

Optimal Regeneration Planning for Old-Growth Forest: Addressing Scientific Uncertainty in Endangered Species Recovery through Adaptive Management

Clinton T. Moore and Michael J. Conroy

Abstract: Stochastic and structural uncertainties about forest dynamics present challenges in the management of ephemeral habitat conditions for endangered forest species. Maintaining critical foraging and breeding habitat for the endangered red-cockaded woodpecker (*Picoides borealis*) requires an uninterrupted supply of old-growth forest. We constructed and optimized a dynamic forest growth model for the Piedmont National Wildlife Refuge (Georgia, USA) with the objective of perpetuating a maximum stream of old-growth forest habitat. Our model accommodates stochastic disturbances and hardwood succession rates, and uncertainty about model structure. We produced a regeneration policy that was indexed by current forest state and by current weight of evidence among alternative model forms. We used adaptive stochastic dynamic programming, which anticipates that model probabilities, as well as forest states, may change through time, with consequent evolution of the optimal decision for any given forest state. In light of considerable uncertainty about forest dynamics, we analyzed a set of competing models incorporating extreme, but plausible, parameter values. Under any of these models, forest silviculture practices currently recommended for the creation of woodpecker habitat are suboptimal. We endorse fully adaptive approaches to the management of endangered species habitats in which predictive modeling, monitoring, and assessment are tightly linked. FOR. SCI. 52(2):155–172.

Key Words: Forest planning, red-cockaded woodpecker, *Picoides borealis*, optimization, stochasticity, decision modeling, habitat.

IN RECENT DECADES, management of public forest lands in the United States has shifted from a primary focus on commodity production to one that provides greater consideration of other resource objectives (Rose and Chapman 2003). In particular, an emphasis on ecosystem management has spurred natural resource agencies to integrate and plan for old-growth conditions in many forest management schemes (Franklin 1997, Thomas 1997). These conditions are requisite for the recovery and population maintenance of several species of special concern, including the endangered red-cockaded woodpecker (*Picoides borealis*, hereafter “woodpecker”), a native cavity-nesting bird of pine (*Pinus* sp.) forests throughout the southeastern United States (Ligon et al. 1986, Jackson 1994).

Thus, on forests that are identified as woodpecker recovery areas, there is a management desire to promote old-growth conditions as rapidly as possible and to sustain suitable amounts of these conditions over perpetuity. A plan for recovery of the woodpecker on forests within the US Fish and Wildlife Service National Wildlife Refuge System (US Fish and Wildlife Service 1998) provides guidelines for regeneration of pine forest stands so that the maintenance of

an old-growth component is assured. This plan proposes the “area control” method of forest regeneration planning for even-aged or two-aged management systems. Under this method, the proportion of forest regenerated each year, either through the clearing of all (clearcutting) or a majority (seed tree, shelterwood) (Smith 1962) of the mature overstory, is a constant ratio of compartment entry cycle to forest rotation length. Thus, 8% of forest in a compartment should be regenerated annually if the forest is managed on a 100-year rotation and compartments are visited every 8 years for treatment.

The area control approach may be overly simplistic for several reasons. First, it assumes that the forest age class distribution is uniform, which may not be realistic, e.g., in naturally regenerated forests. Recognizing this, the Refuge System plan makes the ad hoc recommendation of avoiding harvests of the two oldest age classes (where age classes are defined by the entry cycle) in any year until a uniform age distribution is achieved (US Fish and Wildlife Service 1998). Second, managers may expect to gradually lose (e.g., succession to hardwood) or gain pine acreage over the long term, and anticipation of those trends may affect the optimal

Clinton T. Moore, USGS Patuxent Wildlife Research Center, Warnell School of Forest Resources, University of Georgia, Athens, GA 30602—Phone: 706-542-1609; Fax: 706-542-1235; cmoore@forestry.uga.edu. Michael J. Conroy, USGS Georgia Cooperative Fish and Wildlife Research Unit, Warnell School of Forest Resources, University of Georgia, Athens, GA 30602—mconroy@uga.edu.

Acknowledgments: Funding and support for our work came from the US Environmental Protection Agency STAR Fellowship Program (Fellowship U-915424-01-0), the USGS Georgia Cooperative Fish and Wildlife Research Unit (Work Order #34, Coop. Agreement #14-16-0009-1551), and the University of Georgia Warnell School of Forest Resources. We thank the staff of the Piedmont National Wildlife Refuge for their assistance and advice. Our manuscript was greatly improved by the many constructive suggestions of three anonymous reviewers and the Associate Editor. The Georgia Cooperative Fish and Wildlife Research Unit is jointly sponsored by the US Geological Survey, the University of Georgia, the Georgia Department of Natural Resources, and the Wildlife Management Institute.

choice of current harvest. Third, failure to acknowledge stochastic events such as forest disturbance and hardwood succession may also result in suboptimal decisions (Reed and Errico 1986, Gassmann 1989, van Kooten et al. 1992, Teeter et al. 1993). Fourth, uncertainty exists about the system's mechanistic structure, i.e., the fundamental dynamics of forest growth. As a consequence, management decisions that would be appropriate under the belief in one model of the system may be inappropriate if another model is more realistic (Conroy and Moore 2002). These reasons raise legitimate concerns about the unquestioning application of the area control method (or its ad hoc variants) for optimally managing old-growth habitat.

Optimal decision-making for the creation and maintenance of old-growth conditions or any other transitory forest condition must be predicated on models that explicitly acknowledge these dynamics and sources of uncertainty. Stochastic uncertainty has been addressed in previous applications of dynamic optimization models in forest management, for example, in burning to promote persistence of a woodland shrub (McCarthy et al. 2001) and in holding areas of primary forest to maximize future social nonuse benefits (Bulte et al. 2002), but to our knowledge, none have also taken into account structural uncertainty.

We present a forest succession model for the Piedmont National Wildlife Refuge, Georgia, where the objective is sustenance of a red-cockaded woodpecker population. Thus, forest management at the Refuge is oriented toward the creation and maintenance of old-growth (≥ 80 years) loblolly (*P. taeda*) and shortleaf (*P. echinata*) pine forest conditions. Our model projects the response of the forest to harvest from each of three forest age classes and to stochastic influences of disturbance and forest succession.

For each possible age class composition of the forest, our goal was to find optimal amounts of forest area to regenerate annually for the return of a maximum stream of old-growth habitat over an infinite time frame and under stochastic outcomes of disturbance and succession. Most importantly, we also required this decision policy to account for structural uncertainty in the average rate of conversion from pine to hardwood forests, which are inhospitable to woodpeckers. Different hypotheses about succession imply different optimal harvest policies; however, managers cannot postpone decision-making until this uncertainty is resolved. The value of our policy is that it provides optimal decision values both in cases of complete uncertainty among competing hypotheses and in cases where some knowledge is available, perhaps as a result of management experience. Because the policy is computed by taking into account the expected flow of information returned by each cycle of decision-making, the decision policy is said to be actively adaptive. Decision-making under an actively adaptive policy may sacrifice some gain in the short term to elicit information that ultimately improves management performance over the long term (Walters and Hilborn 1978, Williams 1997, Williams et al. 2002).

Study Area

The Piedmont National Wildlife Refuge, located in central Georgia, is a 14,136-ha unit of the US National Wildlife Refuge System (Figure 1). The site supports a second-growth mixed pine (loblolly and shortleaf) and hardwood (*Quercus* sp., *Carya* sp.) forest that regenerated naturally on severely eroded farmland abandoned before 1940 (Gabrielson 1943, Czuhai and Cushwa 1968). Forest management is directed toward the maintenance of all native flora and fauna, sustenance of important ecosystems, and provision of public recreation, including wildlife viewing and sport harvest of some wildlife species (US Fish and Wildlife Service, Piedmont National Wildlife Refuge, <http://piedmont.fws.gov>. Accessed 23 Sep 2005). The Refuge is also a designated recovery site of the red-cockaded woodpecker (US Fish and Wildlife Service 2003) and hosts the largest woodpecker population (about 40 breeding groups) in the Piedmont physiographic province.

The woodpecker's preferred foraging and breeding habitat consists of pure, open stands of mature (≥ 80 years) pine with a fire-maintained herbaceous understory and sparse hardwood midstory (Loeb et al. 1992). Throughout the southern United States these forest habitats have become highly fragmented or have disappeared altogether, particularly since the early 20th century, as intensification of management on industrial forest lands emphasized shorter timber rotations and as exclusion of fire permitted increased hardwood succession (Ligon et al. 1986). Because the Refuge is identified as a woodpecker recovery site, forest management is oriented toward increasing woodpecker abundance. To this end, forest managers use thinning and regeneration cutting, prescribed burning, and mechanical vegetation removal to maintain the desired habitat conditions (US Fish and Wildlife Service, Piedmont National Wildlife Refuge, Habitat Management Plan, 1982 [unpublished report]. Hereafter, "Refuge Habitat Management Plan").

The Refuge uses retention (irregular) shelterwood silviculture (Refuge Habitat Management Plan; J.D. Metteauer, personal communication, Piedmont National Wildlife Refuge, 2000), an even-aged regeneration system in which some shelter trees are left on-site until the following rotation (Smith 1962). In addition to protecting the developing seedling crop, the retained trees may serve as woodpecker foraging (US Fish and Wildlife Service 1998) or nesting (Conner et al. 1991) habitat. Refuge managers designate pine stands as sapling (P1, < 16 years), poletimber (P2, 16–40 years), or sawtimber (P3, ≥ 40 years).

Until recently (2003), pine stands were managed on an 80-year average rotation length; the specific length depended on site fertility (Refuge Habitat Management Plan). Average rotation length is now 100 years (C. Schmidt, personal communication, Piedmont National Wildlife Refuge, 2005) in keeping with woodpecker recovery guidelines for loblolly pine forests (US Fish and Wildlife Service 1998, 2003). Rotations of length > 100 years are considered infeasible because of limitations imposed by site conditions,

0 100 200 Kilometers

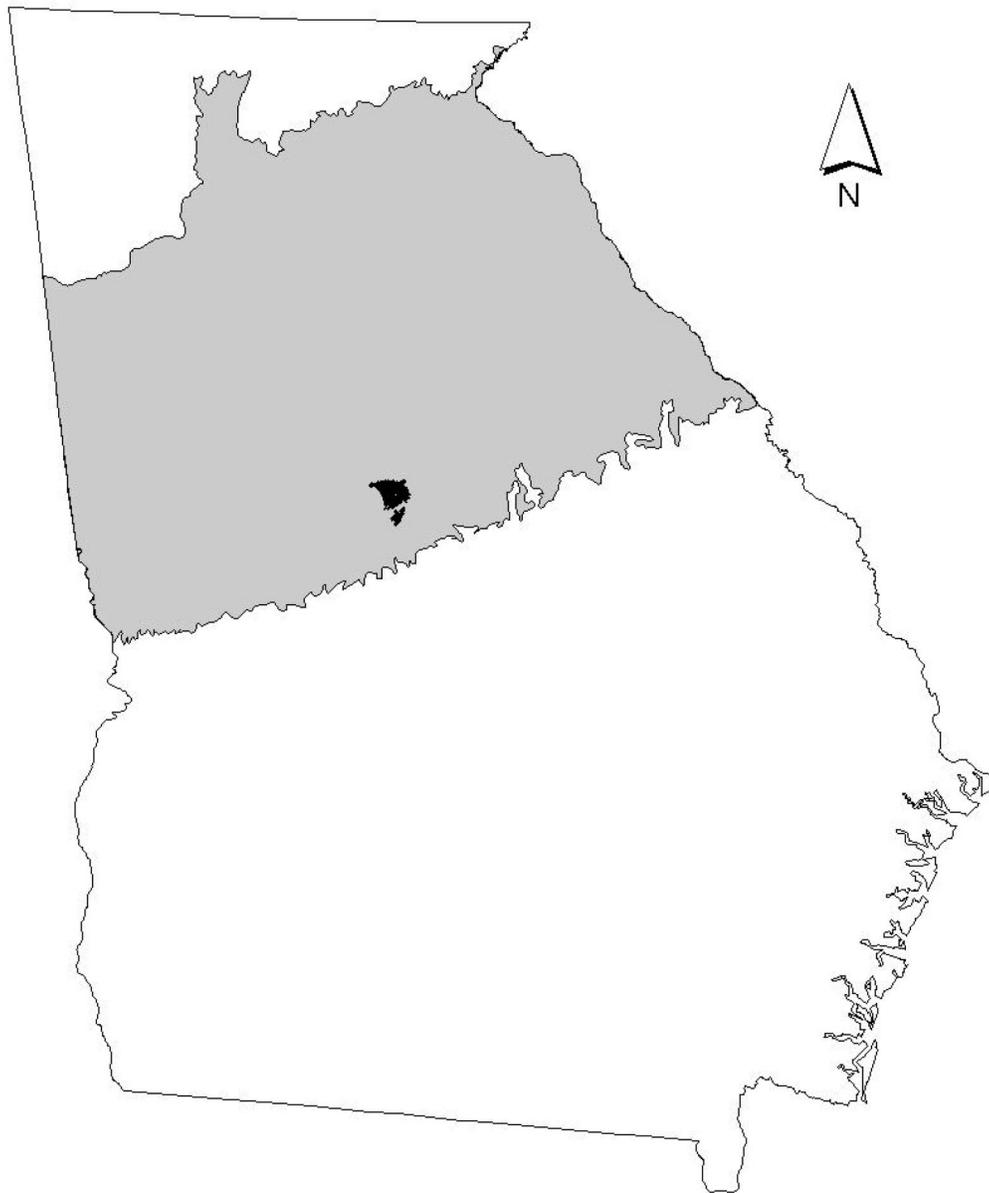


Figure 1. Piedmont National Wildlife Refuge (darkened area), Georgia, USA, and extent of Piedmont physiographic province within Georgia (shaded).

the biological environment, and physiological characteristics of the tree (C. Schmidt, personal communication, Piedmont National Wildlife Refuge, 2005). A series of intermediate thinnings is carried out over the rotation to achieve specified density and crown closure goals for woodpeckers (Refuge Habitat Management Plan).

Pine stands throughout the Refuge are transforming into hardwood stands, often rapidly (Refuge Habitat Management Plan). Refuge managers use fire and mechanical removal to control the hardwood understory in those stands where woodpeckers exist or are desired. However, in an effort to increase Refuge habitat diversity, managers do not

try to impede hardwood succession on many other upland areas (Refuge Habitat Management Plan). Managers forecasted in 1982 that succession will continue until 40% of the forest cover comprises upland and bottomland hardwood types (Refuge Habitat Management Plan).

Methods

Forest Decision Model

To manage woodpeckers and their habitat, Refuge managers must make a series of decisions through space and time. We considered regeneration cutting through time as

the primary regulator of the amount of old-growth forest, and we constructed a forest decision model to analyze cutting decisions. Although spatial arrangement of suitable habitat is a central concern in woodpecker management (Walters 1991, Letcher et al. 1998, Azevedo et al. 2000), decisions about *where* habitat is located could be reasonably made in a secondary step after determining *how much* habitat to create (e.g., Boston and Bettinger 2001). Therefore, we ignored stand-level characteristics, including spatial distribution of stands, and we instead focused on harvest of age classes at the whole-forest level. Our model also did not consider prescribed burning as a means of controlling hardwood succession at the whole-forest level. Placement of prescribed burning at the Refuge occurs mostly in response to current understory vegetation conditions at existing or potential woodpecker nesting sites. Financial constraints and conservation mandates prevent the Refuge from applying fire at the greater scales needed to control hardwood succession at the whole-forest level (J.D. Metteauer, personal communication, Piedmont National Wildlife Refuge, 2000).

The state of the forest at any point in time was represented in our model as the distribution of total forest area into five seral classes: P1 and P2 age classes (as previously defined), P3 (pine 40–80 years), P4 (old-growth pine ≥ 80 years), and UH (upland hardwood). We split the Refuge's single ≥ 40 -year pine class into the classes P3 and P4 because of the need to specifically recognize and manage for very old stands suitable for woodpecker nesting habitat. The Refuge distinguishes upland from bottomland hardwood types, but we ignored this latter type in our model on assumption that these habitats remain largely unchanged in extent over time ($\sim 9.4\%$ of Refuge area).

Our model advanced the state of the forest through time, in annual steps, in response to forest growth, stochastic succession and disturbance events, and cutting decisions. This Markov decision process (Puterman 1994) was represented in a stage-based matrix model (Caswell 2001), in which a set of transition parameters expressed rates of exchange among the pine age classes and the upland hardwood forest type.

Model Structure

We represented forestwide proportions of the four pine age classes and the upland hardwood class occurring at time t by the vector \mathbf{y}_t . The model thus projected the forest state from year t to year $t + 1$ as:

$$\mathbf{y}_{t+1} = \mathbf{A}_t \mathbf{y}_t,$$

where \mathbf{A}_t is the matrix product $\mathbf{C}_t \times \mathbf{G} \times \mathbf{H}_t \times \mathbf{D}_t$. The components of \mathbf{A}_t represent forest disturbance (\mathbf{C}_t), growth (\mathbf{G}), hardwood succession (\mathbf{H}_t), and regeneration decisions (\mathbf{D}_t), respectively, and are described in detail below. All processes except growth (\mathbf{G}) were time-specific and thus are indexed by t . The five forest class proportions (P1, P2, P3, P4, UH) occurred in positions 2–6 of the state vector \mathbf{y}_t . Position 1 was a forest class used to temporarily store the current year's regeneration cut (see below).

Consistent with Refuge management, our model allowed regeneration harvests from P2, P3, and P4. Although in practice small hardwood patches may be cut for regeneration of pine, these cuts are infrequent and chosen opportunistically to augment a neighboring pine harvest: Because cutting of hardwood stands is not part of the scheduled silviculture at the Refuge, we did not consider cuts from the UH class in our model. The first step of the model applied a set of class-specific cutting decisions $\{d_{2t}, d_{3t}, d_{4t}\}$ at time t . These decisions transferred $d_{2t}y_{P2,t} + d_{3t}y_{P3,t} + d_{4t}y_{P4,t}$ amount of pine forest into a temporary "regeneration" class (Figure 2), which was used only for accounting purposes and to prevent newly regenerated forest from immediately growing into the P2 class. After the model calculated the growth transitions (\mathbf{G} , see below), the regeneration class was completely emptied into the P1 class and was thus zeroed out. Matrix \mathbf{D}_t was as follows:

$$\mathbf{D}_t = \begin{bmatrix} 1 & 0 & d_{2t} & d_{3t} & d_{4t} & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 - d_{2t} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 - d_{3t} & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 - d_{4t} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

The model next calculated portions of the pine forest lost to hardwood succession (Figure 2). We assumed that the average rate of loss is persistent each year but is controllable by the amount of harvest taken from each pine class. Thus, through regular cutting, managers may slow or stop the rate of pine cover loss. We assumed that the loss rate, e_{it} , $i = 1, 2, 3, 4$ for the four pine classes, was stochastic. For example, environmental conditions that affect hardwood establishment or the ability of managers to suppress hardwood may vary unpredictably each year. The realized rate of loss, E_{it} , was the positive amount of e_{it} that exceeded d_{it} , i.e.,

$$E_{it} = \max(0, e_{it} - d_{it}).$$

Note that, because $d_{1t} = 0$ (i.e., no harvest in P1), $E_{1t} = e_{1t}$. Each simulation year, we drew a random variate s_t from a beta distribution with mean μ and coefficient of variation v_e . We obtained the individual e_{it} by applying a set of scaling factors $\mu(e_i)$ to s_t . Thus, the e_{it} were random, but they covaried perfectly with each other. The E_{it} were applied in the model as

$$\mathbf{H}_t = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 - E_{1t} & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 - E_{2t} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 - E_{3t} & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 - E_{4t} & 0 \\ 0 & E_{1t} & E_{2t} & E_{3t} & E_{4t} & 1 \end{bmatrix}.$$

Following succession to hardwood, portions of each pine class graduated to the next older class (Figure 2). Transition rates were constant through time and were denoted τ_{ij} , where the transition occurs from the younger class j to the older class i . Thus, τ_{21} , τ_{32} , and τ_{43} describe the rates of transition from P1 to P2, from P2 to P3, and from P3 to P4,

respectively. A portion of P4, denoted τ_{14} , automatically regenerated each year. Because not all sites on the Refuge are of sufficient quality to support the oldest stands of trees (Refuge Habitat Management Plan), only a portion, k , of the P3 class eligible to graduate actually entered the P4 stage. We assumed that the remainder regenerated as P1. Finally, the model emptied the temporary regeneration class into the P1 class. These transitions appear as

$$\mathbf{G} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 - \tau_{21} & 0 & (1 - k)\tau_{43} & \tau_{14} & 0 \\ 0 & \tau_{21} & 1 - \tau_{32} & 0 & 0 & 0 \\ 0 & 0 & \tau_{32} & 1 - \tau_{43} & 0 & 0 \\ 0 & 0 & 0 & k\tau_{43} & 1 - \tau_{14} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

The last component of the model applied an annual stochastic catastrophic disturbance event c_t to the older pine classes and to the UH class (Figure 2). The event converted the affected pine cover into the P1 type. Only a portion, q , of the affected hardwood type regenerated as pine cover, with the rest remaining as type UH. The c_t were drawn from a beta distribution with mean $\mu(c)$ and variance $\sigma^2(c)$. The matrix \mathbf{C} expresses these transitions as

$$\mathbf{C}_t = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & c_t & c_t & c_t & c_t q \\ 0 & 0 & 1 - c_t & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 - c_t & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 - c_t & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 - c_t q \end{bmatrix}$$

Applying these transitions to the vector \mathbf{y}_t in the above order produced the matrix product $\mathbf{\Lambda}_t = \mathbf{C}_t \times \mathbf{G} \times \mathbf{H}_t \times \mathbf{D}_t$. Elements in $\mathbf{\Lambda}_t$ specified rates of transition among all compartments of the model (Figure 3).

Model Parameterization

Thirteen parameters controlled this model. Unfortunately, no data exist from the Refuge to estimate any of them. In some cases (e.g., age transition) we were able to infer reasonable parameter values based on first principles and reasonable, biological assumptions. In other cases the parameter values are purely guesses, but were chosen to produce outcomes in forest composition that are consistent with past experience.

We fixed values for the age class transitions (τ_{21} , τ_{32} , τ_{43} , τ_{14}) at (1/16, 1/24, 1/40, 1/40). In other words, we assumed that rates of transition among age classes occurred in proportion to the time span of the younger age class. This is a reasonable assumption if ages are uniformly distributed within an age class. We fixed the value of the P4 admission rate, k , at 0.5, implying that only half of the forest type leaving the P3 age class enters the P4 class. This value is reasonable for the Refuge, as only approximately half of the Refuge is of sufficient fertility to support the oldest stands of pine (Refuge Habitat Management Plan).

We simulated the proportion of forest destroyed by catastrophic disturbance in year t , c_t , from a beta distribution with mean 0.003394 and variance 0.0003572, providing

98th and 99.8th percentile values of 0.05 and 0.20, respectively, equivalent to destruction of 5% of the forest during a 50-year disturbance event and 20% destruction during a 500-year event. These probabilities appear consistent with estimates of tornado (Peterson 2000) and hurricane (Hooper and McAdie 1995) strike probabilities for the region. Below the 90th percentile (i.e., 10-year and more common events), practically none (<0.002) of the forest is destroyed. The proportion, q , of the hardwood class destroyed by a catastrophic event that regenerates as pine cover was fixed at 0.25.

We had little objective guidance in choosing values for parameters representing mean proportions of succession to hardwood in each class, $\mu(e_i)$, and coefficient of variation of the succession rate, v_e . After much trial and error, we chose the values (0.0006, 0.015, 0.015, 0.03) for the $\mu(e_i)$, $i = 1, 2, 3, 4$, and we fixed the value of v_e at 40%. The v_e value was small enough to provide a nearly symmetric distribution of beta variates yet large enough to realistically reflect considerable ecological variability. We believed that the rate of hardwood succession should be larger in older than in younger stands, thus $\mu(e_i)$ increases with stand age. Simulations of the model with these values and with harvest rates fixed at the values $d_{it} = \mu(e_i)$ projected the UH component to grow to an asymptotic limit of $\sim 50\%$. Although these values may be the most arbitrary of the entire set, they are consistent with available data. The average of the $\mu(e_i)$, weighted by the age class transition rates τ_{ij} , was 0.012, comparable to the 0.0087 annualized rate of hardwood succession in natural loblolly pine stands in Georgia during the period 1961–1972 (Bechtold et al. 1991).

Model Optimization

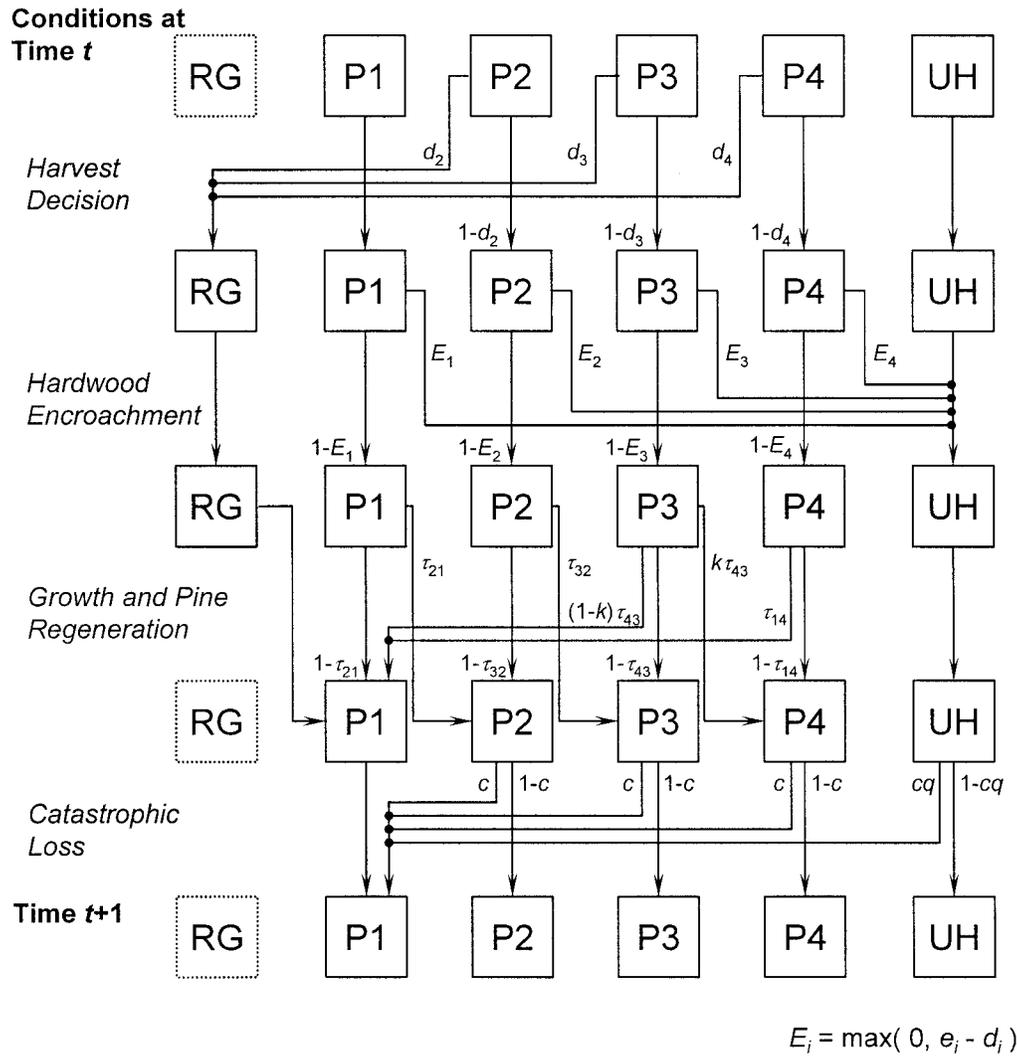
Our goal was to derive a state-specific, stationary (time-independent) regeneration policy that maximizes, over time, the amount of nesting habitat available for the red-cockaded woodpecker. The estimated minimum amount of nesting habitat needed to support the recovery goal of 96 woodpecker groups on the Piedmont National Wildlife Refuge (US Fish and Wildlife Service 1998, 2003) is 388.5 ha, or 0.03096 of the modeled forest cover, assuming a 4.05-ha minimum cluster size (US Fish and Wildlife Service 1998, 2003) and assuming total forest area (pine and upland hardwood) remains constant over time. We defined an objective function

$$v(\mathbf{y}_t, \{d_{it}\}, s_t, c_t) = y_{P4,t+1},$$

where the amount of habitat in the P4 class at the next time period ($y_{P4,t+1}$) is a function of the forest state (\mathbf{y}_t), set of decisions ($\{d_{it}\}$), and stochastic outcomes (s_t, c_t) at time t . For any arbitrary current forest state \mathbf{y}_{t0} , we searched for the series of cutting decisions $\{d_{it}\}$ over a distant time horizon T that maximized

$$V(\mathbf{y}_{t0}) = \sum_{t=t_0}^T v(\mathbf{y}_{t0}, \{d_{it}\}, s_t, c_t),$$

with respect to all realizations of the stochastic variables s_t and c_t . Thus, the objective function directs management to seek



Model Parameters	
Decisions	d_2, d_3, d_4
Succession rates*	e_1, e_2, e_3, e_4 (0.0006, 0.015, 0.015, 0.03)
Age class transitions	$\tau_{21}, \tau_{32}, \tau_{43}, \tau_{14}$ (1/16, 1/24, 1/40, 1/40)
P4 admission rate	k (0.5)
Catastrophic loss rate*	c (0.003394)
UH \rightarrow P1 conversion rate	q (0.25)
*stochastic parameters	

Figure 2. Transitions among cover types of a forest dynamics model, developed for forest management at the Piedmont National Wildlife Refuge, are displayed for successive processes within a single time step. Amounts of forest area are transferred among regeneration (RG), sapling (P1), poletimber (P2), sawtimber (P3), and old-growth (P4) pine age classes and among an upland hardwood (UH) type. Rates of hardwood succession, e_i , are partially controllable by harvest actions (d_2, d_3 , and d_4), yielding realized rates of succession, E_i . Parameter values used in the baseline (F0) dynamics model are indicated.

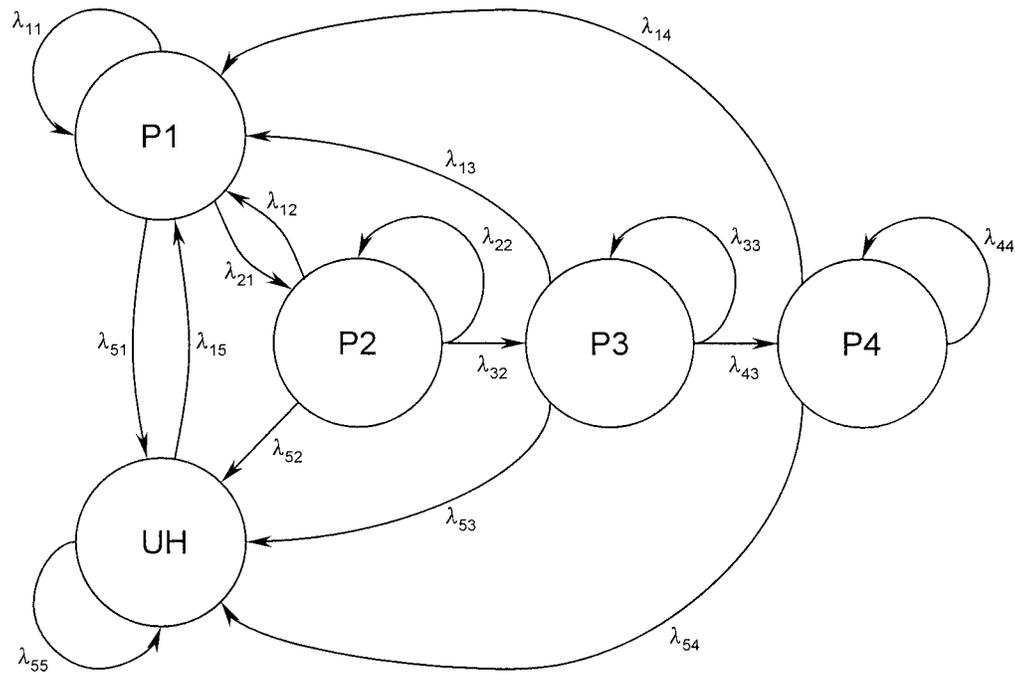
those decisions that yield maximum return of woodpecker nesting habitat as quickly as possible, but in a context in which these returns are sustained over a long time frame.

We used stochastic dynamic programming (Nemhauser 1966, Gluss 1972, Anderson 1975, Dreyfus and Law 1977) implemented in program ASDP (Lubow 1995, 1997) to search for an optimal regeneration decision policy under this model. In general terms, stochastic dynamic program-

ming seeks the decision set δ that solves the recurrence relationship

$$V(\mathbf{x}_t) = \text{Max}_{\delta} \left\{ E_z [v(\mathbf{x}_t, \delta_t, \mathbf{z}_t)] + \sum_{\mathbf{x}_{t+1}} p(\mathbf{x}_{t+1} | \mathbf{x}_t, \delta_t) V(\mathbf{x}_{t+1}) \right\},$$

where \mathbf{x}_t , δ_t , and \mathbf{z}_t are the vectors of possible system states, decisions, and stochastic events at time t , v is the immediate



$$\begin{aligned}
 \lambda_{11} &= (1-\tau_{21}+c\tau_{21})(1-E_1)+cqE_1 \\
 \lambda_{12} &= d_2+(c(1-E_2)+cqE_2)(1-d_2) \\
 \lambda_{13} &= d_3+(((1-k)\tau_{43}+c(1-\tau_{43})+ck\tau_{43})(1-E_3)+cqE_3)(1-d_3) \\
 \lambda_{14} &= d_4+((\tau_{14}+c(1-\tau_{14}))(1-E_4)+cqE_4)(1-d_4) \\
 \lambda_{15} &= cq \\
 \lambda_{21} &= (1-c)\tau_{21}(1-E_1) \\
 \lambda_{22} &= (1-c)(1-\tau_{32})(1-E_2)(1-d_2) \\
 \lambda_{32} &= (1-c)\tau_{32}(1-E_2)(1-d_2) \\
 \lambda_{33} &= (1-c)(1-\tau_{43})(1-E_3)(1-d_3) \\
 \lambda_{43} &= (1-c)k\tau_{43}(1-E_3)(1-d_3) \\
 \lambda_{44} &= (1-c)(1-\tau_{14})(1-E_4)(1-d_4) \\
 \lambda_{51} &= (1-cq)E_1 \\
 \lambda_{52} &= (1-cq)E_2(1-d_2) \\
 \lambda_{53} &= (1-cq)E_3(1-d_3) \\
 \lambda_{54} &= (1-cq)E_4(1-d_4) \\
 \lambda_{55} &= 1-cq
 \end{aligned}$$

Figure 3. Transitions (λ_{ij}) among classes of the forest dynamics model for the Piedmont National Wildlife Refuge, expressed as products of the individual processes of forest disturbance (c , proportion of forest disturbed; q , proportion of destroyed hardwood class regenerating to pine), growth (τ_{ij} , transition rate from pine age class j to age class i ; k , admission rate to P4 pine age class from P3), hardwood succession (E_i , realized rate of hardwood succession in pine class i), and regeneration decisions (d_i).

value returned by carrying out decision δ_t given state \mathbf{x}_t and random outcome \mathbf{z}_t , $p(\mathbf{x}_{t+1}|\mathbf{x}_t, \delta_t)$ is the probability of transition from state \mathbf{x}_t to new state \mathbf{x}_{t+1} (where stochasticity is inherited from \mathbf{z}_t) given δ_t , and $V(\mathbf{x}_{t+1})$ is the expected maximum cumulative return given arrival at future state \mathbf{x}_{t+1} (Nemhauser 1966, Gluss 1972, Williams et al. 2002). In other words, the goal is to find the set of decisions that maximizes the immediate return plus the maximum accumulated return expected over all possible values of the

future state \mathbf{x}_{t+1} , without regard to how state \mathbf{x}_t came to be from whatever decision sequence preceded it. This definition of the problem is an expression of Bellman's (1957) Principle of Optimality, and it suggests a recursive computational strategy that explores the decision space backward through time, facilitated by the representation of forest states, harvest decisions, and stochastic effects as discrete points on a grid (Nemhauser 1966, Gluss 1972). For each step backward in time, t_A , and for each possible forest state,

the dynamic programming procedure provides a decision that maximizes the expected return accumulated over the interval t_A, t_{A+1}, \dots, T . The entire set of optimal decisions over all forest states and time events constitutes the decision policy. The policy is said to be stationary when further steps backward in time do not elicit further changes in the array of state-specific decisions (Anderson 1975); at this point, the decision policy requires no reference to time and can therefore be appropriately applied in the context of decision-making over the infinite time horizon.

We discretized cutting decisions for each harvestable class (values of d_{ij}) in steps of 0.02 over the range 0.0–0.2. We defined 997 distinct forest states (combinations of seral classes), and the policy found by ASDP provided optimal cutting decisions from P2, P3, and P4 for each state. We required ASDP to consider $T = 1,000$ stage (year) iterations. We selected this iteration limit to assure that we obtained a policy that was optimal with respect to a reasonably long time horizon. We counted the number of final iterations over which the solution was unchanged for evidence that the solution was truly stationary.

Parametric Uncertainty

Because the parameter values in our model were established under considerable uncertainty, we investigated the sensitivity of the optimal decision policy to perturbations in the parameters. We used ASDP to calculate optimal decision policies under two alternatives to the default model (model F0): (1) parameter values that projected a lower rate of transition to the P4 type and a higher transition rate to the hardwood type (model F1) relative to model F0, and (2) parameter values that projected a higher rate of transition to the P4 type and a lower transition rate to the hardwood type (model F2) relative to model F0 (Table 1). Most parameter values were simply halved or doubled to produce the desired dynamics in P4 and UH (Table 1). Disturbance events were less likely in model F1 (5% destruction in 100-year event, 20% destruction in 1,000-year event) than in F0 and more likely in model F2 (5% destruction in 25-year event, 20% destruction in 250-year event). We did not alter values for parameters τ_{21} , τ_{32} , and τ_{43} among model alternatives because we assumed that they consistently reflected transitions between age classes under any model (Table 1).

We compared the alternative decision policies with respect to their aggressiveness of cutting actions in each of the pine types. We also extracted from each policy the optimal regeneration decisions that would have been applicable for the estimated state of the Refuge forest in year 2000 (Moore 2002).

Adaptive Optimization under Model Uncertainty

Beyond the obvious differences in details (management objectives, forest states, and system dynamics), our problem formulation to this point is rather indistinct from the optimal timber production models of Lembersky and Johnson

Table 1. Values chosen for parameters in baseline (F0) and alternative (F1, F2) models of forest stage dynamics for the Piedmont National Wildlife Refuge

Parameter ¹	Model F0	Model F1 ²	Model F2 ²
τ_{21}	1/16	1/16	1/16
τ_{32}	1/24	1/24	1/24
τ_{43}	1/40	1/40	1/40
τ_{14}	1/40	1/20	1/60
$\mu(e_1)$	0.0006	0.0012	0.0003
$\mu(e_2)$	0.015	0.03	0.0075
$\mu(e_3)$	0.015	0.03	0.0075
$\mu(e_4)$	0.03	0.06	0.015
v_e	40%	80%	20%
k	0.5	0.25	1.0
q	0.25	0.125	0.5
$\mu(c)$	0.003394	0.001707	0.006711
$\sigma^2(c)$	0.0003572	0.0001816	0.0006911

¹ Parameters: τ_{ij} , transition rate from pine age class j to age class i ; $\mu(e_i)$, mean rate of succession to hardwood in pine age class i ; v_e , coefficient of variation of hardwood succession rate; k , admission rate to P4 pine age class from P3; q , proportion of destroyed hardwood class regenerating to pine; $\mu(c)$, $\sigma^2(c)$, mean and variance of beta distribution of size (area) of catastrophic disturbance.

² Alternative model F1 (F2): lower (higher) rate of transition to P4 age class and higher (lower) rate of transition to hardwood class than in baseline model F0.

(1975), van Kooten et al. (1992), and Teeter et al. (1993). However, uncertainty due to model structure (i.e., existence of a discrete or continuous set of candidate Markovian transition matrices, each a plausible description of system dynamics) was not formally addressed in these previous studies. Under structural uncertainty, optimal decisions are chosen to minimize Bayesian loss, given the current relative degree of confidence that each candidate model is the appropriate model of system dynamics (Ducey 2001, Kangas and Kangas 2004). An actively adaptive approach to optimal decision-making under uncertainty takes the further step of balancing this choice against the need to elicit information (i.e., response by the system) that will allow learning about the system to occur as rapidly as possible.

The adaptive optimization problem generalizes the earlier problem by explicitly recognizing (1) alternative models of state dynamics, (2) model-specific optimal return values, and (3) an “information state” of time-specific model probabilities (relative degrees of confidence that each model is the appropriate representation of system dynamics) (Williams 1996a). In general terms, our aim is to find the set of decisions δ that solve

$$\bar{V}(\mathbf{x}_t, \boldsymbol{\pi}_t) = \text{Max}_{\delta} \left\{ \bar{v}(\mathbf{x}_t, \boldsymbol{\delta}_t, \mathbf{z}_t) + \sum_{x_{j+1}} \bar{p}(\mathbf{x}_{t+1} | \mathbf{x}_t, \boldsymbol{\delta}_t) \bar{V}(\mathbf{x}_{t+1}, \boldsymbol{\pi}_{t+1}) \right\},$$

where

$$\begin{aligned} \bar{V}(\mathbf{x}_t, \boldsymbol{\pi}_t) &= \sum_i \pi_{t,i} V_i(\mathbf{x}_t), \\ \bar{v}(\mathbf{x}_t, \boldsymbol{\delta}_t, \mathbf{z}_t) &= \sum_i \pi_{t,i} E_z[v_i(\mathbf{x}_t, \boldsymbol{\delta}_t, \mathbf{z}_t)], \\ \bar{p}(\mathbf{x}_{t+1} | \mathbf{x}_t, \boldsymbol{\delta}_t) &= \sum_i \pi_{t,i} p_i(\mathbf{x}_{t+1} | \mathbf{x}_t, \boldsymbol{\delta}_t), \end{aligned}$$

and the $\pi_{t,i}$ are model-specific elements of the probability vector $\boldsymbol{\pi}_t$ (Williams et al. 2002). This problem statement resembles the previous one, and as before, the optimal decision at t is chosen in light of the expected accumulated return from the current state of the system \mathbf{x}_t and its future possible states. However, adaptive optimization also takes into account the current state of information about the system, $\boldsymbol{\pi}_t$, and how this measure of system uncertainty may evolve in response to the current decision and all future decisions: This evolution occurs via a Bayesian linkage between $\boldsymbol{\pi}_t$ and $\boldsymbol{\pi}_{t+1}$ embedded in the optimization equation (Williams et al. 2002). Because it explicitly considers the relationship between decisions, reduction in model uncertainty, and expected accumulation of returns, an adaptive decision policy may trade off management performance in the short term for information expected to help identify the appropriate model and ultimately deliver larger gains in the long term (Walters and Hilborn 1978, Walters 1986, Williams 1997, Williams et al. 2002).

The adaptive optimization problem is solved using backward iterative search under dynamic programming (Williams 1996b). The approach is much the same as previously described, but a dimension of computation is added by having to accommodate the information state, which ASDP does. We represented the information state at time t as the vector $\boldsymbol{\pi}_t = (\pi_{t,0}, \pi_{t,1}, 1 - \pi_{t,0} - \pi_{t,1})$, where the components reflect current degree of belief in forest models F0, F1, and F2, respectively. We used ASDP over $T = 1,000$ stage iterations to produce a set of optimal cutting decisions from P2, P3, and P4 for all combinations of 997 forest states and 10 discrete levels of the information state. As we did for the single-model optimization runs, we assessed the iteration history for evidence that the decision policy at the final iteration was stationary.

This policy anticipates that information will be accrued through decision-making and that the information state will evolve over time. For the purpose of estimating the management value of resolving uncertainty in harvest decision-making, we computed a second policy for the special case in which uncertainty among the models is still considered but is assumed to remain fixed through time, i.e., forest responses to management are either never measured or applied to the reduction of uncertainty. Assigning the probabilities $\boldsymbol{\pi} = (1/3, 1/3, 1/3)$ as weights on models F0, F1, and F2, we computed an optimal policy based on model-averaged state predictions and decision returns using the single-model dynamic programming technique described above. We computed the expected value of perfect information (EVPI), an estimate of the value, in units of the resource, that would be gained by resolving structural uncertainty in this problem (Lindley 1985):

$$\text{EVPI}(\mathbf{y}) = [V_{F0}(\mathbf{y}) + V_{F1}(\mathbf{y}) + V_{F2}(\mathbf{y})]/3 - V_F(\mathbf{y}),$$

where $V_i(\mathbf{y})$ is expected maximum cumulative return for system state \mathbf{y} under one of the system models ($i = F0, F1, F2$) or the model-averaged form ($i = F$).

Simulations

We simulated each model under its own nonadaptive optimal policy for each of four initial forest states (i.e., a specific type and age class composition of the forest). The selections were arbitrary but represented a range of forest conditions: a mostly young forest (S1), a forest with balanced age structure (S2), a mostly mature forest (S3), and the estimated forest composition in year 2000 (S4) (Moore 2002). Additionally, we simulated each model under the adaptive decision policy corresponding to complete uncertainty among models, in which we held the information state at a fixed value of $\boldsymbol{\pi} = (1/3, 1/3, 1/3)$ over the simulated time frame. We also simulated the area control policy because of its prominence in red-cockaded woodpecker recovery planning (US Fish and Wildlife Service 1998). Each initial state \times model \times policy combination was simulated 500 times, and mean annual amounts of old-growth forest occurring at 10, 100, and 1,000 years were recorded.

We also used simulation to assess how rapidly uncertainty could be resolved in this system if a forest monitoring program was in place and to discover how observation error may affect learning rate. We simulated each model under a realistic implementation of adaptive decision-making and information updating. A simulation run began with an initial forest state \mathbf{y}_0 expressed in logit form \mathbf{u}_0 , i.e., each vector element $u_k = \log(y_k/y_3)$ for $k = 1, \dots, 4$ (\mathbf{y} simplified here as a five-element vector). Given a fixed SD value $\sigma \geq 0$ to represent size of observation error in measuring forest state, we simulated an observed forest state at time 0 as

$$\tilde{\mathbf{y}}_0 = \text{logit}^{-1}(\tilde{\mathbf{u}}_0) = \text{logit}^{-1}(\mathbf{u}_0 + \boldsymbol{\xi}),$$

where $\boldsymbol{\xi}$ was a random vector drawn from a multivariate normal distribution with mean 0 and variance $\mathbf{I}\sigma^2$, and $\text{logit}^{-1}(\mathbf{x})$ was the inverse logit operation that transformed vector \mathbf{x} with $K - 1$ elements into a K -element vector of proportions. Thus, the simulation kept track of both the true (\mathbf{y}) and observed ($\tilde{\mathbf{y}}$) states of the forest (and were equivalent only when $\sigma = 0$). Using $\tilde{\mathbf{y}}_0$ and the initial value of the information state $\boldsymbol{\pi}_0 = (1/3, 1/3, 1/3)$ as “look-up” terms in the adaptive policy table, the simulation program retrieved the corresponding set of harvest decisions $\{d_2, d_3, d_4\}$. The model then advanced the true system state \mathbf{y}_0 to a new state \mathbf{y}_1 given the set of decisions and stochastic realizations of hardwood succession and forest disturbance. In the same manner as before, a corresponding observed new system state $\tilde{\mathbf{y}}_1$ was simulated from \mathbf{y}_1 and σ .

We updated the information state vector $\boldsymbol{\pi}_0$ to new state $\boldsymbol{\pi}_1$ as follows. Given the observed system state $\tilde{\mathbf{y}}_0$ and set of decisions $\{d_2, d_3, d_4\}$ chosen at time 0, we computed 500 bootstrap predictions of the outcome \mathbf{y}_1 under each system model i . Each set of these $\hat{\mathbf{y}}_{1i}$ constituted a sampling distribution for assessing the likelihood that the new observed system state $\tilde{\mathbf{y}}_1$ could have arisen under model i given $\tilde{\mathbf{y}}_0$. The $\hat{\mathbf{y}}_{1i}$ and $\tilde{\mathbf{y}}_1$ were transformed to the logit scale as vectors $\hat{\mathbf{u}}_{1i}$ and $\tilde{\mathbf{u}}_1$, respectively. Scatterplots of the $\hat{\mathbf{u}}_{1i}$ revealed that the models induced high pairwise correlations among forest components and sharp domain boundaries, making proper

assessment of $\tilde{\mathbf{u}}_1$ in the context of the multivariate distribution of the $\hat{\mathbf{u}}_{1i}$ difficult, at least in terms of the $\hat{\mathbf{u}}_{1i}$ coordinate system. Therefore, for each model i , we computed principal component scores $\hat{\mathbf{u}}_{1i}^*$ through the linear transformation $\mathbf{U}^* = \mathbf{U}\mathbf{E}$, where \mathbf{U} was the 500×4 matrix formed by appending the (mean-corrected) $\hat{\mathbf{u}}_{1i}$ vectors and \mathbf{E} was the matrix of eigenvectors for the product $\mathbf{U}'\mathbf{U}$. Similarly, we projected $\tilde{\mathbf{u}}_1$ into the same coordinate space, yielding $\tilde{\mathbf{u}}_1^*$. For each vector element \tilde{u}_k^* in $\tilde{\mathbf{u}}_1^*$, $k = 1, \dots, 4$, we calculated its rank $r_{k,i}$ among the 500 ordered values of $\hat{u}_{k,i}^*$, the corresponding elements from the vectors $\hat{\mathbf{u}}_{1i}^*$. We computed an empirical likelihood of membership for \tilde{u}_k^* among the $\hat{u}_{k,i}^*$ as

$$p_{k,i} = 1 - |250 - r_{k,i}|/250,$$

or $p_{k,i} = 0.001$, whichever was larger. In other words, likelihood of membership was proportional to rank difference (in either direction) from the median value of $\hat{u}_{k,i}^*$, but no smaller than the arbitrary minimum bound of 0.001. Given $\tilde{\mathbf{u}}_0$ and the set of decisions, we obtained a total expression of likelihood for $\tilde{\mathbf{u}}_1$ under model i as $P_{1i} = \prod_k p_{k,i}^{0.25}$. After performing the bootstrap analysis for the remaining models and obtaining those values of P_{1i} , we updated the information state vector by applying Bayes formula:

$$\pi_{1i} = \pi_{0i}P_{1i} / \sum_{j=1}^3 \pi_{0j}P_{1j}.$$

By this point, the simulation had arrived at decision period 1 with new values of \mathbf{y}_1 , $\tilde{\mathbf{y}}_1$, and $\boldsymbol{\pi}_1$, and it proceeded to find values of \mathbf{y}_2 , $\tilde{\mathbf{y}}_2$, and $\boldsymbol{\pi}_2$ in an identical manner as above.

We simulated adaptive decision-making under uncertainty for each underlying model, for each of the four initial states, and for each value σ in the set $\{0.0, 0.01, 0.02, 0.03, 0.04, 0.05\}$. Each trial simulated a 100-year time span of management, and we ran 500 simulation trials for each initial state \times model \times σ combination. We tracked the information state through time, particularly the probability value corresponding to the underlying model to determine under what conditions that value approached 1.0 (indicating correct identification of the model), and how quickly it occurred.

Results

Program ASDP did not achieve a stationary cutting policy for model F0 within 1,000 stage iterations. However, the decision policy changed only 11 times in the last 100 iterations, and each occurrence was limited to a single state of the 997 possible states. Therefore, it is likely that the policy produced at the final stage iteration very closely resembles a bona fide stationary (time-independent) policy. Optimal harvest values for four representative initial states of the forest reveal that the policy is indeed state-dependent: Harvest decisions are not generally equal across seral stages or unvarying among states (Table 2, Policy F0).

Distribution of cutting amounts within the optimal policy was highly dependent on choice of forest dynamics model (Figure 4). Policy F1 identified more states requiring aggressive harvesting of P2 and P3 forest than did policies F0 or F2. The no-harvest decision was identified more often in

Table 2. Optimal regeneration harvest decisions for each of four sample forest states at the Piedmont National Wildlife Refuge, under certainty with respect to each of three alternative forest dynamics models, and under uncertainty with respect to all models

	Initial forest state (seral stage compositions) ¹					Policy ²	Regeneration amounts (proportion taken) from seral stage		
	P1	P2	P3	P4	UH		P2	P3	P4
S1. Mostly young forest	0.47	0.38	0.03	0.03	0.09	F0	0.04	0.0	0.0
						F1	0.08	0.0	0.0
						F2	0.0	0.0	0.0
						Uncertain	0.08	0.0	0.0
S2. Even-aged forest	0.16	0.23	0.37	0.20	0.04	F0	0.0	0.0	0.0
						F1	0.0	0.0	0.0
						F2	0.02	0.0	0.0
						Uncertain	0.02	0.0	0.0
S3. Mostly mature forest	0.10	0.10	0.40	0.30	0.10	F0	0.0	0.0	0.0
						F1	0.0	0.0	0.0
						F2	0.0	0.0	0.0
						Uncertain	0.10	0.20	0.02
S4. Piedmont NWR, 2000 ³	0.03	0.15	0.57	0.07	0.18	F0	0.02	0.14	0.02
						F1	0.02	0.14	0.02
						F2	0.0	0.0	0.0
						Uncertain	0.0	0.14	0.08

¹ Seral stages are P1 (pine 0–16 yr), P2 (16–40 yr), P3 (40–80 yr), P4 (≥ 80 yr), and UH (upland hardwood). Values are proportions of total forest area.

² F0, decision policy for baseline model; F1, policy for alternative model of high hardwood competition; F2, policy for alternative model of low hardwood competition.

³ Estimated composition of the Refuge forest in year 2000 (Moore 2002).

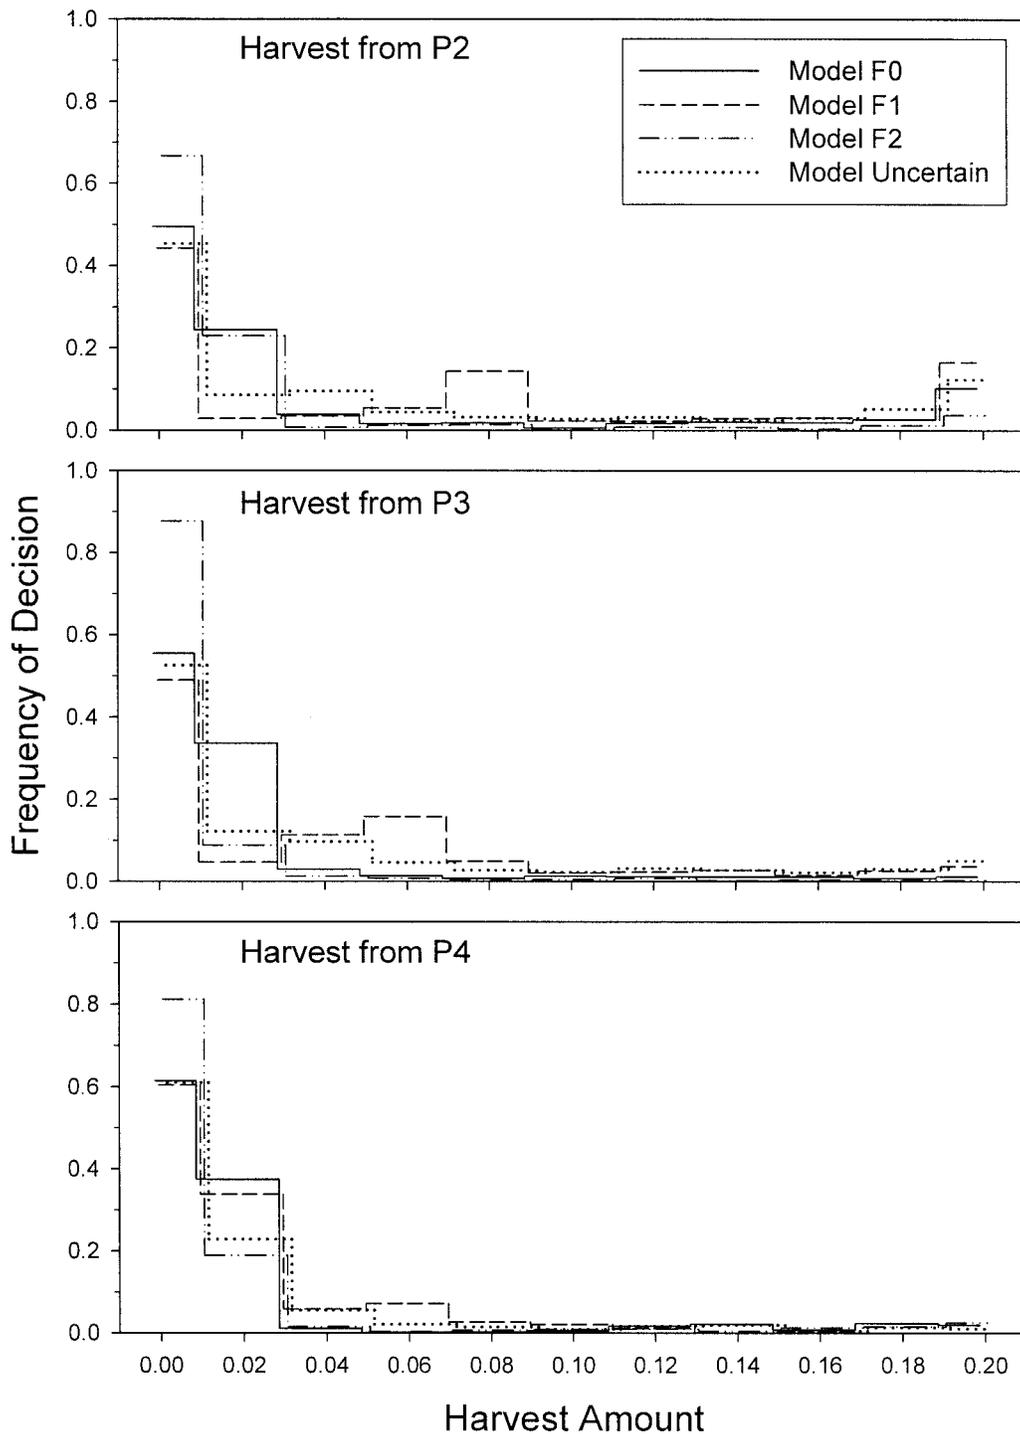


Figure 4. Distribution of age-specific (P2, P3, and P4) forest harvest amounts tabulated from optimal decision policies (997 system states per policy) that correspond to three alternative forest dynamics models for the Piedmont National Wildlife Refuge. Decision frequencies for a fourth policy (9,970 states), that of uncertainty among the three models, is also displayed.

policy F2 than in the others. The sensitivity of the decision to choice of model is also apparent for the four representative initial forest states presented earlier (Table 2). For this sample of states, decisions are least divergent among models when the forest occurs in a predominately mature state and most divergent for the contemporary (year 2000) state of the Refuge forest (Table 2).

The adaptive optimization procedure under uncertainty among the three models also failed to converge on a stationary solution; however, so few decisions changed in the last 100 iterations ($\bar{x} = 3.1$ of 9,970 information-forest states per iteration) that we concluded that the solution closely resembled a stationary policy. Under complete uncertainty with respect to model choice, one uses that part of

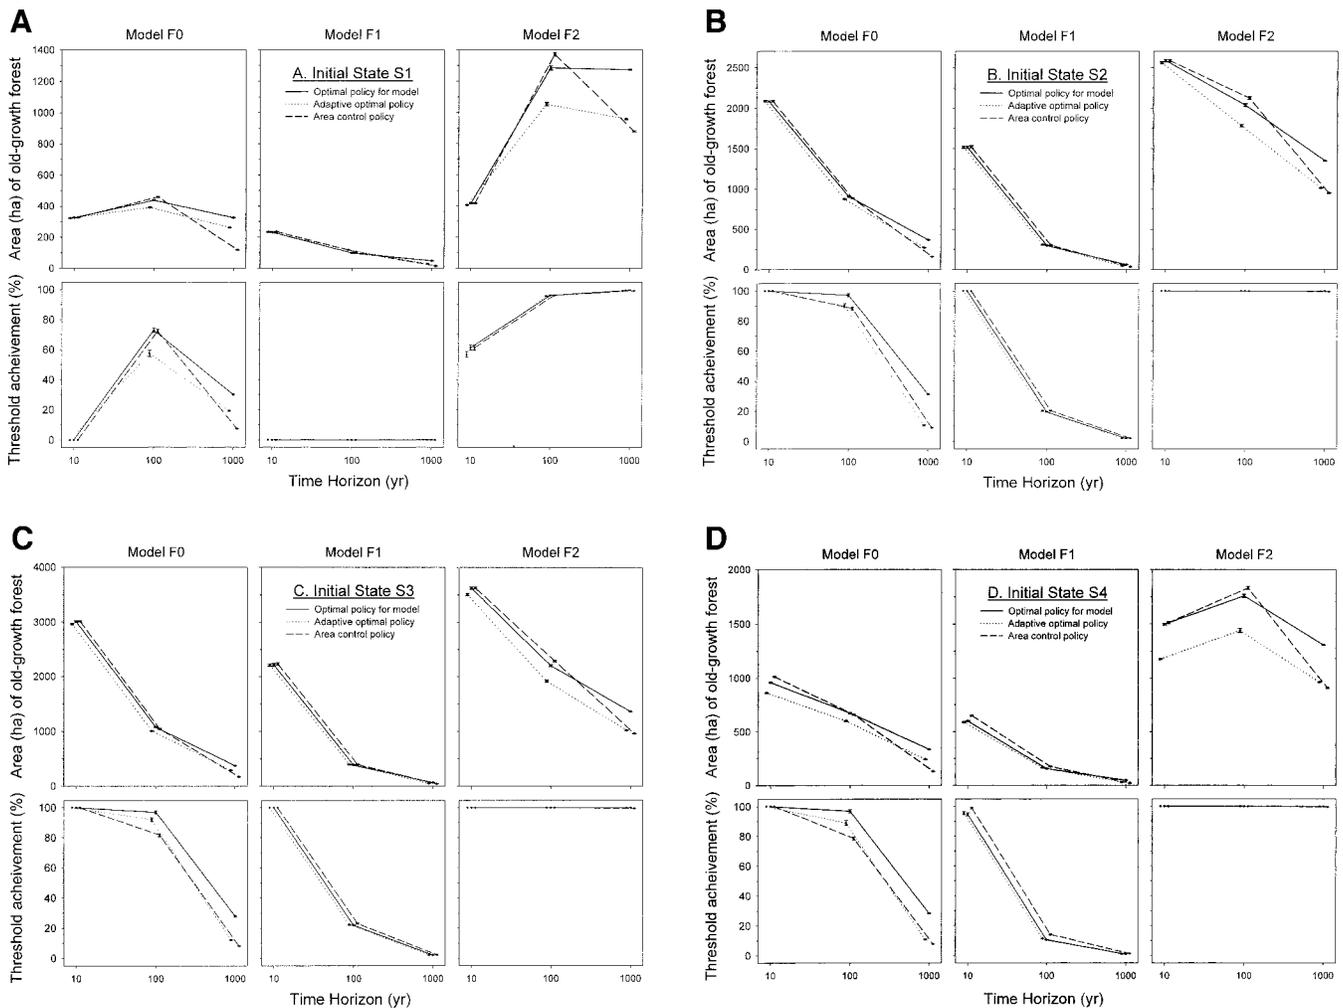


Figure 5. Results for decision policies simulated from four initial forest states (A–D) under three alternative forest dynamics models. Response variables are averages, over three time horizons (10, 100, 1,000 years), of annual area (ha) of old-growth forest habitat that occurs (upper plot panels) and annual success rate of exceeding the Refuge target habitat threshold of 388.5 ha (lower plot panels). Starting states (S1–S4) and models (F0–F2) are described in Table 2. For each response variable, simulation means and 99% confidence intervals (staggered on the time axis for visual clarity) are provided for the optimal policy corresponding to the simulated model (solid line), for the adaptive optimal policy under continuing uncertainty among the models (dotted line), and for the area control policy (dashed line).

the decision policy that corresponds to placing confidence weight $\frac{1}{3}$ on each model.

In general, decisions under complete uncertainty appear to be a compromise action among the competing models, with some notable exceptions. Under complete uncertainty, frequencies of different levels of harvest are intermediate relative to those found in the policies of model certainty (Figure 4). Similarly, for the sample of initial forest states, optimal decisions under uncertainty tend toward “average” decision values found under the certainty policies (Table 2). However, decisions for the initial state of mostly mature forest (Table 2, S3) portray a striking exception to this general pattern. Whereas 0.0 is the optimal cutting amount from all harvest classes under any model, the optimal action when choice of model is uncertain is to cut from all classes; substantially for the younger classes. Because such an action is not an optimal one when the model is known—whatever that model may be—this harvest action suggests a trading off of management performance in the

short term for a rapid gain of information intended to improve management over the long term. The fact that optimal decisions are similar under scenarios of certainty and uncertainty for the other initial states indicates that such aggressive “probing” for information (Walters and Hilborn 1978, Walters 1986) is too costly in terms of the long-term management objective and that resolution of uncertainty must proceed in a more risk-averse manner from these states.

Average (over the 997 discrete states of forest composition) optimal returns (expected annual average of proportion of forest in P4 habitat) were 0.0241, 0.00414, and 0.0788 for models F0, F1, and F2, respectively. These values compare to the model-averaged result of 0.0197, the mean optimal value expected under continuing model uncertainty. Average EVPI was 0.0160 (SE = 1.66×10^{-5} , min/max = 0.0133/0.0171), indicating that failing to resolve model uncertainty implies a resource sacrifice (i.e., an opportunity cost of management) of approximately 0.016 proportion

units, a substantial cost in light of the average optimal values above.

Analysis of each model simulated under its corresponding optimal policy demonstrated that (1) expected annual area in old-growth pine forest (Figure 5A–D, upper plot panels) approached long-term stable values of >1,000 ha for model F2, <100 ha for model F1, and some intermediate value for model F0; (2) average annual frequency of old-growth area exceeding the habitat goal threshold of 388.5 ha (Figure 5A–D, lower plot panels) approached 100% under model F2, 0% under model F1, and some intermediate value under model F0; and (3) direction of these trends depended on starting state.

Simulation of the models under the adaptive decision policy but assuming continuing uncertainty indicated some sacrifice in old-growth habitat (Figure 5A–D, dotted lines) relative to the model-certain policies (solid lines), particularly over the long time frame. This is not surprising, as the models were simulated without performing model likelihood updating, i.e., the system was treated artificially as if uncertainty is never resolved through time. In practice (and as carried out in the fully adaptive simulation below), as management experience is accumulated, evidence toward or against specific models is expressed in the reassignment of model likelihoods over time, and optimal decisions begin to more closely reflect those from one of the model-certain policies. Thus, losses in system performance due to model uncertainty would diminish over time as more is learned about the system. Simulation results for the area control policy (Figure 5A–D, dashed lines) reveal that the policy always provides less habitat over the long term (1,000 years), even relative to the policy under continuing model uncertainty.

Simulations of fully adaptive decision-making (updating of the information state based on forest response monitoring) under uncertainty indicated that the underlying models were quickly identified (belief weight assigned to underlying model ≥ 0.98 after two annual cycles of forest harvesting) if forest composition was measured without error (Figure 6). The simulation model generated little variability in individual age class components of the observed forest state vector (CV often 5% or less), even at the maximum value of 0.05 for the observation error parameter σ . However, ability to correctly identify the underlying models F0 or F2 diminished rapidly as observation error increased (Figure 6), with models most often misidentified as F1.

For the estimated composition of the Refuge in year 2000, optimal total area of regeneration cutting was 0 ha for model F2, 1,056 ha for models F0 and F1, and 1,071 ha under model uncertainty. During the decade 1989–1998, the average area of pine forest annually regenerated was 53.6 ha (J.D. Metteauer, personal communication, Piedmont National Wildlife Refuge, 2001), an amount more consistent with management under belief in model F2 than belief in any other model or uncertainty with respect to all models.

Discussion

Scope and Limitations

The adaptive optimization and informed decision-making framework we have presented is applicable in any forest management setting in which structural uncertainty can be described, an objective of management exists, management occurs as a sequence of decisions with clear alternatives at each stage, and some form of system monitoring provides periodic feedback of information. Furthermore, because information can be obtained and applied to the reduction of uncertainty whether or not decisions happen to be optimal ones, there is no strict requirement for dynamic optimization for carrying out adaptive forest management. Thus, alternative techniques such as simulation, nondynamic (single-step) optimization, or heuristic procedures (e.g., Moore et al. 2000, Conroy and Moore 2001) can be used to provide reasonable approximations to optimal dynamic decisions under uncertainty, yet the steps of decision-making, monitoring, and model updating can nevertheless be followed to reduce uncertainty before selection of any subsequent decision.

The matrix structure of our model is one commonly used for modeling linear dynamics of stage-based systems (Caswell 2001), including many forest management applications (e.g., Lembersky and Johnson 1975, Reed and Errico 1986, Gassmann 1989, van Kooten et al. 1992, Teeter et al. 1993). Often, stages represent the compositional parts of a whole, such as proportions of total forest area in age classes, density classes, species types, or some other categorical attribute. The model can be made spatially explicit by including components representing two or more spatial units and by including parameters representing reclassification of forest area within units conditional on unit membership. With added complexity, however, optimization becomes more challenging.

Of course, our implementation of the model is specific to this application of upland pine forest harvest management in the Piedmont region of central Georgia, in which pine forest dynamics are believed to be governed by harvesting, hardwood succession, and forest disturbance. However, all of the parameter selections (Table 1) were subjective and received only circumstantial support from empirical sources. Therefore, there is little basis to argue why our model would be less appropriate in similar forest and physiographic settings within or outside of the state. But it was this very issue of uncertainty in a data-poor climate that cautioned us to not place full faith in any single model but to instead consider highly distinct, but plausible, alternatives to the baseline model. Our hope is that other environments of uncertainty will prompt forest managers to propose alternative models that try to frame the limits of uncertainty, especially models that challenge existing data or conventional wisdom.

Our approach to optimization of forest habitat is distinctly different from the scheduling approaches used by Bettinger et al. (1997), Moore et al. (2000), and Bettinger et al. (2002). Our approach recognizes that the managed forest

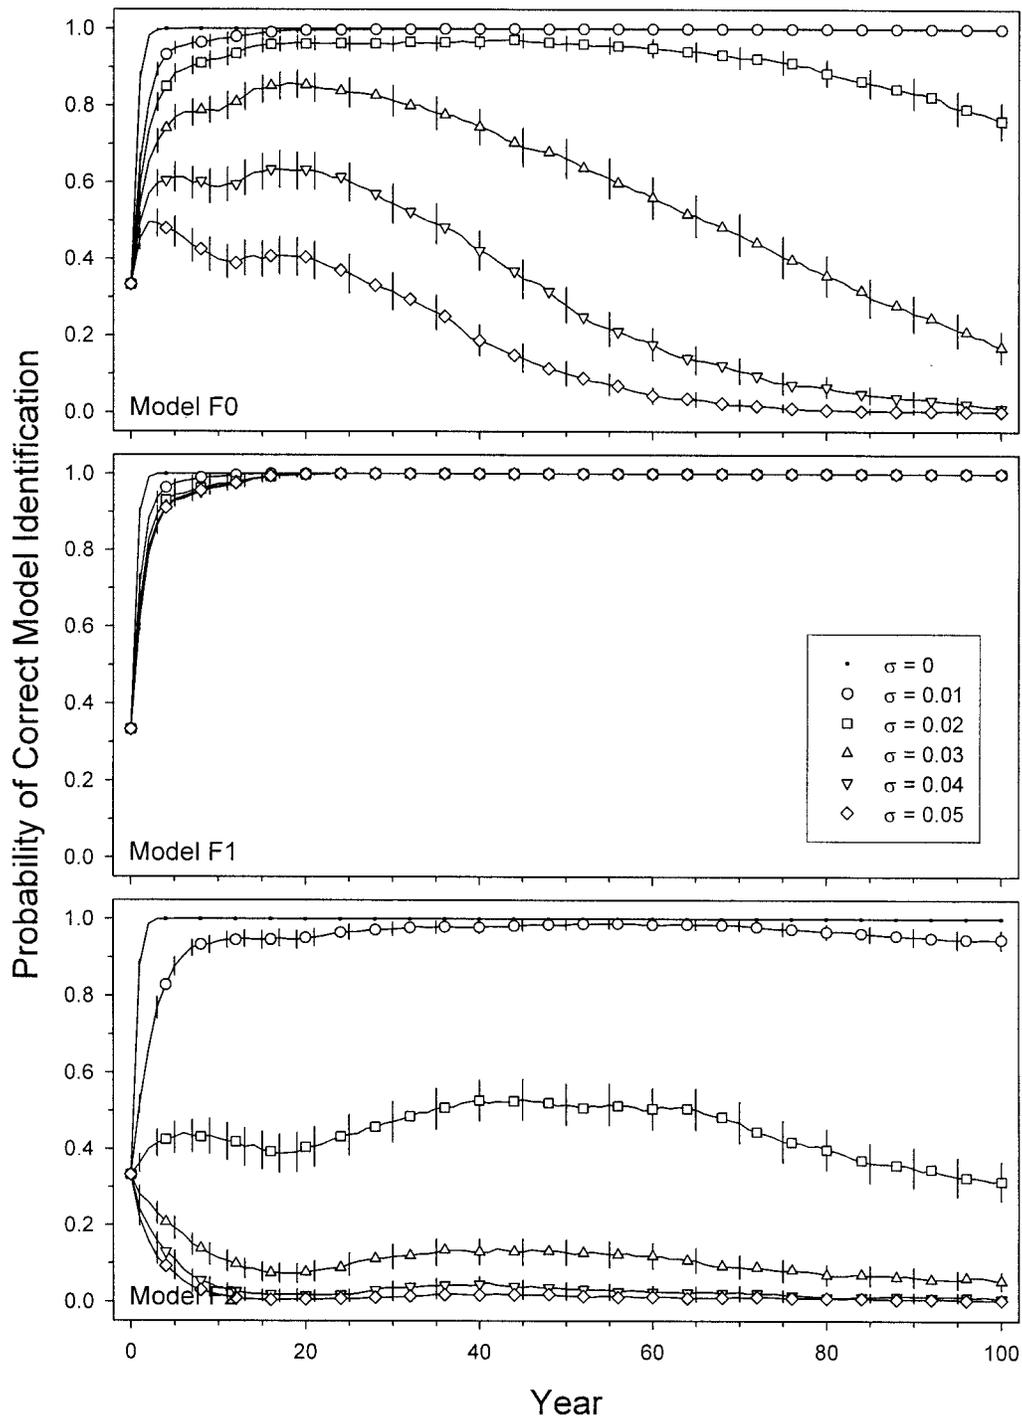


Figure 6. Simulation of adaptive decision-making and updating of the information state under three alternative forest models, starting from an initial forest state of balanced age classes (state S2; similar response patterns were produced for three other initial states). Responses are means and (selected) 99% confidence intervals of probability weight assigned to the underlying model when the forest state is measured under alternative levels of observation error (σ).

system is dynamic and responsive to stochastic perturbations rather than static and deterministic. As such, any future state of the forest can be anticipated as a consequence of its previous state and intervening management decisions and stochastic events. An optimal dynamic policy provides an optimal decision for any arbitrary state of the forest at each decision opportunity, whereas a harvest schedule of-

fers a fixed prescription of management actions regardless of unpredictable events that transpire over the management time frame. Of course, a harvest schedule can be re-computed at any time in response to such events, but this, by definition, constitutes a suboptimal policy, as the course correction was not anticipated at the beginning of the time frame (Dreyfus and Law 1977).

One obvious limitation of our model is the fact that the management objective of maximizing old-growth forest does not, in itself, address the habitat needs of the red-cockaded woodpecker, nor does the model take into account temporal and spatial dynamics of the existing woodpecker population. Old-growth forest patches may become suitable woodpecker nesting patches only if certain additional conditions exist: a sparse overstory of large-diameter trees, almost complete absence of a hardwood midstory, a fire-maintained ground cover, and spatial connectivity of the patch to source patches and foraging areas (Walters 1991, Loeb et al. 1992, Letcher et al. 1998, Azevedo et al. 2000, US Fish and Wildlife Service 2003). However, woodpecker family groups, forest patches, and the fine-scale dynamics among them exist in the context of a whole-forest matrix of age and type classes whose dynamics play out over decades and centuries of management. We believed that no formal approach that simultaneously accommodated all of these processes at their disparate scales, even if one existed, would provide decisions of perceptibly higher quality than would an informal approach that decomposed the problem into hierarchical, conditionally independent pieces; Boston and Bettinger (2001) made such a conclusion in their assessment of a two-stage approach to forest harvest scheduling with spatial habitat constraints for the red-cockaded woodpecker. Therefore, we separated both the management objectives and the modeling and optimization tasks along these lines of scale to find best forest management strategies, under uncertainty, for woodpecker conservation goals. The purpose of our dynamic optimization model described above was to provide annual decisions on total cutting amounts that lead to best old-growth outcomes at the whole-forest level over a very long time frame. A subsequent step (Moore 2002) uses a spatially explicit model of woodpecker clusters and forest growth to allocate those cutting amounts and prescribed fire quotas into specific patches, where the allocations are constrained by fixed selection and adjacency rules, to maximize number of woodpecker clusters over a 100-year time frame. Thus, fine-scale, spatially explicit decisions intended to increase woodpecker population size are predicated on harvest decisions chosen to meet objectives at the aggregated, whole-forest scale over a multirotation time horizon; Bettinger and Sessions (2003) distinguished these as “tactical” and “strategic” planning activities, respectively.

Current Regeneration Approaches

Under our stage dynamics model, the optimal amount of total harvest was dependent on the age class composition of the forest, and amounts harvested were generally unequal among age classes. This contrasts with the approach prescribed by the area control method, which calls for a constant rate of harvest through time, with harvests taken in equal proportion from all age classes. In light of our simulation results, we suggest that the area control method, the regeneration method recommended for even-aged and two-aged systems on federal refuges (US Fish and Wildlife

Service 1998), is suboptimal for the generation of woodpecker old-growth forest habitat.

Of course, this claim depends on how well our forest model represents growth dynamics of the Refuge forest. The structure of this model was simple and integrates generally understood forest succession processes in a straightforward way. Whereas we believe our model structure is a reasonable one, one could (and should) question our choice of parameter values used in the model, as these quantities are highly uncertain. By proposing two alternatives to our baseline model, we acknowledged this uncertainty and were able to carry out a crude sensitivity analysis with respect to the optimal decision policy. These alternative parameterizations represented fairly extreme departures from the baseline model. Yet in each case, we found the area control method was inferior to other regeneration strategies for increasing and maintaining old-growth habitat. There may exist a model of forest growth and a set of forest states for which the area control method is optimal, but in our view the requisite assumptions and applicable situations have not been explicitly spelled out.

One consequence of model differences was disagreement in optimal amount of total forest area to regenerate for the estimated forest conditions that occurred on the Refuge in 2000: 0 ha under model F2 versus 1,056 ha under models F0 or F1. The average amount (53.6 ha) actually harvested each year during the preceding decade more closely corresponded to the optimal harvest amount called for under model F2 than under the other models. In other words, recent harvest practice on the Refuge was consistent with an optimal harvest practice corresponding to a relatively optimistic view of old-growth forest generation and persistence. However, to carry out such actions when a more appropriate model of dynamics is operative can be extremely suboptimal, leading to rapid depletion of the old-growth component and loss of future opportunities for its regeneration. Whether the recent practices are in fact optimal ones requires assessing the relative quality of each of these models (or any other competing model) in predicting forest response to management actions and to the stochastic environment. A program of regular monitoring is therefore essential in making informed decisions about forest management.

Structural Uncertainty and Decision-Making

The optimal policy was highly sensitive to the choice of forest model. Managers' decisions about what and how much to harvest are based not only on current forest state but also on the model thought to best approximate the system. This is true in all forest decision problems, whether the model is written as a formal mathematical construct or exists only conceptually in the manager's mind. Thus structural uncertainty, if unresolved, poses a significant challenge to forest managers. Indeed, we have demonstrated that the recovery goal of 388.5 ha of habitat may not even be achievable under particular assumptions, with strategic implications for rangewide species recovery and for objectives

of the National Wildlife Refuge system. For example, is recovery of the woodpecker in the Piedmont physiographic province a priority if doing so requires forest management to become more species-oriented than ecosystem oriented?

EVPI provides an estimate of the management cost of structural uncertainty, measured in units of the resource. Our estimate of EVPI (averaged over all forest states) was 0.016, or approximately 200 ha of P4 habitat. In other words, if uncertainty among these three models is never resolved, the expected sacrifice of woodpecker nesting habitat is about 200 ha, over half of the habitat goal. This result should help convince managers that direct fiscal savings achieved by not implementing measures that would reduce uncertainty (e.g., monitoring programs, model development and analysis) may be offset by “opportunity” costs in species conservation (Moore and Kendall 2004).

Computing an optimal harvest policy with respect to all candidate models, with models weighted by their relative degree of credibility, appears to be a reasonable way to proceed under uncertainty. There are two general approaches to policy computation, and they differ in how future information from the system is expected to be used in assessing credibility of each model. Passive adaptive optimization ignores the return of such information and proceeds as if model credibility weights will remain fixed through time (Johnson and Williams 1999, Williams et al. 2002). Of course, this does not have to be the case in practice, and a new passive policy can be computed whenever model credibility weights are updated (Williams et al. 2002). Our optimization of model-averaged returns for the computation of EVPI was an example of passive adaptive optimization.

In contrast, we used active adaptive optimization (Johnson and Williams 1999, Williams et al. 2002). Relative to passive optimization, it is a forward-looking approach that exploits the fact that learning usually occurs as decisions are made through time. In a circumstance where two alternative decision actions are expected to provide about the same short-term resource return (e.g., similar amounts of old-growth forest habitat expected under either cutting decision), an actively adaptive procedure will favor the option that produces the most divergent system response with respect to all the candidate models. Thus, when outcomes of the decision are then compared to outcomes predicted by each model, chances have been increased that one or more models will be distinguished from the others in terms of predictive credibility. In other words, the decision process is used fruitfully, and with little sacrifice of returns over the short term, to gain knowledge for making better decisions over the long term. A comparison of optimal decisions under certainty and uncertainty for a forest in a mostly mature state (Table 2, initial state S3) is evidence of an opportunity in which information can be more aggressively pursued for relatively little short-term management sacrifice (for comparison, the optimal decision for this state under passive optimization was 0.0 total harvest, exactly as found for each of the alternative models). Whereas passive optimization seeks to maximize the long-term resource return

with respect to an unvarying state of uncertainty, active optimization anticipates that uncertainty will change through time.

Adaptive Management

Improvements in management occur only when uncertainties about responses to management are reduced, and reductions in uncertainty only come about when system responses are compared to predicted outcomes. System monitoring and informative updating of prediction models are requisite components of adaptive management (Nichols et al. 1995, Williams 1997, Moore et al., in press). Despite the ubiquity of the term “adaptive management” throughout the natural resources literature, we suspect that few systems are truly managed adaptively, as these two critical elements are rarely both in place.

At the Piedmont National Wildlife Refuge, a systematic (spatially and temporally) assessment of forest cover types would serve the purpose of updating our models. Our simulations suggest that uncertainty can be resolved quickly under such a monitoring system. Unfortunately, none has been established, and consequently there is no means of resolving uncertainty about system management or evaluating decision-making.

Four general forms of uncertainty hamper decision-making in natural resource management: stochastic uncertainty, structural uncertainty, partial controllability, and partial observability (Nichols et al. 1995, Williams 1997). Many researchers have previously addressed the topic of stochastic uncertainty in decision-making for dynamic forest systems (Lembersky and Johnson 1975, Reed and Errico 1986, Gassmann 1989, van Kooten et al. 1992, Teeter et al. 1993), but none to our knowledge has formally addressed the problem of dynamic optimization under structural uncertainty, such as the kind we have described. Under partial controllability, the action carried out is not the decision intended. For example, the width of fire lines (breaks designed to control the spread of fire past designated limits) is a decision variable for controlling probability of holding a wildfire, but additional uncertainty in wildfire outcome is introduced because constructed fire lines generally do not match their targeted width (Mees and Strauss 1992).

Our simulations also explored the influence of partial observability, in which the true state of the system is not observable and therefore not fully informative either for decision-making or for updating system models. For example, Eid (2000) noted that ignored sampling errors and biases from forest inventories translated into lost timber revenues. In our analyses, failing to account for errors in measuring forest composition affected the inferences made on competing models and in turn distorted the course of optimal decision-making. Observational errors caused system dynamics under any model to resemble those under model F1, and optimal decisions were then driven according to that model. We did not simulate another common example of partial observability in which forest composition is measured precisely each time but at measurement intervals

longer than decision intervals. Optimal adaptive policies under partial observability may be found, but the computational requirements may be considerable (Williams 1996a,b). Monitoring programs should be designed to reduce observation error as much as practicable, and we suggest the use of simulation to assess the sensitivity of decision-making to expected degrees of observation error. Not all of the above forms of uncertainty may bear on a given decision problem (Kangas and Kangas 2004), but they deserve greater attention as the focus of forest management on public lands continues to broaden beyond the traditional scope of commodity production.

We recommend an adaptive program of decision-making whenever uncertainty accompanies the management of wildlife habitats. That is, models should be used to forecast outcomes of decisions, and those models should be evaluated on the basis of information gathered through system monitoring. We especially advise this approach whenever managers must maintain continuous streams of an ephemeral habitat for the sustenance of an endangered species population. As the recovery of the red-cockaded woodpecker population is focused almost entirely on federal lands (US Department of Defense, US Forest Service, US Fish and Wildlife Service, National Park Service, US Department of Energy), federal management agencies should receive the resources and administrative commitment needed to start and sustain such efforts.

Literature Cited

- ANDERSON, D.R. 1975. Optimal exploitation strategies for an animal population in a Markovian environment: A theory and an example. *Ecology* 56:1281–1297.
- AZEVEDO, J.C.M., S.B. JACK, R.N. COULSON, AND D.F. WUNNEBURGER. 2000. Functional heterogeneity of forest landscapes and the distribution and abundance of the red-cockaded woodpecker. *For. Ecol. Manage.* 127:271–283.
- BECHTOLD, W.A., G.A. RUARK, AND F.T. LLOYD. 1991. Changing stand structure and regional growth reductions in Georgia's natural pine stands. *For. Sci.* 37:703–717.
- BELLMAN, R. 1957. *Dynamic programming*. Princeton University Press, Princeton, NJ. 342 p.
- BETTINGER, P., D. GRAETZ, K. BOSTON, J. SESSIONS, AND W. CHUNG. 2002. Eight heuristic planning techniques applied to three increasingly difficult wildlife planning problems. *Silva Fennica* 36:561–584.
- BETTINGER, P., AND J. SESSIONS. 2003. Spatial forest planning: To adopt, or not to adopt? *J. For.* 101(2):24–29.
- BETTINGER, P., J. SESSIONS, AND K. BOSTON. 1997. Using Tabu search to schedule timber harvests subject to spatial wildlife goals for big game. *Ecol. Model.* 94:111–123.
- BOSTON, K., AND P. BETTINGER. 2001. Development of spatially feasible forest plans: A comparison of two modeling approaches. *Silva Fennica* 35:425–435.
- BULTE, E., D.P. VAN SOEST, G.C. VAN KOOTEN, AND R.A. SCHIPPER. 2002. Forest conservation in Costa Rica when nonuse benefits are uncertain but rising. *Am. J. Agric. Econ.* 84:150–160.
- CASWELL, H. 2001. *Matrix population models*, 2nd ed. Sinauer Associates, Sunderland, MA. 722 p.
- CONNER, R.N., A.E. SNOW, AND K.A. O'HALLORAN. 1991. Red-cockaded woodpecker use of seed-tree/shelterwood cuts in eastern Texas. *Wildl. Soc. Bull.* 19:67–73.
- CONROY, M.J., AND C.T. MOORE. 2001. Simulation models and optimal decision making in natural resource management. P. 91–104 *in* *Modeling in natural resource management: Development, interpretation, and application*, Shenk, T.M., and A.B. Franklin (eds.). Island Press, Washington, DC.
- CONROY, M.J., AND C.T. MOORE. 2002. Wildlife habitat modeling in an adaptive framework: The role of alternative models. P. 205–218 *in* *Predicting species occurrences: Issues of accuracy and scale*. Scott, J.M., P.J. Heglund, M.L. Morrison, J.B. Haufler, M.G. Raphael, W.A. Wall, and F.B. Samson (eds.). Island Press, Covelo, CA.
- CZUHAI, E., AND C.T. CUSHWA. 1968. A resume of prescribed burnings on the Piedmont National Wildlife Refuge. USDA For. Serv. Res. Note SE-86. 4 p.
- DREYFUS, S.E., AND A.M. LAW. 1977. *The art and theory of dynamic programming*. Academic Press, New York. 284 p.
- DUCEY, M.J. 2001. Representing uncertainty in silvicultural decisions: An application of the Dempster-Shafer theory of evidence. *For. Ecol. Manage.* 150:199–211.
- EID, T. 2000. Use of uncertain inventory data in forestry scenario models and consequential incorrect harvest decisions. *Silva Fennica* 34:89–100.
- FRANKLIN, J.F. 1997. Ecosystem management: An overview. P. 21–53 *in* *Ecosystem management*, Boyce, M.S., and A. Haney (eds.). Yale University Press, New Haven, CT.
- GABRIELSON, I.N. 1943. *Wildlife refuges*. Macmillan, New York. 257 p.
- GASSMANN, H.I. 1989. Optimal harvest of a forest in the presence of uncertainty. *Can. J. For. Res.* 19:1267–1274.
- GLUSS, B. 1972. *An elementary introduction to dynamic programming*. Allyn and Bacon, Boston, MA. 402 p.
- HOOPER, R.G., AND C.J. MCADIE. 1995. Hurricanes and the long-term management of the red-cockaded woodpecker. P. 148–166 *in* *Red-cockaded woodpecker: Recovery, ecology and management*, Kulhavy, D.L., R.G. Hooper, and R. Costa (eds.). Center for Applied Studies in Forestry, Stephen F. Austin State University, Nacogdoches, TX.
- JACKSON, J.A. 1994. Red-cockaded woodpecker (*Picoides borealis*). P. 1–20 *in* *Birds of North America* (No. 85), Poole, A., and F. Gill (eds.). Acad. Natur. Sci., Philadelphia, PA, and Am. Ornithol. Union, Washington, DC.
- JOHNSON, F., AND K. WILLIAMS. 1999. Protocol and practice in the adaptive management of waterfowl harvests. *Conserv. Ecol.* 3(1):8 [online URL: <http://www.consecol.org/vol3/iss1/art8>].
- KANGAS, A.S., AND J. KANGAS. 2004. Probability, possibility and evidence: Approaches to consider risk and uncertainty in forestry decision analysis. *For. Policy Econ.* 6:169–188.

- LEMBERSKY, M.R., AND K.N. JOHNSON. 1975. Optimal policies for managed stands: An infinite horizon Markov decision process approach. *For. Sci.* 21:109–122.
- LETCHER, B.H., J.A. PRIDDY, J.R. WALTERS, AND L.B. CROWDER. 1998. An individual-based, spatially-explicit simulation model of the population dynamics of the endangered red-cockaded woodpecker, *Picoides borealis*. *Biol. Cons.* 86:1–14.
- LIGON, J.D., P.B. STACEY, R.N. CONNER, C.E. BOCK, AND C.S. ADKISSON. 1986. Report of the American Ornithologists' Union Committee for the Conservation of the Red-Cockaded Woodpecker. *Auk* 103:848–855.
- LINDLEY, D.V. 1985. Making decisions, 2nd ed. John Wiley and Sons, London, UK. 207 p.
- LOEB, S.C., W.D. PEPPER, AND A.T. DOYLE. 1992. Habitat characteristics of active and abandoned red-cockaded woodpecker colonies. *South. J. Appl. For.* 16:120–125.
- LUBOW, B.C. 1995. SDP: Generalized software for solving stochastic dynamic optimization problems. *Wildl. Soc. Bull.* 23:738–742.
- LUBOW, B.C. 1997. Adaptive stochastic dynamic programming (ASDP): Supplement to SDP user's guide, version 2.0. Colo. Coop. Fish and Wildlife Res. Unit, Colo. State Univ., Fort Collins, CO.
- MCCARTHY, M.A., H.P. POSSINGHAM, AND A.M. GILL. 2001. Using stochastic dynamic programming to determine optimal fire management for *Banksia ornata*. *J. Appl. Ecol.* 38:585–592.
- MEES, R., AND D. STRAUSS. 1992. Allocating resources to large wildland fires—A model with stochastic production-rates. *For. Sci.* 38:842–853.
- MOORE, C.T. 2002. Forest decision making under uncertainty: Adaptive management for the conservation of bird populations on a national wildlife refuge. PhD dissertation, University of Georgia, Athens, GA. 333 p.
- MOORE, C.T., M.J. CONROY, AND K. BOSTON. 2000. Forest management decisions for wildlife objectives: System resolution and optimality. *Comput. Electron. in Agric.* 27:25–39.
- MOORE, C.T., AND W.L. KENDALL. 2004. Costs of detection bias in index-based population monitoring. *Anim. Biodivers. Conserv.* 27(1):287–296.
- MOORE, C.T., W.T. PLUMMER, AND M.J. CONROY. 2005. Forest management under uncertainty for multiple bird population objectives. *Bird Conservation Implementation and Integration in the Americas: Proc. of the Third International Partners in Flight Conference 2002*, Ralph, C.J., and T.D. Rich (eds.) (two volumes). USDA For. Serv., PSW-GTR-191.
- NEMHAUSER, G.L. 1966. Introduction to dynamic programming. John Wiley and Sons, New York. 256 p.
- NICHOLS, J.D., F.A. JOHNSON, AND B.K. WILLIAMS. 1995. Managing North American waterfowl in the face of uncertainty. *Annu. Rev. Ecol. Syst.* 26:177–199.
- PETERSON, C.J. 2000. Catastrophic wind damage to North American forests and the potential impact of climate change. *Sci. Total Environ.* 262:287–311.
- PUTERMAN, M.L. 1994. Markov decision processes. John Wiley and Sons, New York. 649 p.
- REED, W.J., AND D. ERRICO. 1986. Optimal harvest scheduling at the forest level in the presence of the risk of fire. *Can. J. For. Res.* 16:266–278.
- ROSE, S.K., AND D. CHAPMAN. 2003. Timber harvest adjacency economics, hunting, species protection, and old growth value: Seeking the dynamic optimum. *Ecol. Econ.* 44:325–344.
- SMITH, D.M. 1962. The practice of silviculture, 7th ed. John Wiley and Sons, New York. 578 p.
- TEETER, L., G. SOMERS, AND J. SULLIVAN. 1993. Optimal forest harvest decisions: A stochastic dynamic programming approach. *Agric. Syst.* 42:73–84.
- THOMAS, J.W. 1997. Foreword. P. ix-xii in *Creating a forestry for the 21st century*, Kohm, K.A., and J.F. Franklin (eds.). Island Press, Washington, DC.
- US FISH AND WILDLIFE SERVICE. 1998. Strategy and guidelines for the recovery and management of the red-cockaded woodpecker and its habitats on national wildlife refuges. US Dept. Interior, Fish and Wildl. Serv., Atlanta, GA. 51 p.
- US FISH AND WILDLIFE SERVICE. 2003. Recovery plan for the red-cockaded woodpecker (*Picoides borealis*), 2nd revision. US Dept. Interior, Fish and Wildl. Serv., Atlanta, GA. 296 p.
- VAN KOOTEN, G.C., R.E. VAN KOOTEN, AND G.L. BROWN. 1992. Modeling the effect of uncertainty on timber harvest: A suggested approach and empirical example. *J. Agric. Res. Econ.* 17:162–172.
- WALTERS, C.J. 1986. Adaptive management of renewable resources. MacMillan, New York. 374 p.
- WALTERS, C.J., AND R. HILBORN. 1978. Ecological optimization and adaptive management. *Annu. Rev. Ecol. Syst.* 9:157–188.
- WALTERS, J.R. 1991. Application of ecological principles to the management of endangered species: The case of the red-cockaded woodpecker. *Annu. Rev. Ecol. Syst.* 22:505–523.
- WILLIAMS, B.K. 1996a. Adaptive optimization and the harvest of biological populations. *Math. Biosci.* 136:1–20.
- WILLIAMS, B.K. 1996b. Adaptive optimization of renewable natural resources: Solution algorithms and a computer program. *Ecol. Model.* 93:101–111.
- WILLIAMS, B.K. 1997. Approaches to the management of waterfowl under uncertainty. *Wildl. Soc. Bull.* 25:714–720.
- WILLIAMS, B.K., J.D. NICHOLS, AND M.J. CONROY. 2002. Analysis and management of animal populations. Academic Press, San Diego, CA. 817 p.