

AN APPROACH TO MANAGING FISHERIES WHEN WEAK AND STRONG STOCKS MIX

David Tomberlin, NOAA Fisheries, david.tomberlin@noaa.gov

ABSTRACT

When weak and strong fish stocks are caught in the same fishery, managing for the protection of the weak stock may result in foregone economic benefits from harvest of the strong stock, while managing for the strong stock may result in overfishing of the weak stock. A particular complication arises when the stocks are not easily distinguished in the catch, for example when the stocks are of the same or a similar species. This paper demonstrates how the partially observable Markov decision process can be used to explore policy options for managing mixed-stock fisheries when the stocks are imperfectly identified.

Keywords: mixed-stock fisheries, information, partially observable Markov decision process

INTRODUCTION

Distinct stocks of fish are often caught together in assemblages, which poses particular management challenges addressed by the growing literature on multi-species fisheries. This paper considers a particular aspect of the general problem of mixed-stock or mixed-species fisheries, namely the management of the fishery when the degree of mixing among stocks or species is difficult to assess. The purpose of the analysis here is to demonstrate the application of the partially observable Markov decision process (POMDP) to explore the trade-offs between less fishing, which provides better protection of weak stocks, and more fishing, which generates greater revenues and also—importantly—more information.

The analysis here is a response to the situation in the salmon troll fishery along the coasts of Oregon and California in the United States. There, the ocean salmon troll fishery targets primarily chinook salmon, predominantly Central Valley fall run chinook or Klamath River fall run chinook, which mingle in the open ocean. In recent years, the fishery has relied heavily on the Central Valley stock, while the Klamath stock has been generally weak. Prosecution of the fishery has been constrained by the need to protect the Klamath stock. A particular challenge in this type of ‘weak stock management’ is determining how much the weak stock is actually threatened by the fishery targeting the strong stock, and what management actions are likely to strike a successful balance between protecting the weak stock and allowing harvest of the stock stock.

Routledge (2001) analyzes the case in which effort can be directed away from a mixed-stock fishery to a single-stock terminal fishery, in particular demonstrating the disincentives for protecting the weak stock. Here, we consider a somewhat different case, in which a strong and a weak stock may or may not be mixed at a particular time, and the question is whether to allow fishing or to temporarily close the fishery in order to protect the weak stock. The essence of the approach taken here is to identify the optimal policy as a function of managers’ confidence in the state (i.e., the degree of mixing among stocks) of the fishery, given costly and imperfect observations on this status. Available scientific methods can determine the stock of a particular fish with a good deal of confidence. Barnett-Johnson et al. (2006), for example, report being able to distinguish wild from hatchery salmon with 90% accuracy, and natal river of wild salmon with 95% accuracy, based on strontium isotope ratios in otoliths. Genetic Stock Identification (GSI) is another means of determining the mixing of stocks that has been repeatedly applied to Pacific salmon fisheries (Winans et al. 2001). However, there is substantial cost in obtaining estimates of stock mixing, and better estimates will generally entail greater costs.

The POMDP approach to mixed-stock fisheries management may be thought of as a particular formalization of adaptive management, in that decisions are made in a dynamic and stochastic environment based on beliefs that change as new information becomes available. Below, the approach is demonstrated in a stylized example of a fishery in which managers may choose at any time to switch between a fully open fishery, a limited ‘test’ fishery intended to provide some revenue to the fleet along with some information (based on catch composition) to fisheries managers, and a fully closed fishery. The result of the analysis is an optimal policy that maps managers’ beliefs about the status of the mixed-stock fishery into actions that maximize the expected value of the stream of fisheries revenues over time, taking into account the cost of stock information and its evolution over the decision horizon.

MODEL

The POMDP is a collection of sets $\{S, P, A, W, \Theta, R\}$ (Cassandra 1994), where S is the system’s state variables, P represents state dynamics as transition probabilities, A is the actions available to an agent, W is the rewards to taking particular actions in particular states, Θ is a set of possible observations on the state variables, and R is a set of observation probabilities. Observations $\theta \in \Theta$ are the only information the agent has on the unobservable true state, S . The observation model R describes the probabilistic relationship between observations θ and the true state S . The problem solver (here, the fishery manager) uses observations θ and the observation model R to estimate the state S .

The model assumes the manager’s goal is to maximize long-run discounted total value of the fishery, which includes the cost of monitoring stock mixing. The actions that achieve this goal are identified with dynamic programming (Bertsekas 2000) through a recursively defined value function V :

$$V_t(\pi) = \max_a \left[\sum_i \pi_i q_i^a + \beta \sum_{i,j,\theta} \pi_i p_{ij}^a r_{j\theta}^a V_{t+1}[T(\pi | a, \theta)] \right]$$

where

π_i = subjective probability of being in state $i \in S$ at time t

q_i^a = immediate reward for taking action $a \in A$ in state $i \in S$ at time t

β = discount factor

p_{ij}^a = probability of moving from state $i \in S$ at time t to state $j \in S$ at time $t+1$
after taking action $a \in A$

$r_{j\theta}^a$ = probability of observing $\theta \in \Theta$

after taking action $a \in A$ and moving to state $j \in S$

T = function updating beliefs based on prior beliefs and observed θ

V is the greatest expected net benefit that the agent can achieve over time, taking into account that as conditions change in the future, different actions may be warranted. The solution of V yields an optimal policy, which is a mapping from beliefs about the current state, $\pi(S)$, into the optimal action.

In our setting, the state variable S is the degree of mixing between two fish stocks, which for expository purposes takes only two possible values, *Pure Strong* if only the strong stock is present and otherwise *Mixed*. The action set A consists of *Normal* fishing (i.e., allowing the fishery to proceed along customary lines), *Test* fishing (i.e., allow a limited-effort fishery that yields less catch than the *Normal* fishery but still provides some revenue and also some catch-based information), and *Suspend* fishing (i.e., temporarily close the fishery entirely). Each action is available in each period. The observation set consists of the same two possible values as S , *Pure Strong* and *Mixed*, but an observation of $\theta = \textit{Pure Strong}$ does not necessarily mean that the true state $S = \textit{Pure Strong}$. Instead, we define an observation model R as follows:

$$R_{j\theta}^1 = \begin{bmatrix} 0.99 & 0.01 \\ 0.01 & 0.99 \end{bmatrix} \quad R_{j\theta}^2 = \begin{bmatrix} 0.75 & 0.25 \\ 0.25 & 0.75 \end{bmatrix} \quad R_{j\theta}^3 = \begin{bmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{bmatrix}$$

Each matrix, with the state $j \in S$ defined by row and each observation θ defined by column, defines the probabilistic relationship of observation to true state under a different action. $R_{j\theta}^1$, for example, tells us that after taking action $a=1$ (*Normal*) and moving to the unobservable state $j=Pure\ Strong$, we would observe $\theta=Pure\ Strong$ with 99% probability and $\theta=Mixed$ with 1% probability. That is, prosecuting the *Normal* fishery yields a very high-quality signal of the true state of mixing (‘what you see is what you get’). $R_{j\theta}^2$, in contrast, tells us that allowing a limited *Test* fishery ($a=2$), yields a weaker basis for inference on S , because with fewer fish being caught, the probability that the true state reveals itself is less. Finally, $R_{j\theta}^3$ indicates that immediately after choosing to *Suspend* the fishery, observations based on no catch tell us nothing about the true state of mixing in the fishery—each possible observation is equally likely under any given state, so an observation cannot tell us anything about the probability that the fishery is in a particular state.

The stochastic dynamics of the mixing state S are given by transition probability matrices defined as follows:

$$P_{ij}^1 = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix} \quad P_{ij}^2 = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix} \quad P_{ij}^3 = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix}$$

These matrices embody the assumption that the transition probabilities (e.g., probability of moving from *Pure Strong* at time t to *Mixed* at time $t+1$) are invariant to the decision made by the manager. Thus, while there is a 10% chance that a *Mixed* stock becomes *Pure Strong* from t to $t+1$, this chance is the same regardless of what the manager does.

Finally, the reward structure is as follows:

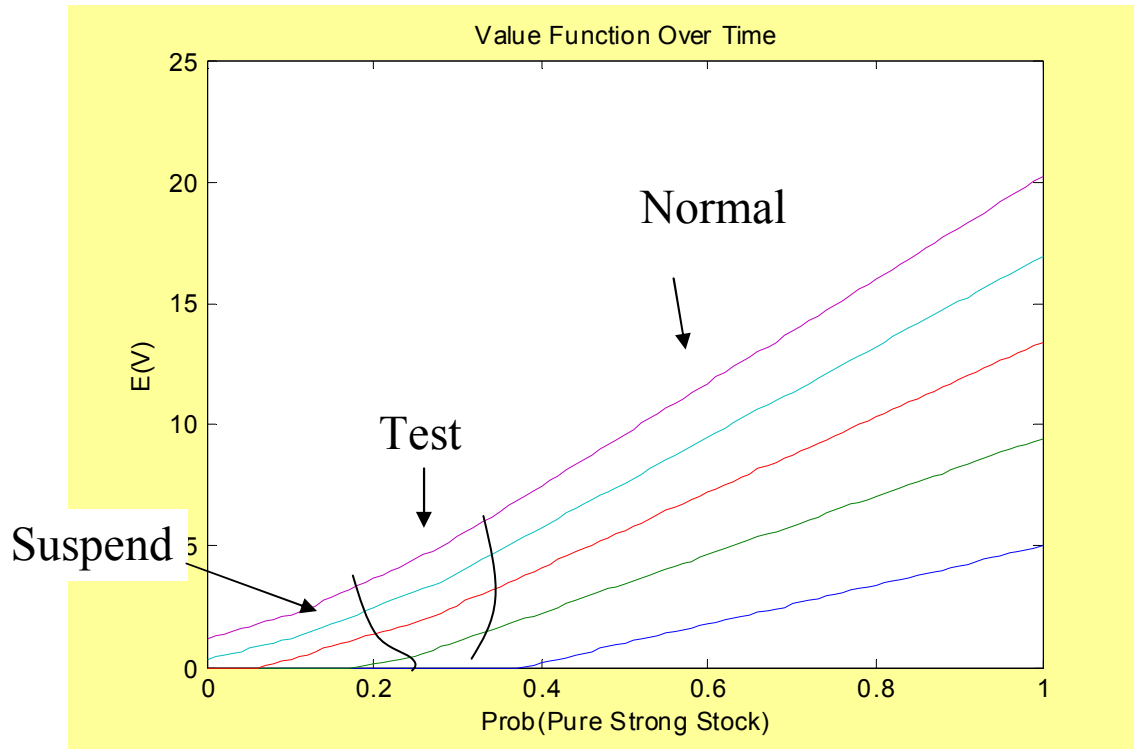
$$W_j^1 = [5 \quad -3] \quad W_j^2 = [1 \quad -1] \quad W_j^3 = [0 \quad 0]$$

Here the columns of each vector represent the rewards (in unspecified units) of taking a particular action in a particular state. W^1 tells us that allowing the *Normal* fishery while in the *Pure Strong* state will generate a payoff of 5, while allowing the *Normal* fishery in the *Mixed* state will lead to a penalty of -3 (which may represent lost social welfare, actual financial penalties imposed on the fleet, or any other sort of loss). W^2 , the payoffs to the *Test* fishery, are smaller in magnitude (both reward and penalty) due to the reduction of catch under this action. Finally, W^3 tells us that suspending the fishery will not result in any reward or penalty. Comparing these payoff vectors, it’s obvious that if the manager knew the true state to be *Pure Strong*, the best choice would be to implement the *Normal* ($a=1$) fishery, and if the manager knew the true state to be *Mixed*, the best thing to do would be to *Suspend* ($a=3$). However, the premise of the model, and the reality that managers generally face, is that the true state is unknown.

RESULTS

Figure 1 shows the value function, V , as it evolves over a 5-period decision horizon (the lowest solid line is V in period T-1, the next down is V in period T-2, etc.). The black lines show the division of the state space into policy regions, i.e., the beliefs for which the actions *Normal*, *Test*, or *Suspend* are optimal. The most salient features of the solution are that 1) as the decision horizon lengthens, *Suspend* is the optimal action for a progressively smaller portion of the belief space, and 2) *Test* enters the optimal policy at T-2

and becomes the optimal action over a larger portion of the belief space as the decision horizon lengthens. This is because information generated under the *Test* fishery has more value in longer-term planning, since it can be used to inform more decisions.



Further iteration on the value function could be used to explore the optimal policy for a decision horizon of any desired length. Our example here has been stylized both for ease of presentation and because POMDPs are known to be computationally intractable. Research on solution techniques for POMDPs is an active field in applied mathematics—more fully developed applications to fisheries management will employ recently developed heuristics to allow for larger state and action sets.

DISCUSSION

While restrictions on fishing effort may help to protect weak fish stocks, in many cases they also impose costs in terms of both foregone revenue and foregone information. The POMDP provides a tool for assessing the information value of a particular policy, which derives from the utility of whatever can be learned from fishing now about the likely current and future states of the fishery. Of particular interest are cases in which information value reverses the optimal choice between two candidate policies, as when a policy that appears to be suboptimal becomes optimal by virtue of the learning opportunity it affords. This is the idea of ‘probing’ policies from the dual control literature, which has been discussed in a fisheries setting by Walters (1986).

In the stylized model presented here, we found that for a 5-period decision horizon, implementing a *Test* fishery is optimal over a range of beliefs concerning the mixing state of the fishery, by virtue of the fact that the *Test* fishery provides more information that is not available when managers choose to Suspend the fishery. In this example, for most beliefs the opportunity cost of suspending the fishery exceed the expected benefit—suspending was preferred only when the manager had a fairly strong a

priori belief that the fishery was *Mixed*. Of course, these results are specific to the parameter values applied in this ad hoc example—a different parameter set may yield very different results.

Where do the beliefs about the degree of mixing come from? They are subjective beliefs that may come from personal experience, field experiments, studies from other areas, or other sources, updated within the model by incorporating new observations by Bayes' rule. That is, the POMDP is a Bayesian decision framework in which to change as the information set changes. Because different individuals will generally have different priors, they will at any given time have different beliefs even though they update their beliefs based on the same observations. While this may lead them to different conclusions regarding the best policy, it will often be the case the people with quite different beliefs about the state of the system can nonetheless agree on a policy. For example, in the example developed above, an individual who believes the probability that the stock is *Pure Strong* is 100% and another who believes this probability is only 50% should still be able to agree that the expected-value maximizing policy is the *Normal* fishery.

Of course, maximizing expected value may not be the managers' objective. While we suspect that imposing a risk-averse utility function in the POMDP will require significant algorithmic enhancement, forward simulation of candidate policies provides a basis for comparing the likely ranges of outcomes and associated statistics (e.g., performance variance). While this simulation approach is not equivalent to solution for a risk-averse objective function, it is easily implemented and provides a more rigorous basis for risk assessment than is usual in fisheries management.

REFERENCES

- Barnett-Johnson, R., F. Ramos 3, C. Grimes 1, and R. MacFarlane. 2006. Use of otoliths to identify river and hatchery of origin of California Central Valley Chinook salmon in the ocean fishery: potential application to Klamath River Chinook salmon. Informational Report to the Pacific Fisheries Management Council, available at http://www.pcouncil.org/bb/2006/0606/inforpt_1.pdf
- Bertsekas, D. 2000. *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific.
- Cassandra, A. 1994. *Optimal Policies for Partially Observable Markov Decision Processes*. Brown University Department of Computer Science Technical Report CS-94-14.
- Walters, C. 1986. *Adaptive Management of Renewable Resources*. New York: MacMillan Publishing.
- Winans, G., D. Viele, A. Grover, M. Palmer-Zwahlen, D. Teel, and D. Van Doornik. 2001. An Update of Genetic Stock Identification of Chinook Salmon in the Pacific Northwest: Test Fisheries in California. *Reviews in Fisheries Science*, 9(4):213-237.