

# MODELING IN NATURAL RESOURCE MANAGEMENT

DEVELOPMENT, INTERPRETATION, AND APPLICATION

EDITED BY  
TANYA M. SHENK AND  
ALAN B. FRANKLIN

**ISLAND PRESS**

Washington • Covelo • London

## Chapter 6

.....

# Simulation Models and Optimal Decision Making in Natural Resource Management

*Michael J. Conroy and Clinton T. Moore*

Renewable resource management typically involves, at its core, an optimization problem of intrinsically dynamic systems. Viewed this way, management involves three elements: quantifying an objective, often involving multiple resources and inherent tradeoffs; defining system dynamics, including components of uncertainty and detail such as spatial resolution; and specifying a decision, including possible dynamic and spatial aspects. Williams (1989) has compared the major features of optimization and simulation, and Conroy and Noon (1996) have discussed them in relation to spatially explicit models.

### THE OPTIMAL CONTROL PROBLEM

Much, if not all, of natural resource management can be described as the formulation of a set of management actions that are expected to result in the achievement of a desired objective. An *objective function* is simply a mathematical statement that describes, in terms of measurable system attributes or *system states* (such as population abundance), values ascribed to possible out-

comes following a decision. *Optimal control* is the problem of finding a set of actions that maximize or minimize the objective function. The resulting set of actions through time is known as an *optimal policy*—frequently termed an *optimal strategy* when the system is stochastic (Dreyfus and Law 1977; Lubow 1995). In most cases, natural resource managers must consider not only the immediate consequences of a decision (such as this year's harvest) but also the decision's impact on the future condition of the resource and its value (future harvest opportunities). Given this long-term view, the objective function must now explicitly incorporate the value of future system states, which of course are not presently observable. To do this we must have a mathematical model for *system dynamics* that describes the relationship between certain management actions and these future system states given the current system state and assumptions about system behavior. Our model may be *deterministic*—that is, each combination of management actions leads to the same trajectory of the system through time—or, more realistically, *stochastic*, in which case intervening random events (such as environmental variation or demographic processes) cause the system to deviate from a completely predictable path. Often decisions are made and resource systems are monitored at discrete points in time—for example, once a year—and thus it may be appropriate to model system dynamics in discrete time as

$$\mathbf{X}(t+1) = \mathbf{X}(t) + f(\mathbf{X}, \mathbf{Z}(t), \mathbf{U}(t), t) \quad [6.1]$$

where  $\mathbf{X}(t)$  is a vector describing the system state (abundance of several species, habitat conditions) at time  $t$ ,  $\mathbf{Z}(t)$  is a vector of random variables, and  $\mathbf{U}(t)$  is a vector of management actions taken at time  $t$ . The variable  $t$  may denote any discrete unit of time. In this discussion we refer to the interval  $(t, t+1)$  as a *time step*. A general form for the objective function is

$$J = \sum_{t=t_0}^T V(\mathbf{X}(t), \mathbf{U}(t), t) + V_1(\mathbf{X}(T)) \quad [6.2]$$

where  $V$  is a function describing the value or return from management actions  $\mathbf{U}(t)$  and system state  $\mathbf{X}(t)$  and where  $V_1$  is a value given by the system state  $\mathbf{X}(T)$  at some "terminal time"  $T$ , usually taken as the distant future (100 years, for example, or even  $T = \infty$ ). The long-term nature of the decision problem is made explicit in that the objective value is obtained by summing from the present decision time ( $t = t_0$ ) to this distant future time. The optimization problem can be then mathematically formulated as

$$\begin{aligned} &\text{maximize } J \\ &\{\mathbf{U}(t)\} \in \mathbf{U} \end{aligned}$$

subject to

$$\begin{aligned} \mathbf{X}(t+1) &= \mathbf{X}(t) + f(\mathbf{X}(t), \mathbf{Z}(t), \mathbf{U}(t), t) \\ \mathbf{X}(t_0) &= \mathbf{X}_0 \end{aligned} \quad [6.3]$$

where for stochastic systems

$$J = E \left[ \sum_{t=t_0}^T V(\mathbf{X}(t), \mathbf{U}(t), t) + V_1(\mathbf{X}(T)) \right] \quad [6.4]$$

that is, the objective is now over the average of the random outcomes influenced by  $\mathbf{Z}(t)$ .

A number of mathematical approaches may be taken to solve the optimal control problem and are reviewed by Williams (1989). Here we focus on applying two methods—dynamic programming (DP) and forward simulation-optimization (FSO)—to finding optimal solutions for a simple natural resource decision problem. Briefly, optimization approaches such as DP have the advantage of guaranteeing that the resulting decision will be optimal with respect to the objective function. Dynamic programming (Bellman 1957; Dreyfus and Law 1977; Mangel and Clark 1988) works by application of the "principle of optimality," which states that "an optimal policy has the property that, whatever the initial state and decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision" (Bellman 1957). The objective form described earlier lends itself to this principle because the objective function is composed of two parts: an "immediate reward" part that starts at  $t_0$  and moves to  $t_1$  and a subsequent part that starts at  $t_1$  and moves to the terminal time  $T$ . The optimality principle says that in order for the overall strategy to be optimal (given initial system conditions  $\mathbf{X}(t_0)$ ), the strategy over this second portion must be optimal as well. Dynamic programming then applies the principle to finding the optimal strategy at  $t_0$  by working backward from  $T$ ; once we arrive at  $t_0$ , the overall strategy must now be optimal by definition. The same approach can be extended to stochastic systems by invoking the stochastic form of the state dynamics and objective function (Dreyfus and Law 1977; Lubow 1995).

Dynamic programming is an extremely powerful means of finding optimal solutions for dynamic systems. Indeed, it is guaranteed to provide optimal decisions given the assumptions of the method. It is nearly impossible, however, to include much complexity in the model of system dynamics—particularly if this involves spatial or individual animal components. And although dynamic programming allows for dynamic and stochastic systems typical of natural resource management, this method works only for systems in which there are just a few state variables and decisions. By contrast, simulation models (such as FSO) allow for virtually unlimited modeling of the details of systems, including spatial resolution and individual animal behavior. These methods may be used to seek optimal solutions, but they provide no guarantee that a solution is optimal. Moreover, complex simulation models are often of dubious reliability because of limited validation and difficulties in parameter estimation (Conroy et al. 1995).

In summary, optimization methods such as DP are known to give optimal solutions for problems, but only for simple systems where everything can be specified. Thus these methods are limited in terms of the detail of the systems that can be modeled. Simulation methods such as FSO can accommodate much more detail, but they are not guaranteed to find an optimal solution. Thus a natural question arises for resource managers: To what degree are the approaches comparable in terms of the optimal decision for a given system state and the value of the objective function given the optimal strategy?

### A COMPARISON OF METHODS

We approached this question by constructing a simple dynamic problem for which we could specify the parameters but which nonetheless contained the essential elements of dynamic decision making and a tradeoff in natural resource objectives. For this problem we applied both a backward-iteration, stochastic optimization approach (Bellman 1957) and a forward-simulation model with optimal one-time-step decision making (FSO). We then compared the optimal decision and resulting objective function—in this case involving the abundance trajectories of two vertebrate species of interest—in each approach.

#### An Example System

Although it was simple, we used a model system that was motivated by natural resource management problems encountered in forestlands of the southeastern United States—in particular our experiences at the Piedmont National Wildlife Refuge in Georgia. Our model system contained a 1000-hectare landscape of forest composed of two successional stages: early (0–40 years) and late (40–120 years). These stages might correspond to early succession following clear-cutting and regeneration of loblolly pine (*Pinus taeda*) and older stands following thinning to remove understory and mid-story vegetation. Figure 6.1a illustrates the habitat dynamics of this system between the two successional stages (that is, habitat states); the proportions indicate the 5-year rates of transition between early and late successional stages. In all of our models, habitat dynamics were completely deterministic. Habitat state dynamics can thus be represented by the expression

$$\mathbf{x}_{t+1} = \mathbf{L}\mathbf{x}_t \quad (6.5)$$

where

$$\mathbf{x}_t = \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix}$$

is a column vector describing the system habitat state, with  $x_i(t)$  representing the number of hectares in forest stage  $i$  at time step  $t$  and

$$\mathbf{L} = \begin{bmatrix} 0.875 & 0.0625 \\ 0.125 & 0.9375 \end{bmatrix}$$

defining the rate of transition between stages. The values in  $\mathbf{L}$  are purely a function of the length of time spent in each stage (40 and 80 years, respectively) and the length of a time step (5 years). Thus every 5 years  $5/40 = 0.125$  of the early stage moves to late with  $35/40 = 0.875$  remaining in the early successional stage. Likewise  $1/16 = 0.0625$  of the late successional stage reverts to early succession with  $15/16 = 0.9375$  remaining behind. The columns always sum to 1 because of area conservation (that is, the total acreage of 1000 hectares does not change through time). The acreage in each stage after each 5-year transition is the total of the acreage remaining in that stage from the previous time step plus the acreage moving into the stage. For example, given a distribution of acreage at time  $t = 4$  of

$$\mathbf{x}_4 = \begin{bmatrix} 750 \\ 250 \end{bmatrix}$$

the state at time  $t = 5$  (that is, in 5 years) will be

$$\mathbf{x}_5 = \mathbf{L}\mathbf{x}_4 = \begin{bmatrix} 0.875 & 0.0625 \\ 0.125 & 0.9375 \end{bmatrix} \begin{bmatrix} 750 \\ 250 \end{bmatrix} = \begin{bmatrix} 0.875(750) + 0.0625(250) \\ 0.125(750) + 0.9375(250) \end{bmatrix} = \begin{bmatrix} 671.875 \\ 328.125 \end{bmatrix}$$

with  $671.875 + 328.125 = 1000$ .

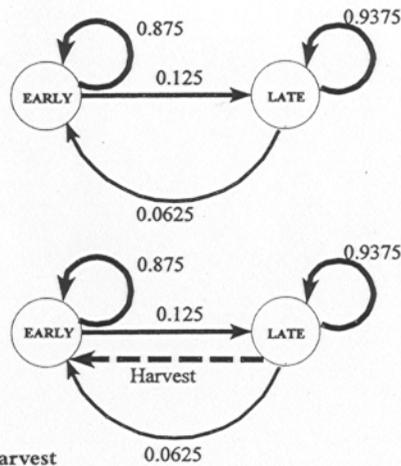
The single decision to be made, every 5 years, is what proportion of the late successional stage to harvest in five evenly spaced levels from 0 to 1; there was no harvest in the early stage. The effect of harvest on habitat dynamics is to immediately transfer the indicated proportion of late-stage habitat to the early stage with succession rates continuing as before (Figure 6.1b). Thus the system of equations representing succession dynamics is modified to include harvest

$$\mathbf{x}_{t+1} = \mathbf{L}(\mathbf{H}_t\mathbf{x}_t)$$

where

$$\mathbf{H}_t = \begin{bmatrix} 1 & h_t \\ 0 & (1 - h_t) \end{bmatrix}$$

and  $h_t$  is the proportion of the later stage that is harvested, again with no harvest occurring in the early stage. For example, if  $h_t = 0.50$  and the ini-



a. Succession only

b. Succession plus harvest

Figure 6.1. Schematic of habitat dynamics for a simple forest management problem. (a) Model incorporating succession only. (b) Model incorporating succession plus harvest of older stage.

tial stage is defined as in the previous example, the distribution of acreage at  $t = 5$  is

$$\begin{aligned} \mathbf{x}_5 &= \mathbf{L}(\mathbf{H}_4 \mathbf{x}_4) = \begin{bmatrix} 0.875 & 0.0625 \\ 0.125 & 0.9375 \end{bmatrix} \begin{bmatrix} 1 & 0.5 \\ 0 & 0.5 \end{bmatrix} \begin{bmatrix} 750 \\ 250 \end{bmatrix} \\ &= \begin{bmatrix} 0.875 & 0.0625 \\ 0.125 & 0.9375 \end{bmatrix} \begin{bmatrix} 750 + 0.5(250) \\ 0.5(250) \end{bmatrix} \\ &= \begin{bmatrix} 0.875 & 0.0625 \\ 0.125 & 0.9375 \end{bmatrix} \begin{bmatrix} 875 \\ 125 \end{bmatrix} \\ &= \begin{bmatrix} 773.4375 \\ 226.5625 \end{bmatrix} \end{aligned}$$

We represented population response for the two hypothetical species by species-specific source/sink models (Pulliam 1988; Conroy and Noon 1996). For species 1, early forest was source habitat; for species 2, late forest was source. Upper limits in the source and sink habitats were 0.1 animal/hectare and 0.5 animal/hectare, respectively. Species that could be associated with these respective source and sink habitats for a southeastern pine system might include Carolina wren (*Thryothorus ludovicianus*) for early successional stage and red-cockaded woodpecker (*Picoides borealis*) for late successional stage. The population abundance for each species is determined both by the amount of suitable habitat (that is, acreage in each successional

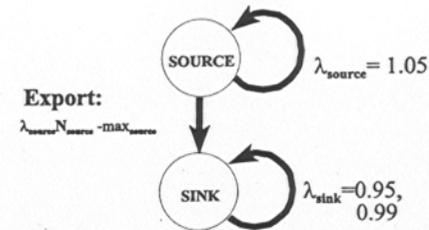


Figure 6.2. Schematic model of source/sink population dynamics. Each species increases at the rate of  $\lambda = 1.05$  in its source habitat (defined as early forest stage for species 1, late stage for species 2) with a ceiling density specified as 0.1 animal/hectare of source. When density reaches this ceiling, all surplus animals are assumed to emigrate to the sink habitat—up to the ceiling density in the sink of 0.5 animal/hectare of sink (defined as late forest stage for species 1, early stage for species 2). Species decrease at the rate of  $\lambda = 0.95$  for species 1 and  $\lambda = 0.99$  for species 2 in sink habitats.

stage) present and by the ability of each species to export excess animals from source habitats to sink habitats. For a single species (assuming no density limitation in the sink habitat), population abundance in the source and sink habitats at time  $t + 1$  depends on whether the size of the source population,  $N_{\text{source}}(t)$ , is expected to exceed the saturation abundance  $N_{\text{source}}^*$  at time  $t + 1$  given the finite rate of growth of the source population,  $\lambda_{\text{source}}$ . If  $\lambda_{\text{source}} N_{\text{source}}(t) < N_{\text{source}}^*$  (future population does not exceed saturation abundance), then  $N_{\text{source}}(t + 1) = \lambda_{\text{source}} N_{\text{source}}(t)$  and  $N_{\text{sink}}(t + 1) = \lambda_{\text{sink}} N_{\text{sink}}(t)$ , where  $\lambda_{\text{sink}}$  and  $N_{\text{sink}}(t)$  have analogous meanings to  $\lambda_{\text{source}}$  and  $N_{\text{source}}$ . Otherwise (future population exceeds saturation abundance), then  $N_{\text{source}}(t + 1) = N_{\text{source}}^*$  and  $N_{\text{sink}}(t + 1) = \lambda_{\text{sink}} N_{\text{sink}}(t) + \lambda_{\text{source}} N_{\text{source}}(t) - N_{\text{source}}^*$ . However, because we also enforce a destiny limit in the sink habitats (0.5 animal/hectare),  $N_{\text{sink}}(t + 1)$  can never exceed a saturation abundance  $N_{\text{sink}}^*$ . We assume that each species' growth rates ( $\lambda$ ) are uniform within the source or sink habitats. For both species under consideration we set  $\lambda_{\text{source}} = 1.05$ ,  $\lambda_{\text{sink}} = 0.95$  for species 1 (affinity for early forest stages), and  $\lambda_{\text{sink}} = 0.99$  for species 2 (later forest stages), resulting in greater source-to-sink dispersal rates for the latter species than the former (Figure 6.2).

We introduced random variation into this model by using a discrete log-normal distribution to model variation in  $\lambda$ , as

$$\lambda_i = \bar{\lambda} (1 + z_i \text{CV}) \quad [6.6]$$

where  $\lambda_i$  is a value of  $\bar{\lambda}$  (as in source habitats),  $\bar{\lambda}$  is the mean (1.05) of the distribution from which  $\lambda_i$  is taken,  $z_i$  is a standard normal deviate, and CV is the coefficient of variation on the logarithmic scale. We took discrete values of  $z_i$  as  $\{-2.367, -1.15, 0, 1.15, 2.367\}$  corresponding to approximate

1, 25, 50, 75, and 99 percent quantiles of the standard normal distribution, and  $CV = 0.10$  for all runs. Values of  $z_i$  were not different between species for source or sink growth. That is, a value of  $z_i$  drawn for species 1 source growth was also used as source growth for species 2, and a different  $z_i$  drawn for sink growth was applied to both species.

### Optimization

We defined the objective for optimization as the 100-year abundances for each species, weighted by a value for each species, and penalized for quasi-extinction of either species:

$$J_T = [N_{1T}(v_1 + \frac{N_{2T} - N_2^{\min}}{N_2^{\min}}) + N_{2T}(v_2 + \frac{N_{1T} - N_1^{\min}}{N_1^{\min}})] \quad [6.7]$$

where  $T$  is the terminal time period under investigation,  $v_1$  and  $v_2$  are relative values for each species, and  $N_1^{\min}$  and  $N_2^{\min}$  are quasi-extinction levels for each species. In our analyses we set the values equal for both species ( $v_1 = v_2 = 1.0$ ) and set quasi-extinction levels to 10 ( $N_1^{\min} = N_2^{\min} = 10$ ). Given Equation (6.7), if either species declined to extinction the resulting objective value is zero. We evaluated the objective function (Equation 6.7) at  $T = 20$  five-year time steps—that is, after 100 years.

For both optimization approaches, we evaluated the expected value of the objective for each state with expectations taken across the random  $\lambda_i$  outcomes. The optimal decision was the one that maximized this expected value. With DP, the optimal decision was obtained by backward iteration from the optimal terminal value to the present—that is, 20 iterations or 100 years. Five state variables described the status of the modeled system at any point in time: amount of early forest habitat (amount of late forest is known by subtraction) and abundances of both species in respective source and sink habitats. We discretized each of the state variables into 11 levels; thus the strategy yielded a decision for each of the  $11^5$  state combinations. We used Program ASDP, a successor to SDP (Lubow 1995) to perform these computations.

To illustrate the backward-iteration procedure, start with the terminal state of the system. In our case, this is the terminal distribution of habitats between successional stages and the abundance of animals of each species in the sources and sinks. The objective function is then evaluated for each combination of the random  $\lambda$ —in this case illustrated (Figure 6.3) by the five possible values for the source  $\lambda$ . The expected value for each decision is evaluated and the optimal decision selected. This process is repeated for the previous time step and continued until the present time is reached—resulting in an optimal state-specific and time-specific decision strategy (Figure 6.3).

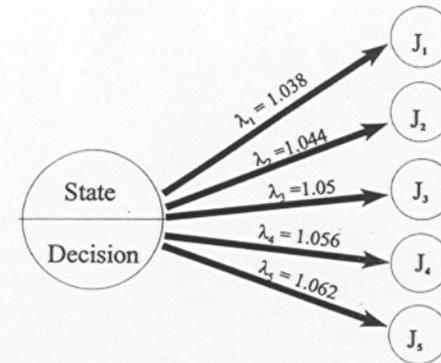


Figure 6.3. Schematic representation of decision making under dynamic programming (DP). The return ( $J$ ) on the objective function is evaluated from the state at time  $t = T$  (the terminal state, 100 years in the example) for each combination of random outcomes (in this case illustrated by five possible  $\lambda_i$  for a source population). The optimal state-specific decision is taken as the decision that maximizes the expected value of the objective. Conditioned on this decision, the process is repeated at the previous time step for all possible combinations of states and decisions. The resulting decisions at each previous time, including the current time ( $t = 0$ ), are optimal with respect to the state dynamics and objective function.

With FSO, we began at the present time with a given combination of initial habitat and population states. We evaluated the expected value of the decision for all possible decisions and selected the decision resulting in the optimum for one time step—that is, the best value of the objective over the next 5 years. We then selected random values of  $\lambda$  for the source and sink populations and, implementing the optimal decision, used forward simulation to move to the next time. Figure 6.4 illustrates the basic steps in the FSO algorithm. Given an initial state, then: the objective values for each possible decision are averaged across the random  $\lambda$ ; the optimal decision is selected, resulting in the highest expected value of the objective function; random  $\lambda$  are selected to move the system forward one time step; and the process is repeated at the next time step. We used a stochastic forward-optimization procedure written in SAS PROC IML (SAS Institute 1990) to perform these computations.

### Simulation of Strategies

We compared the FSO and DP strategies by simulation starting from a range of initial state conditions (243 combinations of habitat and initial population size for each species in each habitat). For each initial state com-

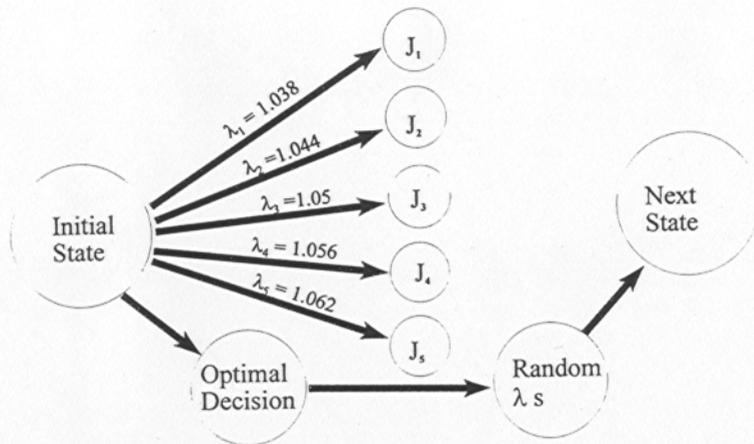


Figure 6.4. Schematic representation of decision making under forward simulation-optimization (FSO). Beginning at an initial state, the objective function return ( $J$ ) is evaluated for each combination of random outcomes (again illustrated by five possible  $\lambda_i$  for a source population) and the decision that is optimal (provides the maximum expected value of the objective) over one time step is selected. Given this decision, random values representing a source  $\lambda$  and a sink  $\lambda$  are selected, which are used with the model of system dynamics to simulate a new state of the system. The process is repeated until the terminal time ( $t = 100$  years) is reached.

bination we used identical realizations of random outcomes for  $\lambda_i$  in order to specify the population's trajectory to the next time step conditional on the habitat conditions (influenced by the strategy selected) and the initial population state. We continued this process for 20 iterations, or 100 years, and recorded the value of the objective at the terminal time. For each simulation we kept track of the initial ( $t = 0$ ) harvest decision as being relevant to initiating an "optimal" 100-year strategy as well as the objective values at  $T = 20$  resulting from each strategy. We then compared the strategies by computing

$$\frac{J_{20}(\text{DP}) - J_{20}(\text{FSO})}{J_{20}(\text{DP})} \quad [6.8]$$

to represent the relative difference of the DP (presumed optimal) strategy compared to the FSO strategy, where  $J_{20}(\text{DP})$  and  $J_{20}(\text{FSO})$  were obtained from Equation (6.7) evaluated at  $T = 20$  under the simulated DP and FSO strategies, respectively. Finally, we replicated this procedure 10,000 times for each of the 243 starting states.

As expected, FSO and DP yielded identical optimal strategies, with

Equation (6.8) equal to zero, for all state combinations at  $t = 19$ , that is, one time step before the terminal time  $T = 20$  (100 years). Although this comparison was not of particular interest and is not reported here, it confirmed that the two procedures were solving the same one-time-step problem. That is, from the perspective of a decision maker at  $t = 19$ , the one-time-step, "myopic" strategy yielded by FSO is equivalent to a long-term, "far-sighted" strategy provided by DP.

As the decision horizon lengthened to 100 years (that is, time starting at  $t = 0$  rather than  $t = 19$ ), the two strategies diverged. The 100-year strategies were compared with respect to whether the  $t = 0$  harvest decisions differed and the direction of difference if any (Figure 6.5). Under a broad range of initial conditions, FSO yields an optimal strategy that is more aggressive (involves more cutting) than does DP—which makes sense given the myopic nature of FSO. That these myopic strategies are suboptimal with respect to the long-term objective is seen by comparing the simulated 100-year ( $t = 20$ ) objective outcomes (Equation 6.8; Figure 6.6). The results confirm the superiority of the DP strategy. For all state combinations, the average over 10,000 simulations always resulted in positive values for Equation (6.8), although FSO happened to outperform DP in some simulations. The degree of suboptimality of the FSO strategy depended on initial habitat conditions: poorer performance occurred

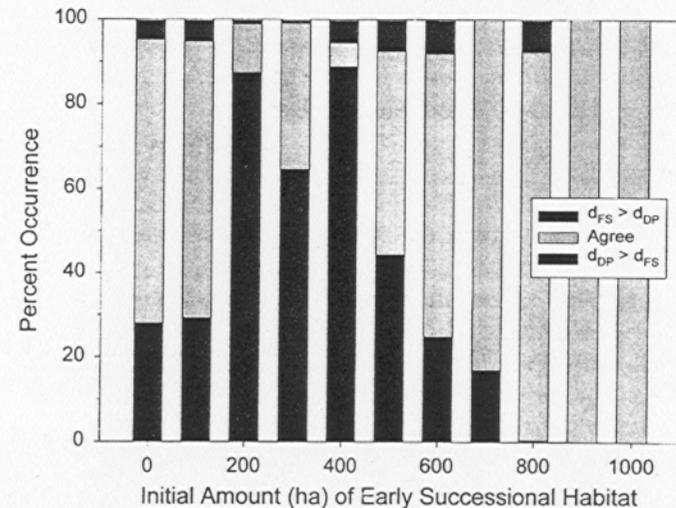


Figure 6.5. Differences in  $t = 0$  harvest decision between DP and FSO over state combinations displayed over levels of the habitat state variable. Darkest and intermediate-shaded bar sections represent cases of more and less aggressive cutting strategies, respectively, under FSO compared to DP.

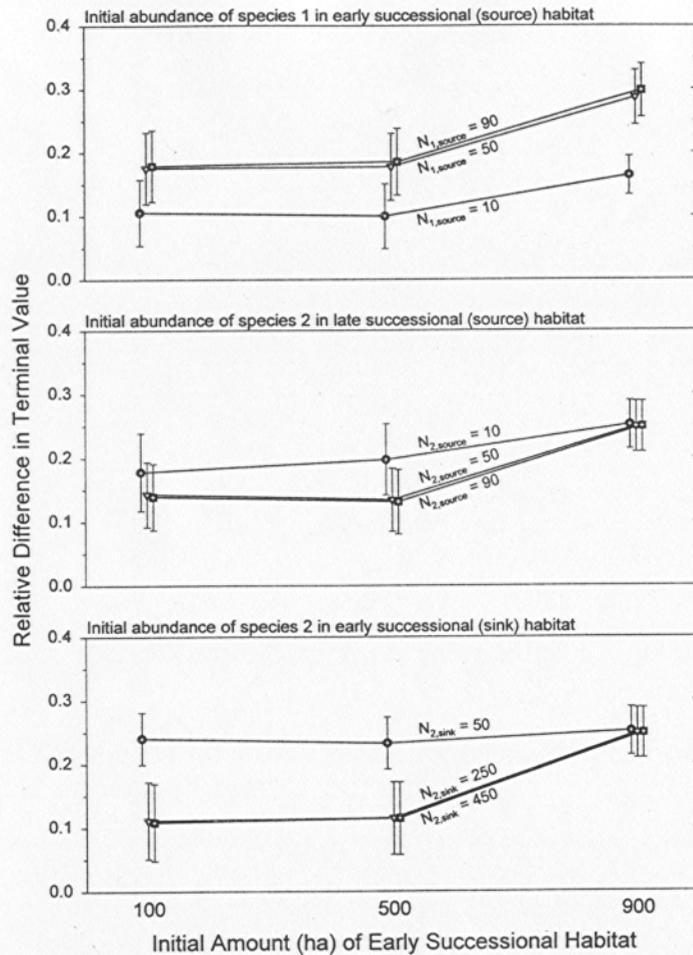


Figure 6.6. Relative differences in expected performance (objective value) between DP and FSO (Equation 6.8) over 10,000 simulation trials displayed over three two-way combinations of initial system states. Error bars around means represent one standard deviation of simulated difference measures over all other initial system states.

when more versus less of the landscape is in early successional habitat. We note, however, that under the worst-case scenario FSO only approached a 30 percent loss in optimality and for a wide range of initial state conditions FSO provided decisions that approached 90 percent of the optimal DP results.

### CONSIDERATIONS IN USING OPTIMIZATION VERSUS SIMULATION

Our example illustrates the benefits of optimization approaches such as DP, which are designed to efficiently solve dynamic problems in which the objective is specified over a long time frame. Because of its backward, iterative approach, DP is able to decompose the time frame into stages and then solve the problem one stage at a time. In contrast, FSO (and simulation in general), by working forward through time, deals with problems that increase geometrically in complexity at each time step. A forward-looking algorithm like FSO is simply incapable of anticipating all the possible pathways that decisions could take through time and therefore becomes increasingly inferior to DP as time horizons lengthen. For the problem just described, FSO would have had to consider decisions over  $(5 \times 25)^{20} = 8.7 \times 10^{41}$  decision pathways (5 decision levels, 25 random source and sink  $\lambda$  outcomes, 20 decision periods) from a single starting state of the system—and most problems are more complex. Nonetheless, forward-looking but myopic simulation procedures such as FSO can be useful for exploring complex problems and, as illustrated here, may yield results that are not substantially inferior to DP or other optimization procedures. For complicated systems involving multiple species and spatial components, DP may be incapable of providing solutions at all—due to the “curse of dimensionality” (Bellman 1957)—in which case managers are forced to use simulation or other heuristic approaches.

There is, however, a troubling conceptual issue: we may never be able to assess the cost of moving from truly optimal solutions to solutions obtained by incorporating simulation. Despite the simplicity of our model, stripped of all but the most fundamental elements of long-term, multispecies management, we still faced computational challenges. For a more realistic problem there will not be a benchmark optimal strategy—otherwise we would never consider simulation—so it may never be possible to evaluate the relative performance of simulation-optimization strategies.

*Adaptation* is the explicit incorporation of information from monitoring, management, and experimentation into decision making. At present, formal coupling of adaptation to optimization is very complex—even for simple systems (Williams 1996). Combining simulation and optimization and incorporating adaptation for complex systems such as ours will require innovative applications of hardware and software and will no doubt require the development of new methods. Our future work will explore how adaptation can best be incorporated into decision problems like those described here. The results of this work should have important implications for the role of monitoring programs in adaptive resource management.

## ACKNOWLEDGMENTS

We thank T. M. Shenk, A. B. Franklin, and two anonymous referees for critical comments on an earlier draft. The Georgia Cooperative Fish and Wildlife Research Unit is jointly sponsored by the USGS Biological Resources Division, the University of Georgia, the Georgia Department of Natural Resources, and the Wildlife Management Institute.

## REFERENCES

- Bellman, R. E. 1957. *Dynamic Programming*. Princeton: Princeton University Press.
- Conroy, M. J., and B. R. Noon. 1996. Mapping of species richness for conservation of biological diversity: Conceptual and methodological issues. *Ecological Applications* 6:763-773.
- Conroy, M. J., Y. Cohen, F. C. James, Y. G. Matsinos, and B. A. Maurer. 1995. Parameter estimation, reliability, and model improvement for spatially-explicit models of animal populations. *Ecological Applications* 5:17-19.
- Dreyfus, S. E., and A. M. Law. 1977. *The Art and Theory of Dynamic Programming*. New York: Academic Press.
- Lubow, B. 1995. SDP: Generalized software for solving stochastic dynamics optimization problems. *Wildlife Society Bulletin* 23:738-742.
- Mangel, M., and C. W. Clark. 1988. *Dynamic Modeling in Behavioral Ecology*. Princeton: Princeton University Press.
- Pulliam, H. R. 1988. Sources, sinks, and population regulation. *American Naturalist* 132:652-661.
- SAS Institute. 1990. *SAS/IML Software: Usage and Reference*. Version 6. 1st ed. Cary, N.C.: SAS Institute.
- Williams, B. K. 1989. Review of dynamic optimization methods in renewable natural resources management. *Natural Resource Modeling* 3:137-216.
- . 1996. Adaptive optimization and the harvest of biological populations. *Mathematical Biosciences* 136:1-20.