

Systems identification and the adaptive management of waterfowl in the United States

Authors: Williams, Byron K., and Nichols, James D.

Source: Wildlife Biology, 7(3) : 223-236

Published By: Nordic Board for Wildlife Research

URL: <https://doi.org/10.2981/wlb.2001.027>

BioOne Complete (complete.BioOne.org) is a full-text database of 200 subscribed and open-access titles in the biological, ecological, and environmental sciences published by nonprofit societies, associations, museums, institutions, and presses.

Your use of this PDF, the BioOne Complete website, and all posted and associated content indicates your acceptance of BioOne's Terms of Use, available at www.bioone.org/terms-of-use.

Usage of BioOne Complete content is strictly limited to personal, educational, and non - commercial use. Commercial inquiries or rights and permissions requests should be directed to the individual publisher as copyright holder.

BioOne sees sustainable scholarly publishing as an inherently collaborative enterprise connecting authors, nonprofit publishers, academic institutions, research libraries, and research funders in the common goal of maximizing access to critical research.

Systems identification and the adaptive management of waterfowl in the United States

Byron K. Williams & James D. Nichols

Williams, B.K. & Nichols, J.D. 2001: Systems identification and the adaptive management of waterfowl in the United States. - *Wildl. Biol.* 7: 223-236.

Waterfowl management in the United States is one of the more visible conservation success stories in the United States. It is authorized and supported by appropriate legislative authorities, based on large-scale monitoring programs, and widely accepted by the public. The process is one of only a limited number of large-scale examples of effective collaboration between research and management, integrating scientific information with management in a coherent framework for regulatory decision-making. However, harvest management continues to face some serious technical problems, many of which focus on sequential identification of the resource system in a context of optimal decision-making. The objective of this paper is to provide a theoretical foundation of adaptive harvest management, the approach currently in use in the United States for regulatory decision-making. We lay out the legal and institutional framework for adaptive harvest management and provide a formal description of regulatory decision-making in terms of adaptive optimization. We discuss some technical and institutional challenges in applying adaptive harvest management and focus specifically on methods of estimating resource states for linear resource systems.

Key words: adaptive management, dynamic optimization, harvest, mallards, population models, sport hunting, waterfowl

Byron K. Williams, U.S. Geological Survey, Cooperative Research Units, 12201 Sunrise Valley Drive, Reston, Virginia 22092, USA - e-mail: Byron_Ken_Williams@usgs.gov

James D. Nichols, U.S. Geological Survey, Patuxent Wildlife Research Center, 11510 American Holly Drive, Laurel, Maryland 20708-4017, USA

The regulation of waterfowl harvests in the United States involves public announcements, deliberations, and joint decision making by the federal and state governments. The federal government derives its responsibility for establishing sport-hunting regulations from the Migratory Bird Treat Act of 1918 (as amended), which implements provisions of international treaties for migratory bird conservation. The Act directs the Secretary of Agriculture to periodically adopt hunting regulations for migratory birds, "having due regard to the zones of temperature and to the distribution, abundance, economic value, breeding habits, and times and lines of migratory flight of such birds" (U.S. Department of the Interior 1975). The responsibility for

managing migratory bird harvests has since been passed to the Secretary of the Interior and the U.S. Fish and Wildlife Service. Other legislative acts, such as the National Environmental Policy Act, the Endangered Species Act, the Administrative Procedures Act, the Freedom of Information Act and the Regulatory Flexibility Act, provide additional responsibilities in the development of hunting regulations, and help define the nature of the regulatory process (Blohm 1989).

An essential element of harvest regulations is the annual collection and analysis of data on breeding population status, harvest levels, survival, production, migration and other population characteristics (Smith, Blohm, Kelly & Reynolds 1989). Long-term databases

of monitoring data are used to estimate key population parameters such as survivorship and reproduction, and to predict harvest impacts on population dynamics (Nichols, Conroy, Anderson & Burnham 1984, Johnson, Sparling & Cowardin 1987, Johnson, Nichols & Schwartz 1992). The information thus accumulated is folded into waterfowl population 'models', which in turn are used to inform the regulations process (Cowardin & Johnson 1979, Johnson, Nichols, Conroy & Cowardin 1988, Williams & Nichols 1990).

Though biologists have long recognized a need for informative harvest management (Anderson & Burnham 1976, Nichols et al. 1984, Montalbano, Johnson, Miller & Rusch 1988, Williams & Nichols 1990), traditionally the regulation of waterfowl harvests has not focused on uncertainty about regulations impacts. Nor has the opportunity to use regulations to reduce uncertainty been exploited. An unfortunate result is an unnecessarily slow rate of learning about population dynamics, and a correspondingly slow rate of improvement in management over time.

The harvest regulations process

Most waterfowl hunting regulations are established annually, within a timetable that is constrained by the timing of biological surveys and the need to give states and the public an opportunity to influence regulations. The annual regulatory cycle includes analysis and interpretation of biological data, development of regulatory proposals and solicitation of public comment, leading in turn to the promulgation and publication of hunting regulations in the autumn of each year.

A key component of the regulatory process consists of data collected each year on population status, habitat conditions, production, harvest levels and other system attributes of management interest (Smith et al. 1989). Population and habitat monitoring is essential for discerning resource status and modifying hunting regulations in response to changes in environmental conditions. The system of waterfowl monitoring in North America is unparalleled in its scope and is made possible only by the cooperative efforts of the U.S. Fish and Wildlife Service, the Canadian Wildlife Service, state and provincial wildlife agencies and various research institutions.

Each year monitoring data are used to estimate key population parameters such as survival and reproductive rates, and to associate levels of harvest with various regulatory scenarios (Martin, Pospahala & Nichols 1979). These and other estimators are combined to

produce and refine dynamic population models, which describe how waterfowl abundance varies in response to harvest and uncontrolled environmental factors (Williams & Nichols 1990). These models in turn are used to inform the regulations process, on assumption that population status is directly related to harvest and harvest can be predicted as a function of hunting regulations (Johnson, Williams, Nichols, Hines, Kendall, Smith & Caithamer 1993). By building on accumulated monitoring data, these models constantly evolve to reflect a growing understanding of waterfowl population dynamics and the impacts of harvest.

Unfortunately, the modeling of waterfowl populations and their harvest continues to be characterized by great uncertainty. In many cases, the sheer number and complexity of hunting regulations, combined with inadequate replication and experimental controls, has precluded reliable inference about the relationship between regulations and harvests (Nichols & Johnson 1989). Managers know even less about the impact of harvest on subsequent waterfowl population size. Particularly problematic in this regard are questions about the nature of density-dependent population regulation, which provides the theoretical basis for sustainable exploitation (Hilborn, Walters & Ludwig 1995). Uncertainties about the relationships among hunting regulations, harvest and population size constitute a principal source of controversy in the regulations-setting process.

In response to frustrations about the continuing lack of the biological understanding needed to inform harvest regulations, in 1995 the U.S. Fish and Wildlife Service implemented a new system to harvest regulation under the name of adaptive harvest management (AHM). This system is a rather formalized example of Adaptive Resource Management (Holling 1978, Walters 1986), which often is described in terms of 'management by experiment' or 'learning by doing' (Walters & Holling 1990). An appropriate definition for AHM is 'management in the face of uncertainty, with a focus on its reduction' (Williams & Johnson 1995). An adaptive approach to harvest regulations emphasizes uncertainty about regulatory effects, and incorporates uncertainty as a factor guiding management actions (Johnson et al. 1993). The goal is to reduce uncertainty over time, and thereby improve long-term management.

Operational elements of adaptive harvest management

Adaptive harvest management explicitly accounts for

uncertainty and the value of information in the regulatory process, recognizing that there are at least four identifiable sources (Williams 1997). The first and most obvious is uncontrollable (and sometimes unrecognized) environmental variation, which influences biological processes and induces stochasticity in population dynamics. The second is uncertainty about resource status, called partial observability to indicate limitations in one's ability to observe a resource system through monitoring. A third source of uncertainty is referred to as partial controllability, to emphasize the limited influence of management decisions on harvest and other actions. Finally, structural uncertainty concerns the lack of understanding (or lack of agreement) about the structure of biological relationships, such as the influence of harvest on survivorship. Environmental variation, structural uncertainty, partial observability and partial controllability all limit a manager's ability to make informed regulatory decisions (Nichols, Johnson & Williams 1995).

Along with an institutional framework and appropriate monitoring programs as described above, four elements are definitive of the adaptive process for setting waterfowl regulations:

- 1) An array of regulatory options that are available for the regulation of waterfowl harvest. These options include various combinations of regulations representing, e.g. 'restrictive', 'liberal', and 'moderate' regulations, with possible constraints on allowable fluctuations from year to year. The set of feasible regulatory options can be limited or expanded as the need and desirability to do so is recognized by management.
- 2) An objective function by which to evaluate and compare these options. The general form of the objective function is a weighted sum of harvests (or harvest utilities) over some recognized time frame. This is in keeping with traditional goals for waterfowl harvest management, and ensures that the focus is on harvest and harvest opportunity. A long time frame and harvest utilities that devalue harvest at low population sizes provide a conservation perspective, by preventing excessive harvest in the short term and thereby ensuring long-term sustainability.
- 3) A set of waterfowl models representing an array of meaningful hypotheses about the impact of regulations on waterfowl populations. For example, the set currently in use includes models that incorporate the hypothesis of additive hunting mortality, and others that incorporate the hypoth-

esis of completely compensatory hunting mortality. These models are used to gauge the consequences of different regulations. At present four models are used, each developed from data bases that have accrued as a result of waterfowl monitoring and research programs.

- 4) Measures of reliability for the models used in selecting harvest regulations. Reliability measures are used to weight the model outputs and are updated each year as additional data about resource status and the impacts of regulation become available. The notion of reliability is included in the process as an acknowledgment that the 'correct' or best approximating model for use in evaluating regulatory options is not known with certainty, and this uncertainty should be incorporated somehow in the procedure for evaluating and selecting regulations.

Adaptive harvest management is framed in terms of sequential decision making under uncertainty, in which one annually observes the state of the resource system (e.g. population size and relevant environmental features) and takes some management action (e.g. hunting regulations). An immediate return accrues as a result, which is expressed as a function of the benefits and costs that are relevant to the stated objectives of management. In response to the combined influence of management actions and uncontrolled environmental variation, the resource system subsequently evolves to a new state. The manager then observes the new system state, makes a new decision, accumulates additional returns, and the system evolves to yet another state (Fig. 1). And so on. The goal of management is to make a sequence of such decisions, each based on information about current system status, so as to maximize management returns over an extended time frame.

A major advantage of AHM is the explicit acknowledgement of alternative hypotheses describing the effects of regulations and other environmental factors on population dynamics. These hypotheses are codified in a set of system models, each of which has an associated weight reflecting its ability to describe system dynamics. Each year the weights are updated by comparing the model-specific prediction of changes in population size against the actual change observed from the monitoring program. By iteratively updating model weights and optimizing regulatory choices, the process eventually should identify which model is most appropriate to describe the dynamics of the managed population, and thereby should allow more effective, because better informed, management.




Figure 1. A sequential decision-making process in which management decisions (i.e. hunting regulations) made over time (t) produce an immediate return, expressed as a function of the benefits and costs that are relevant to the stated objectives of management. In response to the combined influence of management actions and uncontrolled environmental variation, the resource system subsequently evolves to a new state. The manager then observes the new system state, makes a new decision, accumulates additional returns and the system evolves to yet another state.

Population dynamics and the value of harvest

In what follows we describe AHM in a context of optimal adaptive control theory. Structural uncertainty is characterized here with multiple models of population dynamics over a discrete time frame, along with model-specific measures of uncertainty about which model is most appropriate. Without loss of generality as to optimal management, we combine environmental variation and partial controllability into a single stochastic factor z_t affecting population dynamics. Resource status is characterized by \underline{x}_t , recognizing that \underline{x}_t includes attributes that are definitive of resource state. Management action at time t is designated by a_t , and policies describing actions over the remainder of the time frame are designated by A_t .

With these notational conventions, consider a biological population that annually is subjected to harvest, with management actions that are based on resource status \underline{x}_t and the projected effects on future resource states. Models depicting population responses play prominently in the assessment of impacts. Several models of the form

$$\underline{x}_t = \underline{G}_i(\underline{x}_{t-1}, \underline{a}_{t-1}, z_{t-1})$$

are assumed to be available, where \underline{a}_t and z_t represent management controls and random variation, respec-

tively. Initially one does not know which model most appropriately represents population change in response to management. This uncertainty is captured in a set $p_i(t)$ of likelihoods that express one's confidence in the models at time t . The notation $p_i(t)$ allows for evolving likelihood values in response to accumulating information about management controls and population responses. By affecting population dynamics, management can influence the evolution of the likelihoods, and thereby promote learning.

Benefits and costs attend the implementation of harvest controls over time, and these can be captured in a utility function that itself may be model-specific. For simplicity we describe utilities as functions of the current resource state and action, recognizing that the utility function might also

represent an average of utilities across potential outcome states. Thus, $R_i(a_t | \underline{x}_t)$ is the utility for model i if the resource status is \underline{x}_t and action a_t is taken. In the case of waterfowl harvest, the utility is given in terms of harvest yield, recognizing an aversion to harvest decisions that would result in an expected population size below the goal of the North American Waterfowl Management Plan (NAWMP; U.S. Department of the Interior, Environment Canada, and Secretario de Desarrollo Social Mexico 1994). The capacity of available breeding habitat to promote population growth is considered in determining an optimal regulatory decision for \underline{x}_t . Thus, liberal hunting regulations could be appropriate even if the population is below the NAWMP goal, if current habitat conditions are expected to result in good production of young. On the other hand, restrictive regulations may be appropriate when reproductive success is expected to be low, even if populations are at or above the NAWMP goal.

In general, an overall value for harvest utility that accounts for model uncertainty is the average

$$\bar{R}(a_t | \underline{x}_t, p_t) = \sum_i p_i(t) R_i(a_t | \underline{x}_t),$$

based on model-specific utilities $R_i(a_t | \underline{x}_t)$ and model likelihoods $p_i(t)$. If there is only a single model under consideration, or if the likelihood is assumed to be $p_i(t)=1$ for model i , the utility corresponding to action a_t simplifies to

$$\bar{R}(\underline{a}_t|\underline{x}_t, \underline{p}_t) = R_i(\underline{a}_t|\underline{x}_t).$$

Each of the population models characterizes transitions of the population over time, as influenced by such factors as survivorship, recruitment and migration, along with the controls affecting them. These factors always are subject to environmental variation and other stochastic factors, including randomness in the effects of controls. Thus, the projected resource status \underline{x}_{t+1} for model i inherits a probability distribution $p_i(\underline{x}_{t+1}|\underline{x}_t, \underline{a}_t)$ from environmental and other sources of variation. The challenge is to choose harvest controls that will maximize aggregate harvest utility in the face of stochastic effects, while also accounting for uncertainties about the biological processes that drive population dynamics.

Dynamic programming

Adaptive harvest management (AHM) utilizes a variant of stochastic dynamic programming, an iterative procedure for identifying optimal state-specific actions for dynamic systems. Stochastic dynamic programming can be described in terms of an observable system

$$\underline{x}_t = \underline{G}(\underline{x}_{t-1}, \underline{a}_{t-1}, z_{t-1}),$$

with \underline{a}_t and z_t representing management controls and random variation, respectively, at time t . Let policy A_t specify a state-specific control \underline{a}_t for every state \underline{x}_t at every time in a time frame $\{t, t+1, \dots, T\}$. A value $V(A_t|\underline{x}_t)$ can be associated with A_t by accumulating utilities over the remainder of the time frame:

$$V(A_t|\underline{x}_t) = E \left\{ \sum_{\tau=t}^T R(\underline{a}_\tau|\underline{x}_\tau) \mid \underline{x}_t \right\},$$

where the expectation is with respect to environmental variation and partial controllability over the time frame. The notation $V(A_t|\underline{x}_t)$ indicates that the accumulation of utilities begins at time t , the start of the time frame for A_t . It also expresses the fact that accumulated utilities are conditional on the population state \underline{x}_t , in that $V(A_t|\underline{x}_t)$ can (and usually does) vary for different population states \underline{x}_t .

Decomposing the sum above into current and future utilities, we have

$$V(A_t|\underline{x}_t) = E \left\{ R(\underline{a}_t|\underline{x}_t) + \sum_{\tau=t+1}^T R(\underline{a}_\tau|\underline{x}_\tau) \mid \underline{x}_t \right\},$$

$$\begin{aligned} &= R(\underline{a}_t|\underline{x}_t) + \sum_{\underline{x}_{t+1}} p(\underline{x}_{t+1}|\underline{x}_t, \underline{a}_t) E \left\{ \sum_{\tau=t+1}^T R(\underline{a}_\tau|\underline{x}_\tau) \mid \underline{x}_{t+1} \right\} \\ &= R(\underline{a}_t|\underline{x}_t) + \sum_{\underline{x}_{t+1}} p(\underline{x}_{t+1}|\underline{x}_t, \underline{a}_t) V(A_{t+1}|\underline{x}_{t+1}), \end{aligned}$$

with the transition probabilities $p(\underline{x}_{t+1}|\underline{x}_t, \underline{a}_t)$ capturing environmental variation and partial controllability. These probabilities often are assumed to be stationary, in that they change through time only as a result of controls. However, stationarity is not a theoretical requirement, and non-stationary transition probabilities can be denoted in the above expression simply by conditioning explicitly on time, as in $p(\underline{x}_{t+1}|\underline{x}_t, \underline{a}_t, t)$.

Values for the aggregate utilities $V(A_t|\underline{x}_t)$ can be obtained for every possible policy A_t over the time frame. Thus, by proper choice of A_t these values can be optimized. A backward iteration algorithm to determine an optimal policy is given by the Hamilton-Jacobi-Bellman (HJB) equation

$$V^*(\underline{x}_t) = \max_{\underline{a}_t} \left\{ R(\underline{a}_t|\underline{x}_t) + \sum_{\underline{x}_{t+1}} p(\underline{x}_{t+1}|\underline{x}_t, \underline{a}_t) V^*(\underline{x}_{t+1}) \right\}$$

(Stengel 1994), with $V^*(\underline{x}_t)$ the optimal value of aggregate utility corresponding to state \underline{x}_t at time t .

The iterative application of the HJB equation starting at the end of the time frame is known as stochastic dynamic programming (Bellman & Dreyfus 1962, Dreyfus & Law 1977). The optimal policy $A^*(\underline{x}_t)$ thus described identifies optimal actions for all population states at all times in the time frame, along with a field of optimal values $V^*(\underline{x}_t)$ for all population states and times. An optimal policy having been found with dynamic programming one need only identify the population state \underline{x}_t at a particular time, and then apply the control specified by the policy for that state at that time.

Optimization in adaptive harvest management

Now consider the control of a population for which several models describing population dynamics are available, but the most appropriate model is not known with certainty, i.e. $p_i(t) \neq 1$. Policy value is given in terms of accumulated harvest utilities, averaged over all models based on the model likelihoods:

$$\bar{V}(A_t|\underline{x}_t, \underline{p}_t) = E \left\{ \sum_i p_i(t) \sum_{\tau=t}^T R_i(\underline{a}_\tau|\underline{x}_\tau) \mid \underline{x}_t \right\}$$

$$= \sum_i p_i(t) V_i(A_t | \underline{x}_t).$$

This expression can be further decomposed into current and future utilities by

$$\begin{aligned} \bar{V}(A_t | \underline{x}_t, \underline{p}_t) &= \sum_i p_i(t) E \left\{ \sum_{\tau=t}^T R_i(\underline{a}_\tau | \underline{x}_\tau) \mid \underline{x}_t \right\} \\ &= \sum_i p_i(t) \left\{ R_i(\underline{a}_t | \underline{x}_t) + \sum_{\underline{x}_{t+1}} p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) V_i(A_{t+1} | \underline{x}_{t+1}) \right\} \\ &= \bar{R}(\underline{a}_t | \underline{x}_t, \underline{p}_t) + \sum_i \sum_{\underline{x}_{t+1}} p_i(t) p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) V_i(A_{t+1} | \underline{x}_{t+1}). \end{aligned}$$

The term $p_i(t) p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t)$ in the latter expression can be replaced by $p_i(t+1) \bar{p}(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t)$ via Bayes' theorem

$$\begin{aligned} p_i(t+1) &= \frac{p_i(t) p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t)}{\sum_i p_i(t) p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t)} \\ &= \frac{p_i(t) p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t)}{\bar{p}(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t)} \end{aligned}$$

(Lee 1992), so that

$$\begin{aligned} \bar{V}(A_t | \underline{x}_t, \underline{p}_t) &= \bar{R}(\underline{a}_t | \underline{x}_t, \underline{p}_t) + \sum_i \sum_{\underline{x}_{t+1}} p_i(t+1) \bar{p}(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) V_i(A_{t+1} | \underline{x}_{t+1}) \\ &= \bar{R}(\underline{a}_t | \underline{x}_t, \underline{p}_t) + \sum_{\underline{x}_{t+1}} \bar{p}(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) \sum_i p_i(t+1) V_i(A_{t+1} | \underline{x}_{t+1}) \\ &= \bar{R}(\underline{a}_t | \underline{x}_t, \underline{p}_t) + \sum_{\underline{x}_{t+1}} \bar{p}(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) \bar{V}(A_{t+1} | \underline{x}_{t+1}, \underline{p}_{t+1}). \end{aligned}$$

A value $\bar{V}(A_t | \underline{x}_t, \underline{p}_t)$ for the average accumulated utility can be obtained for every possible policy A_t over the time frame, starting at any particular time t and any combination $(\underline{x}_t, \underline{p}_t)$. By proper choice of A_t these values can be optimized, with a solution algorithm that is based on

$$\bar{V}^*(\underline{x}_t, \underline{p}_t) = \max_{\underline{a}_t} \left\{ \bar{R}(\underline{a}_t | \underline{x}_t, \underline{p}_t) + \sum_{\underline{x}_{t+1}} \bar{p}(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) \bar{V}^*(\underline{x}_{t+1}, \underline{p}_{t+1}) \right\} \quad (1).$$

This is a stochastic dynamic programming problem, though complicated somewhat by the characterization of system state by $(\underline{x}_t, \underline{p}_t)$. Transitions via Bayes' theorem are required for \underline{p}_t , and transitions for \underline{x}_t are given in terms of the transition probabilities $\bar{p}(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t)$.

The optimization problem can be solved by iterative application of Equation 1, starting at the end of the time frame and proceeding backward in time. An optimal solution consists of a policy $A^*(\underline{x}_t, \underline{p}_t)$ that identifies a specific action for every combination $(\underline{x}_t, \underline{p}_t)$ of resource state \underline{x}_t and likelihood state \underline{p}_t , along with a field of optimal values $\bar{V}^*(\underline{x}_t, \underline{p}_t)$ for all resource states and model likelihoods at all times in the time frame. To implement the optimal policy, at each time one must (i) determine the resource status, (ii) update the likelihoods with Bayes' theorem, and (iii) apply the regulatory control specified by the optimal policy for the resource state and set of updated likelihoods (Table 1).

We note that when $p_i(t)=1$ the optimal policy and values for $(\underline{x}_t, \underline{p}_t)$ are $A_i^*(\underline{x}_t)$ and $\bar{V}_i^*(\underline{x}_t)$ respectively for a single model i . This intuitive result follows from the fact that if $p_i(t)=1$, $\bar{p}(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) = p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t)$ throughout the remainder of the time frame, so the computing algorithm (1) reduces to the HJB equation for model i :

$$\bar{V}_i^*(\underline{x}_t) = \max_{\underline{a}_t} \left\{ R_i(\underline{a}_t | \underline{x}_t) + \sum_{\underline{x}_{t+1}} p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) V_i^*(\underline{x}_{t+1}) \right\} \quad (2).$$

This formula describes a straightforward stochastic dynamic programming problem, which can be solved by iterative application of Equation 2 as described above.

Optimal decision-making with partial observability

In almost all management applications, resource status is not known with certainty, and instead must be estimated at each time with field data. An estimate $\hat{\underline{x}}_t$ inherits a distribution from data collected in the field, conditional on the field sampling design and the actual population size \underline{x}_t . Let \underline{y}_t represent field data collected at time t , and Y_t represent the accumulation of data up to t . Each year's monitoring effort adds to the accumulation of data, by $Y_{t+1} = \{Y_t, \underline{y}_{t+1}\}$. Assume that an estimate $\hat{\underline{x}}_t$ of resource status can be obtained as a function $\hat{\underline{x}}_t = \hat{\underline{x}}_t(\underline{y}_t | Y_{t-1})$ of the data accumulated up to time t . Since \underline{y}_t is conditional on \underline{x}_t , the estimate $\hat{\underline{x}}_t$ inherits conditional distributions $f_1(\underline{x}_t | \hat{\underline{x}}_t)$ and $f_2(\hat{\underline{x}}_t | \underline{x}_t)$ from \underline{y}_t , and the transition from $\hat{\underline{x}}_t = \hat{\underline{x}}_t(\underline{y}_t | Y_{t-1})$ to $\hat{\underline{x}}_{t+1} = \hat{\underline{x}}_{t+1}(\underline{y}_{t+1} | Y_t)$ is given in terms of the model-specific probabilities

$$p_i(\hat{\underline{x}}_{t+1} | \underline{x}_t, \underline{a}_t) = \sum_{\underline{x}_{t+1}} \sum_{\hat{\underline{x}}_t} f_1(\underline{x}_t | \hat{\underline{x}}_t) p_i(\underline{x}_{t+1} | \underline{x}_t, \underline{a}_t) f_2(\hat{\underline{x}}_{t+1} | \underline{x}_{t+1}) \quad (3).$$

Under these conditions a solution of the optimization problem for model i is obtained by iterative application

Table 1. Optimal regulatory choices for midcontinent mallards during the 1999 hunting season. This strategy is based on very restrictive (VR), restrictive (R), moderate (M) and liberal (L) regulatory alternatives, along with current model weights, based on a dual objective of maximizing long-term cumulative harvest and achieving a population goal of 8.7 million. An appropriate regulatory action is identified for each combination of population and habitat conditions, based on information from resource surveys. Table cells with no regulatory entries correspond to season closure.

Mallards ^b	Ponds ^a									
	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0
<5.0										
5.0										
5.5							VR	VR	VR	VR
6.0		VR	VR	VR	VR	VR	R	R	R	R
6.5	VR	VR	VR	VR	R	R	M	M	M	L
7.0	R	R	R	R	M	M	M	L	L	L
7.5	R	M	M	M	M	L	L	L	L	L
8.0	M	M	M	L	L	L	L	L	L	L
8.5	M	L	L	L	L	L	L	L	L	L
≥9.0	L	L	L	L	L	L	L	L	L	L

^a Estimated number in millions of ponds in Prairie Canada in May.

^b Estimated number in millions of midcontinent mallards during May.

of the HJB equation, expressed in terms of estimated rather than actual system state:

$$\bar{V}_i^*(\hat{x}_t) = \max_{a_t} \left\{ \sum_{x_t} f_1(x_t | \hat{x}_t) R_t(a_t | x_t) + \sum_{\hat{x}_{t+1}} p_1(\hat{x}_{t+1} | x_t, a_t) V_i^*(\hat{x}_{t+1}) \right\} \quad (4)$$

(Williams 1996a).

A key feature in incorporating partial observability is the statistical association between x_t and \hat{x}_t , from which are derived the conditional distributions $f_1(x_t | \hat{x}_t)$ and $f_2(\hat{x}_t | x_t)$. These distributions derive from the stochastic structure of the statistic $\hat{x}_t = \hat{x}_t(y_t | Y_{t-1})$, which is parameterized by x_t . The distribution $f_2(\hat{x}_t | x_t)$ arises naturally from $\hat{x}_t = \hat{x}_t(y_t | Y_{t-1})$, based on sampling variation in y_t . On the other hand, the derivation of $f_1(x_t | \hat{x}_t)$ can be quite difficult to derive and is a subject of considerable theoretical interest (see below).

The extension to adaptive optimization with multiple models is straightforward. Then the HJB equation 4 becomes

$$\bar{V}^*(\hat{x}_t, p_t) = \max_{a_t} \left\{ \sum_{x_t} f_1(x_t | \hat{x}_t) \bar{R}(a_t | x_t, p_t) + \sum_{\hat{x}_{t+1}} \bar{p}(\hat{x}_{t+1} | x_t, a_t) \bar{V}^*(\hat{x}_{t+1}, p_{t+1}) \right\} \quad (5)$$

where $\bar{V}^*(\hat{x}_t, p_t)$ and $\bar{V}^*(\hat{x}_{t+1}, p_{t+1})$ are defined as before and

$$\bar{p}(\hat{x}_{t+1} | \hat{x}_t, a_t) = \sum_{x_t} \sum_{\hat{x}_t} f_1(x_t | \hat{x}_t) \bar{p}(x_{t+1} | x_t, a_t) f_2(\hat{x}_{t+1} | x_{t+1}).$$

To identify the transition probabilities $\bar{p}(\hat{x}_{t+1} | \hat{x}_t, a_t)$ in Equation 5 one must determine the transition probability $\bar{p}(x_{t+1} | x_t, a_t)$ for every state x_t for which there is a non-zero probability $f_1(x_t | \hat{x}_t)$, and calculate the average utility $\bar{R}(a_t | x_t, p_t)$ for all states with non-zero probabilities.

These requirements result in a substantial increase in computations, well beyond what is required for a solution with a single model.

Sequential identification of resource state

In the development above we assumed an estimator $\hat{x}_t = \hat{x}_t(y_t | Y_{t-1})$ and described the system transitions in terms of \hat{x}_t . In what follows we generalize this situation somewhat by describing system transitions directly in terms of the accumulated observations. Thus, at each time in the time frame the system is observed, a management action a_t then is taken and the system evolves to a new state x_{t+1} at time $t+1$. The system is assumed to be only partially observable, so that x_t cannot be observed directly, and information about the system state must be obtained from y_t . System dynamics are recorded in terms of the transition of observations from y_t to y_{t+1} , rather than transitions from x_t to x_{t+1} . Transition probabilities, expressing the stochastic influence of environmental uncertainty, partial controllability and partial observability, can be represented in terms of

$$p_{\beta}(y_t | Y_{t-1}, a_{t-1}) = \sum_{x_{t-1}} \sum_{\hat{x}_t} f_1(x_{t-1} | Y_{t-1}) f_{\beta}(x_t | x_{t-1}, a_{t-1}) f_2(y_t | x_t) \quad (6)$$

where $Y_t = \{Y_{t-1}, y_t\}$ accumulates observations up to time t . Here we use f_1 and f_2 to denote distributions based on observation data y_t , in contrast to Equation 3 in which they represent distributions for the estimators \hat{x}_t . We also characterize the transition probabilities by $f_{\beta}(x_t | x_{t-1}, a_{t-1})$ with structural uncertainty now represented by the parameter β .

Of the three probability distributions in Equation 6, the distribution

$$f_{\underline{\beta}}(\underline{x}_t | \underline{x}_{t-1}, \underline{a}_{t-1}) \quad (7)$$

captures the notion of a sequential linkage of systems states, whereby the transition from \underline{x}_{t-1} to \underline{x}_t is influenced by management control \underline{a}_{t-1} and is parameterized by $\underline{\beta}$ (recall that the subscript i was used earlier to denote a limited suite of models). This distribution is based on stochastic models of population dynamics which themselves express hypothesized relationships among system, environmental and control variables. This formulation allows structural uncertainty to be expressed in terms of variation in the parameter $\underline{\beta}$.

The component $f_{2(Y_t | \underline{x}_t)}$ of Equation 6 essentially specifies that system state informs system observations, in that variation in the observations \underline{y}_t is conditional on state \underline{x}_t . The idea is that observations are tied to the system state, with stochastic variation in \underline{y}_t as a result of random sampling. This variation can be modelled with field data, based on an assumed form for $f_{2(Y_t | \underline{x}_t)}$.

The lead distribution in Equation 6

$$f_1(\underline{x}_{t-1} | Y_{t-1}) \quad (8),$$

reflects the fact that the actual system state \underline{x}_{t-1} is conditionally associated with accumulated observations Y_{t-1} . Identifying distribution (8) can be especially problematic, in large part because the conditioning variable Y_{t-1} is itself subject to stochastic and time-varying influences (i.e. sampling variability) that are not present in \underline{x}_{t-1} . The stochastic identification of \underline{x}_{t-1} is an example of statistical 'calibration' (Graybill 1976), whereby the value of a conditional predictor variable is sought given one or more values of a stochastic response variable. The difficulty here is that the predictor variable itself evolves stochastically according to the process equation

$$\underline{x}_t = \underline{G}_{\underline{\beta}}(\underline{x}_{t-1}, \underline{a}_{t-1}, \underline{e}_t) \quad (9),$$

where \underline{e}_t can be thought of as a general environmental white noise process with an assumed mean of $\underline{0}$ and dispersion \underline{W}_t . Equation 9 provides the biological basis for distribution (6). To simplify notation, we suppress the subscript $\underline{\beta}$ in the argument below, recognizing, however, the potential for structural uncertainty.

In what follows, we focus on the stochastic prediction of \underline{x}_t based on the accumulated data Y_t up to time t , and in particular we seek the distribution $f_{2(\underline{x}_t | Y_t)}$, as represented by estimates of the conditional mean $E(\underline{x}_t | Y_t)$ and conditional dispersion $\sum_{\underline{x}_t | Y_t}$. To begin, consider initially a linear dynamic system, with \underline{x}_t a vector of time-specific state variables (in our case popu-

lation and/or habitat status) and \underline{y}_t a vector of time-specific observation variables (from e.g. breeding grounds surveys). The idea is to say something about the (actual) state \underline{x}_t at each point in time, given a record of observations Y_t up to t . State transition equations for a linear system can be expressed as

$$\underline{x}_t = \underline{F}_t \underline{x}_{t-1} + \underline{e}_t \quad (10),$$

where \underline{F}_t is a full rank $n \times n$ matrix of potentially time-varying but non-random parameters. A key structural assumption is that the environmental vector \underline{e}_t adds to the process components $\underline{F}_t \underline{x}_{t-1}$. Then the first and second system moments are given by

$$E(\underline{x}_t) = \underline{F}_t E(\underline{x}_{t-1}) \quad (11)$$

and

$$\begin{aligned} \sum_{\underline{x}_t} &= \underline{F}_t \sum_{\underline{x}_{t-1}} \underline{F}_t' + \underline{W}_t \\ &= \underline{R}_t \end{aligned} \quad (12).$$

Note that the notation for system control has been suppressed; on including controls, see below.

The system is assumed here to be partially observable, with observation equations

$$\begin{aligned} \underline{y}_t &= \underline{H}_t \underline{x}_t + \underline{\epsilon}_t \\ &= \underline{H}_t (\underline{F}_t \underline{x}_{t-1} + \underline{e}_t) + \underline{\epsilon}_t \end{aligned} \quad (13),$$

where \underline{H}_t is a $k \times n$ matrix of potentially time-varying but non-random parameters and $r(\underline{H}_t) = \min\{k, n\}$. The vector $\underline{\epsilon}_t$ represents sampling variation, and is assumed to have $\underline{0}$ mean and dispersion \underline{V}_t . Then

$$E(\underline{y}_t) = \underline{H}_t E(\underline{x}_t) \quad (14)$$

and

$$\sum_{\underline{y}_t} = \underline{H}_t \underline{R}_t \underline{H}_t' + \underline{V}_t \quad (15).$$

The linkage between system state and observation state in Equation 13 is controlled by the magnitude of dispersion \underline{V}_t , and also by the rank of the transform \underline{H}_t . A full rank transform with $\underline{0}$ dispersion is tantamount to complete observability.

The problem of estimating \underline{x}_t is non-trivial because the transition to \underline{x}_t each time is conditional on the previous state variable value \underline{x}_{t-1} , which itself is unobservable. An iterative procedure utilizes the condi-

tional distribution of \underline{x}_t , given the current data \underline{y}_t and an unbiased estimate $\hat{\underline{x}}_{t-1}$ based on previous data. The procedure is based on the joint distribution

$$\begin{bmatrix} \underline{x}_t \\ \underline{y}_t \end{bmatrix} \sim \left(\begin{bmatrix} \underline{F}_t \underline{E}(\underline{x}_{t-1}) \\ \underline{H}_t \underline{F}_t \underline{E}(\underline{x}_{t-1}) \end{bmatrix}, \begin{bmatrix} \underline{R}_t & \sum_{\underline{x}_t, \underline{y}_t} \\ \sum_{\underline{x}_t, \underline{y}_t} \underline{H}_t \underline{R}_t \underline{H}'_t + \underline{V}_t \end{bmatrix} \right)$$

of \underline{x}_t and \underline{y}_t , where $\underline{E}(\underline{x}_t)$ and $\sum_{\underline{x}_t}$ are given as in Equations 11 and 12, $\underline{E}(\underline{y}_t)$ and $\sum_{\underline{y}_t}$ are given as in Equations 14 and 15 and

$$\begin{aligned} \sum_{\underline{x}_t, \underline{y}_t} &= \underline{E}(\underline{x}_t - \underline{E}(\underline{x}_t))(\underline{y}_t - \underline{E}(\underline{y}_t))' \\ &= \underline{E}(\underline{x}_t \underline{y}_t) - \underline{E}(\underline{x}_t) \underline{E}(\underline{y}_t)' \\ &= \underline{E}(\underline{F}_t \underline{x}_{t-1} + \underline{e}_t)(\underline{H}_t [\underline{F}_t \underline{x}_{t-1} + \underline{e}_t] + \underline{e}_t)' - \underline{E}(\underline{x}_t) \underline{E}(\underline{y}_t)' \\ &= \underline{R}_t \underline{H}'_t. \end{aligned}$$

Thus,

$$\begin{bmatrix} \underline{x}_t \\ \underline{y}_t \end{bmatrix} \sim \left(\begin{bmatrix} \underline{F}_t \underline{E}(\underline{x}_{t-1}) \\ \underline{H}_t \underline{F}_t \underline{E}(\underline{x}_{t-1}) \end{bmatrix}, \begin{bmatrix} \underline{R}_t & \underline{R}_t \underline{H}'_t \\ \underline{H}_t \underline{R}'_t & \underline{H}_t \underline{R}_t \underline{H}'_t + \underline{V}_t \end{bmatrix} \right).$$

From this distribution it follows that, conditional on observations \underline{y}_t , the distribution of system state \underline{x}_t has the mean

$$\underline{E}(\underline{x}_t | \underline{y}_t) = \underline{F}_t \underline{E}(\underline{x}_{t-1}) + \underline{R}_t \underline{H}'_t (\underline{H}_t \underline{R}_t \underline{H}'_t + \underline{V}_t)^{-1} [\underline{y}_t - \underline{H}_t \underline{F}_t \underline{E}(\underline{x}_{t-1})] \quad (16)$$

and the dispersion

$$\sum_{\underline{x}_t | \underline{y}_t} = \underline{R}_t - \underline{R}_t \underline{H}'_t (\underline{H}_t \underline{R}_t \underline{H}'_t + \underline{V}_t)^{-1} \underline{H}_t \underline{R}_t \quad (17)$$

(Graybill 1976: 106). Substituting $\hat{\underline{x}}_{t-1}$ and $\hat{\underline{x}}_t$ for the expected values in Equation 16 yields the iterative algorithm

$$\hat{\underline{x}}_t = \underline{F}_t \hat{\underline{x}}_{t-1} + \underline{R}_t \underline{H}'_t (\underline{H}_t \underline{R}_t \underline{H}'_t + \underline{V}_t)^{-1} [\underline{y}_t - \underline{H}_t \underline{F}_t \hat{\underline{x}}_{t-1}] \quad (18)$$

for updating the estimate $\hat{\underline{x}}_{t-1}$ at time $t-1$ to $\hat{\underline{x}}_t$ at t based on the data \underline{y}_t . A number of interesting points are noteworthy in Equations 16-18:

- The estimate $\hat{\underline{x}}_t$ in expression (18) is a linear combination of two components, the first of which is simply the propagation of $\hat{\underline{x}}_{t-1}$ with the system transition matrix \underline{F}_t , absent any updating with data. The second adjusts this propagation with a factor that accounts for the components \underline{y}_t .

- Likewise, the conditional dispersion in expression (17) is a linear combination of two components, the first of which is the dispersion of \underline{x}_t based on the system transition Equation 10, absent environmental and observability considerations for time t . The second component adjusts this dispersion with a factor that accounts for the latter sources of variation.
- If \underline{F}_t , \underline{H}_t , \underline{V}_t , and \underline{W}_t are stationary the evolution of moments for $(\underline{x}_t | \underline{y}_t)$ display the following patterns:
 - Both the estimate $\hat{\underline{x}}_t$ and the dispersion tend to increase or decrease through time depending on the eigenvalues of \underline{F} .
 - Both environmental dispersion \underline{W} and observability dispersion \underline{V} influence the evolution $\hat{\underline{x}}_t$, through the second term in expression (18).
 - Environmental variation and partial observability also influence the evolution of the conditional dispersion, through the second term in expression (17).
- The following conditions contribute to a large adjustment in the mean of $(\underline{x}_t | \underline{y}_t)$ through time:
 - $\underline{y}_t - \underline{H}_t \underline{F}_t \hat{\underline{x}}_{t-1}$ is large;
 - \underline{V}_t is small;
 - \underline{H}_t is $n \times n$ and full rank.
- The following conditions contribute to the reduction in dispersion of $(\underline{x}_t | \underline{y}_t)$ through time:
 - \underline{V}_t is small,
 - \underline{H}_t is $n \times n$ and full rank.

With this formulation it is easy to see why adaptive decision-making is so much easier in the presence of complete observability. If $\underline{V}_t = \underline{0}$ and \underline{H}_t is dimension n and full rank, the estimate $\hat{\underline{x}}_t$ in expression (14) becomes

$$\begin{aligned} \underline{F}_t \hat{\underline{x}}_{t-1} + \underline{R}_t \underline{H}'_t (\underline{H}_t \underline{R}_t \underline{H}'_t + \underline{V}_t)^{-1} (\underline{y}_t - \underline{H}_t \underline{F}_t \hat{\underline{x}}_{t-1}) &= \underline{F}_t \hat{\underline{x}}_{t-1} + \underline{H}_t^{-1} (\underline{y}_t - \underline{H}_t \underline{F}_t \hat{\underline{x}}_{t-1}) \\ &= \underline{H}_t^{-1} \underline{y}_t \\ &= \underline{x}_t \end{aligned}$$

and the dispersion of $(\underline{x}_t | \underline{y}_t)$ in expression (17) becomes

$$\begin{aligned} \underline{R}_t - \underline{R}_t \underline{H}'_t (\underline{H}_t \underline{R}_t \underline{H}'_t + \underline{V}_t)^{-1} \underline{H}_t \underline{R}_t &= \underline{R}_t - \underline{R}_t \\ &= \underline{0}. \end{aligned}$$

In essence, complete observability means that \underline{x}_t and \underline{y}_t encode the same information, so that the system state \underline{x}_t at time t is known with certainty once the data \underline{y}_t are available. The adaptive management job then becomes one of controlling an observable process with temporally varying dispersion \underline{R}_t that is known at each point in time.

Equation 18 is a version of the Kalman filter (Kalman 1960, Kalman & Bucy 1961), with filter gain $\mathbf{R}_t \mathbf{H}'_t (\mathbf{H}_t \mathbf{R}_t \mathbf{H}'_t + \mathbf{V}_t)^{-1}$ influencing both the conditional means and the conditional dispersions. On assumption of multivariate normality for \mathbf{y}_t and \mathbf{w}_t , it can be shown that the Kalman filter provides a minimum variance estimator of system state. Note that data-based adjustments occur in the estimation of $\hat{\mathbf{x}}_t$ but not in the temporal updating of dispersion, since the latter is influenced by sampling variation only through the sampling dispersion \mathbf{V}_t in the filter gain. A readable description of the Kalman filter is given by Meinhold & Singpurwalla (1983), though their approach differs from the development here. A theoretically comprehensive description is given by Stengel (1994).

Incorporating harvest controls

One of the nice things about this formulation is that additional complexity can be accommodated without much effort. For example, additive controls can be included simply by incorporating an additional term \mathbf{a}_{t-1} in transition Equation 10:

$$\mathbf{x}_t = \mathbf{F}_t \mathbf{x}_{t-1} + \mathbf{G}_t \mathbf{a}_{t-1} + \mathbf{e}_t.$$

Then Equation 11 for the mean $E(\mathbf{x}_t)$ becomes

$$\underline{E}(\mathbf{x}_t) = \mathbf{F}_t \underline{E}(\mathbf{x}_{t-1}) + \mathbf{G}_t \underline{a}_{t-1},$$

Equation 16 for the conditional mean becomes

$$E(\mathbf{x}_t | \mathbf{y}_t) = \mathbf{F}_t E(\mathbf{x}_{t-1}) + \mathbf{G}_t \mathbf{a}_{t-1} + \mathbf{R}_t \mathbf{H}'_t (\mathbf{H}_t \mathbf{R}_t \mathbf{H}'_t + \mathbf{V}_t)^{-1} [\mathbf{y}_t - \mathbf{H}_t (\mathbf{F}_t E(\mathbf{x}_{t-1}) + \mathbf{G}_t \mathbf{a}_{t-1})],$$

and the updating algorithm in Equation 13 is

$$\hat{\mathbf{x}}_t = \mathbf{F}_t \hat{\mathbf{x}}_{t-1} + \mathbf{G}_t \mathbf{a}_{t-1} + \mathbf{R}_t \mathbf{H}'_t (\mathbf{H}_t \mathbf{R}_t \mathbf{H}'_t + \mathbf{V}_t)^{-1} [\mathbf{y}_t - \mathbf{H}_t (\mathbf{F}_t \hat{\mathbf{x}}_{t-1} + \mathbf{G}_t \mathbf{a}_{t-1})].$$

Note that because \mathbf{a}_t is assumed to be non-random, Equation 13 for the dispersion of $(\mathbf{x}_t | \mathbf{y}_t)$ remains unchanged. It follows that controls in a linear system can influence the system state (and thus the data-based estimate $\hat{\mathbf{x}}_t$ of system state), but not the corresponding dispersion. This is not surprising, since \mathbf{a}_t essentially constitutes a linear random influence on the system mean. Equation 16 still describes a (relatively!) simple algorithm for calculating the dispersion trajectory, one that is especially easy to use if transfer matrices are constant over time. Then the eigenstructure of the gain will be diagnostic of change in dispersion, with patterns that

exhibit either explosive growth or exponential decay over the long term depending on the eigenvalues.

One can introduce partial controllability into the problem by assuming that the control at is random,

$$\mathbf{a}_t \sim (\hat{\mathbf{a}}_t, \underline{\Sigma}_{\mathbf{a}_t}),$$

where the distribution mean $\hat{\mathbf{a}}_t$ represents an 'intended' control at each point in time and the dispersion represents uncontrolled variation about the mean (as with the other stochastic model components, we assume here that there is no autocorrelation). Then the system mean becomes

$$E(\mathbf{x}_t) = \mathbf{F}_t E(\mathbf{x}_{t-1}) + \mathbf{G}_t \hat{\mathbf{a}}_{t-1},$$

the system dispersion becomes

$$\begin{aligned} \underline{\Sigma}_{\mathbf{x}_t} &= \mathbf{F}_t \underline{\Sigma}_{\mathbf{x}_{t-1}} \mathbf{F}'_t + \mathbf{G}_t \underline{\Sigma}_{\mathbf{a}_{t-1}} \mathbf{G}'_t + \mathbf{W}_t \\ &= \mathbf{R}_t^*, \end{aligned}$$

and the updating algorithm becomes

$$\hat{\mathbf{x}}_t = (\mathbf{F}_t \hat{\mathbf{x}}_{t-1} + \mathbf{G}_t \hat{\mathbf{a}}_{t-1}) + \mathbf{R}_t^* \mathbf{H}'_t (\mathbf{H}_t \mathbf{R}_t^* \mathbf{H}'_t + \mathbf{V}_t)^{-1} [\mathbf{y}_t - \mathbf{H}_t (\mathbf{F}_t \hat{\mathbf{x}}_{t-1} + \mathbf{G}_t \hat{\mathbf{a}}_{t-1})]$$

with

$$\underline{\Sigma}_{\mathbf{x}_t | \mathbf{y}_t} = \mathbf{R}_t^* - \mathbf{R}_t^* \mathbf{H}'_t (\mathbf{H}_t \mathbf{R}_t^* \mathbf{H}'_t + \mathbf{V}_t)^{-1} \mathbf{H}_t \mathbf{R}_t^*.$$

Note that stochastic controls, like deterministic controls, influence the system state through the mean value $\hat{\mathbf{a}}_t$. However, unlike deterministic controls they also inflate system dispersion. Again this is not surprising, since stochastic control simply adds another random element to the system transitions.

Special cases

There are several interesting special cases that correspond to restrictions on the sources of system variation. Examples include:

Constant system state: On assumption that $\mathbf{x}_t = \mathbf{x}_{t-1}$ over the time frame, system 'dynamics' really consist of variation in the system observations through time:

$$\mathbf{y}_t = \mathbf{H}_t \mathbf{x} + \mathbf{e}_t.$$

This formulation is simply the general linear model of classical estimation theory, with \mathbf{x} playing the role of the vector $\underline{\beta}$ to be estimated and \mathbf{H}_t playing the role of the predictor data matrix $\underline{\mathbf{X}}$. Estimation in this case pro-

ceeds in a straightforward manner, according to the usual computing algorithms for general linear models (Graybill 1976).

We note that a filter-like form can be used for recursive updating of a least-squares estimator. To illustrate this, consider a univariate response y_t at each point in time, with $y'_t = [y'_{t-1}, y_t]$ and $H'_t = [H'_{t-1}, h'_t]$ representing the accumulation of data up to time t . If \hat{x}_{t-1} and \hat{x}_{t-1} are time-specific least-squares estimates, then

$$\hat{x}_t = \hat{x}_{t-1} + K_t (y_t - h'_t \hat{x}_{t-1})$$

with

$$K_t = \frac{\sum_{\hat{x}_{t-1}} h'_t}{h_t \sum_{\hat{x}_{t-1}} h'_t + \sigma^2}$$

The dispersion of \hat{x}_t also can be expressed recursively as

$$\sum_{\hat{x}_t}^{-1} = \sum_{\hat{x}_{t-1}}^{-1} + h'_t h_t / \sigma^2,$$

or in terms of the estimator gain K_t as

$$\begin{aligned} \sum_{\hat{x}_t} &= \sum_{\hat{x}_{t-1}} - \frac{\sum_{\hat{x}_{t-1}} h'_t h_t \sum_{\hat{x}_{t-1}}}{h_t \sum_{\hat{x}_{t-1}} h'_t + \sigma^2} \\ &= \sum_{\hat{x}_{t-1}} - K_t h_t \sum_{\hat{x}_{t-1}} \end{aligned}$$

The recursive approach to least-squares estimation is described in Walters (1986).

Observation error only: In this case the system state varies through time according to Equation 10, except that process variation is assumed to be captured in the transition matrix F_t . In effect, there are no stochastic components in the state transitions, and the state trajectory is completely determined by the initial state x_0 . The effect of this assumption is registered in a simplified dispersion for \hat{x}_t in Equation 12, which no longer contains a term \underline{W}_t for the dispersion of e_t . All other computing forms are operative, and the overall effect is to increase system gain: basically, with less noise in the system there is more information in the accumulated data about \hat{x}_0 .

Process variation only: In this case the system contains random 'process variation' e_t , but it is observed perfectly. Without loss of generality, the observation Equa-

tion 13 is simplified to $y_t = x_t$, with $H_t = I$ and $V_t = 0$. As mentioned above, the estimation issue then vanishes, and the problem reduces to the optimal control of an observable stochastic system.

Discussion

It is worth keeping in mind that the theoretical developments of the previous section are really about the (apparently) small problem of identifying $f_1(x_{t-1}|Y_{t-1})$ in the decomposition

$$f_{\underline{B}}(y_t|Y_{t-1}, a_{t-1}) = f_1(x_{t-1}|Y_{t-1}) f_{\underline{B}}(x_t|x_{t-1}, a_{t-1}) f_2(y_t|x_t)$$

of transition probabilities in the Hamilton-Jacobi-Bellman equation. By imposing rather stringent linearity conditions on the system transitions as in Equation 10, it is possible to derive efficient algorithms for estimating or 'projecting' the system state x_t at each point in the time frame. On the assumption that stochasticities are normally distributed, the Kalman filter can be shown to produce an optimal estimator \hat{x}_t .

It also is worth keeping in mind the limitations of the Kalman filter, especially for biological systems with severe non-linearities and non-additive variance components. The linear system transitions in Equation 10 can be generalized naturally as in Equation 9, to include mathematical non-linearities as well as non-additivity in the stochastic environmental influences. Either or both features require a different approach to the derivation of $f_1(x_{t-1}|Y_{t-1})$. A possible attack on non-linear system dynamics is by way of a neighbourhood analysis around a system equilibrium utilizing a Taylor series expansion to produce a quasi-linearized approximation of the system equations. This approach leads into the realm of the extended Kalman filters for non-linear systems, a mathematically complex issue that nevertheless is worth exploring in the context of adaptive harvest management.

The focus of the development presented here is on the identification of resource state in the case of partial observability. Together with other recent technical developments (Williams 1996a,b, Williams, Nichols & Conroy in press, B. Lubow, pers. comm.), the preceding material provides the framework needed to actually carry out adaptive optimization. Despite the recent advances on this particular topic, technical challenges nevertheless remain in the implementation of adaptive management. Some of these challenges involve optimization computations, others involve possible changes

in the nature of the objective function and still others involve more fundamental issues associated with the development of the *a priori* model set.

There are very substantial computational requirements with the approach presented here for projecting future resource state in the face of partial observability, and presented elsewhere (Williams 1996a,b, Williams et al. in press) for projecting future information state. Software development for this purpose is proceeding rapidly, however, and should permit application of this approach in the very near future (B. Lubow, pers. comm.). The stochastic dynamic programming approach to computing optimal state-specific policies (Lubow 1994, 1995) has been applied successfully to the adaptive harvest management of midcontinent mallard ducks *Anas platyrhynchos* in North America (Williams & Johnson 1995, Williams, Johnson & Wilkens 1996, Johnson, Moore, Kendall, Dubovsky, Caithamer, Kelley & Williams 1997, Johnson & Williams 1999), a problem involving a single population and four competing models of system dynamics. Recent decisions to consider the use of adaptive management for additional populations of mallards and perhaps other waterfowl species could increase the dimension of the optimization problem, adding substantially to the computational load. Although most of our practical experience is based on the management of hunting regulations, recent attention has focused on habitat acquisition and management (e.g. Johnson, Williams & Schmidt 1996, Johnson, Anderson, Baydack, Nelson, Ringelman, Koneff, Bailey, Martin & Rubec 1997). The simultaneous incorporation of both habitat and harvest actions into an adaptive framework is certainly possible, but the existence of different temporal scales of the two classes of action (e.g. annual hunting regulations, planting of nesting cover at multi-year periods, one-time acquisition of nesting or wintering habitat) may complicate both the modeling and the optimization algorithms.

The objective function currently used in mallard management focuses on the size of the harvest. Because waterfowl hunting in the United States is not a commercial enterprise, it can be argued that the objective function should incorporate measures of hunter 'satisfaction' that include factors in addition to harvest (Ringelman 1997, Johnson & Case 2000). Development of appropriate metrics for hunter satisfaction, monitoring programs for these metrics, and models relating satisfaction to hunting regulations represent an important technical challenge (also see Johnson & Case 2000).

The optimal policies and learning identified via the outlined approach to adaptive management are conditional on the members of the model set. This conditional

nature of the process encourages substantial care and effort in the development of the competing system models. If none of the models closely approximates system dynamics in response to management actions, then the information state of the process will not evolve in the desired manner to decrease uncertainty. Because even models that are well-grounded empirically may not perform well when conditions change, we recommend that biologists focus on mechanistic models to the degree practicable (e.g. Johnson et al. 1993, Williams et al. in press). If the modelled processes themselves change over time, then the information state can be expected to evolve in the direction of the models best approximating reality. The success of the adaptive process in this situation should be a function of the relative rates of change in the underlying process versus the information state.

In addition to these technical challenges, the continuation and expansion of adaptive management of waterfowl harvest will involve a number of institutional challenges. Perhaps the most important of these challenges are political in nature. Adaptive harvest management was implemented by the U.S. Fish and Wildlife Service in 1995 and habitat conditions have remained favourable for mallard production. The combination of reasonable mallard populations and good habitat conditions has led to liberal hunting regulations during each of the years under adaptive management. Political pressure for changes in regulatory decisions tend to be much greater in years of poor populations and/or habitat conditions. It is hoped that harvest managers can withstand political pressures when these conditions occur, and continue to establish hunting regulations according to the policies developed through the adaptive optimization process.

We believe that political input to the adaptive management process should largely focus on the development of objective functions and the identification of regulatory alternatives. Political influence is to be expected with these two components of the adaptive management process, but politically motivated changes should be restricted to periods of program reassessment and should not occur frequently. Frequent changes in the regulatory packages, for example, limit one's ability to predict harvest rates from regulations, and thereby retard the reduction of a principal source of uncertainty (partial controllability).

Another institutional challenge involves the technical and relatively mathematical nature of this formal approach to adaptive management. Work on this topic requires a level of technical expertise that is not common among wildlife managers. The group of sci-

entists and managers responsible for the initial efforts with mallards is relatively small, and it is essential to maintain this group and to train other technicians in this methodology.

Despite these and other challenges (e.g. Johnson & Case 2000), we believe that adaptive processes of the type described in this paper will become increasingly important in North American waterfowl harvest management. The political and scientific selective pressures that resulted in the consideration and adoption of this process in 1995 are even stronger and more compelling today (Nichols 2000). We thus believe that the approach merits the additional research to deal with the remaining technical challenges and to expand the approach to other populations and species.

References

- Anderson, D.R. & Burnham, K.P. 1976: Population ecology of the mallard. VI. The effect of exploitation on survival. - U.S. Fish and Wildlife Service Resource Publication 128, U.S. Fish and Wildlife Service, Washington, DC., USA, 66 pp.
- Bellman, R. & Dreyfus, S.E. 1962: Applied dynamic programming. - Princeton University Press, Princeton, New Jersey, USA, 363 pp.
- Blohm, R.J. 1989: Introduction to harvest - understanding surveys and season setting. - In: Beattie, K.H. (Ed.); Proceedings of the Sixth International Waterfowl Symposium. Ducks Unlimited, Inc., Memphis, Tennessee, USA, pp. 118-133.
- Cowardin, L.M. & Johnson, D.H. 1979: Mathematics and mallard management. - *Journal of Wildlife Management* 43: 18-35.
- Dreyfus, S.E. & Law, A.M. 1977: The art of dynamic programming. - Academic Press, San Diego, California, USA, 284 pp.
- Graybill, F.A. 1976: Theory and application of the linear model. - Duxbury Press, North Scituate, Massachusetts, USA, 704 pp.
- Hilborn, R., Walters, C.J. & Ludwig, D. 1995: Sustainable exploitation of renewable resources. - *Annual Review of Ecology and Systematics* 26: 45-67.
- Holling, C.S. (Ed.) 1978: Adaptive environmental assessment and management. - John Wiley and Sons, London, Great Britain, 377 pp.
- Johnson, D.H., Nichols, J.D., Conroy, M.J. & Cowardin, L.M. 1988: Some considerations in modeling the mallard life cycle. - In: Weller, M.W. (Ed.); Waterfowl in winter. University of Minnesota Press, Minneapolis, Minnesota, USA, pp. 9-20.
- Johnson, D.H., Nichols, J.D. & Schwartz, M.D. 1992: Population dynamics of breeding waterfowl. - In: Batt, B.D.J., Afton, A.D., Anderson, M.G., Ankney, C.D., Johnson, D.H., Kadlec, J.A. & Krapu, G.L. (Eds.); Ecology and management of breeding waterfowl. University of Minnesota Press, Minneapolis, Minnesota, USA, pp. 446-485.
- Johnson, D.H., Sparling, D.W. & Cowardin, L.M. 1987: A model of the productivity of the mallard duck. - *Ecological Modelling* 38: 257-275.
- Johnson, F.A., Anderson, M.G., Baydack, R.K., Nelson, J.W., Ringelman, J.K., Koneff, M.D., Bailey, R.O., Martin, T.E. & Rubec, C. 1997: Enhancing biological performance of the North American Waterfowl Management Plan. - *Transactions of the North American Wildlife and Natural Resources Conference* 62: 377-385.
- Johnson, F.A. & Case, D.J. 2000: Adaptive regulation of waterfowl harvests: lessons learned and prospects for the future. - *Transactions of the North American Wildlife and Natural Resources Conference* 65: 94-108.
- Johnson, F.A., Moore, C.T., Kendall, W.L., Dubovsky, J.A., Cathamer, D.F., Kelley, J.R., Jr. & Williams, B.K. 1997: Uncertainty and the management of mallard harvests. - *Journal of Wildlife Management* 61: 202-216.
- Johnson, F.A. & Williams, B.K. 1999: Protocol and practice in the adaptive management of waterfowl harvests. - *Conservation Ecology* 3: 8.
- Johnson, F.A., Williams, B.K., Nichols, J.D., Hines, J.E., Kendall, W.L., Smith, G.W. & Caithamer D.F. 1993: Developing an adaptive management strategy for harvesting waterfowl in North America. - *Transactions of the North American Wildlife and Natural Resources Conference* 58: 565-583.
- Johnson, F.A., Williams, B.K. & Schmidt, P.R. 1996: Adaptive decision making in waterfowl harvest and habitat management. - In: Ratti, J.T. (Ed.); Proceedings of the Seventh International Waterfowl Symposium. Ducks Unlimited, Inc., Memphis, Tennessee, USA, pp. 27-33.
- Kalman, R.E. 1960: A new approach to linear filtering and prediction theory. - *Journal of Basic Engineering* 82: 34-45.
- Kalman, R.E. & Bucy, R.S. 1961: New results in linear filtering and prediction theory. - *Journal of Basic Engineering* 83: 95-108.
- Lee, P.M. 1992: Bayesian statistics: an introduction. - Halsted, New York, USA, 294 pp.
- Lubow, B. 1994: Stochastic dynamic programming (SDP) user's guide, Version 1.06. - Colorado Cooperative Fish and Wildlife Research Unit, Colorado State University, Fort Collins, Colorado, USA, 138 pp.
- Lubow, B. 1995: SDP: generalized software for solving stochastic dynamic optimization problems. - *Wildlife Society Bulletin* 23: 738-742.
- Martin, F.W., Pospahala, R.S. & Nichols, J.D. 1979: Assessment and population management of North American migratory birds. - In: Cairns, J., Jr., Patil, G.P. & Walters, W.E. (Eds.); Environmental biomonitoring, assessment, prediction, and management - certain case studies and related quantitative issues. International Cooperative Publishing House, Fairland, Maryland, USA., pp. 187-239.
- Meinhold, R.J. & Singpurwalla, N.D. 1983: Understanding the Kalman filter. - *American Statistician* 37: 123-127.

- Montalbano, F.H., III, Johnson, D.H., Miller, M.R. & Rusch, D.H. 1998: Workshop summary: hunting vulnerability and mortality. - In: Weller, M.W. (Ed.); *Waterfowl in winter*. University of Minnesota Press, Minneapolis, Minnesota, USA, pp. 575-579.
- Nichols, J.D. 2000: Evolution of harvest management for North American waterfowl: selective pressures and preadaptations for adaptive harvest management. - *Transactions of the North American Wildlife and Natural Resources Conference* 65: 65-77.
- Nichols, J.D., Conroy, M.J., Anderson, D.R. & Burnham, K.P. 1984: Compensatory mortality in waterfowl populations: a review of the evidence and implications for research and management. - *Transactions of the North American Wildlife and Natural Resources Conference* 49: 535-554.
- Nichols, J.D. & Johnson, F.A. 1989: Evaluation and experimentation with duck management strategies. - *Transactions of the North American Wildlife and Natural Resources Conference* 54: 566-593.
- Nichols, J.D., Johnson, F.A. & Williams, B.K.. 1995: Managing North American waterfowl in the face of uncertainty. - *Annual Review of Ecology and Systematics* 26: 177-199.
- Ringelman, J.K. 1997: Effects of regulations and duck abundance on duck hunter participation and satisfaction. - *Transactions of the North American Wildlife and Natural Resources Conference* 62: 361-376.
- Smith, R.I., Blohm, R.J., Kelly, S.T. & Reynolds, R.E. 1989: Review of data bases for managing duck harvests. - *Transactions of the North American Wildlife and Natural Resources Conference* 54: 537-544.
- Stengel, R.F. 1994: *Optimal control and estimation*. - Dover Press, New York, USA, 639 pp.
- U.S. Department of the Interior 1975: Issuance of annual regulations permitting the sport hunting of migratory birds. U.S. Fish and Wildlife Service Final Environmental Impact Statement. - U.S. Fish and Wildlife Service, Washington, DC., USA, 276 pp.
- U.S. Department of the Interior, Environment Canada, and Secretario de Desarrollo Social MEXico 1994: 1994 update to the North American Waterfowl Management Plan: expanding the commitment. - Washington, D.C., USA, 30 pp.
- Walters, C.J. 1986: *Adaptive management of renewable resources*. - Macmillan, New York, USA, 374 pp.
- Walters, C.J. & Holling, C.S. 1990: Large-scale management experiments and learning by doing. - *Ecology* 71: 2060-2068.
- Williams, B.K. 1996a: Adaptive optimization and the harvest of biological populations. - *Mathematical Biosciences* 136: 1-20.
- Williams, B.K. 1996b: Adaptive optimization of renewable natural resources: solution algorithms and a computer program. - *Ecological Modelling* 93: 101-111.
- Williams, B.K. 1997: Approaches to the management of waterfowl under uncertainty. - *Wildlife Society Bulletin* 25: 714-720.
- Williams, B.K. & Johnson F.A. 1995: Adaptive management and the regulation of waterfowl harvests. - *Wildlife Society Bulletin* 23: 430-436.
- Williams, B.K., Johnson, F.A. & Wilkins, K. 1996: Uncertainty and the adaptive management of waterfowl harvests. - *Journal of Wildlife Management* 60: 223-232.
- Williams, B.K. & Nichols, J.D. 1990: Modeling and the management of migratory birds. - *Natural Resource Modeling* 4: 273-311.
- Williams, B.K., Nichols, J.D. & Conroy, M.J. in press: The analysis and management of animal populations. - Academic Press, San Diego, California, USA.