively small and the fish do not delay in their upstream migrations. Exceptions occur occasionally during extreme tides when the fish may enter and leave the fishery more than once before proceeding upstream. A similar exception would occur if a fishing district were of sufficient size that individual fish must necessarily spend more than a single time period in it. In these more general cases we must include all prior allocations as conditioning quanitities, i.e., the appropriate density would be of the form $g_i(\hat{\eta}_i|\hat{\eta}_1, \hat{\eta}_2, \ldots \hat{\eta}_{i-1}, \delta_i, \underset{\sim}{\theta})$. An equivalent but more concise notation would be to condition the distribution of $\hat{\eta}_i$ by δ_i and by the state of the system at the start of the ith time period, S_i, i.e., $g_i\,(\hat{\eta}_i|\delta_i, S_i)$. That this is an equivalent conditioning follows from our previous definition of the state of the system as reflecting the true state of nature as well as the effects of all previous policy decisions. If we retain the assumption of independence of the allocations, the joint probability density of $(\hat{\eta}_i, \hat{\eta}_2, \ldots \hat{\eta}_m)$ may be written in factored form as

$$g\,(\hat{\eta}_i, \hat{\eta}_2, \ldots \hat{\eta}_m | \delta_1, \delta_2, \ldots \delta_m; \underset{\sim}{\theta}) = \prod_{i=1}^{m} g_i\,(\hat{\eta}_i | \delta_i, \underset{\sim}{\theta}). \tag{5}$$

It is now possible to construct the risk functions appropriate for the start of the kth time period where, as usual, k is arbitrary. The experimental cost, $C(\zeta)$, has been assumed fixed in advance in which case it is equal to its expected value. Thus the first term of the risk, $R_{1,k}(\zeta)$, is given simply by

$$R_{1,k}\,(\zeta) = C(\zeta) \tag{6}$$

for all k.

The risk function associated with the catch may be thought of as consisting of two parts. The first part is the risk corresponding to the loss already accrued through the kth time period for which the management decisions have already been made. The second is the risk over the remainder of the run for which the decisions $\delta_{k+1}, \delta_{k+2}, \ldots \delta_m$ remain to be made. From (2) and (5) we obtain

$$R_{2,k}(\delta_1, \delta_2, \ldots \delta_k | \underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1}; \zeta) =$$
$$\sum_{i=1}^{k} v_i \int d\underset{\sim}{\theta}\, f_k(\underset{\sim}{\theta} | \underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1}; \zeta)$$
$$\int d\hat{\eta}_i\, (\eta_i\,(\underset{\sim}{\theta}) - \hat{\eta}_i)\, n_i\,\,(\underset{\sim}{\theta})\, g_i\,(\hat{\eta}_i \mid \delta_i, \underset{\sim}{\theta}) \tag{7}$$

for the risk through the kth time period. Although we are considering the start of the kth time period, $\underset{\sim}{Y}_{k-1}$ will have been acquired so that δ_k may be chosen on the basis of this and all previous observations. Similarly, the risk over the remainder of the run is given by

$$R_{3,k}\,(\delta_{k+1}, \delta_{k+2}, \ldots \delta_m \,|\, \underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1}; \zeta) =$$

$$\sum_{i=k}^{m} \, {}_1 \; v_i \int d\underset{\sim}{\theta}\, f_k\,(\underset{\sim}{\theta} \,|\, \underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1}; \zeta)$$

$$\int d\hat{\eta}_i\,(\eta_i\,(\underset{\sim}{\theta}) - \hat{\eta}_i)\, n_i\,(\underset{\sim}{\theta})\, g_i\,(\hat{\eta}_i | \delta_i; \underset{\sim}{\theta}) \tag{8}$$

where, in an implicit sense, the decisions δ_{k+1}, $\delta_{k+2}, \ldots \delta_m$ will be conditional upon the prior decisions $\delta_1, \delta_2, \ldots \delta_k$ as well as upon the observations $\underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1}$. Equations (7) and (8) have a similar structure although in Equation (7) we are weighting past decisions by our present knowledge of the state of nature while in Equation (8) the future decisions must necessarily reflect only the information obtained through the (k-1)st time period.

The risk for the escapement is assumed to be nonseparable so that the entire run must be considered at once. The general expression for this portion of the risk is then given by

$$R_{4,k}\,(\delta_1, \delta_2, \ldots \delta_m \,|\, \underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1}; \zeta) =$$

$$\int d\underset{\sim}{\theta}\, f_k\,(\underset{\sim}{\theta} \,|\, \underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1}; \zeta) \int d\hat{\eta}_1 \cdots$$

$$\int d\hat{\eta}_m \; \phi\,[E(\underset{\sim}{\theta}), \hat{E}\,(\hat{\eta}_1, \hat{\eta}_2, \ldots \hat{\eta}_m)] \prod_{i=1}^{m} g_i\,(\hat{\eta}_i | \delta_i; \underset{\sim}{\theta}). \tag{9}$$

The total Bayes risk is then the sum of the risks given by (6), (7), (8), and (9), i.e.,

$$R_k = \sum_{i=1}^{4} R_{i,k}. \tag{10}$$

These equations have been derived under very general conditions and assumptions with little effort towards characterizing any of the functions indicated. It is interesting to note, however, that loss functions of the same general form as those appearing in the integrands of (7), (8), and (9) may be obtained if we assume 1) a steady state spawner-return relation of the type proposed by Ricker (1958) and 2) the economic loss (or gain) is proportional to the catch. The steady state Ricker spawner-return relation is

$$\hat{N}_{tot} = \theta_1 \hat{E} e^{-\theta_2 \hat{E}} \tag{11}$$

where the parameter vector $\underset{\sim}{\theta} = (\theta_1, \theta_2)$ describes the run and $\hat{N}_{tot}$ and $\hat{E}$ denote the actual run size and the escapement respectively. The corresponding catch is given by

$$\hat{X} = \theta_1 \hat{E} e^{-\theta_2 \hat{E}} - \hat{E} \tag{12}$$

from which it follows that the optimum escapement, E, is given by the solution of

$$\frac{d\hat{X}}{d\hat{E}} = 0$$

or

$$\theta_1\,(1 - \theta_2 E)\, e^{-\theta_2 E} = 1. \tag{13}$$

The appropriate loss function in terms of the catch is

$$L\,(X, \hat{X}) = \bar{v}\,(X - \hat{X}) \tag{14}$$

where $\bar{v}$ is the average unit value of the fish. The equivalent expression

$$L'\,(E, \hat{E}) = \bar{v}\,[\theta_1\,(Ee^{-\theta_2 E} - \hat{E}e^{-\theta_2 \hat{E}}) - (E - \hat{E})] \tag{15}$$

in terms of the escapement is easily obtained by substituting (12) in (14). The term $\bar{v}\theta_1$ $[\exp(-\theta_2 E) - \exp(-\theta_2 \hat{E})]$ is nonlinear while $\bar{v}\,(E - \hat{E})$ is a linear and additive function of the daily escapements. The substitution of (15) into (9) would then result in an integral of the same general form as (9) and two additional integrals corresponding to (7) and (8).

As indicated earlier, the experimental design,

ζ, is usually fixed prior to the start of fishing and, except for minor variations, remains essentially unchanged during the run. We define the optimum experimental design, ζ^*, by

$$\min_{\zeta} R_0 (\bar{\delta}_1, \bar{\delta}_2, \ldots \bar{\delta}_m ; \zeta) = R_0 (\bar{\delta}_1, \bar{\delta}_2, \ldots \bar{\delta}_m ; \zeta^*) \tag{16}$$

where the overbars on the $\{\delta_i\}$ denote averaging overall allowable decision rules δ in the set D and the risks are determined prior to the taking of any observations. Similarly, the optimum decision rules $(\delta_1^*, \delta_2^*, \ldots \delta_m^*)$ are that set of decisions that minimizes the average risk over the duration of the run.

DISCUSSION

A mathematical description of a salmon fishery that includes both stochastic and dynamic elements has been formulated although the final result is relatively general and somewhat abstract. Indeed, the mathematics was formulated specifically to simulate the actual assessment and management of the typical Bristol Bay sockeye salmon fishery. The statistician or management biologist periodically acquires additional data, such as catch reports and test fishing results, from which he can make repeatedly more refined estimates of the true state of nature. Also, although perhaps quite unconsciously, he attempts to estimate the losses (or, if an optimist, the gains) associated with any course of action and the relative probability of occurrence of these losses. Then, based on all data obtained to date, including all past decisions and outcomes, he attempts to formulate a future strategy that will minimize his risk. The analysis of the preceding section attempted to express this sequence of events in a more formalized and quantitative manner.

The apparent fidelity of statistical decision theory to the real world suggests that it provides a very general theoretical tool for the description of such processes. However, the implementation of such a theory may give rise to some practical problems of considerable difficulty, some of which were discussed in the Introduction. In particular it was emphasized there that the ability to formulate a problem as a dynamic program or, almost equivalently, as a problem in sequential statistical decision theory, by no means assures that a solution may be obtained. In this section some additional general features of dynamic programming, as they apply to the fisheries management problem just formulated, will be further elaborated.

The set of optimum decision rules has been defined as that set that minimizes the Bayes risk over the duration of the run. From this it follows that the kth decision must be chosen optimally as a function of the set of prior observations $(\underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1})$ and as a function of all prior decisions $(\delta_1, \delta_2, \ldots \delta_{k-1})$. In other words, δ must be chosen at each stage in an optimal manner taking into account all prior observations and decisions. At this point we continue Hillier and Lieberman's (1967) characterization of dynamic programming, the first three principles of which were presented in the Introduction. Their principle number four states that: Given the current state of the system, an optimal policy for the remaining stages is independent of the policies adopted in the previous stages. This is a paraphrase of the fundamental "Principle of Optimality" of Bellman (1957, p. 83) which states that: "An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision." The principle of optimality thus assures that the policy we have specified is indeed an optimal policy.

The contradiction, which is more apparent than real, will now be resolved between the principle of optimality just stated and our previous contention that the choice of an optimal decision δ_k depends not only on $(\underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1})$ but also on the previous decisions $(\delta_1, \delta_2, \ldots \delta_{k-1})$. Recall that the state of the system at time k, say, S_k, is assumed to be uniquely determined by $(\underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1})$, $(\delta_1, \delta_2, \ldots \delta_{k-1})$, and $\underset{\sim}{\theta}$, the true state of nature. The converse is not true, however, since a multiplicity of different decisions and observations may lead to the same S_k, i.e., there is generally no unique path to S_k. Thus, while it is perhaps more appropriate to state that the optimum δ_k is a function only of S_k, it should be borne in mind that S_k has been uniquely determined by past decisions and observations so that, in an implicit sense, the $\{\delta_i\}(i = 1, \ldots m)$ are not generally mutually independent.

Consider now a typical Bristol Bay salmon fishery. The usual allowable decision rules consist of either opening or closing the fishery. In addition the management biologists also have the option of allowing fishing over an increased or decreased area depending on whether the run is larger or smaller than normal. Thus, in the most general case, a total of four distinct strategies is available although it is unlikely that both increased and decreased fishing areas would be used during a single season.[5] In the usual case, then, a total of three distinct strategies is available each day from which it follows that a total of 3^m separate courses of action may be pursued during a fishing season m days long. Typically m is equal to about 20 days in Bristol Bay so that the total number of allowable strategies is of the order of 10^9. This is not a number to be taken lightly and is an example of what Bellman (1957, p. 6) refers to as "The Curse of Dimensionality."

The principle of optimality, which is particularly useful in multistage allocation processes, may be invoked in an effort to reduce this problem in many dimensions to a sequence of problems in one dimension. Assume that we are at the beginning of the mth time period where the state of the system is characterized by S_m where S_m reflects the observations $(\underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{m-1})$ as well as past decisions $(\delta_1, \delta_2, \ldots \delta_{m-1})$. Thus the only decision at our disposal is δ_m and presumably an optimal δ_m may be chosen as a function of S_m. Consider next the beginning of the $(m-1)$st time period for which the system is characterized by S_{m-1}. For every δ_{m-1} selected and executed the system is transformed, after the acquisition of the data $\underset{\sim}{Y}_{m-1}$, into the state S_m for which an optimal decision has already been obtained. Thus at this stage we need optimize only with respect to δ_{m-1}. In this manner we can proceed backward to the beginning of the first time period, characterized by an S_1 depending on Y_0 at which point an optimal δ_1 is selected.

This backward recursive scheme is typical of the method of attack on dynamic programming problems. For a concise but elegant example of this technique applied to a simulated, but numerical, fishery problem see Rothschild (1970). The particular example he used had discrete stages in time with a finite number of strategies available for each stage. The desired solution described the optimum sequence, or path, in time of visiting various fisheries, for a fixed total number of time periods, so that the total catch was maximized. While highly idealized, this problem constitutes a true dynamic program. However, it lacks the stochastic features that are an essential feature of the present discussion.

Verbally this method of solution appears to be most attractive since we have apparently overcome the problem of excessive dimensionality by the recursive consideration of a sequence of problems of lower dimension. This feature is emphasized in the previously cited example presented by Rothschild. However, in problems of larger scale, either in terms of the number of stages or the number of possible states per stage, a rather more subtle problem of dimensionality appears. This concerns the successive specifications of the states of the system $\{S_k\}(k = 1, \ldots m)$. We recall that S_k is characterized not only by $\underset{\sim}{\theta}$, the true state of nature, but all prior observations $(\underset{\sim}{Y}_0, \underset{\sim}{Y}_1, \ldots \underset{\sim}{Y}_{k-1})$ and all prior decisions $(\delta_1, \delta_2, \ldots \delta_{k-1})$. The observation vectors are necessarily multidimensional, each component of which represents a particular piece of data or the observation of a particular entity. To make matters worse, the dimension of the $\underset{\sim}{Y}_i$ will generally increase with i since new forms of data, such as catch reports, tower counts, catch per unit of effort, etc., will become successively available. Thus, the $\{\underset{\sim}{Y}_i\}(i = 0, \ldots k - 1)$ required to specify S_k will bring with them their own dimensionality which will soon become overwhelming unless S_k can be described adequately by relatively few parameters. Obviously, aside from possible computer core limitations, the most useful algorithms are those that can accommodate the requisite dimensionality. In a more practical

[5] It is also possible to impose waiting periods for the entry of gear into selected fisheries but this will not be considered here.

sense, however, cognizance must be taken of dimensionality problems to avoid foundering in details, many of which may be irrelevant.

Possible alternative procedures to decrease the number of strategies to consider would be to: 1) combine various contiguous time periods into longer units and thus effectively decrease m and 2) reject those strategies that, based on the value of Y_0 available, have a sufficiently high a priori probability of being nonoptimal. The combining of time periods would be particularly appropriate near the beginning and end of the run when relatively few fish are entering the fishery and the consequences of nonoptional procedures are not so serious. Also, the necessity to consider increased or decreased fishing areas can often be resolved prior to the start of fishing by consideration of the preseason information contained in $\underset{\sim}{Y}_0$.

LITERATURE CITED

BELLMAN, R.
1957. Dynamic programming. Princeton Univ. Press, 340 p.

HILLIER, F. S., AND G. J. LIEBERMAN.
1967. Introduction to operations research. Holden-Day, San Franc., 639 p.

MATHEWS, S. B.
1967. The economic consequences of forecasting sockeye salmon (*Oncorhynchus nerka* Walbaum) runs to Bristol Bay, Alaska: A computer simulation study of the potential benefits to a salmon canning industry from accurate forecasts of the runs. Ph.D. Thesis, Univ. Washington, Seattle, 238 p.

RICKER, W. E.
1958. Handbook of computations for biological statistics of fish populations. Fish. Res. Board Can., Bull. 119, 300 p.

ROTHSCHILD, B. J.
1970. A systems view of fishery management with some notes on the tuna fisheries. Quantitative Science Paper No. 14. Center for Quantitative Science in Forestry, Fisheries and Wildlife and Fisheries Research Institute, Univ. Washington, Seattle, 78 p.

ROTHSCHILD, B. J., AND J. W. BALSIGER.
1971. A linear-programming solution to salmon management. Fish. Bull, U.S. 69:117-140.

ROYCE, W. F.
1965. Almanac of Bristol Bay sockeye salmon. Univ. Wash. Fish. Res. Inst., Circ. 235, 48 p.