# Fine-scale fishing strategies of factory trawlers in a midwater trawl fishery for Pacific hake (*Merluccius productus*)

**Martin W. Dorn**

**Abstract**: A Markov decision process model is developed to investigate the optimal scheduling of haul setting and retrievals on a factory trawler in the Pacific hake (*Merluccius productus*) fishery. The model was used to investigate changes in fishing behavior during a temporary ban on night fishing in 1992. The cycle of setting the net, fishing, and retrieving the net is modeled with different costs for each activity. The optimal controls generally consist of a bin threshold that signals the vessel to start fishing and a catch threshold that signals the vessel to stop fishing. A range of simple "rule of thumb" strategies generated nearly as much net revenue as the optimal control, indicating that the reward surface is flat in the region of the optimal control. A model with diel variation in catch rates caused the vessel to adjust bin and catch thresholds through the day to stockpile fish during the daylight hours and then to cut back on fishing at night when the catch rate is lower. With a ban on night fishing, vessels accumulated fish to a greater extent during the day, but daily net revenue did not decline significantly. Operational models of fishing have important applications in evaluating the consequences of management actions on the fishing industry. Regulations that ignore the constraints and trade-offs under which fishing vessels operate may fail to achieve their intended purposes, or may have unforseen adverse consequences.

**Résumé** : Nous décrivons un modèle décisionnel de Markov mis au point pour déterminer la planification optimale de la mise à l'eau et de la remontée du chalut d'un chalutier-usine dans la pêche au merlu du Pacifique (*Merluccius productus*). Ce modèle a servi à l'étude du changement dans les comportements de pêche durant une période d'interdiction provisoire de la pêche nocturne, en 1992. Nous avons modélisé l'enchaînement des opérations de mise à l'eau, de pêche et de remontée du chalut en appliquant un coût différent à chacune. En général, les régulateurs optimaux sont le seuil de pêche, le bateau commençant à pêcher lorsqu'il est atteint, et le seuil d'arrêt, où la pêche cesse. Avec diverses stratégies basées sur une simple règle empirique, nous avons obtenu des revenus nets presque aussi élevés qu'avec le régulateur optimal, ce qui indique que la surface de gain est plate dans la région du régulateur optimal. Dans un modèle supposant des variations de prises sur 24 h, le bateau a ajusté toute la journée les seuils de pêche et d'arrêt de façon à accumuler le poisson durant le jour, la pêche cessant la nuit, période où le taux de capture est faible. Durant la période où la pêche a été interdite la nuit, les bateaux ont accumulé plus de poisson le jour, mais les revenus nets quotidiens n'ont pas baissé dans une mesure significative. La modélisation des opérations de pêche a des applications importantes dans l'évaluation des effets des mesures de gestion sur l'industrie halieutique. Si les autorités de réglementation ne tiennent pas compte des contraintes et des arrangements qui caractérisent l'exploitation des bateaux de pêche, leurs objectifs pourraient n'être pas réalisés ou leurs interventions pourraient avoir des effets nuisibles inattendus.

[Traduit par la Rédaction]

## Introduction

Although the use of catch per unit effort (CPUE) to monitor the abundance of fish stocks has a long history (Smith 1988), researchers have only recently begun to consider fishing behavior as a legitimate subject for investigation (Hilborn and Ledbetter 1979; Gillis et al. 1995*a*). This emerging field of study is not as concerned with the time-honored objective of indexing abundance as it is with gaining a better understanding of the principles that govern fishing behavior. Much of this research has been guided by ecological models of optimal patch selection by animal foragers (Gillis et al. 1993), although economic aspects of fishing strategy have occasionally been studied (Lane 1989). Markov decision process models (or stochastic dynamic programming models) have proven to be useful tools for studying fishing behavior as a particular case of a general class of optimal foraging problems in behavioral ecology (Mangel and Clark 1988). Ultimately, this research has the promise of developing new techniques for interpreting fishing experience in a less restrictive way to monitor the abundance of fish populations. Operational models of fishing also have immediate practical applications in evaluating the consequences of management actions on the fishing industry. Management actions that ignore the constraints and trade-offs under which fishing vessels operate may fail to achieve their intended purposes, or may have unforseen adverse consequences (Gillis et al. 1995*a*).

The choices of fishing strategy available to fishers depend on the spatial and temporal scales of their assessment of the environment. Models and analyses of fishing strategy have mainly focused on the problem of patch selection by fishing vessels (Hilborn and Ledbetter 1979; Lane 1989; Gillis et al.

Received August 6, 1996. Accepted June 6, 1997.
J13590

**M.W. Dorn.** Alaska Fisheries Science Center, National Marine Fisheries Service, National Oceanic and Atmospheric Administration, 7600 Sand Point Way NE, BIN C15700, Seattle, WA 98115-0070, U.S.A.
e-mail: Martin.Dorn@noaa.gov

**Fig. 1.** Temporal and spatial scales in a hierarchy of decisions made by factory trawlers (adapted from Holling 1992).

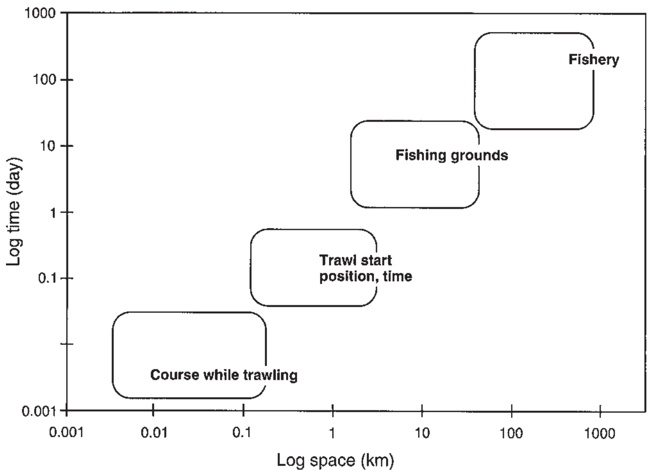


1993), leaving other important aspects of fishing behavior unexamined. Recent ecological models of animal foraging have emphasized the hierarchical character of decisions made by foraging animals (Holling 1992). A decision hierarchy appropriate for trawlers consists of at least four levels: (1) the decision to participate in a fishery, (2) the selection of a patch or fishing grounds on which to operate, (3) the scheduling of fishing operations within that patch, and (4) decisions associated with fishing depth, trawling speed, and compass bearing while actively fishing. Each level has a characteristic spatial and temporal scale. Decisions on which fishery to participate in are made annually with a spatial scale of ~1000 km; patch selections are made for periods ranging from several days to several weeks with a spatial scale of ~10–50 km (Fig. 1). Decision-making within a patch affects behavior over shorter intervals and smaller spatial scales.

Models with different temporal and spatial scales can shed light on different aspects of fishing behavior. This paper adopts a fine-scale focus to develop a model for decision-making over short periods of time by fishing vessels operating within a large-scale aggregation of fish. At this level of detail, it is reasonable to suppress the spatial aspects of foraging and focus exclusively on time scheduling problems. On trawlers, an important component of decision-making at this scale is the temporal scheduling of two events: setting the net, which initiates fishing, and retrieving the net, which ends fishing. A factory trawler will seek to maintain the flow of fish from holding bins into the factory by starting to fish when the amount of fish in the bins becomes low. If the vessel waits too long to start fishing, there is a risk that the factory will run out of fish before the next haul can be landed. If the vessel starts to fish too soon, it may catch fish for which there is no space in the holding bins. While fishing, the vessel must decide whether to continue fishing or to retrieve the net with the fish already caught. If the vessel continues to fish for too long, it may capture so many fish that retrieval is difficult and fish are damaged during retrieval, lowering product quality.

This research is motivated by a series of events that occurred in 1992. Based on an analysis that suggested that chinook salmon (*Oncorhynchus tshawytscha*) bycatch was higher at night, the Pacific Fishery Management Council (PFMC) established a night closure for the 1992 hake fishery extending from midnight to 1 h after official sunrise. At-sea processors objected to the night fishing ban, contending that the night closure disrupted fishing and processing operations. In response to these concerns and to a reanalysis of the salmon bycatch that suggested that bycatch rates were not significantly higher at night, the PFMC limited the night closure to south of 42°N latitude, leaving most of the fishing grounds open to fishing at night. Evidently, factory trawlers considered night fishing important, despite lower catch rates (the catch rate at night is approximately one third of the catch rate during the day) and possibly higher bycatch rates. The reasons for this are unclear. One possibility is that the factory trawlers were unable to catch enough fish during daylight hours to keep the factory operating at night, forcing shutdowns that reduced efficiency. However, not enough was known about daily fishing operations on a factory trawler to evaluate whether fishing at night was necessary for a profitable operation. To explore these hypotheses required the development of new analytical tools.

Here, I present a Markov decision process model for the scheduling of net setting and retrievals for a factory trawler. The model will be used to study how the state of the factory trawler influences decision-making, as measured by the tons of unprocessed catch in holding bins and the amount of fish in the net when the vessel is fishing. Since factory trawlers process their catch onboard, optimization models where the fishing vessel maximizes the landed value of catch are not appropriate for studying the fine-scale aspects of decision-making. Such vessels are constrained by their daily processing capacity, so an increase in daily catch above a certain level would not result in a corresponding increase in daily production. Here, I assume that the objective of the factory trawler is to maximize the daily net revenue, which implies that they are risk-neutral (Squires and Kirkley 1991).

Although the major focus in this paper is on optimal strategies, recent studies on foraging behavior of animals suggest that animals do not generally act according to predictions of optimal foraging models and are as likely to use simple "rules of thumb" to guide their foraging activities (Kareiva et al. 1989). In addition, the reward surface is often flat, suggesting that a wide range of strategies can result in close to the maximum reward. To explore these possibilities, a section of this paper uses forward simulation of the fishing process to test simple rules of thumb against the optimal strategies.

The following section of this paper provides an overview of the factory trawler fleet fishing for Pacific hake (*Merluccius productus*). In addition, the spatial features of the Pacific hake population as they relate to fishing strategy are briefly described using data from the 1992 National Marine Fisheries Service (NMFS) acoustic survey. Additional sections develop the Markov decision process model and describe the dynamic programming algorithm used to obtain the optimal controls. An analysis of a simple prototype is presented, and the results from simulations of a vessel following the optimal control are compared with at-sea observer data and some simple rules of thumb that mimic the more complex optimal controls. The simple prototype assumes that the random catch increment is an independent draw from a single probability distribution. Subsequent sections consider more complex and realistic models with correlated catch increments and diel patterns in catch increments. The question of whether a ban on night fishing significantly reduces net revenue is examined using the model with a diel trend in the catch increments. Finally, the implications of the model results are discussed and the major conclusions are briefly summarized.

**Factory trawlers and Pacific hake**

U.S. factory trawlers are a recent and influential component of the fishing fleet off the west coast of North America. Since 1990, these vessels have accounted for more than 1.0 million metric tons (tonnes, t) of groundfish catch per year from off the West Coast and in Alaskan waters (NMFS 1996). The at-sea processors in the Pacific hake fishery have onboard surimi (minced fish flesh) production capacity and were designed to fish primarily for walleye pollock (*Theragra chalcogramma*) in Alaska fisheries. Generally, these vessels also have the capacity to produce frozen fillet blocks and have a fish meal plant that processes the waste from the making of surimi. Although some fishing companies have several boats participating in the hake fishery, most vessels are from different companies. There are two classes of at-sea processors in the Pacific hake fishery: factory trawlers, which catch and process their own fish, and motherships, which process deliveries from catcherboats, but do not fish on their own. During 1991–1995, an average of 14 factory trawlers and five motherships participated in the fishery, with an average aggregate catch of 150 000 t. The fishery is managed using an annual harvest guideline (i.e., quota), with a separate allocation reserved for the shore-based processing sector. The at-sea fishery operates as a "derby" fishery in which all vessels compete for the fleet-wide quota. The fishery extends for 3–4 weeks after the annual opening date of April 15.

Factory trawlers are large vessels (mean length ~92 m), carry a crew of 70–100, and are capable of trips lasting several months. Between 50 and 70% of the crew is engaged in surimi production. Several shifts keep the factory in operation 24 h/day. Midwater trawls (mean trawl opening $90 \times 55$ m) are used exclusively for Pacific hake. As is usual with targeted midwater trawling, catches in the hake fishery are almost pure, with bycatch typically amounting to less than 3% of the total catch by weight. The most common bycatch species are pelagic rockfishes (yellowtail rockfish (*Sebastes flavidus*), widow rockfish (*Sebastes entomelas*), and Pacific ocean perch (*Sebastes alutus*)) and two species of mackerel (jack mackerel (*Trachurus symmetricus*) and chub mackerel (*Scomber japonicus*)). Most of the bycatch is discarded at sea. Although the bycatch of chinook salmon is low (4000–6000 fish/year), it has become an important concern since the listing of several west coast salmon runs as endangered under the *Endangered Species Act*.

Pacific hake form aggregations of sufficient density to support fishing activity along the continental shelf break from northern California to Vancouver Island. The 1992 NMFS acoustic survey of Pacific hake provides a synoptic view of hake spatial patterns (Fig. 2). In Fig. 2, the 0.5-nautical-mile (nmi) acoustic transect segments were marked with a vertical bar if the density of hake was greater than twice the density that would produce the mean catch rate in the 1993 fishery. The purpose here is to obtain a rough approximation of the spatial pattern and total area of the region where fishing would be successful. Acoustic densities were converted to expected catch rates by assuming that a midwater trawl captures all the fish in the water column and calculating the area swept by a net with a 90-m horizontal opening.

**Fig. 2.** Spatial distribution of Pacific hake in the 1992 NMFS west coast acoustic survey. The vertical bars mark the 0.5-nmi trackline segments where the acoustic density was more than twice the density that would produce a catch rate equal to the mean fishery CPUE. The 200-m isobath is also shown.

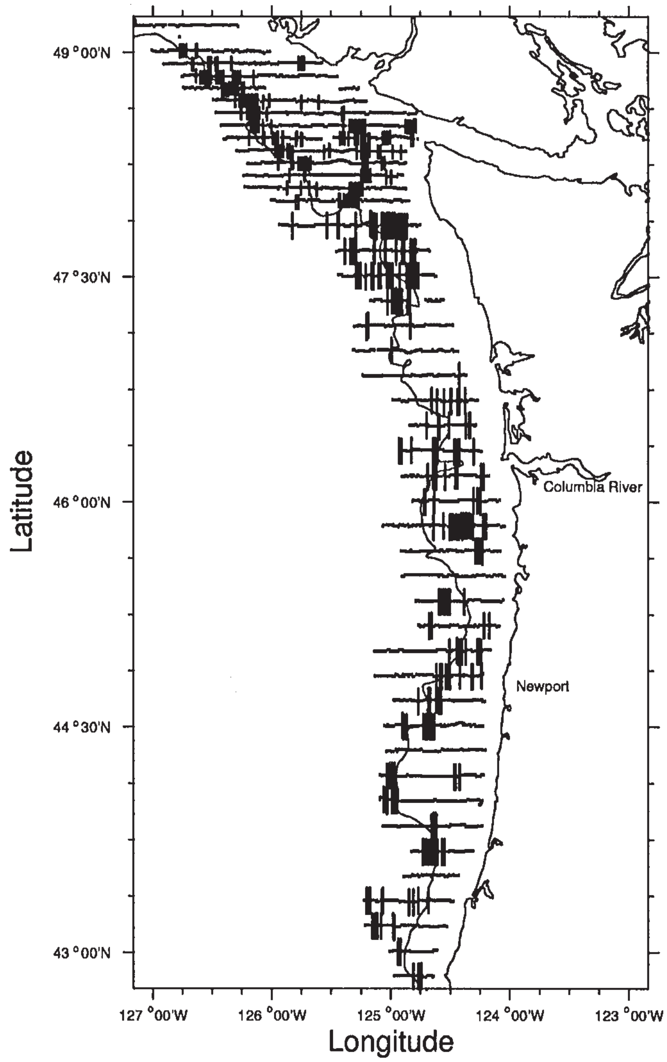


The affinity of Pacific hake for the shelf break habitat produces a fishing region that is much narrower in its east–west dimension than its north–south dimension. The north–south range of fishing is 600 km, while the fishery is conducted over bottom depths ranging from 150 to 600 m, an area that is between 10 and 30 km wide. Interspersed regions of high and low density extend along the entire coast. Certain features along the shelf break support higher densities of Pacific hake (e.g., Heceta Bank off central Oregon, Willapa and Guide canyons off southwest Washington, and Juan de Fuca Canyon off Cape Flattery), but Pacific hake aggregations are not confined to these areas. Between 15 and 25% of the total area between 150 and 600 m contains fishable densities of Pacific hake. These areas of high density are not persistent features, so they will have to be detected by the vessel before they can be exploited. From the perspective of the fishers, then, the key

**Table 1.** State dynamics on a factory trawler.

| Ship status, $s_k$ | Description | Decision, $u_k$ | Updated state | Per period cost, $\kappa(s_k, u_k)$ |
|---|---|---|---|---|
| 1 | Haul back | 0 | $s_{k+1} = 2$<br>$x_{k+1} = \max(0, x_k - p)$<br>$c_{k+1} = c_k$ | $\kappa_{sr}$ |
| 2 | Haul back (catch added to fish in bins) | 0 | $s_{k+1} = 3$<br>$x_{k+1} = \max(0, x_k - p) + c_k$<br>$c_{k+1} = 0$ | $\kappa_{sr}$ |
| 3 | Vessel able to fish | 0 | $s_{k+1} = 3$<br>$x_{k+1} = \max(0, x_k - p)$<br>$c_{k+1} = 0$ | $\kappa_{nf}$ |
| | | 1 | $s_{k+1} = 4$<br>$x_{k+1} = \max(0, x_k - p)$<br>$c_{k+1} = 0$ | $\kappa_{sr}$ |
| 4 | Vessel fishing | 0 | $s_{k+1} = 1$<br>$x_{k+1} = \max(0, x_k - p)$<br>$c_{k+1} = c_k$ | $\kappa_{sr}$ |
| | | 1 | $s_{k+1} = 4$<br>$x_{k+1} = \max(0, x_k - p)$<br>$c_{k+1} = c_k + w_k$ | $\kappa_f$ |

spatial characteristics of the hake population are (*i*) a narrow elongated region of potential occurrence and (*ii*) transient fishable aggregations of 15–30 km in size that can be fished multiple times.

## Methods

### Model development

Markov decision process models are discrete time representations of continuous time processes. The sequence of events that occur at each time step is as follows. First, the state of the vessel is observed. Based on the observed state, a decision is made whether to fish or not. Next, a random catch increment is generated according to a given probability distribution. The reward is then calculated and added to previous rewards. Finally, the state equations update the state to the next step. If the next step is the final period, then the terminal reward is added to the previous rewards and the process stops; otherwise, the same sequence of events is repeated for the next time step. The model was configured to increment time in 15-min steps, which gave the model sufficient resolution to capture the pattern of net setting, fishing, and net retrieval that make up the daily schedule on a factory trawler.

The state of the fishing vessel at each time step $k$ consists of three variables: the weight of unprocessed fish on board, $x_k$, the weight of fish currently in the net, $c_k$, and a ship status indicator variable, $s_k$, which keeps track of the current activity of the vessel. States $s_k = 1, 2$ occur in sequence after the decision to retrieve the net. The second state in the sequence, $s_k = 2$, is needed because once the vessel begins hauling back, it is committed to completing this activity during the following time step. In state $s_k = 3$, the vessel is able to fish, but not yet doing so. In state $s_k = 4$, the vessel is actively fishing. With these dynamics, the catch is added to the fish in the bins 45 min after the decision to stop fishing. Fish begin entering the net 15 min following the decision to start fishing. The state variables $x_k$ and $c_k$ are observable at each time step, implying that fish bins can be monitored periodically and that the vessel can monitor the catch already in the net while fishing. This assumption is reasonable because net telemetry on a factory trawler typically includes headrope-mounted side-scan sonar, which detects fish entering the net, and net tension indicators, which trigger as the net fills with fish.

In each time step, the weight of fish in the bins, $x_k$, decreases by the quantity the factory processes in a time step, $p$. In vessel state $s_k = 2$, the catch is added to the fish in the bins at the end of the time step. The state dynamics for the weight of fish in the net, $c_k$, depend on the ship status. If the vessel is fishing, the weight of fish in the net increases by a catch increment, $w_k$, drawn from a known nonnegative probability distribution. During a haul-back sequence ($s_k = 1, 2$), the weight of fish in the net remains constant, $c_{k+1}$. When the vessel is not fishing ($s_k = 3$), $c_{k+1} = 0$.

The decision set is $u_k \in \{0, 1\}$, where $u_k = 0$ is the decision not to fish and $u_k = 1$ is the decision to fish. During the haul-back sequence ($s_k = 1, 2$), the vessel is committed to completing the sequence, reducing the decision set to the single element $u_k \in \{0\}$. Although this model is conceptually straightforward, considerable detail is needed to treat the haul-back and net-setting sequences realistically. The state transitions may be easier to follow as listed in Table 1.

The state equations (Table 1) are subject to the constraints $0 \le x_k \le x_{\max}$ and $0 \le c_k \le c_{\max}$, where $x_{\max}$ is the maximum capacity of the fish bins and $c_{\max}$ is the maximum capacity of the net. The need for a constraint on the weight of fish in the bins is obvious and could be achieved by spilling fish from a net that has been retrieved or by discarding some of the fish already in the bins to make more space. A constraint on the maximum amount of fish in the net is a reasonable requirement, but awkward to model. If a vessel continues to fish when $c_k = c_{\max}$ and catch does not increase, then a potentially viable strategy might be to fill the net with fish and then delay the decision to haul back until some later time. In reality, this strategy would not be possible. Consequently, when $c_k = c_{\max}$ the decision set was reduced to the single element $u_k \in \{0\}$so that the vessel would not have the option of continuing to fish. Vessels avoid overfilling the net, and the model ought to include this behavior. This was done by adding a penalty for oversize catches which will be described shortly.

The random catch increments were modeled with a Δ-distribution, which specifies the probability of a zero catch and models the nonzero values with a lognormal distribution (Pennington 1983). Although the robustness of the use of this distribution to estimate abundance indices from survey data has been questioned (Smith 1990), the Δ-distribution was used here only for modeling. The mean of Δ-distribution was set equal to the mean fishery catch rate in the at-sea fishery in 1993, while the probability of a zero catch increment and the coefficient of variation (CV) of the Δ-distribution were obtained from the acoustically measured density of hake during the 1992 NMFS acoustic survey (Dorn et al. 1994). The purpose of this two-step process was to obtain a parametric probability distribution that approximated the catches per 15-min time step that would be encountered by a factory trawler in the Pacific hake fishery. For the optimization model, the Δ-distribution was discretized so that possible catch increments were even multiples of the processing rate ($p$ = 4.0 t/15-min time step) from zero to 200 t. The state variables $x_k$ and $c_k$ were discretized in the same way so that no interpolation was needed to update the state to the next time step.

The reward per time step is the value of surimi and fish meal produced during that time step, minus costs and penalties. The positive part of the reward is the processing rate multiplied by a conversion factor, ρ, which converts raw fish weight to product value. It is assumed that product value does not depend on the intrinsic characteristics of the fish (i.e., size or sex) or on how long the fish have been held before processing. Although Pacific hake are highly perishable, a vessel that processes its catches sequentially would hold fish for a maximum of $x_{\max} p^{-1}$ h (about 12 h for most factory trawlers). Declines in product quality are relatively slight over this interval. The activity costs of the vessel, $\kappa(s_k, u_k)$ depend on vessel status and on the decision, $u_k$. There are three different costs, a per time step cost while fishing, $\kappa_f$, a per time step cost while not fishing, $\kappa_{nf}$, and a per time step cost while the vessel is setting or retrieving the net, $\kappa_{sr}$.

Penalties can be incurred for discarded fish, $\nu_1$, and for oversize catches, $\nu_2$. Discards occur only when the net is retrieved and the total catch is greater than the remaining space in the bins. Under current management policy, no financial penalties are incurred by a vessel for discarded fish. However, discard is subtracted from the quota so that the total amount of fish available to the fishery is reduced by a corresponding amount. Furthermore, fishermen may have other motives for avoiding discard, such as ethical concerns about fully utilizing the catch and a desire to gain advantage in political settings where decisions are made about how to allocate the catch between user groups. The penalty for oversize catches is imposed for catches above a threshold, $c_{\text{thresh}}$. Large catches are hazardous to the deck crew, damage fish, and increase the likelihood of equipment failure. These would impose additional costs or reduce revenue. The net reward during a time step, $g_k$, is given by

$$g_k(x_k,c_k,s_k,u_k) = \begin{cases} \rho \min(p,x_k) - \kappa(s_k,u_k) - (\nu_1 + \nu_2) & \text{if } s_k = 2 \\ \rho \min(p,x_k) - \kappa(s_k,u_k) & \text{otherwise} \end{cases}$$

except for the final step, where $g_T(x_T) = \rho x_T$. The penalties $\nu_1$ and $\nu_2$ are incurred only when the catch is added to the fish already in the bins ($s_k = 2$) and are given by

$$\nu_1 = \begin{cases} \gamma_1(x_{k+1} - x_{\max}) & \text{if } x_{k+1} > x_{\max} \\ 0 & \text{otherwise} \end{cases}$$

$$\nu_2 = \begin{cases} \gamma_2(c_k - c_{\text{thresh}})^2 & \text{if } c_k > c_{\text{thresh}} \\ 0 & \text{otherwise} \end{cases}$$

where $\gamma_1$ and $\gamma_2$ are penalty weights that specify the relative importance of the penalties for discard and oversize catches, respectively. The influence of these weights on decision-making will be examined later. The penalty for discard, $\nu_1$, increases linearly with an increase in discard. The penalty for oversize catches, $\nu_2$, is quadratic above the threshold, modeling the escalating difficulties of handling larger catches. The use of a quadratic function to model these difficulties was based on interviews with observers and limited personal observations on fishing vessels. The most significant problem with large tows reported by fishers is the decline in product quality due to compression of fish in codends as they are hauled out of the water.

## Solving the optimization problem

Assume that the vessel fishes for $T$ time steps, with $T$ arbitrarily large. The objective of the vessel is to maximize the total expected reward over this period. The total expected reward, $J$, as a function of the initial states $x_0$, $c_0$, and $s_0$ is given by

$$J(x_0,c_0,s_0|u_0,u_1,u_2,\ldots,u_{T-1}) = \underset{w_k,\,k=0,\ldots,T-1}{E} \left\{ g_T(x_T) + \sum_{k=0}^{T-1} g_k(x_k,c_k,s_k,u_k) \right\}.$$

The expected reward will depend on the sequence of decisions $u_0$, $u_1$, $u_2$,..., $u_{T-1}$ made over this period. To solve the optimization problem, we seek a control law, a sequence of functions $\pi = \{\mu_0,\mu_1,\mu_2,\ldots,\mu_{T-1}\}$, where $\mu_k$ defines the decision $u_k$ for any vessel state at step $k$: $u_k = \mu_k(x_k,c_k,s_k)$. A control law is a plan of action, determined in advance, that specifies the decision for every possible state of the vessel. The optimal control law $\pi^*$ is the element of the set of feasible control laws $\Pi$ that maximizes expected reward:

$$J_0^*(x_0,c_0,s_0) = J(x_0,c_0,s_0|\pi^*)$$

$$= \max_{\pi \in \Pi} \underset{w_k,\,k=0,\ldots,T-1}{E} \left\{ g_T(x_T) \sum_{k=0}^{T-1} g_k(x_k,c_k,s_k,\mu_k) \right\},$$

where $J_0^*(x_0,c_0,s_0)$ is the maximum expected reward for the entire period. To solve this problem the dynamic programming algorithm is used (Bertsekas 1987). The dynamic programming algorithm is based on the Bellman optimality principle, which states that a control law that is optimal for the entire period must also be optimal starting from some intermediate step. It is therefore possible to solve the optimization problem for the next to last time step and then work backwards in time toward the initial time step to obtain the optimal control law for the entire period. The recursive equations are given by

$$J_T^* = g_T(x_T),$$

$$J_k^* = (x_k,c_k,s_k) = \max_{u_k(s_k)} \underset{w_k}{E} \{g_k(x_k,c_k,s_k,u_k) +$$

$$J_{k+1}^*[x_{k+1}(x_k,c_k,s_k,u_k),c_{k+1}(c_k,s_k,u_k,w_k),s_{k+1}(s_k,u_k)]\},$$

where $J_k^*(x_k,c_k,s_k)$ is the maximum expected reward starting at the $k$th time step.

For the models considered in this paper the optimal control law becomes stationary as the recursive equations get further away from the final step. In the results, I focus on stationary control laws rather than time-dependent controls, which become important only when little time remains before the closure of the fishery. A stationary control law was approached within 100–200 time steps (~1–2 days) moving backwards in time before the terminal time step. However, time-dependent controls may be important in assessing regulations, particularly fishing closures.

## Parameter values for a simple prototype

To investigate fishing strategy in the at-sea Pacific hake fishery, the optimization model was configured with parameters representative of a typical factory trawler. For this simple prototype, the random catch increments were generated from a single probability distribution with no serial correlation or trend. Model parameters were estimated from observer data obtained during the 1993 at-sea fishery and from cost and production surveys (Freese et al. 1995). Since all at-sea processors carry observers, the observer database is a complete record of the position, set and retrieval time, and total catch of all hauls by all vessels in the fishery. For some parameters, however, the available data permitted only a rough estimate of the appropriate value; in such cases the sensitivity of the optimal control to the parameter value was assessed. The following parameters were used in the prototype.

(1) Total bin capacity: 200 t. Estimates of bin capacity are not available for all factory trawlers. Two hundred tonnes is the mean of five bin capacities recorded by observers in the 1993 hake fishery.

(2) Maximum net capacity: 160 t. The maximum of observed catches in 1993 was 161 t. Observer logbooks record instances where hauls 150–170 t could not be brought up the stern ramp, and some of the catch had to be discarded.

(3) Processing rate: 4.0 t/15-min time step. At this processing rate the vessel could process 384 t/day, which is the 84th percentile of the catch per day during the 1993 fishery. Since daily catches should not greatly exceed the daily processing capacity, this processing rate is a reasonable approximation. No direct observations of factory processing rates are available.

(4) Revenue per tonne of hake catch (includes both surimi and fish meal): \$260. (Freese et al. 1995.)

(5) Costs. An estimate of the cost of operating a factory trawler is \$80 000/day (\$833/time step) (Freese et al. 1995). The model provides for three different per time step costs: a cost incurred while fishing, $\kappa_f$, a cost incurred while not fishing, $\kappa_{nf}$, and a cost incurred while setting or retrieving the net, $\kappa_{sr}$. Based on fuel consumption, the potential for equipment failure or damage, and the crew's activity level, a ranking of the costs $\kappa_{nf} < \kappa_f < \kappa_{sr}$ is reasonable. For the simple prototype, the values $\kappa_{nf} = 800$, $\kappa_f = 840$, and $\kappa_{sr} = 880$ were used.

(6) Penalty weight for discards: \$130. With this penalty weight, a tonne of discard incurs a penalty equal to one half the revenue that it would have produced. If Pacific hake were managed using an individual transferrable quota system (ITQ) where discards are counted against the quota, the penalty for discard would be equal to the

**Fig. 3.** SOC for the simple prototype. The black area indicates the part of the state space where the optimal control is to continue fishing and the open area indicates where the optimal control is to begin retrieving the net. At the lower edge of the figure, the optimal control for $s_k = 3$ (able to fish but not currently fishing) is indicated.

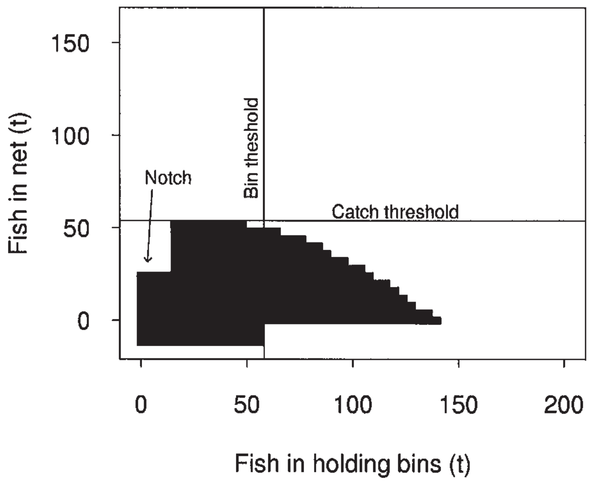


revenue that it would have produced. Here, some intent to avoid discard is assumed, but less than would be expected under an ITQ system.

(7) Penalty weight for oversize catches: \$0.722. With this penalty weight, the penalty incurred by a catch equal to the net capacity (160 t) is equal to the revenue generated by 10 t of unprocessed catch. A catch smaller than 100 t incurs no penalty. This threshold is the 93rd percentile of the observed haul weight for factory trawlers in the 1993 at-sea fishery.

Parameter values for the Δ-distribution of random catch increments were obtained from two sources: the catch rates in the 1993 at-sea fishery and acoustically measured densities of hake along transects during the 1992 NMFS acoustic survey. The best way to estimate these parameters would be to use the catch by 15-min time steps in the fishery. Collecting this information would require the installation of a calibrated echosounder on a fishing vessel. With the fisheries data that are available, catch rates can be calculated only for entire tows (total catch per haul duration). The mean of these catch rates is appropriate to use for the mean of the Δ-distribution, but the CV and the probability of a zero catch would be higher for 15-min time steps than the catch rate for entire tows. The mean catch per 15-min time step in the at-sea fishery in 1993 during the day (06:00–20:00) (all times are in Pacific Daylight Time) was 13.9 t/15 min (CV = 1.18), while at night the mean catch rate was 5.0 t/15 min (CV = 1.00). The day catch rate of 14.0 t/15 min was used as the mean of the probability distribution of catch increments for the basic prototype.

The acoustically measured densities of hake along transects during the 1992 NMFS acoustic survey were used to obtain the CV and the probability of a zero catch increment. These data consisted of mean acoustic densities for 1-nmi transect segments with bottom depths between 150 and 600 m from 43°00′N to 48°15′N ($n$ = 377). This region corresponded to the area where the at-sea fleet fished in 1993. Because mean towing speed in the at-sea hake fishery is 4.0 kn, a vessel transits ~1 nmi in 15 min. Mean acoustic densities were converted to nominal catches for 15-min time steps by assuming that a midwater trawl captures all the fish in the water column and calculating the area swept by a net with 90-m horizontal opening in 1 nmi of trawling. Because fishing is not random within the area where the fishery operates, this approach may misrepresent the distribution of potential catch increments. However, out of the limited alternatives available, this method was considered to be the most reliable.

The empirical distribution of acoustic hake densities was right-skewed and contained 7.7% zero observations (mean = 9.4 t/15-min interval, CV = 1.5). A Kolmogorov–Smirnov test (using the critical values by Lilliefors (1967)) indicated that the logarithm of the nonzero observations differed significantly from a normal distribution ($\alpha$ = 0.05). This significant result was due to the presence of seven small nonzero values. The Kolmogorov–Smirnov test for the logarithm of values >0.135 was not significant, and quantile–quantile plots against a normal distribution indicated that a lognormal distribution was appropriate for catch increments >0.135 t/15 min. Because rounding these smaller values down to zero does not substantially change the mean and the CV, a Δ-distribution was considered an appropriate parametric distribution to model these data. The probability of a zero catch increment was estimated as 36/377 = 0.095. Conversion formulas in Pennington (1983) were used with the CV of the acoustic data (1.5) and the mean fishery catch rate (14.0 t/15-min time step) to obtain the mean and standard deviation of the lognormal distribution for the nonzero part of the distribution.

The final step consisted of obtaining a discrete probability distribution for catches of 0, 4, 8, ...,160 t as required by the dynamic programming algorithm. This was obtained by summing the probability mass between the bin limits. The bin limits were located at the midpoints of adjacent values of 0, 4, 8, ...,160 t on a log scale. The bin boundary between zero and 4 t was arbitrarily set at 2 t. Two additional stochastic models for the catch increments are developed in subsequent sections, a model with first-order serial correlation and a model with diel trends. In both cases, the parameters for the model were obtained in a similar way to the method described above, where the expected catch increment was obtained from fishery catch rate data, and other parameters (i.e., variance, probability of zero catches) were estimated with acoustic survey data.

## Results

### Description of the stationary optimal control for a simple prototype

The dynamic programming algorithm for the simple prototype reached a stationary optimal control (SOC) in 86 time steps (~22 h). Decisions are possible in two ship states, $s_k = 3$, when the vessel is not fishing, and $s_k = 4$, when the vessel is fishing. For $s_k = 3$, the optimal control indicates the appropriate decision of whether to start fishing or not for each of the possible bin levels. For $s_k = 4$, the SOC divides the state space into two regions: a region where the optimal control is to continue fishing and a region where the optimal control is to begin retrieving the net. The fishing region occupies the quadrant of the state space where both the catch in the net ($c_k$) and weight of fish in the bins ($x_k$) are low (Fig. 3). The important characteristics of the SOC are as follows. First, there is a bin threshold, such that when the weight of fish in the bins drops to the threshold (56 t for the simple prototype), the vessel should start fishing. Second, the upper margin of fishing region below the bin threshold is flat, indicating that when the catch in the net increases above a catch threshold (52 t for the simple prototype), the optimal control is to cease fishing independent of the weight of fish in the bins. Finally, there is a notch in the fishing region where the weight of fish in bins is low and the catch in the net is between 24 and 52 t. The notch causes the vessel to retrieve the net before reaching the catch threshold and provides additional protection against running out of fish to process.

The sensitivity of the SOC to model parameters was examined by obtaining the SOC for a range of values for a particular parameter with all the other parameters held constant. Sensitivity to the following parameters was examined: the penalty

**Fig. 4.** Sensitivity of SOC to the model parameters. The lower edge of each figure indicates the optimal control for $s_k = 3$; the main part of the figure indicates the optimal control for $s_k = 4$. The top panels show the SOC for a range of penalties on oversize catches (>100 t) with all other model parameters held constant. The middle panels show the SOC for a range of values for the costs of fishing, not fishing, and setting or retrieving the net. The bottom panels show the SOC for a range of discard penalties. Parameter values used to obtain these figures are given in Table 2.

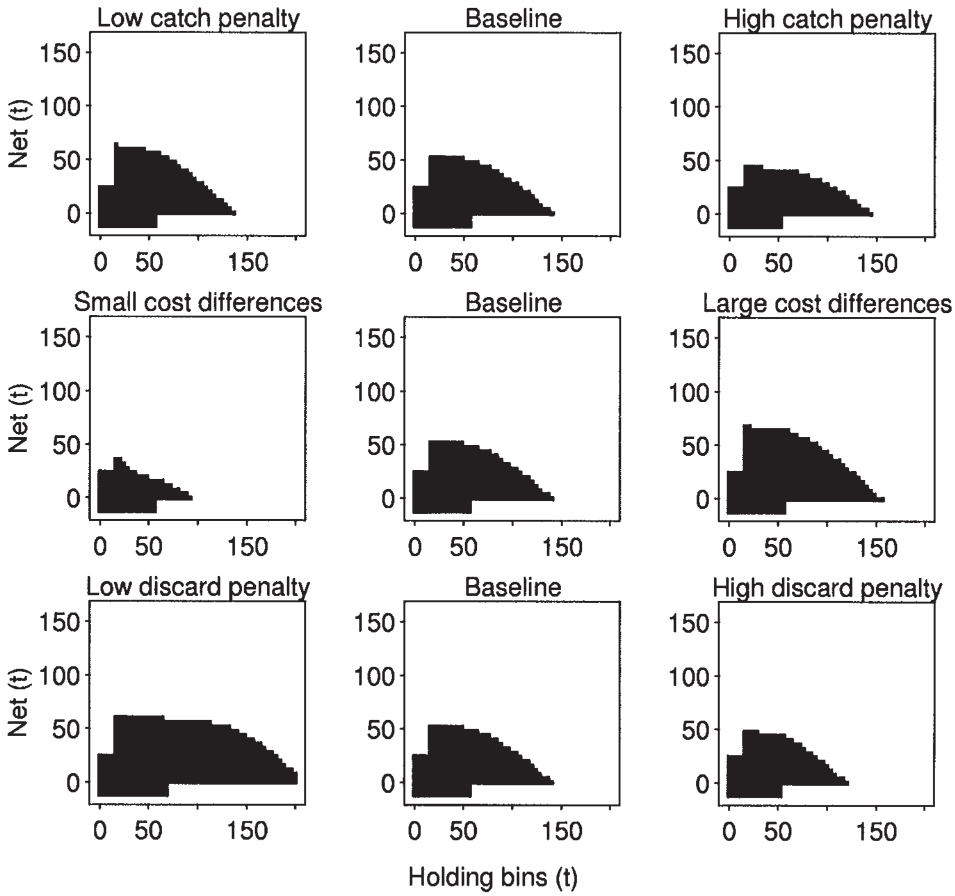


for oversize catches, the cost differences between not fishing, fishing, and setting/retrieval, and the penalty for discard (Table 2). The general characteristics of the SOC are robust to variation in these parameters (Fig. 4). In particular, the bin threshold is extremely robust to variation in parameter values. Changes in the penalty for oversize catches primarily affect the catch threshold. With a large oversize catch penalty, the vessel adopts a more cautious strategy and does not fill the net as full before retrieval (top panels in Fig. 4). However, the change in catch threshold is relatively modest, with a difference of about 20 t between the low penalty and the high penalty. When cost differences are small, the vessel incurs little additional costs for setting and retrieving the net, resulting in a very low catch threshold. Change in the relative costs of different activities has a substantial effect on the catch threshold (middle panels in Fig. 4). As the cost differences increase, the vessel waits until the net is fuller before retrieving it. This would reduce amount of time spent setting and retrieving the net, and thus avoid the higher cost of these activities. The range of discard penalties between zero and the monetary value of the discarded fish had little influence on the SOC. A low discard penalty expands the fishing region, but mostly in a part of the state space that would not be visited in a forward simulation of the state dynamics. With a low discard penalty, the

**Table 2.** Ranges of parameter values used to examine the sensitivity of the optimal control to key parameters.

| | Low | Baseline | High |
|---|---|---|---|
| Oversize catch penalty | 0.361 | 0.722 | 1.444 |
| Cost per time step contrast | $\kappa_{sr} = 810$ | $\kappa_{sr} = 880$ | $\kappa_{sr} = 960$ |
| | $\kappa_{nf} = 800$ | $\kappa_{nf} = 800$ | $\kappa_{nf} = 800$ |
| | $\kappa_f = 805$ | $\kappa_f = 840$ | $\kappa_f = 880$ |
| Discard penalty | 0.0 | 130.0 | 260.0 |

vessels have a higher bin threshold and fill their nets fuller before retrieving, but only slightly (bottom panels in Fig. 4). These results suggest that concern over discarding has a minor effect on the SOC.

Next, the response of a vessel following the SOC to changes in the mean and variance of catch rates was examined. As the mean catch rate increases, the region of the state space where the optimal control is to continue fishing shrinks in size. The bin threshold decreases to less than 50 t for a catch rate of 21 t/interval (top panels in Fig. 5). The catch threshold also decreases as the mean catch rate increases. As a result, the vessel spends less time actively fishing when the catch rate is high. In a more realistic setting, a vessel could use this

**Fig. 5.** Effect of changes in the mean and CV of the catch rate on the SOC. The top panels show the SOC for a range of mean catch rates with the CV fixed at 1.5. The bottom panels show the SOC for a range of CVs with the mean catch rates fixed at 14.0 t/15 min.

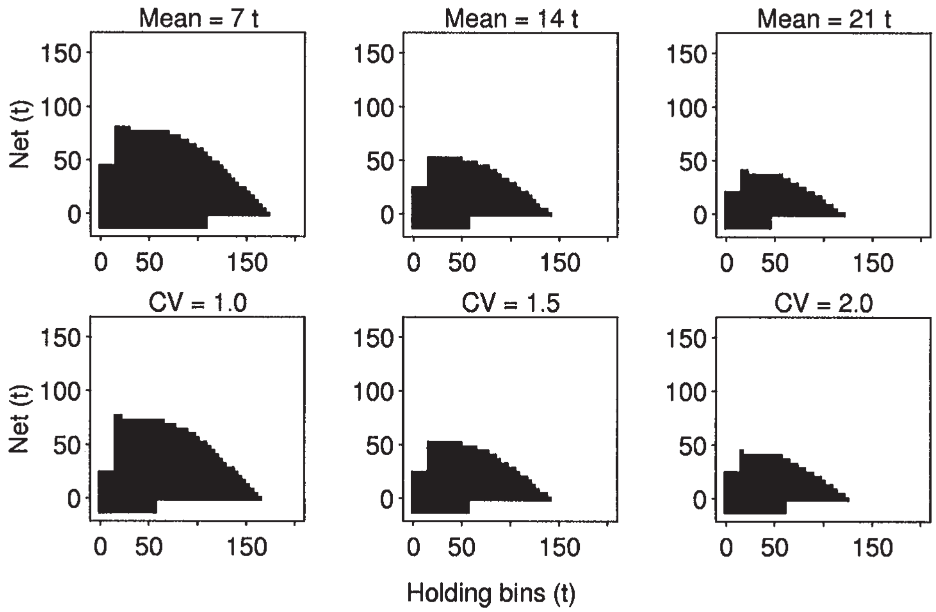


additional time by searching for other high-density aggregations so that the catch rate would continue to be high in subsequent hauls. When catch rates are low, the vessel must spend all its time fishing and would have little time to seek out higher density aggregations. Changes in catch rate variance strongly affect the catch threshold (bottom panels in Fig. 5). Higher variance (as measured by the CV of the catch rate) produces a decrease in the catch threshold. The bin threshold remains stable as variance in the catch rate changes, suggesting that it is determined by the expected catch rate and not by its variance.

**Forward simulation of the optimal control**

A forward simulation model of the state dynamics was constructed to examine the behavior of a vessel following the optimal control (Mangel and Clark 1988). The state dynamics and parameters were the same as those used in the optimization model. In addition, the optimal decision table as a function of vessel state was provided to the model. The model was started with $x_0 = 0$, $c_0 = 0$, $s_0 = 3$ and run forward with the vessel following the SOC. To describe the typical path through the state space, consider a vessel that is not fishing, but has 100 t in its bins. Its path will be to the left along the bottom margin of the state space (Fig. 3). When the bin level drops to 56 t, the vessel starts fishing and starts moving upwards in the state space as catch increments are added to the catch already in the net. When the catch in the net exceeds ~52 t, the vessel retrieves the net over two time steps and returns to the bottom margin of the state space. The simulated vessel cannot catch 52 t exactly because the decision to continue fishing must be made before the random catch increment is added to the catch in the net. A vessel may have more control over its catch than results from a discrete time model with 15-min time steps. Nevertheless, the Markov decision process approach requires the use of a discrete time model, and the use of a 15-min time step in the model is a reasonable level of accuracy to study decisions relating to setting and retrieving the net. The use of a smaller time step in the model might allow the vessel more control over its catch, but the vessel would still have to decide whether to continue fishing before knowing what it would catch in the next time increment.

Frequency distributions and pairwise scatterplots of haul weight, haul duration, and haul interval (time between successive tows) for a 20 000 time step simulation of the SOC were examined to determine the haul characteristics of a vessel following the optimal control (Fig. 6). The haul weight frequency distribution shows a peak just above the catch threshold of 52 t and an extended right tail to the distribution (Fig. 6A). Very few tows less than the catch threshold occur in the forward simulation. Haul durations for the forward simulation ranged between 0.25 and 2.75 h and showed a symmetrical distribution (Fig. 6B). The haul intervals (time between tows) ranged between 1.0 and 9.8 h and had a mean of 3.2 h (Fig. 6C). The distribution of haul intervals is also right-skewed, but does not show a sharp minimum. A haul interval of 1 h is the minimum possible time, and it occurs when the vessel starts fishing immediately. This occurred in ~4% of all the hauls made in the forward simulation. Above the catch threshold of 52 t, the haul intervals increased approximately linearly with increasing haul weight (Fig. 6D). As haul weight increases, more time is required to reduce the amount of fish in the bins to the bin threshold. Because the vessel uses a catch threshold to decide when to retrieve the net, haul weights appear independent of haul duration (Fig. 6E). Since the basic stochastic assumption of the model is that the catch per time increment is a random variable, without the optimal controller adjusting effort (i.e., haul duration), there would be a positive relationship between catch and effort. Haul intervals are inversely related to haul duration (Fig. 6F). When a short-duration haul is emptied into the holding bins, the total amount of fish in the bins will be considerably above the level that triggers fishing, so the vessel does not begin fishing immediately. However, if haul duration is long, the bin level will be substantially below the

**Fig. 6.** Summary graphs for a 20 000 step forward simulation of the SOC: (A) frequency distribution of haul weight; (B) frequency distribution of haul duration; (C) frequency distribution of haul interval; (D) scatterplot of haul interval versus haul weight; (E) scatterplot of haul weight versus haul duration; (F) scatterplot of haul interval versus haul duration. The values for the scatterplots were jittered.

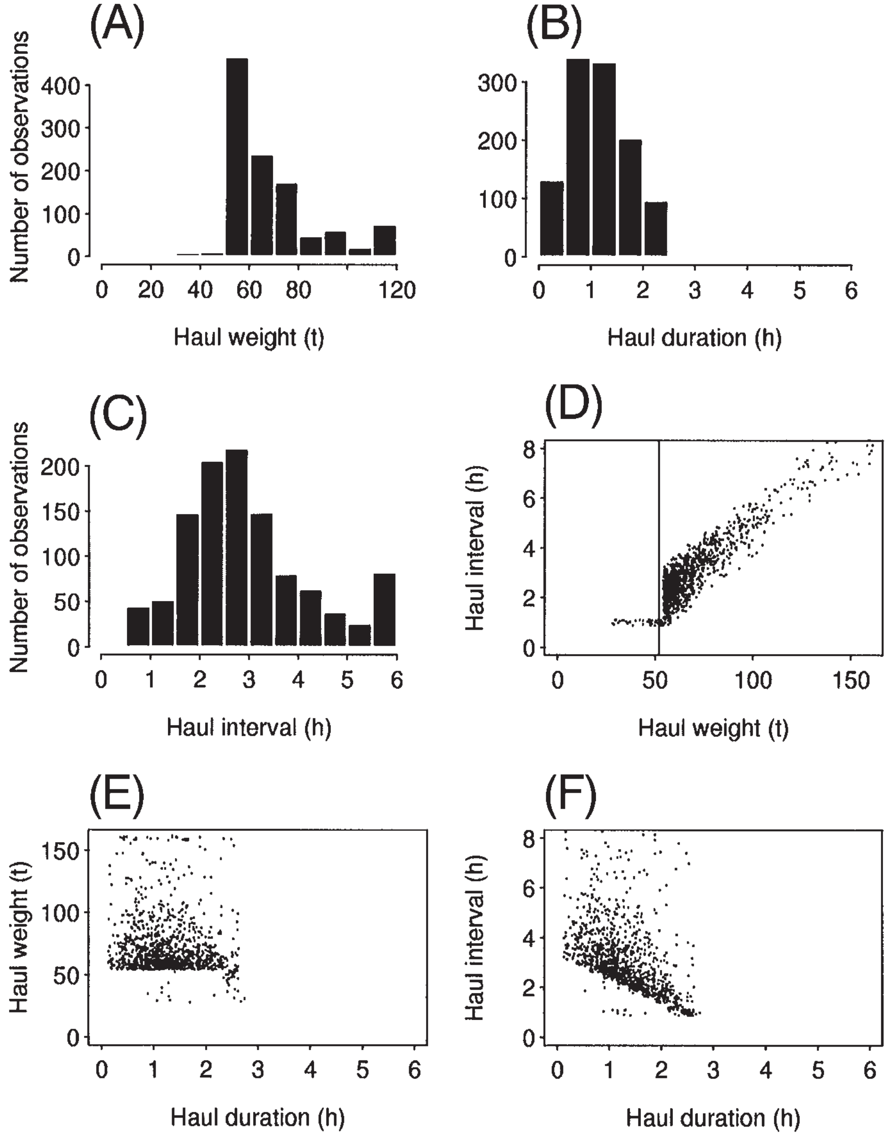


threshold when the net is emptied. As a result, the total amount of fish in the bins will not greatly exceed the bin threshold, and so the vessel will quickly begin fishing again.

## Comparison with at-sea observer data

There are considerable obstacles to comparing the simulated optimal control with the statistics of fishing vessels. During a fishing season, a vessel will adjust its fishing strategy to changing circumstances. A vessel may encounter fish at different mean densities rather than at a single mean density, as was assumed by the optimization model. The higher level aspects of fishing strategy (patch selection and exploratory fishing) will affect the vessel catch statistics in ways that are difficult to predict. Fishing vessels have different maximum bin capacities and processing rates and use different nets. Their fishing strategy would likely vary because of these differences. Nevertheless, if the model captures some fundamental aspects of fishing behavior, there should be some correspondence between a simulation of the optimal control strategy and the qualitative fishing patterns of fishing vessels.

One of the most important predictions of the model that can be assessed with at-sea observer data is the presence of a catch threshold. Frequency distributions were compiled for haul weights in the 1993 fishery during the daylight hours (06:00–20:00) when the vessel traveled less than 50 km before starting the next haul, and the time between hauls was less than 8 h. These criteria were applied to exclude hauls whose properties might reflect diel changes in fishing strategy or changes in behavior related to movement to new fishing grounds. Frequency distributions of haul weight for individual factory trawlers usually showed a left-truncated distribution with few small values (Fig. 7). The range of catch thresholds was from 40 to 70 t and corresponded well to the model predictions. Most haul weight distributions showed a scattering of points

**Fig. 7.** Frequency distributions of haul weight for six factory trawlers in the 1993 at-sea Pacific hake fishery. Vessels were given random codes to preserve confidentiality.

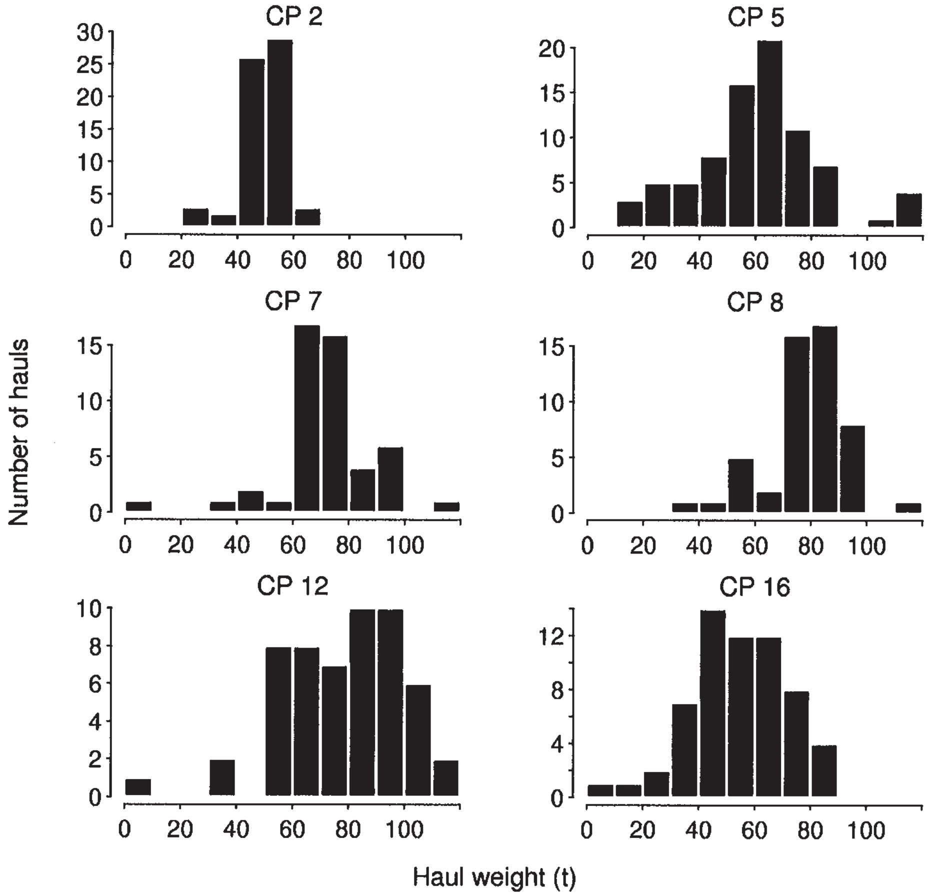


below the catch threshold. These hauls would not have occurred if the vessel had followed a strict catch threshold strategy. One potential explanation is that vessels adjusted their catch threshold in response to changes in fish density. Later in this paper, I examine whether serial correlation in catch sequences would produce simulated haul weight distributions with this characteristic.

Since observers do not record the amount of fish in holding bins, it is not possible to examine directly whether fishing vessels use a bin threshold to decide when to start fishing. However, it is possible to address this issue in an indirect way by assuming a nominal processing rate, keeping track of the derived amount of fish in the holding bins by adding in the catches and subtracting the nominal processing rate, and recording the bin levels when the vessel starts to fish. This procedure was done using data from six factory trawlers for the first 10 days of the 1993 fishery. At the start of the season, no vessel would have fish onboard, so the initial state was known. Catches were added to the fish in the bins 45 min after the recorded start of net retrieval. An estimate of the nominal processing rate was obtained by stepping through a range of processing rates and selecting the processing rate that satisfied the following criteria. First, the maximum tonnes onboard should not greatly exceed 200 t, since most vessels have a bin capacity of ~200 t. Second, the nominal processing rate should not result in the factory being out of fish to process a high proportion of time (i.e., greater than 15% of the time). The nominal processing rates obtained in this way ranged from 2.9 to 4.9 t/15 min for the six vessels. The frequency distributions of nominal bin levels when the vessel decided to start fishing showed a right-truncated distribution (Fig. 8). This suggests that there is an upper limit to the amount of fish in the bins when the vessel starts fishing. These thresholds ranged from 60 to 100 t, higher than the bin threshold of 56 t for the SOC. In forward simulations, very few hauls were started when the bin level was below the threshold, yet apparently this occurs frequently to fishing vessels, suggesting that vessels have more difficulty keeping bin levels above the threshold than occurs in the basic prototype. One possibility for this discrepancy is that the probability model for catch increments does not adequately represent the distribution of catch increments in the fishery.

Frequency distributions and pairwise scatterplots of haul weight, haul duration, and haul interval (time between successive tows) for all factory trawlers in the 1993 at-sea fishery (Fig. 9) correspond to those for the SOC simulation (Fig. 6). As with Fig. 7, only the daytime catches where the vessel traveled less than 50 km to the next tow were selected as appropriate for comparison with the SOC simulation.

Although combining the catch statistics for different vessels obscures the catch thresholds for the individual vessels, a tendency for most of the hauls to be greater than 40 t is apparent in the frequency distribution for all factory trawlers (Fig. 9A). The haul weight distribution has a long right tail of large

**Fig. 8.** Frequency distributions of bin levels when starting to fish for six factory trawlers in the 1993 at-sea Pacific hake fishery. Vessels were given random codes to preserve confidentiality.

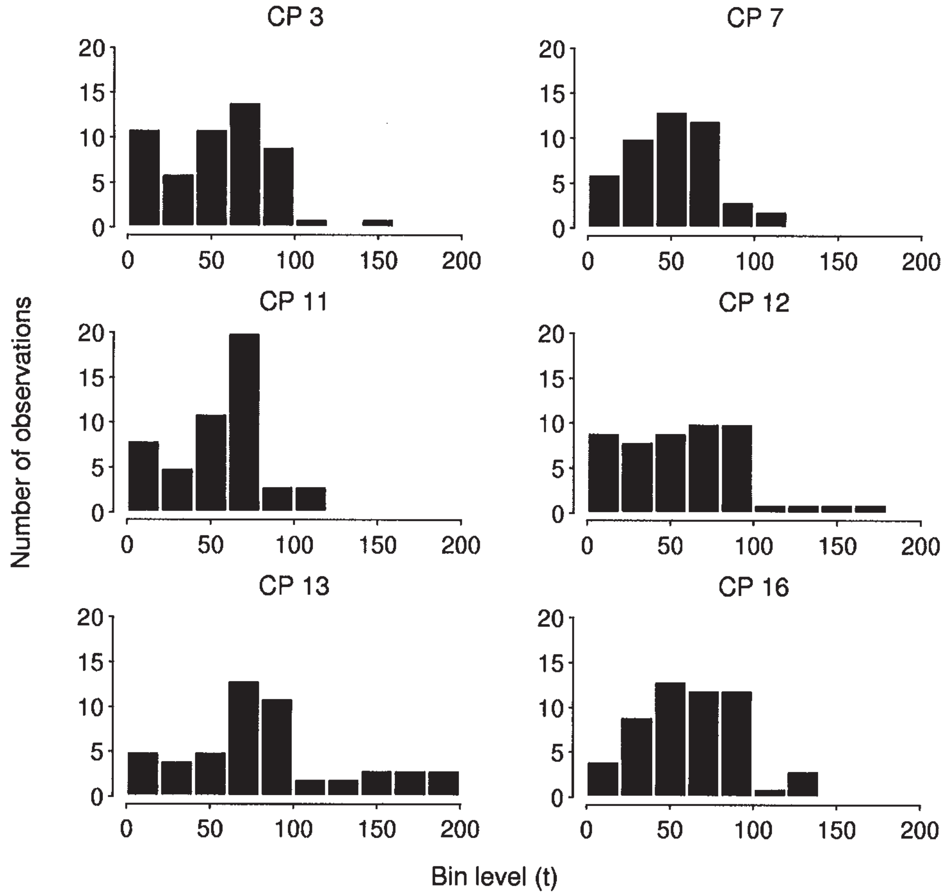


catches similar to the haul weight distribution for the simulated SOC. Haul duration averaged about 2.0 h with a scattering of tows of longer length (Fig. 9B). The haul intervals showed a different frequency distribution than the SOC (Fig. 9C). There is a sharp peak in the distribution at 1.0–1.5 h, indicating that vessels frequently started fishing again soon after retrieving a net. This does not occur as often in the SOC simulations. A reviewer suggested that fishers may be genuinely risk-averse about running out of fish to process and begin fishing sooner than is optimal under assumption of risk-neutrality. Another possibility is that the cost of running out of fish is higher than is assumed in the model.

The pairwise scatterplots generally show similar patterns to the scatterplots for the simulated SOC. Haul interval increases with haul weight above ~50 t (Fig. 9D). This pattern suggests that vessels base the decision when to begin fishing on the amount of fish in the bins, since otherwise, there would be no relationship between haul interval and haul weight. The haul duration versus haul weight plot (Fig. 9E) shows no increase in haul weight with haul duration, similar to the statistics for the simulated SOC. This is a strong indication of decision-making based on the amount of catch in the net, since haul weight would be expected to increase with haul duration with a random distribution for the catch increments. There is an inverse relationship between haul duration and haul interval (Fig. 9F), such that longer tows are followed by shorter haul intervals and vice versa. This also corresponds to the statistics for the simulated SOC.

## Rule of thumb decision algorithms and the shape of the reward surface

The SOC is a bin and catch threshold strategy plus some refinements. A bin and catch threshold strategy is the following: start fishing when the weight of fish in the bins is $\leq x$ tonnes and stop fishing when the weight of fish in the net is $\geq y$ tonnes. Kareiva et al. (1989) used the term "rule of thumb" for decision-making algorithms that foragers might plausibly use. Such algorithms may perform nearly as well as an optimal strategy and may have other advantages such as simplicity, ease of calculation, and robustness. For the foraging problem examined in this paper, a bin and catch threshold strategy could be called a rule of thumb fishing strategy. Note that this strategy is obtained by abstracting the most important characteristics of the optimal control into a simple set of rules. Simulation of these rule of thumb strategies can address two related questions: (1) how well do these algorithms perform relative to the SOC and (2) how sensitive is the reward (daily net revenue) to the decision-making algorithm. For example, it may be possible to obtain most of the potential reward with many different decision algorithms. If the reward is not particularly sensitive to the decision-making algorithm, this could explain differences in fishing strategy between vessels.

**Fig. 9.** Summary graphs for observer data collected on factory trawlers during the 1993 Pacific hake fishery: (A) frequency distribution of haul weight; (B) frequency distribution of haul duration; (C) frequency distribution of haul interval; (D) scatterplot of haul interval versus haul weight; (E) scatterplot of haul weight versus haul duration; (F) scatterplot of haul interval versus haul duration. A lowess smooth (Chambers and Hastie 1992) was fit to each scatterplot to bring out the trend in the data.

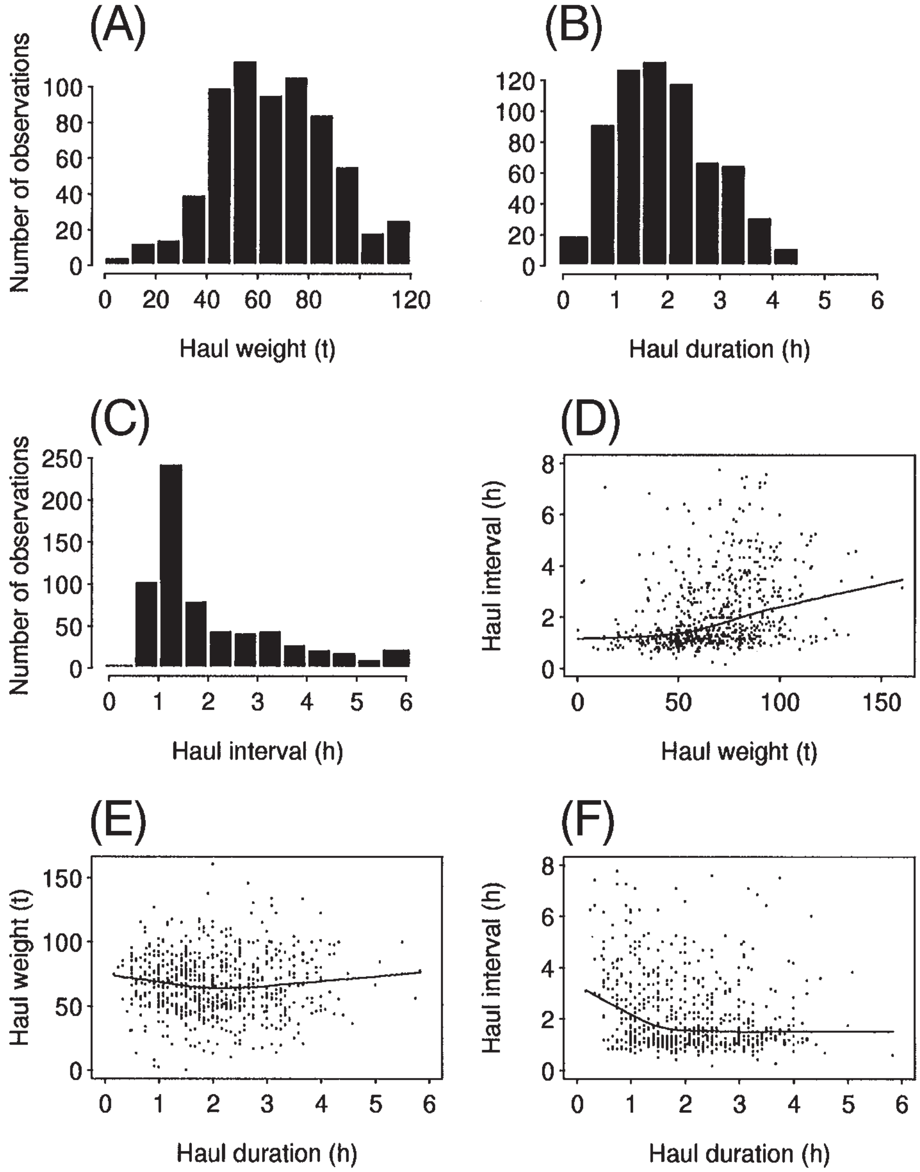


Forward simulations of rule of thumb strategies were used to generate contour plots of net revenue for three different mean catch rates (7.0, 14.0, and 21.0 t/15 min). A grid of all possible combinations of bin and catch threshold strategies in increments of 4 t ($n$ = 41 catch levels × 51 bin levels = 2091) was used. For each bin and catch threshold pair, the mean revenue was obtained from 20 simulations of 2000 time steps, so that each point in the grid was based on simulations roughly equivalent to 20 vessels following that strategy for a 20-day fishery opening for Pacific hake.

The region of the reward surface where bin and catch threshold strategies attain greater than 90% of the net revenue of the SOC is fairly broad for each of the mean catch rates (Fig. 10). The flat-topped portion of the reward surface is roughly the same size for each of the three catch rates, but its location depends on the catch rate. When the catch rate is low (left panel of Fig. 10), maximum revenues occur at high catch and bin thresholds. When the catch rate is high (right panel of Fig. 10), higher revenues occur at lower catch and bin thresholds. This pattern is similar to that for the SOC. A bin and catch threshold strategy that yields close to maximum revenue at one catch rate may not perform as well at other catch rates. Other bin and catch threshold strategies perform well regardless of the catch rate. Strategies that are robust to mean catch rates are favorable strategies because fishers typically will not know the mean catch rate before starting to fish. The region of overlap between the flat-topped portions of the reward surfaces identifies strategies that are robust to changes in mean catch rate.

**Fig. 10.** Reward surfaces for rule of thumb decisions consisting of combinations of bin and catch thresholds for three different mean catch rates. The lowest contour line is where mean daily net revenue is zero. Succeeding contours show the net revenue for the rule of thumb strategy as a proportion of the mean net revenue for the SOC. The highest contour line is where the daily net revenue is 90% of the mean net revenue for the SOC.

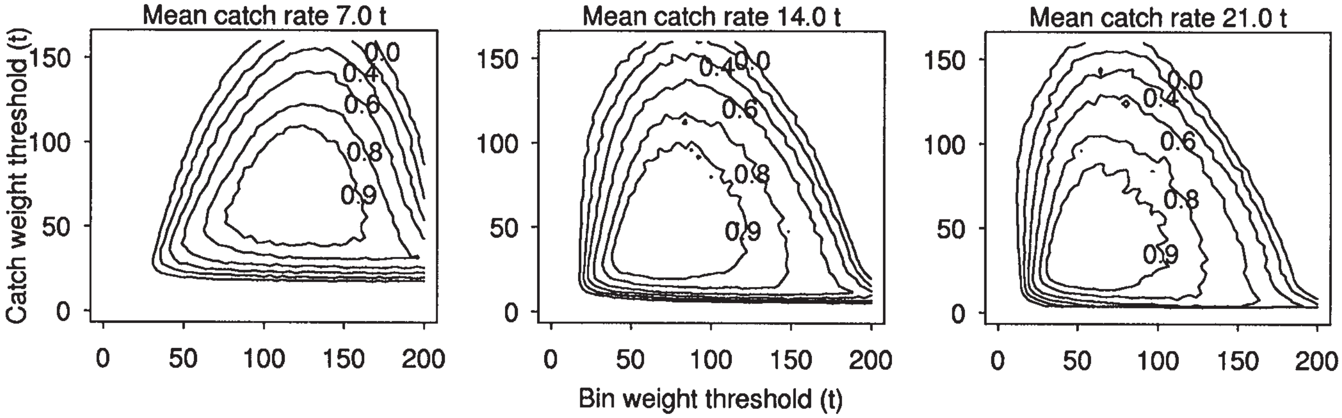


For example, a bin and catch threshold strategy of a bin threshold of 100 t and a catch threshold of 50 t performs well regardless of the mean catch rate.

## A model with correlated catch increments

To add serial correlation to the sequence of catch increments, the technique of state augmentation was used (Bertsekas 1987). In a Markov model for serial correlation, the additional state variable required is the random catch increment in the previous time step, $w_{k-1}$. This variable is assumed to be observable by the fishing vessel, which should be possible using net telemetry. For the simple prototype presented earlier, the stochastic part of the model was a single discrete probability distribution, $\mathrm{pr}(w_k)$, giving the probability of catch increments from zero to 160 t in 4-t increments. For the serial correlation model, the stochastic component is a probability transition matrix, $\mathrm{pr}(w_k|w_{k-1})$, that gives the probability of catch increments from zero to 160 t given the catch increment in the previous time step from zero to 160 t.

Since factory trawlers are unable to make sharp turns while trawling, serial correlation in the sequence of catch increments should be similar to the spatial correlation of fish densities along straightline transects. Consequently, the transition probabilities for the correlated catch increment model were derived from the density of hake along acoustic transects during the 1992 NMFS acoustic survey. The objective in obtaining a probability transition matrix was to keep the correlation structure as simple as possible, yet still capture the most important characteristics of the hake spatial pattern as it relates to fishing strategy. To estimate the transition probabilities, hake densities for 1-nmi transect segments (15 min of towing by a factory trawler) were considered as either a part of a low-density region (<3.0 t/nmi) or a part of a hake aggregation (>3.0 t/nmi) (3.0 t was used because the mean catch rate in the fishery is about 1.5 times the mean acoustic density, and 2.0 t was the boundary between catch increments of zero and 4 t used in obtaining the discrete probability distribution). The transition probabilities were $p_{00} = 0.71$, $p_{01} = 0.29$, $p_{10} = 0.17$, and $p_{11} = 0.83$, where, for example, $p_{01}$ is the probability of a positive catch increment given a zero catch increment in the previous time step.

A Δ-distribution with mean = 14 and CV = 1.5 was generated for each row in the probability transition matrix. Then the probabilities of zero and nonzero catch were adjusted so that the transition probabilities matched those estimated from the acoustic data. Simulations of the correlated catch increment model show a random pattern of high-density aggregations surrounded by low-density regions. After the vessel decides to begin fishing, the catch increment for the first time step does not enter the net because of the time required to set the net. In the serial correlation model, the probability distribution for this initial catch increment is strictly positive (i.e., no zero catch increments) because a fishing vessel would be unlikely to initiate fishing activities without first detecting fish with the echosounder.

Because of the simple structure to the transition probabilities, the SOC has two basic patterns, a strategy when the vessel is in the low-density region and a strategy when the vessel is fishing in an aggregation. When serial correlation is incorporated into the state dynamics, the SOC changes in the following ways. First, the bin threshold increases from 64 t for a noncorrelated model (based on the stationary probability density function to the probability transition matrix) to a threshold of 76 t for the correlated catch increment model (Fig. 11). Second, when the vessel is in the low-density state, the region of the state space where the optimal control is to fish becomes smaller, so that the vessel ceases fishing operations sooner. For example, if the vessel has between 36 and 56 t in the net, and the vessel encounters a low-density region, it immediately retrieves the net. If there is less than 36 t in the net, the vessel will fish longer before retrieving the net (until ~24 t remains in the bins) because the vessel has a probability of 0.29 of encountering another aggregation. At higher levels of serial correlation, the SOC is always to cease fishing immediately if the vessel leaves the aggregation.

Forward simulations of the SOC for correlated catch increments show a catch threshold in the frequency distribution of haul weights similar to that for the random catch increment model (Fig. 12). However, a scattering of tows smaller than the catch threshold occurred when the vessel moved from an aggregation to a low-density region. Thus, serial correlation in catch increments could explain the occurrence of small tows

**Fig. 11.** SOCs for models with random and serially correlated catch increments: (A) SOC for random catch increment model; (B) SOC for correlated catch increment model when the vessel is in the low-density region; (C) SOC for correlated catch increment model when the vessel is in the high-density region.

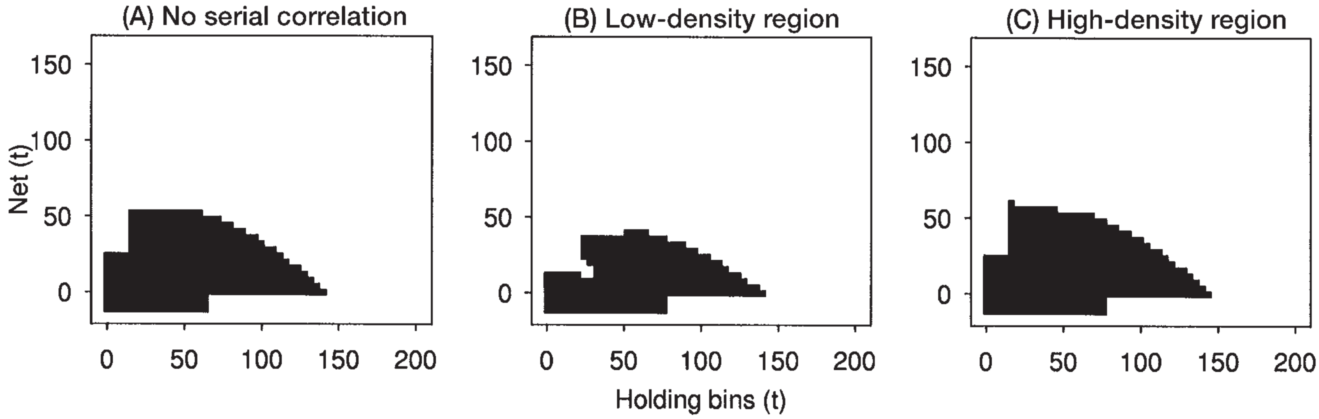


**Fig. 12.** Frequency distribution of haul weight in a forward simulation of the SOC for models with random catch increments and correlated catch increments.

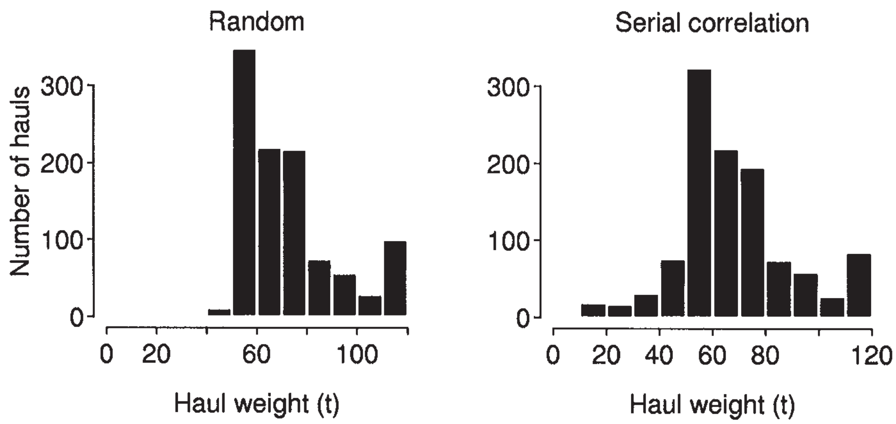


in the observer data, which did not occur in the simple prototype with random catches. A reviewer suggested that the small haul weights in the observer data may also be due to uncertainty in the estimation of the amount of fish that have accumulated in the net, so that the net is occasionally retrieved with less fish than is expected. There were also more short haul intervals with the correlated catch increment model. The vessel began fishing in the shortest possible time (60 min) in 11% of all tows in the forward simulation of correlated catch increment model, but in only 4% of tows for the model without serial correlation. This also agrees more closely with at-sea observer data, where haul intervals ≤1.5 h were common.

### Diel trends in catch rate and the effect of a ban on night fishing

The diel pattern in fishery catch rates is caused by the behavior of Pacific hake. During the day, Pacific hake form aggregations that occupy a distinct layer 10–20 m thick. This layer ranges in depth from 100 to 300 m and may be close to the bottom on the continental shelf or in midwater off the shelf. At dusk, these aggregations break up as hake disperse to begin actively feeding. Usually there is a vertical component to their movement. At night, hake can be found from the surface to 300 m with no discernable pattern of aggregation. The breakdown of hake aggregations coincides with the dusk movement of macrozooplankton, including euphausiids, the primary prey of Pacific hake. The transitions between day and night distribution patterns occur in about an hour at dawn and dusk.

To include this diel pattern in the model, a variable for the time of day (discretized into 96 states of 15-min duration) was included in the state space. To complete the model, the state dynamics and a probability distribution of catch increments for each of these possible times are needed. To model the state dynamics, the counter for the time of day is simply incremented by 1 from 1 to 96 and then returns to 1. Catch distributions from factory trawlers in the 1993 hake fishery were used to obtain the mean catch rate by time of day. From 06:30 to 19:30, a mean catch increment of 14.0 t/15 min was used, and from 20:30 to 05:30, a mean catch increment of 4.0 t/15 min was used. The dawn and dusk transitions occurred linearly in 1 h between the mean day and night catch rates. The Δ-distributions were generated using the methods described in the simple prototype section for each period. The CV and the probability of a zero catch were the same as the basic prototype and did not vary through the day.

The SOC was obtained for two models, a model where night fishing was allowed and a model where fishing was prohibited between midnight and 07:00. A ban on night fishing was easily

**Fig. 13.** Bin and catch thresholds for Markov decision process models with and without night fishing (between midnight and 07:00). The catch threshold was obtained by averaging the catch thresholds for bin levels from 16 t to the bin threshold.

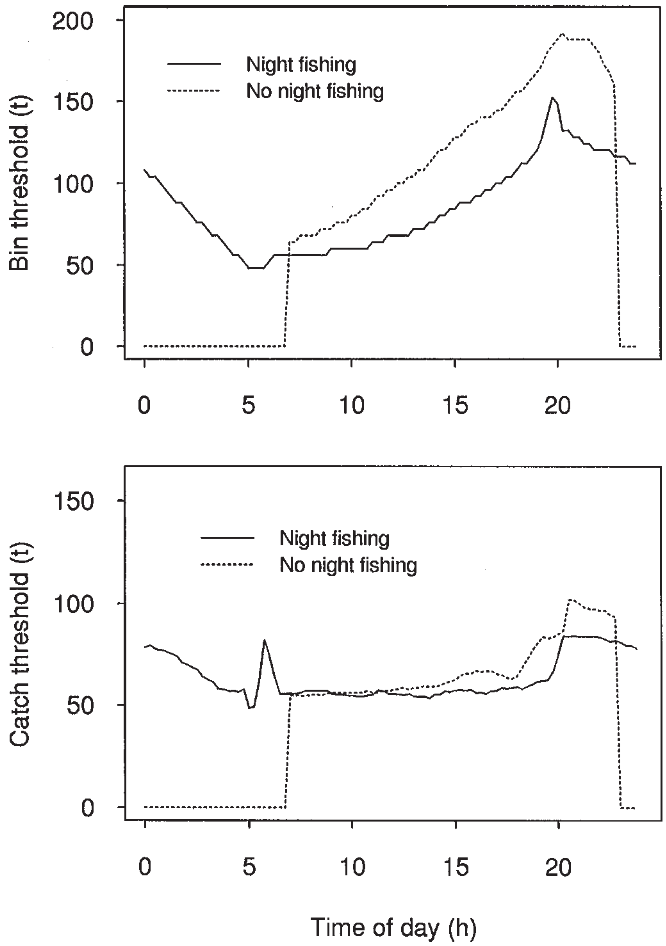


**Fig. 14.** Mean tonnes in holding bins in a forward simulation of the SOC with and without night fishing.

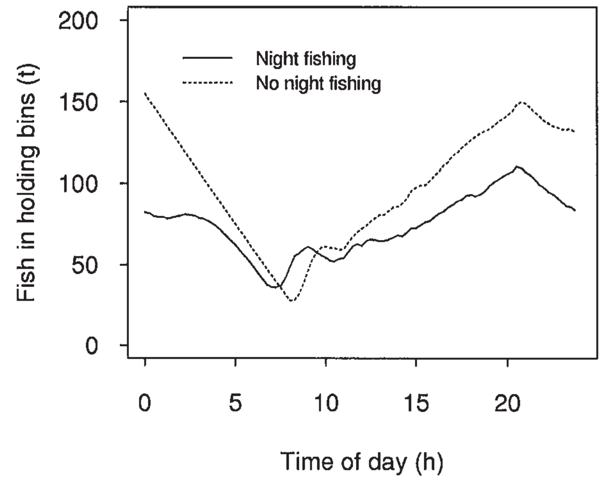


modeled by reducing the decision set to $u_k = 0$ from 23:15 to 07:00 (the additional 45 min before midnight gave the vessel time to get the net on deck). The SOC for the nonfishing state ($s_k = 3$) specifies the appropriate bin threshold by time of day. The SOC for the fishing state ($s_k = 4$) is three-dimensional, with bin level, catch level, and time of day as the three dimensions. The general pattern of the SOC at each time of day is similar to simpler models already analyzed. The top edge of the fishing region is usually flat-topped, resulting in a catch threshold that is independent of the amount of fish in holding bins. As with the simpler models, there is a notch in the fishing region where catch is high and bin level is low that provides additional protection against running out of fish in the factory. A plot of bin and catch threshold by time of day (Fig. 13) captures most of the pattern of the SOC and avoids having to show many plots to characterize the SOC. For the SOC with night fishing, the bin threshold peaks at ~150 t in the early evening (around 20:00) and then declines to a minimum at dawn. When night fishing is closed, the bin threshold is higher throughout the day and peaks at ~190 t in the early evening. The catch threshold for the model with night fishing is between 50 and 60 t during the day and then increases to ~80 t when the mean catch increment decreases at nightfall. The catch threshold without night fishing shows a similar pattern, except for a larger increase (to ~100 t) at nightfall.

In forward simulations of the SOC with and without night fishing, the changing catch and bin thresholds through the day cause an increasing trend in the mean amount of fish onboard starting from a minimum at dawn and reaching a maximum at dusk (Fig. 14). During the night, the mean amount of fish in holding bins decreases. The rate of decrease is more rapid and sustained when night fishing is prohibited. One interesting result is that even when fishing is allowed through the night, the tendency is for the fishing vessel to curtail its fishing activities during the period when the catch rate is low. Apparently, vessels should fish just enough at night to ensure that the fish in holding bins will supply the factory through the night to daybreak when catch rates increase again. Since more time is needed to catch an equivalent amount of fish at night, the cost per tonne of catch is higher at night than during the day.

The forward simulations of the SOC with and without night fishing and a rule of thumb strategy of a 100-t bin threshold and a 50-t catch threshold were compared with observer data from the 1992 fishery (when night fishing was prohibited) and the 1993 fishery (when night fishing was allowed). This rule of thumb strategy was also compared because in previous simulations without a diel trend in catch rates, this strategy produced a large percentage of the potential yield regardless of the mean catch rate. Models with and without night fishing result in a peak of fishing activity in the morning, a period of lower activity during the day, followed by an increase at dusk (Fig. 15A). An examination of hourly vessel activity for the 1992 and 1993 fisheries also indicates that a peak in vessel activity occurs in the morning when fishing opens (Fig. 15D). However, the peak is much higher in both SOCs. In the 1993 fishery data, the proportion of time spent actively fishing is highest in the early hours of the morning whereas the SOC predicts a decrease in fishing effort early in the morning. This suggests that vessels may not vary their bin threshold through the day as predicted by the SOC. The rule of thumb strategy followed throughout the day shows a closer match to the proportion of time spent fishing in the 1993 fishery data (Fig. 15A).

**Fig. 15.** (A) Proportion of the time actively fishing for forward simulations of the SOC with and without night fishing and a rule of thumb (ROT) strategy; (D) proportion of the time actively fishing for factory trawlers in the 1992 fishery (night fishing banned) and 1993 fishery (night fishing allowed); (B and C) mean haul weight and haul duration by time of day for forward simulations of the SOC with and without night fishing and an ROT strategy; (E and F) mean haul weight and haul duration by time of day for factory trawlers in the 1992 and 1993 at-sea fisheries.

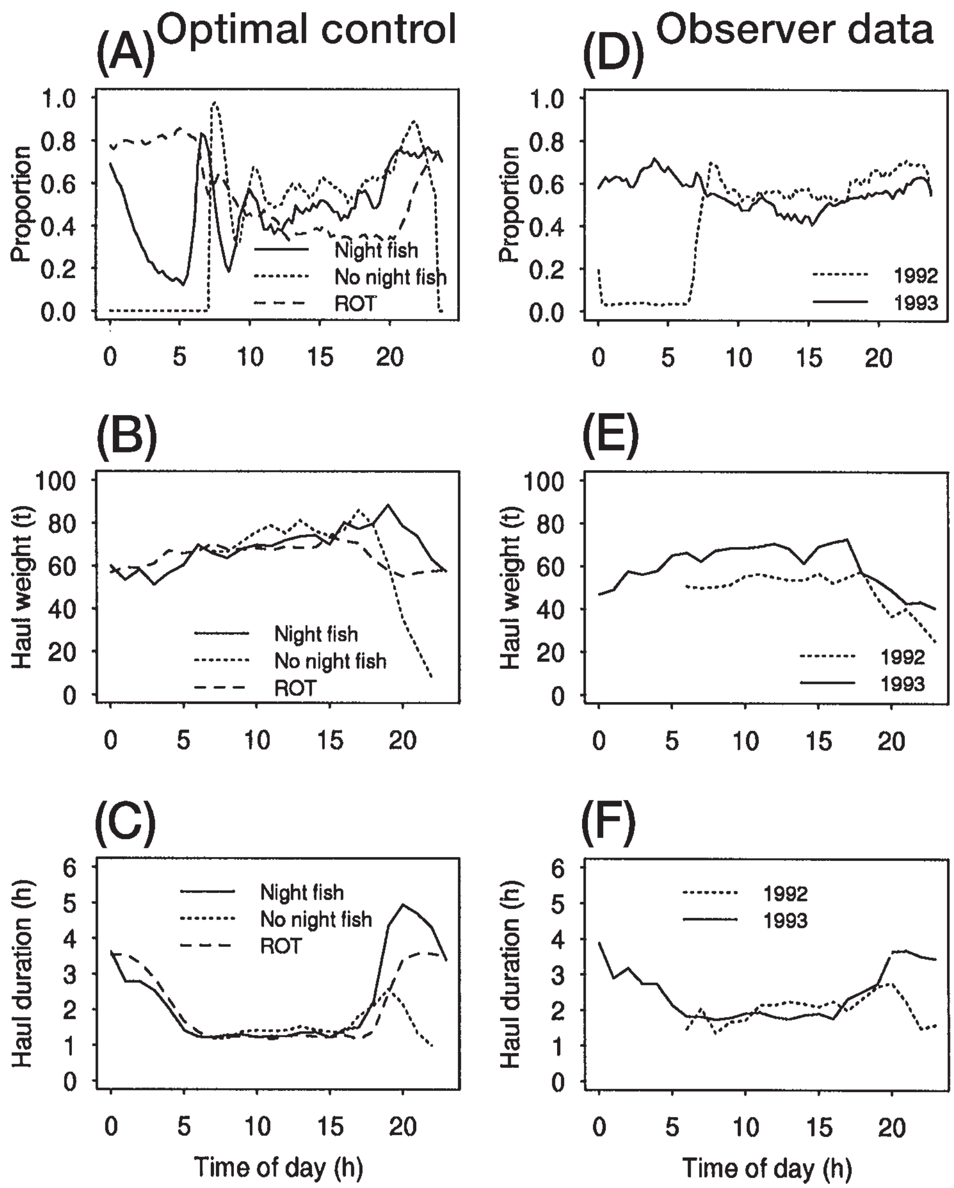


The diel trends in haul weight for the SOC (with and without night fishing) are generally similar to the fishery data for the corresponding conditions (Figs. 15B and 15E). Mean haul weight is highest immediately before dusk because it is advantageous to wait until catch rates decline before retrieving the net. When night fishing is allowed, both the SOC and the fishery data show a decline in haul weight from dusk to midnight and then a gradual increase through the day. When night fishing is banned, the pattern is similar for both the SOC and the fishery data, except for a more rapid decline in haul weight after dusk due to vessels retrieving their nets to comply with the midnight closure. Both the simulated SOCs and the observer data show an increase in haul duration at night (Figs. 15C and 15F). When night fishing is allowed, haul duration for the SOC increases substantially at night whereas the corresponding fishery data show less of an increase. The 1993 fishery data more closely resemble the rule of thumb strategy than the SOC. When night fishing is banned, haul duration begins to increase at dusk but then declines as vessels retrieve their nets to comply with the midnight closure. This pattern is seen in both the SOC and the 1992 fishery data.

The mean revenue per day in simulations of vessels following different strategies depicts the change in profitability of a factory trawler when night fishing is prohibited. When night fishing is prohibited, a vessel following the SOC can still attain a large fraction (~98%) of the potential daily net revenue (Table 3). The most significant difference when night fishing is banned is that the mean discard increases from 1.7 to

**Table 3.** Mean daily costs and revenue for a 50 000 step forward simulation of the SOC when the mean catch increment follows a diel pattern. (The rule of thumb (ROT) strategy is to use a 100-t bin threshold and a 50-t catch threshold throughout the day.)

| | SOC with night fishing | SOC without night fishing | ROT with night fishing | ROT without night fishing |
|---|---|---|---|---|
| Total catch per day (t) | 383.8 | 381.5 | 380.2 | 321.5 |
| Gross revenue per day ($) | 99 780 | 99 199 | 98 840 | 83 584 |
| Costs per day ($) | 80 087 | 79 936 | 80 557 | 79 572 |
| Discard penalty | –278.6 | –1 094.2 | –376.4 | –184.7 |
| Discard per day (t) | 1.7 | 6.8 | 2.4 | 1.2 |
| % discard | 0.45 | 1.76 | 0.62 | 0.36 |
| Oversize catch penalty | –813.4 | –998.7 | –648.6 | –568.5 |
| Net revenue per day ($) | 19 693 | 19 263 | 18 282 | 4 012 |

6.8 t/day. This is because the vessel must fill its holding bins closer to capacity before fishing is closed for the day, running the risk of catching too many fish. However, discard rates still remain low, less than 2%, and a manager might be willing to accept this higher level of discard to reduce salmon bycatch rates. The rule of thumb strategy works well with night fishing (93% of the mean daily revenue of the SOC), but the performance of this rule of thumb strategy degrades substantially when night fishing is closed (21% of mean daily revenue of the SOC). This suggests that the simpler rule of thumb strategy may be adequate when night fishing is allowed, but that the vessel may be compelled to adopt a diel strategy of stockpiling fish during the day when night fishing is banned.

## Discussion and conclusions

The analysis of the simple prototype, where the random catch increment was generated from a single probability distribution, showed that the optimal controls generally consist of a bin threshold and a catch threshold. The bin threshold signals the vessel to start fishing and the catch threshold signals the vessel to stop fishing. Although the optimal controls displayed some sensitivity to key parameters of the model, basic characteristics of the optimal controls were unchanged. Further analyses showed how vessels adjust these controls in response to changes in the mean and variance of the sequence of catch increments entering a net in the water. An exploration of rule of thumb strategies showed that the reward surface is flat in the region of the optimal control. As a result, fishing vessels should have considerable flexibility in selecting simple strategies that generate nearly as much net revenue as the optimal controls. Observer data on haul size, haul duration, and haul interval recorded during the 1993 at-sea fishery were consistent with model results showing that fishing vessels should manage their fishing operations using bin and catch thresholds to maximize net profits. The wide range of catch and bin thresholds evident in the data is consistent with the conclusion obtained by modeling rule of thumb strategies that vessels have considerable flexibility in selecting strategies.

Two elaborations of the simple prototype, a model with serially correlated catch increments and a model with diel variation in catch rates, presented the vessel with a more complex and realistic environment. The resulting optimal controls accommodated that complexity in intuitively reasonable ways. The model with serial correlation showed that the vessel should have a lower catch threshold in the low-density region, such that the net would be retrieved with fewer fish in it than if the vessel remained in a high-density region throughout the tow. The model with diel variation in catch rates caused the vessel to adjust bin and catch thresholds through the day to stockpile fish during the daylight hours when the catch rate is high and then to cut back on fishing at night when the catch rate is low and fishing is less profitable. With a ban on night fishing, vessels accumulated fish to a greater extent during the day, but daily net revenue did not decline significantly.

Diel patterns in the fishery data for 1992 and 1993 were most consistent with a rule of thumb strategy consisting of a fixed bin and catch threshold rather than the optimal strategy of stockpiling fish during the daylight hours. However, the performance of this rule of thumb strategy degraded substantially in simulations when night fishing was closed. The ability to accommodate a daily shutdown of fishing operations would depend on the bin capacity of the vessel. Vessels with smaller bin capacities may have been be more adversely affected by the night closure than vessels with larger bin capacities. Issues of fairness and equity should be addressed when regulations such as the night closure in 1992 are proposed. These issues were not obvious prior to conducting this research.

Since the model was intentionally kept simple to emphasize the essential characteristics of the time-allocation problem on factory trawlers, fishing vessels may have more behavioral flexibility than the model allows. One example of this flexibility is “short wiring,” where a vessel partially retrieves the net so that it is no longer actively fishing. The vessel maintains trawling speed so the fish already captured cannot escape. In the context of the Markov decision process model, allowing vessels to short wire expands the decision set for $s_k = 1$, when the vessel is retrieving the net. Rather than being compelled to move to state $s_k = 2$, as the model now requires, the vessel would have the option of remaining in this state. Limited personal observations suggest that short wiring is a fairly common strategy among factory trawlers in the Pacific hake fishery. Short wiring is helpful because it allows the vessel to reserve its catch until there is room in holding bins to receive the fish. Experimentation with a model where fishing vessels were allowed to short wire suggested that vessels would only short wire when retrieving the net would result in discard, i.e., when $x_k + c_k > x_{max}$. This suggests that short wiring may help vessels manage their discard.

The model developed in this paper also assumes that the factory processing rate is constant. In reality, a vessel can adjust its processing rate by shutting down a production line or reducing the number of factory workers. There may also be a random component to the processing rate due to unpredictable equipment breakdowns. Machinery for surimi processing operates most efficiently when supplied with a continuous flow of fish, and there may be significant startup costs when restarting a production line after a long shutdown. A strategy to reduce the processing rate would only be used when catch rates are low and there is some risk of running out of fish. Because vessels in the at-sea hake fishery are all competing for the fishery-wide quota, vessels with low catch rates would be strongly motivated to search for higher densities of fish rather than adjust their processing rate downwards. Consequently, the ability of the vessel to modify its processing rate is unlikely to have a large influence on the general results obtained in this paper.

An additional assumption of the model is that product value does not depend on the length of time the fish have been held. To model changes in product value over time would require additional state variables and decision options and would substantially increase the complexity of the SOC. For example, the vessel may decide to discard fish held too long to produce top-grade surimi, a decision that is not now an option in the model. Although holding time is an important factor in determining the grade of surimi, potential decreases in product value probably do not play a significant role in decision-making at the time scales addressed in the model. Surimi grades are based on color and gel strength. The gel strength of Pacific hake increases for up to 6 h after landing due to stiffening of the muscle tissue, and in pollock fisheries, the fish are typically “aged” for several hours before processing. Price differences for surimi of different grades fluctuate depending on the market, but typically the price range between top-grade surimi and second-grade surimi is 5–10%.

Pacific hake held for less than about 10 h in holding bins, or for less than 24 h in refrigerated seawater tanks, produce top-grade surimi (Gregory Peters, Oregon State University Seafood Laboratory, Astoria, Oreg., personal communication). About one third of the at-sea fleet has refrigerated seawater tanks; the rest of the fleet stores unprocessed fish in holding bins. For the simple prototype, 1.0% of the fish were held longer than 10 h before processing in forward simulations of the optimal control (mean holding time = 3.6 h). For the diel model with night fishing, 1.4% of the fish where held longer than 10 h (mean holding time = 4.3 h). However, with a night closure on fishing, 9.5% of the fish are held longer than 10 h before processing. This suggests that under normal conditions (i.e., without a night closure or an equipment breakdown), a vessel would be unlikely to lose revenue due to a decrease in product value. Although a ban on night fishing could reduce product value, it is not obvious how a vessel could change its strategy to improve product value, since all the fish processed during the 7-h night closure must be caught earlier in the day.

Gillis et al. (1995*b*) described three types of discarding on commercial fishing vessels: capacity-discarding, exclusion-discarding, and high-grading. Capacity-discarding is discarding because the vessel has no room in holding bins for the catch and is the kind of discard that occurs in the model developed in this paper. Exclusion-discarding is discarding when the species or the size of the fish has no economic value. Since juvenile hake occur mostly off California, south of where the fishery operates, exclusion-discarding of undersized hake is rare. However, all nonhake species are typically discarded. High-grading is the discard of marketable fish to leave space for higher valued fish. In the Pacific hake fishery, high-grading would be unlikely to occur because all marketable hake can be processed into surimi, which has the same value no matter the size of the fish used to produce it.

Gillis et al. (1995*b*) contended that only with high-grading does discard occur at the discretion of the fishers. However, they are concerned only with decision-making that occurs after a haul is brought onboard. The operational models developed in this paper suggest that discarding should be viewed within the context of the entire set of decisions that leads to the discard. The analysis of these models indicate that capacity-discarding can occur as an indirect result of decisions made before the discard occurs, but that influence the probability that discard will occur. Factory trawlers can maximize net revenue by keeping the factory continually supplied with fish. If they start to fish when they already have a substantial amount of fish in bins (i.e., a strategy with a large bin threshold), there is a high probability that they will be able catch enough fish to refill the bins before they become empty. However, vessels following this strategy would have a higher risk of capacity-discarding, since they may occasionally catch so many fish that they have no space in the bins to hold them. The strategy adopted by the fishing vessel can put it in a situation where discarding is unavoidable. If the penalty for discard were to be increased, as would occur if the fishery were managed with ITQs and discard was subtracted from the vessel’s quota, vessels would adopt a strategy to avoid those situations. However, some discard during the fishing season may be unavoidable due to the stochastic nature of fishing and the spatial and temporal overlap between economically valuable fish and those species with little or no economic value.

Although observers in the hake fishery do not record the type of discard, they do estimate the total discard by a vessel. Discard is difficult to monitor by observers because it is episodic in nature and can occur in more than one location on the vessel. The estimated discard of hake in the factory trawler fleet declined from 5.5% in 1992 when night fishing was banned to 3.8% in 1993 when night fishing was allowed. This decline in discard corresponds to the model prediction of lower discard when night fishing is allowed. However, this decrease in discard from 1992 to 1993 occurred within the context of a trend of decreasing discard since the large-scale fishery started in 1991. Consequently, it is impossible to separate the effect of the ban on night fishing on discard from the trend of increasing efficiency as the factory trawler fleet becomes more proficient at fishing for hake. An intriguing conclusion from the night fishing analysis is that a management action designed to control one problem, the bycatch of salmon, can have consequences that are difficult to predict and may be as undesirable as the problem the management action was intended to correct. A potential benefit of the modeling techniques developed in this paper is the ability to predict consequences of management actions before implementing them. However, the data needed to parameterize and test operational models of fishing are not routinely collected. As the regulatory environment under which fisheries operate becomes increasingly restrictive, there is an ongoing need to monitor fishing behavior, develop new models, and test the predictions of those models.

The Markov decision process model developed in this paper is a projection of the general foraging problem of a fishing vessel into one-dimensional time. While the fishing vessel contends with time-allocation problems, it also has to contend with the spatial aspects of foraging, as in deciding where to start the next haul. When the vessel has found a large-scale aggregation of hake and can obtain an acceptable catch rate, the time-allocation concerns are probably in the forefront, since the vessel can quickly reach any location within the aggregation to start the next haul. Approximately 90% of all tows in the 1993 at-sea fishery were located within 18 nmi of the immediately preceding tow, 1–1.5 h of running time by the vessel, suggesting that over short periods the spatial aspects of foraging do not act as constraints. When longer periods are involved, the spatial aspects of foraging cannot be ignored. Decisions associated with patch selection (i.e., spatial aspects of foraging) are a step higher in the hierarchy of decisions from the operational models developed in this paper. At this level, the time scale ranges from 1 day to 1 week, and a new set of issues must be addressed. Fine-scale strategy (the operational aspects of fishing) affects decision-making at larger scales by linking fish abundance, as experienced by the fishing vessel as a catch rate, to factory production.

Decision-making at the patch selection level can be separated into two questions: “when to leave” and “where to go.” The “when to leave” question evaluates whether the abundance of fish within the patch can support the processing capacity of the vessel. The model developed in this paper indicates that vessels saturate rather quickly when catch rates increase above the processing rate. When the catch rate is equal to the factory processing rate, the net revenue of a vessel following the SOC is close to zero. When the catch rate is 1.5 times the factory processing rate, the vessel can obtain 94% of its potential net revenue by following the SOC. At catch rates greater than twice the factory processing rate, the vessel can increase its daily surimi production by very little, suggesting that there would be little incentive for a vessel to seek higher catch rates. A catch rate between 1 and 1.5 times the factory processing rate may prompt the vessel to assess whether to continue fishing in an aggregation. This assessment could take the form of a threshold catch rate that would initiate movement to a new aggregation. The “where to go” aspect of patch selection was not addressed in this paper but is a promising area for research. Informal discussions with fishers and research on other fisheries (Eales and Wilen 1986) suggest that these decisions are based on the fishers’ knowledge of areas historically high in fish density, anecdotal information from other fishing vessels, the location of other vessels, and information gained from searching. Exploring the relative importance and interplay between these various factors is a challenging task.

Further progress in modeling fishing behavior may depend on developing alternative models of decision-making or using the techniques of other scientific disciplines. By viewing decision-making as a hierarchial set of choices, it may be possible to address problems at different spatial and temporal scales, yet keep the problems to a manageable size and relevant to actual decision-making processes. A useful alternative perspective could be gained from an anthropological investigation of how fishers view the environment and their opportunities for decision-making. The technique developed in this paper of modeling fishing behavior using rule of thumb strategies and comparing these strategies with the optimal controls is a tool for simplifying and generalizing successful strategies. Consideration of rule of thumb strategies provides a broader perspective than can be obtained by studying only the optimal solutions to decision process models.

## Acknowledgments

A draft of this paper was reviewed by Gordon Swartzman, Ana Parma, Joseph Terry, Anne Hollowed, and Martin Loefflad. Ray Hilborn and Robert Francis made many helpful suggestions. I thank Marianne Vignaux and an anonymous reviewer for their perceptive comments.

## References

Bertsekas, D.P. 1987. Dynamic programming: deterministic and stochastic models. Prentice-Hall, Inc., Englewood Cliffs, N.J.

Chambers, J.M., and Hastie, T.J. 1992. Statistical models in S. Wadsworth and Brooks, Pacific Grove, Calif.

Dorn, M.W., Nunnallee, E.P., Wilson, C.D., and Wilkins, M.E. 1994. Status of the coastal Pacific whiting resource in 1993. U.S. Dep. Commer. NOAA Tech. Memo. NMFS-AFSC-47.

Eales, J., and Wilen, J.E. 1986. An examination of fishing location choice in the pink shrimp fishery. Mar. Rec. Econ. **2**: 331–351.

Freese, S., Glock, J., and Squires, D. 1995. Direct allocation of resources and cost–benefit analyses in fisheries, an application to Pacific whiting. Mar. Policy, **19**: 199–211.

Gillis, D.M., Peterman, R.M., and Tyler, A.V. 1993. Movement dynamics in a fishery: application of the ideal free distribution to spatial allocation of effort. Can. J. Fish. Aquat. Sci. **50**: 323–333.

Gillis, D.M., Peterman, R.M., and Pikitch, E.K. 1995*a*. Implications of trip regulations for high-grading: a model of the behavior of fishermen. Can. J. Fish. Aquat. Sci. **52**: 402–415.

Gillis, D.M., Pikitch, E.K., and Peterman, R.M. 1995*b*. Dynamic discarding decisions: foraging theory for high-grading in a trawl fishery. Behav. Ecol. **6**: 146–154.

Hilborn, R., and Ledbetter, M. 1979. Analysis of the British Columbia salmon purse-seine fleet: dynamics of movement. J. Fish. Res. Board Can. **36**: 384–391.

Holling, C.S. 1992. Cross-scale morphology, geometry, and dynamics of ecosystems. Ecol. Monogr. **62**: 447–502.

Kareiva, P., Morse, D.H., and Ecceleston, J. 1989. Stochastic prey arrivals and crab spider giving-up times: simulations of spider performance using two simple “rules of thumb.” Oecologica, **78**: 542–549.

Lane, D.E. 1989. A partially observable model of decision making by fishermen. Oper. Res. **37**: 240–254.

Lilliefors, H.W. 1967. On the Kolmogorov–Smirnov test for normality with mean and variance unknown. J. Am. Stat. Assoc. **62**: 399–402.

Mangel, M., and Clark, C.W. 1988. Dynamic modeling in behavioral ecology. Princeton University Press, Princeton, N.J.

National Marine Fisheries Service (NMFS). 1996. Our living oceans. Report on the status of U.S. living marine resources, 1995. U.S. Dep. Commer. NOAA Tech. Memo. NMFS-F/SPO-19.

Pennington, M. 1983. Efficient estimators of abundance for fish and plankton surveys. Biometrics, **39**: 281–286.

Smith, T.D. 1988. Stock assessment methods: the first fifty years. *In* Fish population dynamics. *Edited by* J.A. Gulland. John Wiley & Sons, Inc., New York. pp. 1–33.

Smith, S.J. 1990. Use of statistical models for the estimation of abundance from groundfish trawl survey data. Can. J. Fish. Aquat. Sci. **47**: 894–903.

Squires, D., and Kirkley, J. 1991. Production quota in multiproduct Pacific fisheries. J. Environ. Econ. Manage. **21**: 109–126.