Cycles in Agricultural Production: The Case of Aquaculture

Larry Karp, Arye Sadeh, and Wade L. Griffin

The problem of determining optimal harvest and restocking time and levels is considered. A continuous time deterministic control problem is used to study the case where production occurs in a controlled environment. A stochastic control problem is then used to determine rules for the cultivation of *P. stylirostris* which occurs in a stochastic environment. The deterministic analog of the problem is also solved. The two solutions are used to develop a measure for the value of a controlled environment and for the value of information about the stochastic environment.

Key words: aquaculture, stochastic control, value of information.

Many agricultural and resource economics problems involve cycles. In some cases the length of growing season permits only one crop; often, however, there is the possibility of two or more crops. In these cases the date of harvest constrains the date at which the next cycle begins; the manager must balance the opportunity cost of delaying harvest with the effect of this delay on the payoff of the current cycle.

It may be possible to control the environment, as in a greenhouse or with poultry production, or the environment may be relatively constant, as with mariculture in tropical climates. In these cases it is possible to operate continuously, and cycles of harvest and restocking (or replanting) should be approximately constant. If the environment cannot be controlled, so that production takes place during only part of the year, it is important to determine how many crops to attempt to harvest in the growing season (Talpaz and Tsur). These two situations provide a very general taxonomy of agricultural production. This paper focuses on problems in aquaculture, but the techniques used and the insights obtained are applicable to a wide variety of agricultural problems.

The next section considers the case where production occurs continuously, i.e., where the environment is "controlled." Abstracting from price uncertainty, this can be modeled as a deterministic, continuous time, autonomous control problem. A harvest and subsequent restocking is described as a "jump" in the biomass. The optimality conditions are interpreted and used to place bounds on optimal harvest and restocking levels. This type of problem is usually characterized using the Faustman formula, which was developed to determine the optimal cycle in timber production (Clark, chap. 8). The innovation here is that the optimality conditions determine the restocking level as well as the harvest level.

The second situation, where the environment is uncontrolled, is modeled as a stochastic control problem and solved using dynamic programming. The specific problem studied involves shrimp farming in Texas. This is a relatively new industry, and one of the principal questions confronting it is whether to attempt two crops in a year. Stocking early in the season or leaving shrimp in the ponds past a certain date exposes the farmer to the risk of a sudden fall in temperature and a killoff of the stock. Optimal behavior under certain and stochastic weather is compared, and sensitivity studies are used to test the robustness of the results.

Larry Karp is an assistant professor, Department of Agricultural and Resource Economics, University of California, Berkeley; Arye Sadeh and Wade L. Griffin are, respectively, a research associate and a professor, Department of Agricultural Economics, Texas A&M University.

This work is a result of a research program sponsored by the Texas A&M University Sea Grant College Program supported by the National Oceanic and Atmospheric Administration, Office of Sea Grant, U.S. Department of Commerce under Grant No. 04-7-158-44108 with the Texas Agricultural Experiment Station (Technical Article No. TA-21394).

The authors would like to thank two anonymous *Journal* referees for helpful suggestions on earlier draft.

Review was coordinated by Alan Randall, associate editor.

The question of optimal stocking and harvest levels is a variation of the classical "machine replacement problem" in the dynamic programming literature (Dreyfuss and Law). This technique is well known but underexploited in production problems. The study described below is part of an attempt to develop decision-making tools for aquaculturists in the Southwest. These tools can be very useful for new industries where there has been little opportunity to learn through previous mistakes.

Both the continuous time control problem and the dynamic programming problem concentrate on the stocking/harvesting decision. Inclusion of other types of controls, e.g., feeding levels, is mentioned briefly. Both problems assume prices are known.

A Controlled Environment and Continuous Production

This section provides a simple model of harvest and stock replacement in a controlled environment under continuous production. It derives an intuitive optimality condition and characterizes the solution. The model is deterministic and assumes that the stock of fish can be described by a single state variable, biomass. For many aquaculture problems, the value of a unit of biomass depends on characteristics such as average size of the fish or shrimp. The assumption that a decision can be made on biomass alone is a serious abstraction, but the model is sufficient for the present purposes.

The biomass is given by x, which obeys the growth equation

$$(1) \qquad dx/dt = \dot{x} = f(x, u),$$

where u is a vector of controls, such as feeding rates and water temperature. In most of what follows u is suppressed because the chief concern here is the harvest and stocking decisions and not behavior during the cycle when the farmer incurs maintenance cost at the rate $c(u, x)$. The decision to harvest at the time t_i, $i = 1, 2, \ldots, n$ is modeled as a "jump" in x at t_i; equation (1) is valid only for $t \neq t_i$. At the instant before harvest, t_i^-, the biomass is x_i^-; at the instant after harvest and restocking, t_i^+, it is x_i^+. Where no ambiguity results, the subscript on x is suppressed. The present value of harvesting and restocking at t_i is

$$e^{-rt_i}\phi(x^-, x^+),$$

where r is the discount rate and ϕ is the net revenue function, excluding the maintenance cost $\int c(\ )dt$.

The value of the problem, given the initial biomass x_o at t_o is

$$(2) \quad J(x_o, t_o) = \max \sum_{i=1}^{n} \left[\int_{t_{i-1}^+}^{t_i^-} -e^{-rt}c(x, u)dt + e^{-rt_i}\phi(x^-, x^+) \right],$$

subject to (1) for $t \neq t_i$. The maximization is performed with respect to $u(t)$, t_i, x^- and x^+. Assume n is large enough so that behavior in the first cycle is insensitive to changes in n. Define the Hamiltonian

$$(3) \qquad H = -e^{-rt}c(x, u) + \lambda f(x, u),$$

where $\lambda(t)$, the costate variable, gives $\partial J[x(t), t]/\partial x$, the shadow value of the stock. Standard variational techniques yield the necessary conditions at a jump (Bryson and Ho, 3.7).

$$(4a) \qquad \lambda(t_i^-) = e^{-rt_i}\partial\phi(t_i)/\partial x^-$$

$$(4b) \qquad \lambda(t_i^+) = -e^{-rt_i}\partial\phi(t_i)/\partial x^+$$

$$(4c) \qquad -re^{-rt_i}\phi(t_i) + H(t_i^-) = H(t_i^+).$$

The notation $\phi(t_i)$, $H(t_i^+)$, and $H(t_i^-)$ refers to those functions evaluated at their arguments at t_i, t_i^+, and t_i^-. Equations (4a) and (4b) state that the shadow value of the stock before harvesting and after restocking should equal, respectively, the marginal net revenue of an additional unit at harvest and the marginal cost of an additional unit at restocking. To interpret (4c), use $\partial J/\partial t = -H(t)$, where the Hamiltonian is evaluated at the optimal control at t. Then (4c) states that the foregone interest payment on invested net revenue plus the value of delaying harvest should equal the cost of delaying the beginning of the next cycle.

Equations (3), (4a), and (4b) can be used in (4c) to obtain

$$(5) \quad r\phi(t_i) + c(t_i^-) - c(t_i^+) = \phi_{x^-}(t_i)f(t_i^-) + \phi_{x^+}(t_i)f(t_i^+).$$

The left of (5) is the "single-period" interest payments on the revenue from a cycle plus the difference between costs in the last period and the first period of a cycle. This should equal the change in net revenue due to a unit increase in stock before harvest times the single period increase in stock before harvest plus the change in net revenue due to an increase in restocking times the single-period increase in stock after restocking.

In order to proceed, restrictions are placed on ϕ and c. The most obvious restriction is that at the optimum $\phi_{x^-} \leq |\phi_{x^+}|$. For example, suppose that ϕ is linear: $\phi(\ ) = px^- - kx^+$; p is the unit price of harvested biomass and k is the unit cost of restocking. Clearly $p \leq k$, or profits are infinite. Another example is $\phi(\ ) = \phi[\delta(x^+, x^-)]$, where $\delta(\ ) = x^- - x^+$; net revenue depends only on the difference between the preharvest and postrestocking levels of biomass. This describes the situation where a percentage of the stock is left in the pond to start the next cycle.

A reasonable assumption for $c(\ )$ is $c(t^-) \geq c(t^+)$: the cost of maintaining the pond in the period before harvest is at least as great as the cost after restocking. A sufficient condition for this inequality to hold is $c = c(u)$ and $f(\ ) = g(u)h(x)$ with c increasing and convex in the scalar u and g concave. The first equality states, for example, that the cost of maintenance depends on the amount of food put into the pond, not the number of fish who eat it. The second equality states that it is possible to shift up the natural growth rate; this may be reasonable over some interval of x. The assumptions on $c(\ )$ and $f(\ )$ imply that u is increasing over the cycle, so $c(t^-) > c(t^+)$.

These two assumptions on ϕ and c, together with (5), imply that $f(t_i^-) > f(t_i^+)$: the growth rate immediately before harvest must be greater than the growth rate immediately after harvest. A realistic assumption is that $f[x, u^*(x)]$, where $u^*(x)$ is the optimal control, has a unique maximum; designate this point as x_m. Define x^* as the value $x \neq x^+$ that satisfies $f(x) = f(x^+)$. The conclusion is that $x^+ < x_m$ and $x^- < x^*$.

In order to examine the forces at work, suppose that optimal harvest takes place past the point where the growth rate is maximized: $x^- > x_m$. This assumption is made to simplify the explanation; if the assumption does not hold, the explanation is modified in an obvious way. The situation is shown in figure 1. There are three considerations that tend to make it optimal for the growth rate to be less at stocking than at harvest. The first is the positive discount rate and the fact that stocking costs are incurred at the beginning of the cycle. This encourages x^+ to be lower, which emphasizes the difference between $f(t_i^+)$ and $f(t_i^-)$. The opportunity cost of delaying harvest, which increases with the discount rate, makes it optimal to harvest sooner, decreasing x^-. For $x^- > x_m$, this also emphasizes the difference between $f(t_i^+)$ and $f(t_i^-)$. The second consideration involves $|\phi_{x^+}| - \phi_{x^-} \geq 0$; increasing the difference encourages either or both x^+ and x^- to be lower. If $x^- > x_m$, this unambiguously increases $f(t_i^-) - f(t_i^+)$. The third consideration is that $c(t_i^+) > c(t_i^-)$ means that a unit of time is more expensive at the end of the cycle than at the beginning. In order to compensate for this, the marginal productivity of a unit of time must be greater at the end of the cycle.

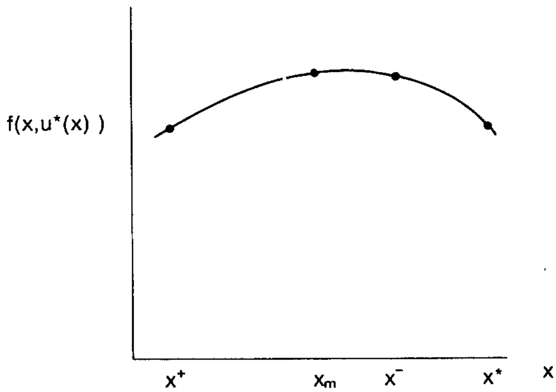

Figure 1. Optimal harvest and stocking levels with continuous production

It is worth emphasizing that $c(t_i^-) > c(t_i^+)$ is a sufficient, not a necessary, condition for $f(t_i^-) > f(t_i^+)$. It is possible that the marginal cost of a unit of time is greater at stocking than harvest $[c(t_i^+) > c(t_i^-)]$ due for example to high nursery costs, but that $f(t_i^-) > f(t_i^+)$ still holds. This is because of the operation of the first two forces mentioned above, those involving r and $|\phi_{x^+}| - \phi_{x^-}$; $f(\ )$ is only one component of the marginal product of time.

The above discussion is heuristic, rather than a rigorous comparative static analysis of r and the parameters of $c(\ )$ and $\phi(\ )$. Changing any parameter may change the entire control trajectory and thus the shape of $f[x, u^*(x)]$. However, (5) involves estimable functions and provides a testable hypothesis. The conjecture that $f(t^-) > f(t_i^+)$ can also be tested; failure of the inequality should imply high initial maintenance costs.

It is possible to write down conditions such as (4) for the case where the growth equation is stochastic using methods described by, e.g., Mangel or Brock. However, these conditions involve second-order partials of the value function J, and have no obvious interpretation. In order to include stochastics, it is convenient to use numerical methods, as described in the next section.

Aquaculture in an Uncontrolled Environment

This section studies optimal cultivation of *P. stylirostris*, a species of shrimp, in an uncontrolled environment. Shrimp mariculture is a rapidly growing new industry in Central and South America and Asia (U.S. Department of Commerce 1984). The first shrimp farm has been established in Texas and others are being planned. Since there is virtually no practical experience in Texas, there is a lack of general agreement about management practices. The principal controversy concerns whether to attempt to grow two crops a year. The shrimp are cultivated in outdoor ponds, where poor weather may result in decreased growth rates or loss of the entire stock. An analogous situation is faced by farmers in general; however, in most cases farmers have the benefit of generations of experience, so it is plausible that a near-optimal strategy has evolved. One goal of this study is to help reduce the learning cost for cultivators of *P. stylirostris* by suggesting optimal practices.

As mentioned above, harvest price depends on the quality of biomass, which depends on factors such as average weight of the shrimp. The two variables (states) used to describe the stock level are weight of the shrimp and number of shrimp. In the following, "stock level" always refers to an ordered pair, giving weight and number.

Cultivators are faced with a sequence of decisions. First, they decide what time of year to begin stocking and at what level to stock. Because weather is stochastic and consequently growth is stochastic, it is not optimal to decide in advance at what date to harvest, i.e., to follow an open loop policy. It is preferable to give harvest rules as a function of the current stock level and time of year, i.e., to follow a closed-loop policy. After the manager has decided to harvest, he decides whether to restock and, if so, at what level. A new cycle begins.

The following three sections discuss, respectively, the biological model, the economic problem, and the summary of results.

The Biological Model

The data to estimate the parameters of the following functions comes from experiments performed at TAMU Mariculture Research Center at Corpus Christi during years 1981–83. The basic equations are

$$\text{(6a)} \quad X = X_0(1 - \alpha X_0 t)^{-1}, \quad \alpha < 0$$

$$\text{(6b)} \quad W = (a_0 + a_1X + a_2X^2)(1 - e^{-kt})^3, \quad k > 0,$$

where X is number of fish, X_0 is the number at stocking, W is average weight (in grams), and t is biological time. Equation (6a) implies the differential equation $\dot{X} = \alpha X^2$. Fish are stocked in large numbers at small weight; the rate at which they die decreases as their number decreases (Pardy et al.). Equation (6b) is a generalization of the von Bertalanffy function (Clark, p. 271). The coefficient of $(1 - e^{-kt})^3$ gives the weight in the limit as $t \to \infty$. In this model, the limiting weight depends on the number of fish in the pond. This reflects the crowding that occurs in a confined environment. If the density of the fish is high, their growth is diminished.

Weekly data was available for W, but only initial and final observations were available for X. The parameters of (6b) were estimated using nonlinear least squares assuming X fixed at the harvest level. The estimated values of a_0, a_1, a_2, all of which were statistically significant, were used in a simulation of system (6). A search over possible values of α and k was performed; the parameter values at which the simulation most nearly duplicated the set of observed boundary conditions were used. It was not possible to estimate all of the parameters of (6) jointly, using, for example, pooled time-series and cross-sectional data because of the absence of weekly observations for X.[1]

The discrete approximations to the differential equations implied by (6a, b) are

$$\text{(7a)} \quad X_t = \alpha X_{t-1}^2 + X_{t-1}, \text{ and}$$

$$\text{(7b)} \quad W_t = (\dot{f}/f + 3\dot{g}/g + 1)W_{t-1},$$

where $f = a_0 + a_1X + a_2X^2$ and $g = 1 - e^{-kt}$. Equation (7b) contains t, biological time. This can be eliminated by solving (6b) for t and

[1] The parameter values used in the control model are $a_0 = 69.13$, $a_1 = -2.67 \times 10^{-3}$, $a_2 = 3.79 \times 10^{-8}$, $k = -1.09 \times 10^{-1}$, $\alpha = -1.3 \times 10^{-6}$.

substituting into (7b). The result is a system of autonomous difference equations. It is preferable to work with autonomous equations because this reduces the dimensions in the optimization problem.

Recent tank experiments (Rubino) measured the effect of pond temperature on growth rate. These results were used to choose four intervals of pond temperature. Temperatures lower than a threshold of 4° C result in killoff of the entire stock. In the other three intervals the growth parameter, k, is set at 0, 50%, and 100% of its estimated value. The observations used for the estimation occurred at (approximately) biologically optimal pond temperatures.

Pond temperature was assumed proportional to air temperature, with the proportionality factor varying over the year. The relationship between air and pond temperature was estimated using 1983 data. Data on air temperature was available over 1949–83 (U.S. Department of Commerce 1949–83). This was used to construct a probability distribution of air temperature for each week in the year. The proportionality factors were used to convert these into probability distributions of pond temperature for each week. This completes the stochastic biological model.

The Economic Problem

The essentials of the problem were discussed above. The manager is allowed to choose the number and weight to stock. Farmers can choose to stock postlarvae (approximately .01-gram [g.] animals) directly into a grow-out pond or stock them in a nursery and then transfer them to the pond as juvenile shrimp (between 1 g. and 5 g. animals). Growth in the nurseries is controlled, and the cost of raising a certain number of shrimp of a certain weight can be estimated. The results of Johns were used to extrapolate the unit costs of raising juveniles of different weights. These internal prices were treated as market prices, and the problem was modeled as if the farmer were able to purchase the shrimp he puts in the pond.

The farmer's selling price was determined using the ex-vessel price for shrimp reported by the National Marine Fisheries Service in the "Fishery Market News Reports" for 1977–81 (U.S. Department of Commerce 1977–81). The average prices were adjusted to 1983 levels. The prices depend on the shrimp weight.

The feeding rate is assumed to be 25% of body weight per day for weights less than 1.5 g. It decreases as a percentage of body weight, reaching 3% for weights over 18 g. Feeding costs were fixed at \$.53 per kilogram (Pardy et al.).

The state is an ordered pair giving weight and number of shrimp in the pond. The number of shrimp was allowed to take 70 possible values, and the weight 50 possible values, giving 3,500 combinations. These combinations are represented by Z_i, $i = 1, 2, \ldots 3{,}500$. An additional state, Z_{3501}, represents an empty pond. The transition from one state to another is governed by the autonomous version of (7a, b) with the parameter k stochastic; the state Z_{3501} is reached by harvesting or as a result of a fall in temperature which causes a killoff.

The following definitions are used: $J_n(Z_h)$ is the expected value of a one-acre pond with stock level Z_h when there are n weeks left in the year and optimal rules are used; $R(Z_h)$ is the revenue of harvesting stock Z_h; $C(Z_h)$ is cost of maintaining Z_h for one week; $S(Z_h)$ is cost of stocking an empty pond at Z_h; $P_n(h, j)$ is the probability that a pond containing stock Z_h when there are n weeks to go will contain Z_j in the next period. (This involves only probabilities resulting from weather conditions.)

The recurrence relation is

$$\text{(8a)}\quad J_n(Z_h) = \max\begin{bmatrix}\text{don't harvest: } -C(Z_h) + \sum_{j\in j^*} P_n(h,j)J_{n-1}(Z_j)\\ \text{harvest: } R(Z_h) + J_{n-1}(Z_{3501})\end{bmatrix}$$

for $h = 1, 2, \ldots, 3{,}500$, where j^* is the set of j for which $P_n(h, j)$ is greater than 0, and

$$\text{(8b)}\quad J_n(Z_{3501}) = \max_{Z_h}\left[-S(Z_h) + \sum_{j\in j^*} P_n(h,j)J_{n-1}(Z_j)\right].$$

The boundary condition is

$$\text{(8c)}\quad J_0(Z_h) = R(Z_h).$$

Equation (8c) states that in the last period, the value of stock level Z_h is the revenue obtained from harvesting Z_h. The assumption that there is a "last period" is reasonable because it would never be optimal to stock during winter when killoff would certainly occur. Equation (8b) states that the value of an empty pond with n weeks to go is equal to the maximum over Z_h of minus the cost of stocking Z_h plus the sum of the expected value of having stock size Z_j next week times the probability that the stock will be Z_j next week given that Z_h was stocked. This equation determines whether to stock an empty pond and, if so, at what level.

Equation (8a) gives rules for harvesting. If the farmer does not harvest, his expected returns consist of two parts. The first is (minus) the cost of feeding the shrimp. The second is the expectation, taken over all possible next-period stock levels, of the expected value of future returns, conditioned on the stock level in the next period. If he does harvest he receives the revenue from the sale of the current stock plus the expected value of having an empty pond in the next period. The rules for harvesting depend on weight and number of shrimp. Neither of these variables is perfectly observed, but it is considerably more difficult to obtain estimates of the number of shrimp. This type of control rule will be more useful when sampling and estimation techniques improve. When a reliable estimate of the number is not available, the history of observations on weather and weight, and equations (7a, b), together with the known initial conditions can be used to infer the current number. Methods for the control of an imperfectly observed Markov chain can also be used (Astrom). These methods involve straightforward extensions of the techniques used here. Their disadvantage is that they severely increase the dimensions of the problem, since the "state" becomes the entire history of observations.

This model has two important limitations. The assumption that there is an internal price for juveniles which can be interpreted as a market price, and used to construct the function $S(\ )$, ignores the fact that production of juveniles is not instantaneous. A more complete model would include the interaction of stochastic demand for juveniles on optimal production of juveniles. The current model is a step in this direction. The second limitation is that harvest is assumed to take one period. Manpower and equipment constraints may make this impossible in a large farm. In this case, it becomes important to stagger harvests for the various ponds. Relaxing these assumptions will result in a much larger problem. Because the state already contains 3,501 elements, the problem may grow unmanageable. It is useful to begin with a simpler situation and then add complexity.

Summary of Results

A deterministic and stochastic version of the model was used to determine the importance of uncertainty. In the deterministic version the probability density for pond temperature in each week was concentrated in the interval with the highest probability. There are several ways of comparing the results of the two versions. One method compares the certainty equivalent path in the stochastic model with the known path in the deterministic model. The certainty equivalent[2] path is obtained by using the control rules derived from the stochastic model and setting the pond temperature for each week in the interval with the highest probability. The two paths describe the behavior under identical conditions of a manager who knows what the future weather is and one who only knows the probability distribution of weather.

In the deterministic model the manager stocks 22,000 2-gram shrimp the first week of April and harvests 13,250 29-gram shrimp around the beginning of August. In the stochastic model the manager stocks 23,000 2-gram shrimp in the last week of April and harvests 14,500 26-gram shrimp in the beginning of August. In both models the managers restock 19,500 3-gram shrimp. In the stochastic model 13,250 27-gram shrimp are harvested fourteen weeks later; in the deterministic model, 12,750 29-gram shrimp are harvested sixteen weeks later. With the weather uncertain, the manager stops cultivation as soon as there is a positive probability of kill-off. At the beginning of the season he postpones initial stocking past the time where probability of

[2] This is not the standard use of "certainty equivalent." The term usually refers to the result of setting the random term equal to its expected value in the optimization process, i.e., of replacing a stochastic problem with a deterministic one. In this case, the complete stochastic problem is solved. The control rules from this problem are used in the state system which describes the evolution of the stock; this depends on stocking and harvesting decisions. Then the random element (the parameter k) is replaced by a specific realization, the one used to solve the deterministic problem.

kill-off falls to zero. During these weeks there is a positive probability that the growth parameter $k = 0$, so weight gain is minimal but mortality still occurs.

The most important result is that both models indicate that two crops are optimal. The first harvest for both models results in greater biomass than the second; for both harvests the biomass is greater in the deterministic model than the stochastic model. Stocking for the second crop uses a smaller number of larger shrimp.

Another method of comparing the two models uses the rules from (8b), which gives the decision whether to harvest. These rules are presented in matrix form. For each week there is a matrix with all possible weight/number combinations. A zero in the i, j position for n weeks to go means do not harvest if the weight is W_i and the number X_j in that week; a one in that position means harvest. This form of control rule is convenient for managers, but it is difficult to summarize and compare so much data. An attempt at this was made using a three-dimensional plot of weight, number, and weeks to go. In each matrix the positions which contain zeros tend to be clustered together, as expected. The boundary of this cluster in the weight, number plane was graphed against weeks to go, resulting in a three-dimensional surface. An example for the deterministic model is given in figure 2. The combinations of size and weight below the surface gives the region of zeros where harvest does not take place. The graph for the stochastic model is very similar. The principal difference is that the entire surface is shifted down, so it is optimal to harvest at lower weight and number combinations. This indicates increased caution in the presence of uncertainty.

In figure 2 flat regions near the beginning and end of the year indicate that it is not worth feeding the shrimp during those weeks because a kill-off is certain. The figure shows two valleys, around the twentieth and thirty-seventh week. These indicate the importance of timing. In these valleys it may be optimal to harvest at particular number-weight combinations for which it is optimal not to harvest a few weeks earlier or later. In general, the height of the surface is decreasing in number; that is, for a particular time-weight-number combination it may be optimal not to harvest; increasing the number makes it optimal to harvest. The explanation is that the death-rate is

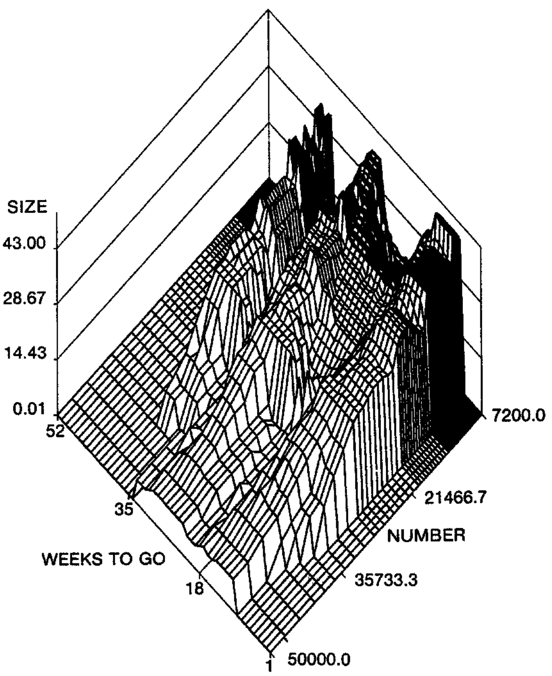

Figure 2. Region beneath surface given weight-number combination where harvest does not occur

higher when there are more shrimp in the pond, so it is less attractive to postpone harvest.

The surface in the figure is not smooth, because there are occasional stray zeros in the midst of clusters of ones. The opposite also occurs but does not affect the graph. This indicates a marginal decision. The manager is nearly indifferent between harvesting and not harvesting; the criterion selects the most profitable alternative, but an incremental change in size or weight may reverse the decision. The matrix of zeros and ones can be replaced by a matrix whose elements give the exact advantage of one decision over the other. This would be useful in sensitivity studies but is not pursued here.

The results can also be used to determine the value of a controlled environment and the value of information about the uncertain environment. Define the following three functions: J_n^{DD}, the value of an empty pond with n weeks to go given that the environment is deterministic and the optimal (for the deterministic problem) rules are followed; J_n^{SS} the expected value of an empty pond with n weeks to go given that the environment is stochastic and the optimal (for the stochastic problem) rules are followed; J_n^{DS}, the expected value of an empty pond with n weeks to go given that the environment is stochastic but the rules from the deterministic problem are used. The functions J_n^{DD} and J_n^{SS} are found by solving equations (8a, b, c) for the appropriate problems; J_n^{DS} is obtained from solving (8a, b, c) with the "max" operator replaced by the operator, "choose the alternative given by rules of the deterministic problem." Since the "deterministic" rules are suboptimal for the stochastic problem, $J_n^{DS} \le J_n^{SS}$ for all n.

The functions J_n^{DD} and J_n^{SS} reach their maximum at, respectively, $n = 39$ (the first week of April) and $n = 36$ (the last week of April), with $J_{39}^{DD} = \$3{,}930$ and $J_{36}^{SS} = \$3{,}060$. That is, being able to maintain the pond with certainty at its median temperature results in an expected increase in value of almost 30%.

The functions J_n^{DD} and J_n^{SS} are both nondecreasing in n, since it is possible to leave a pond empty for another week. The function J_n^{DS} is not (weakly) monotonic in n, since the deterministic rules prescribe stocking too early for the stochastic environment; J_n^{DS} reaches its maximum at $n = 36$, with $J_{36}^{DS} = \$2{,}530$. If the initial stocking occurs at the optimal time (the last week in April) but proceeds according to the deterministic rules, there is an expected loss of almost 17% of the optimal program. If the initial stocking occurs at the time prescribed by the deterministic rules, there is an additional expected loss of 7% ($J_{39}^{DS} = \$2{,}320$). This decomposes the loss due to using a deterministic model to describe a stochastic environment into a loss due to beginning the season too early and to using the wrong stocking/harvesting rules.

The yearly expected value of an empty pond can be used to calculate the capitalized value of an empty pond, which determines the profitability of investment. Suppose that the true environment is stochastic, but the deterministic problem is used to determine stocking/harvesting rules and to approximate the value of the pond. The approximation is upwardly biased by almost 70% ($3{,}930/2{,}320 \simeq 1.7$). In this case, a deterministic approximation would lead to considerable overinvestment in aquaculture.

The model assumed that the threshold temperature below which kill-off occurs is 4° C. To test the importance of this estimate, the threshold was increased to 9° C. This changes the probability distribution for the growth parameter. The stochastic model was re-solved, and the certainty equivalent path was obtained. The stocking date and level is unchanged, but it is optimal to leave the first crop in the pond five weeks longer. At harvest the stock consists of 13,250 29-gram animals; restocking occurs at 8,000 3-gram animals. Harvest occurs two weeks earlier than with the lower threshold, with a stock of 7,400 16-gram animals. The important result is that two harvests are still optimal, but the second crop is very small. This means that very little crowding occurs, so the shrimp grow rapidly. The expected value of the pond falls approximately 15%.

The effect of a change in consumer tastes was tested by increasing the price of large shrimp. The price of 29-gram shrimp was increased by 3%, and for each additional gram the price was increased by a further 3%, so the price of 45-gram shrimp was increased by 51%. The certainty equivalent path stocks only one crop, of 13,750 2-gram animals. These are left in the pond for twenty-eight weeks, and harvested when there are 8,200 44-gram animals. Not surprisingly, the emphasis is on quality. The expected value of a pond is \$3,170, an increase of only about 4%. The model is also insensitive to small changes in the cost of feeding. Feeding costs were increased by 10% and the stochastic model was rerun with both the original sales price of shrimp and the hypothesized higher price schedule. The certainty equivalent path was unchanged in both cases; the expected value of the pond fell by about 2%.

Conclusion

Two approaches to modeling stocking/harvesting problems in aquaculture were presented. The problem of continuous production in a controlled environment was discussed first. This differs from the standard "optimal stopping problem," since harvest and stocking levels are jointly determined. The optimality conditions were interpreted and used to place bounds on the biomass at stocking and harvesting. These bounds involve the growth rate of the biomass. Plausible sufficient conditions were given which insure that the growth rate is greater at harvest than at stocking. These results appear to be new. Given information on the maintenance function, $c(\,)$, and the net revenue function, $\phi(\,)$, equation (5) can be used as a test for optimality. The principal use of the analysis, however, is as an aid in developing intuition about the problem.

The second approach studied the cultivation of *P. stylirostris* in the Southwest, using dynamic programming. This is a new industry, still in the early stages of the learning curve. The results of this study suggest optimal stocking and harvest rates and may decrease the cost of learning.

Several important points emerged. First, the optimality of two harvests was robust to changes in the stochastic specification and in the threshold temperature of kill-off. The optimality of two harvests increases the importance of timing. Second, the time interval during which harvest is most likely to occur is approximately the same in the stochastic and deterministic models. This is evident from inspection of figure 2 and its stochastic analogue and also from the comparison of the deterministic and the certainty equivalent paths. This information suggests a way of reducing the size of the problem: allow the manager the option of harvesting only when he is likely to exercise that option. The size of the model is important because future research will increase the detail of the model by including imperfect state observations and constraints implied by the absence of a market for juveniles. Third, although the control rules in the deterministic and stochastic models are similar (compare comments above regarding fig. 2), the expected values of profits differ considerably. This suggests that a principal use of the stochastic model will be as an aid in determining investment levels in aquaculture. Deterministic models are likely to encourage excessive investment. Fourth, the model provides a convenient way of measuring the value of a controlled environment and the value of information about the uncontrolled environment. The first quantity is the difference between the expected values of an empty pond in the stochastic and deterministic problems, at the beginning of the season. The second quantity is the difference between the expected value of an empty pond in the stochastic problem and the expected value when the (suboptimal) deterministic rules are used in a stochastic environment. This quantity can be decomposed into the value of information about when to start the season and information about how to behave once the season has begun.

The results are conditioned on the accuracy of the biological model. Data problems, particularly the absence of time series on the number of shrimp, make it difficult to obtain reliable parameter estimates. Continued experiments at the TAMU Mariculture Research Center and the cooperation of economists and fish biologists should lead to improved models. The results also depend on the stochastic process for weather. Data to estimate this process can be obtained for other parts of the country, and the methods used here applied elsewhere. The longer-run goal of this study is to develop management tools that can be used in a variety of circumstances. The computer program used to solve this problem can easily be adapted by changing parameter values.

[*Received September 1984; final revision received April 1985.*]

References

Astrom, K. J. "Optimal Control of Markov Processes with Incomplete State Information." *J. Math. Anal. and Applications* 10(1965):174–205.

Brock, William A. "Introduction to Stochastic Control: A User's Manual." Unpublished notes.

Bryson, Arthur E., and Y. Ho. *Applied Optimal Control.* New York: John Wiley & Sons, 1975.

Clark, C. W. *Mathematical Bioeconomics: The Optimal Management of Renewable Resources.* New York: John Wiley & Sons, 1976.

Dreyfuss, S. E., and A. M. Law. *The Art and Theory of Dynamic Programming.* New York: Academic Press, 1977.

Johns, A. J. "Cost and Return Budgets for Penaeid Shrimp Culture Operations." M.S. thesis, Texas A&M University, 1983.

Mangel, Marc. *Decision and Control in Uncertain Resource Systems.* New York: Academic Press, forthcoming.

Pardy, C. R., W. L. Griffin, M. A. Johns, and A. L. Lawrence. "A Preliminary Economic Analysis of Stocking Strategies for Penaeid Shrimp Culture." *J. World Mariculture Soc.* 14(1983):49–63.

Rubino, A. J. "The Polyculture of Penaeus Stylirostris Stimson with Penaeus Aztecus Ives and Penaeus Accidentalis Steets." M.S. thesis, Corpus Christi State University, 1983.

Talpaz, Hovav, and Yacov Tsur. "Optimizing Aquaculture Management of a Single-Species Fish Population." *Agr. Systems* 9(1982):127–42.

U.S. Department of Commerce, National Oceanic Atmospheric Administration, National Marine Fisheries Service. Unpublished landings data. Washington DC, 1949–83.

———. *Fishery Market News Reports.* New Orleans LA, 1977–81.

———. *Shrimp Situation and Outlook.* Fisheries Development Analysis Branch, St. Petersburg FL, June 1984.