Application of integrated Bayesian modeling and Markov chain Monte Carlo methods to the conservation of a harvested species

C. J. Fonnesbeck & M. J. Conroy

Abstract

Application of integrated Bayesian modeling and Markov chain Monte Carlo methods to the conservation of a harvested species.— When endeavoring to make informed decisions, conservation biologists must frequently contend with disparate sources of data and competing hypotheses about the likely impacts of proposed decisions on the resource status. Frequently, statistical analyses, modeling (*e.g.*, for population projection) and optimization or simulation are conducted as separate exercises. For example, a population model might be constructed, whose parameters are then estimated from data (*e.g.*, ringing studies, population surveys). This model might then be used to predict future population states, from current population estimates, under a particular management regime. Finally, the parameterized model might also be used to evaluate alternative candidate management decisions, via simulation, optimization, or both. This approach, while effective, does not take full advantage of the integration of data and model components for prediction and updating; we propose a hierarchical Bayesian context for this integration. In the case of American black ducks (*Anas rubripes*), managers are simultaneously faced with trying to extract a sustainable harvest from the species, while maintaining individual stocks above acceptable thresholds. The problem is complicated by spatial heterogeneity in the growth rates and carrying capacity of black ducks stocks, movement between stocks, regional differences in the intensity of harvest pressure, and heterogeneity in the degree of competition from a close congener, mallards (*Anas platyrynchos*) among stocks. We have constructed a population life cycle model that takes these components into account and simultaneously performs parameter estimation and population prediction in a Bayesian framework. Ringing data are used to develop posterior predictive distributions for harvest mortality rates, given as input decisions about harvest regulations. Population surveys of black ducks and mallards are used to obtain stock–specific estimates of population size for both species, for inputs into the population life–cycle model. These estimates are combined with the posterior distributions for harvest mortality, to obtain posterior predictive distributions of future population status for candidate sets of regional harvest regulations, under alternative biological hypotheses for black duck population dynamics. These distributions might then be used for both the exploration of optimal harvest policies and for sequential updating of model posteriors, via comparison of predictive distributions to future survey estimates of stock–specific abundance. Our approach illustrates advantages of MCMC for integrating disparate data sources into a common predictive framework, for use in conservation decision making.

Key words: Bayesian analysis, Integrated model, Hierarchical model, Harvest, MCMC, Waterfowl.

Resumen

Aplicación de la modelación integrada bayesiana y de los métodos Monte Carlo basados en cadenas de Markov para la conservación de una especie recolectada.— En el momento de tomar decisiones bien fundamentadas, es habitual que los biólogos conservacionistas deban enfrentarse a fuentes de datos dispares e hipótesis alternativas acerca de los impactos probables que tendrán las decisiones propuestas en el estado del recurso. A menudo, tanto los análisis estadísticos, como la modelación (para la proyección poblacional, por ejemplo) y la optimización o simulación, se llevan a cabo como ejercicios independientes.

Así, es posible que se construya un modelo poblacional, cuyos parámetros se estimen a partir de datos (como estudios de anillamiento y estudios poblacionales). Posteriormente, cabe la posibilidad de que este mismo modelo se emplee para predecir situaciones demográficas futuras a partir de las estimaciones de población actuales, utilizando para ello un sistema de gestión determinado. Por último, el modelo parametrizado también puede emplearse para evaluar posibles decisiones de gestión alternativas, a través de la simulación, la optimización, o ambos procedimientos. Si bien este enfoque resulta eficaz, no aprovecha al máximo la integración de datos y los componentes de los modelos para la predicción y actualización. En este estudio proponemos un contexto bayesiano jerárquico que permite efectuar dicha integración. En el caso del ánade sombrío americano (*Anas rubripes*), los gestores deben enfrentarse a la labor de intentar extraer una recolección sostenible de la especie, al tiempo que mantienen los stocks de individuos por encima de umbrales aceptables. El problema se ve agravado por la heterogeneidad espacial que presentan las tasas de crecimiento y la carga cinegética de los stocks de ánades sombríos, el movimiento entre los stocks, las diferencias regionales en la intensidad de la presión recolectora y la heterogeneidad en el grado de competencia por parte de un congénere cercano —el ánade real (*Anas platyrynchos*)— entre los stocks. Hemos formulado un modelo del ciclo vital de la población que toma en consideración estos componentes, al tiempo que permite llevar a cabo una estimación de los parámetros y una predicción de la población en un marco bayesiano. Los datos de anillamiento se emplean para desarrollar distribuciones predictivas posteriores para las tasas de mortalidad durante la recolección, expresadas como decisiones de entrada acerca de la normativa sobre recolecciones. Los estudios poblacionales del ánade sombrío y del ánade real se emplean para obtener estimaciones sobre el tamaño poblacional específicas de los stocks de ambas especies, que se emplearán como entradas para el modelo del ciclo vital de la población. Dichas estimaciones se combinan con las distribuciones posteriores para la mortalidad durante la recolección, con el propósito de obtener distribuciones predictivas posteriores de la situación demográfica futura para posibles conjuntos de normativas regionales acerca de la recolección, de acuerdo con hipótesis biológicas alternativas relativas a la dinámica poblacional del ánade sombrío. En una fase posterior, tales distribuciones pueden utilizarse tanto para la investigación de políticas óptimas en materia de recolección, como para la actualización secuencial de distribuciones posteriores del modelo mediante la comparación de distribuciones predictivas para estimaciones en estudios futuros acerca de la abundancia poblacional presente de forma específica en los stocks. Nuestro enfoque ilustra las ventajas que presentan las técnicas de Montecarlo basadas en cadenas de Markov (MCMC) para integrar fuentes de datos dispares en un marco predictivo común, con vistas a su utilización en la toma de decisiones sobre conservación.

Palabras clave: Análisis bayesiano, Modelo integrado, Modelo jerárquico, Recolección, MCMC, Aves acuáticas.

Christopher J. Fonnesbeck & Michael J. Conroy, Cooperative Fish and Wildlife Research Unit, D.B. Warnell School of Forest Resources, Univ. of Georgia, Athens, GA 30602 U.S.A.

Introduction

Dynamic models frequently are used in conservation biology to aid in the evaluation of alternative conservation decisions, with respect to obtaining some desired outcome. Such approaches ordinarily employ, whether explicitly or not, several elements: First, a *process model* is used to describe how observed system states, such as population size and composition, change through time. Second, system states and relevant parameters such as survival and recruitment rates must be estimated using data, via one or more *statistical models*. Third, one or more control variables must be defined; these are thought to influence the system in such a way as to lead to gains in the management objective. Examples of decisions that are frequently made, and to which decision modeling may be applied, include determination of optimal harvest regulations, setting of forest cutting policies, the timing and intensity of restocking efforts as components of endangered species conservation, and decisions about land acquisition for conservation. Fourth, we need an explicit way of describing the relative value or utility of potential management outcomes. That is, either explicitly or implicitly, there is some overarching resource goal and quantifiable resource objective in any decision–making process. For harvest decisions the objective typically is the maximization of long–term harvest yield; for forest cutting the objective may be the gain of revenue, perhaps subject to constraints on the avoidance of loss of biodiversity; for restocking efforts, perhaps the maximization of the expected time to extinction for some species; for land acquisition, perhaps the maximization of biodiversity conservation under budgetary constraints. Finally, we need a procedure that seeks some optimum combination of decisions and system conditions. Formal procedures exist for all these elements; in ecology, there is a particularly rich literature focused on dynamic modeling and statistical estimation methods, much of which is summarized in Williams et al. (2002). Likewise, there exists an extensive literature on decision theory and, dynamic optimization methods and optimal control theory, much of it also summarized by Williams et al. (2002). However, in our experience, process modeling, statistical estimation, and decision analysis are often considered as distinct enterprises. Thus, statistical models are frequently used to estimate population states and other parameters; these results may subsequently be incorporated into an existing or newly constructed population model. The parameterized model then may be applied to a decision problem, for instance, by exploratory simulation or formal optimization procedures. Although the sequence of events differs from cases to case (*e.g.*, model constructed first, followed by parameter estimation and optimization), the idea is the same: component elements are treated separately, and integration (to the extent it occurs at all) is usually post–hoc, and often ad–hoc. Philosophically, there are close linkages among these elements; pragmatically, there are also strong arguments for integration. For instance, dynamic optimization models include both state dynamics and an objective function. In turn, data are required to estimate system states and state dynamics, and to assess model comportment to reality, which in turn should influence decision making. Because these elements typically depend upon a common data structure, and involve modeling, there is practical motivation for an integrated approach that leverages shared information. In this paper, we first construct a conceptual framework for integrating process modeling, parameter estimation, and model prediction, based on principles of conditional hierarchical modeling. We then demonstrate the approach using a real decision problem, involving optimal harvest management of multiple stocks of American black ducks.

Methods

A generic decision model

We clarify these concepts by means of a generic decision model. To begin with, consider a dynamic system in which the state $\mathbf{X}_t$ (possibly vector–valued) evolves through time according to a specified process model f $(\mathbf{X}_t|\mathbf{Z}_t)$, which includes both the endogenous effect of the state, as exogenous factors $\mathbf{Z}_t$, such as weather; the latter are frequently modeled as random variables (fig. 1). Add to this model inputs from the vector of decisions $\mathbf{d}_t$, which potentially affect both transitions to future states and the utility gained from present and future states (possibly altered under management). Before proceeding further, we wish to use this generic model to reinforce our earlier points. First, figure 2 makes clear that the modeling of decision influences (both on the system itself and our objective gain from the system) are inextricably linked to the process model. In addition, any algorithm that seeks to find decisions which optimize the objective functional is constrained by system dynamics. That is, it is impossible to obtain the maximum of a dynamic decision problem without taking into account system dynamics. Finally, dynamic decisions are often subject to a finite time horizon. In natural resource management it is usually appropriate that this horizon is relatively distant. Decisions are made, usually with feedback from the current system state, in order to seek an optimal result over an appropriately long time horizon. Although simple conceptually, this integration of dynamic modeling with decision making is complex in practice. Further, decisions ordinarily cannot be based directly on the system states and a model, but rather on statistical estimates of the states, and of the parameters of the model. Thus, the actual system state being modeled evolves through time according to some (assumed) model. The observed system state is related to the actual system according to a statis-
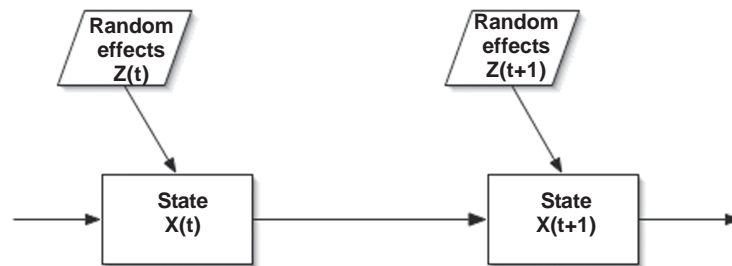
Fig. 1. Generical model of a dynamic process with random effects.

Fig. 1. Modelo genérico de un proceso dinámico con efectos aleatorios.

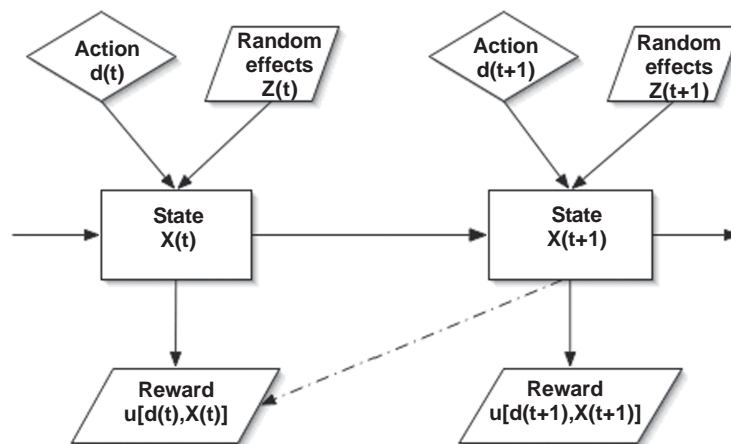

Fig. 2. Incorporation of decision variables and objective function into the stochastic process, showing feedback of future state upon the objective.

Fig. 2. Incorporación al proceso estocástico de variables de decisiones y de la función objetiva, indicando el feedback *del estado futuro con respecto al objetivo.*

tical sampling model. Parameters of the process and decision models are now themselves based on estimates from one or more statistical sampling models, providing a prediction given the current (observed) state (fig. 3):

$$\tilde{\mathbf{X}}_{t+1} = f\left(\mathbf{Y}_t \left| \hat{\theta}, \mathbf{d}_t, \mathbf{Z}_t \right.\right) \qquad (1)$$

Predictions about the future state of the system now inherit statistical uncertainties due these statistical models, as do predictions about utility under each candidate management decision.

The development so far has assumed that the mathematical form and parametric structure of the process model are known. Usually this will not be

the case, therefore it will be important to consider alternative process models. In the context of decision making, these alternative models become important to the extent that the utility of decisions is dependent on belief in the alternate models. Suppose we entertain a single alternate model, denoted as Model 2 (fig. 4). Identical observed system states and candidate decisions induce two sets of values for the predicted utility and predicted system state, one set under each model. Given these alternate predictions, the decision maker must first reconcile the fact that different models may lead to different utility for each decision; thus, the optimal decision may be different for each model. One approach is to form an
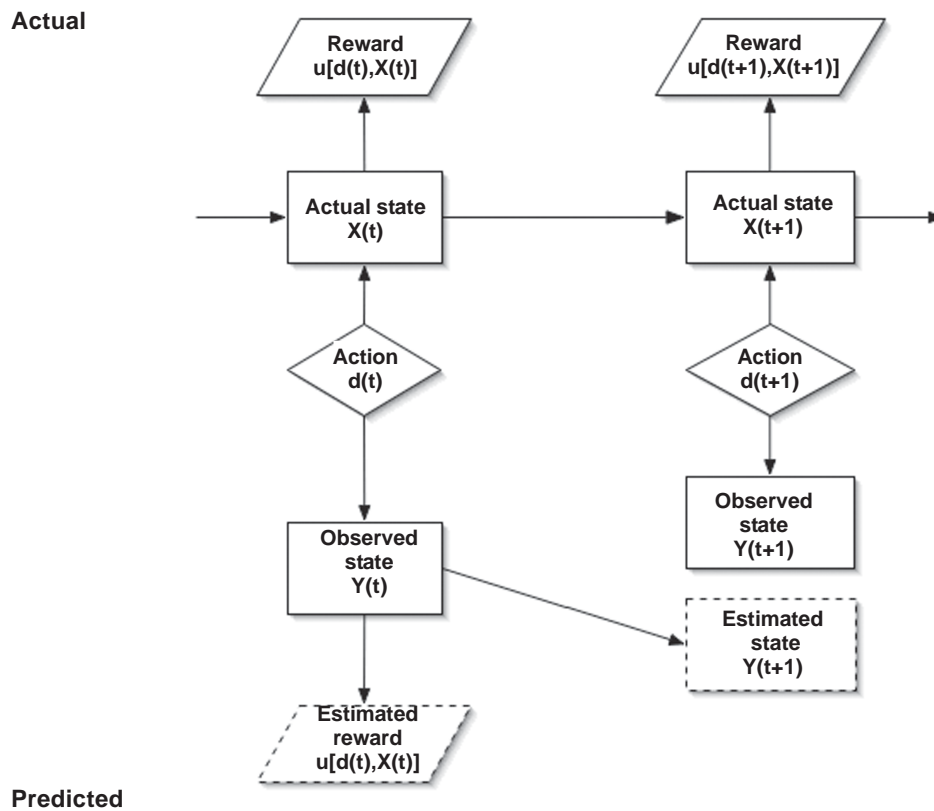
Actual

Fig. 3. Relationship between actual and observed/estimated processes in a dyanmic system.

Fig. 3. Relación entre los procesos reales y los observados/estimados en un sistema dinámico.

expectation of utility across the models, which in turn depends on model probabilities quantifying relative belief in each model. Initially, model probabilities may be impartial, attributing equal weight to each candidate model.

Finally, following decision making, monitoring data are used to compare the predictions under each model for observed future states (fig. 3). These predictions and observations have at least two potential uses: (1) they provide an obvious means of evaluating the degree to which any of the predictive models performs (*i.e.*, validation), as well as a means of discriminating among competing models. As we will describe, this leads to a natural way to update relative belief in the alternate models, using Bayes factors; (2) as noted above, optimal decisions and their utility will generally differ among alternate process models. Thus, to the degree that uncertainty exists with respect to which model best describes and predicts the process, prescribed decisions will be suboptimal; conversely, reduction of process uncertainty through time will result in improved decision making at future decision–making epochs.

Sequential conditioning as a tool for integration

Our approach to this problem exploits well–known and related principles of probability, Bayesian inference and conditional modeling. In Bayesian inference, parameter values θ and observations $\mathbf{Y}$ are both modeled with probability density functions, so it makes sense to consider their joint probability $Pr(\theta, \mathbf{Y})$

$$Pr(\theta, \mathbf{Y}) = Pr(\theta | \mathbf{Y}) Pr(\mathbf{Y}) = Pr(\mathbf{Y} | \theta) Pr(\theta) \quad (2)$$

Because for any sample outcome the probability of the data is a constant, Bayes' theorem states that the posterior (that is, following sampling) probability of θ is proportional to the product of the sampling distribution of the data assuming θ and the prior (unconditional) probability of θ:

$$Pr(\theta, \mathbf{Y}) \propto Pr(\mathbf{Y} | \theta) Pr(\theta) \quad (3)$$

This formulation is readily generalized to alternative model forms. Conditional modeling has increasingly been recognized as a powerful tool for
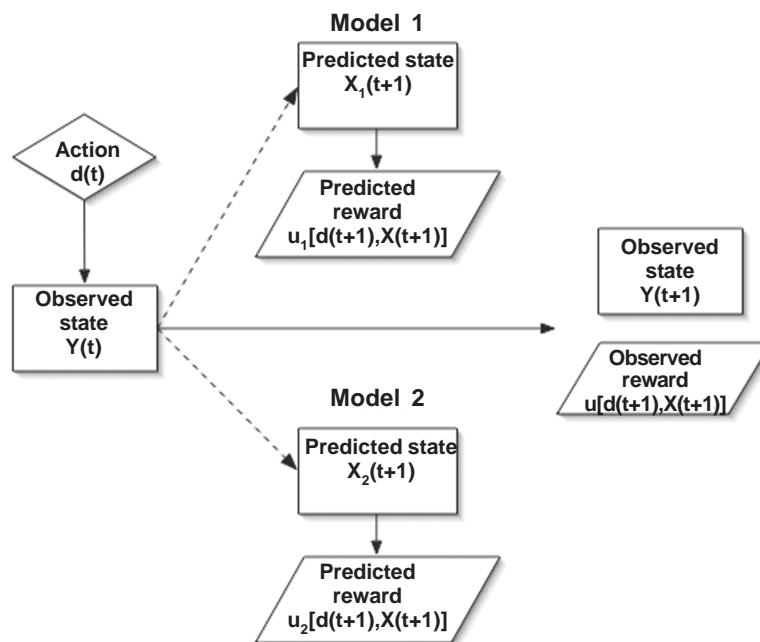
Fig. 4. Alternative process models and predictions. Following decision making, monitoring data are obtained for comparison of the predictions under each model to observed future states. These predictions may then be used to compute Bayes factors for updating the relative belief in the alternate models, for use in the next decision–making epoch.

Fig. 4. Modelos de procesos alternativos y predicciones. Tras la toma de decisiones, se obtienen datos de control para comparar las predicciones facilitadas por cada modelo con los estados futuros observados. Estas predicciones pueden utilizarse para calcular factores de Bayes y actualizar la creencia relativa en los modelos alternativos, con vistas a su utilización en el siguiente período de toma de decisiones.

modeling complex ecological relationships by decomposition as simpler element that are related in a conditional, often hierarchical, manner (*e.g.*, Wikle, 2003). In addition, hierarchical processes naturally lend themselves to conditional sequencing.

These ideas can be effectively combined to solve our generic decision problem in an adaptive, sequential updating model, by expressing the problem as:

$$Pr\left(M,\theta,\tilde{\mathbf{X}}|\mathbf{Y}\right) \alpha \, Pr\left(\mathbf{Y}|M,\theta,\tilde{\mathbf{X}}\right) Pr\left(M,\theta,\tilde{\mathbf{X}}\right) \quad (4)$$

where M represents model structure, θ are parameter values, $\tilde{\mathbf{X}}$ is the predicted and $\mathbf{Y}$ the observed system state. By sequential conditioning we can see that:

$$Pr\left(\tilde{\mathbf{X}},\theta,M|\mathbf{Y}\right) \alpha \, Pr\left(\mathbf{Y}|\tilde{\mathbf{X}},\theta,M\right) Pr\left(\tilde{\mathbf{X}}|\theta,M\right) Pr\left(\theta|M\right) Pr\left(M\right) \quad (5)$$

This decomposition provides a natural, sequential way of dealing with the complexity of simultaneously modeling uncertainty in data, models, and processes. Beginning with the right–most term in (5), conditioned on a distribution of parameter values $Pr(\theta)$, we have predicted values $\tilde{\mathbf{X}}$. Next in

this sequence is the probability of model M, and finally the likelihood of the data, given the model, parameter values, and predictions. We use Markov chain Monte Carlo (MCMC) methods to sample from the conditional posterior distributions of the quantities we seek (*e.g.*, models, predictions). Although not explicitly explored in this paper, conditional decomposition is readily extended to decision making problems, by incorporating a utility function.

Integrated modeling with Markov chain Monte Carlo methods

Markov chain Monte Carlo is a class of general simulation techniques used primarily to solve problems of Bayesian inference (Gamerman, 1997). Specifically, these methods are used to generate samples that are distributed according to some target posterior form $\pi(\theta)$ without having to directly sample from the posterior itself. This is most useful when $\pi(\theta)$ is extremely complex, or otherwise difficult to analyze. For example, the joint posterior characterizing our problem in (5) can be analytical intractable for many biological problems

All MCMC algorithms produce samples from a set of densities (hence, Monte Carlo simulation); these distributions are derived from $\pi(\theta)$, according to the conditional probability of each θ_i. The current set of parameter values $\{\theta_i^{(t)}\}$ is dependent on the previous values $\{\theta_i^{(t-1)}\}$, thereby generating a Markov chain. The Markov chain is constructed in such a way that its limiting distribution, $\pi(\theta)$, is the distribution of interest.

The Metropolis–Hastings algorithm (Metropolis et al., 1953; Hastings, 1970) is the most general MCMC procedure, and therefore the most widely applicable. The Metropolis–Hastings algorithm estimates the posterior density using a form of rejection sampling (see Robert & Casella, 1999). The proposal function $q(\phi|\theta)$ generates candidate values for $\pi(\theta)$ which are accepted or rejected according to each value's probability under the target . Provided that the full support of $\pi(\theta)$ may be sampled, the choice of q can otherwise be arbitrary. In any case, the Metropolis–Hastings algorithm estimates the posterior form directly from a subset of filtered samples, rather than relying on Bayes' rule and conjugate distributions.

Case study: integrated estimation and prediction for American black ducks

To illustrate our approach we use an example of an adaptive decision–making model for American black ducks. The black duck problem involves a process of observation, predictive modeling, and optimization, with the following major elements: (1) historical data have been used to fit empirical relationships between population parameters and key hypothesized factors, under alternative models; (2) a population projection model incorporates key relationships into a discrete–time projection model; (3) population surveys are used each year to infer the state of the system; (4) based on the surveys, parameter estimates, and the projection model, a forecast of system state and of expected utility is obtained; (5) the forecast of system state is compared to observed system state at the next time epoch in order to evaluate the relative predictive ability of the alternative models and compute relative weights for each model; and (6) the parameter estimates, models, and model weights are used in dynamic optimization procedures (*e.g.* Lubow, 1995) to obtain optimal, state–specific decisions for maximizing expected utility.

The models and data structures used are an extension of those described in Conroy et al. (2002) to multiple stocks, and is illustrated in schematic form in figure 5. The elements of this model are summarized as follows: (1) the observed system state is the number of black ducks (N_i) and mallards (M_i) in each of three geographic strata (western, central, and eastern portions of the range in Canada), as estimated from surveys conducted by the Canadian Wildlife Service; (2) historical data involving band recoveries, hunter surveys of wings collected from shot ducks, and population surveys, are used to estimate the relationship between stratum–specific fall age ratios (A_i) and black duck and mallard abundance, under alternative models of density–dependence and density–dependence with competition (Zimpfer, 2004; Conroy et al., 2002); (3) historical data involving band recoveries and population surveys are used to estimate the relationship between non–harvest survival (S^b_j) and black duck abundance, under alternative models of density–dependent (compensatory) and density–independent (additive) mortality (Conroy et al., 2002); (4) the fall age ratios calculated above, together with spring–to–summer survival (S^b, assumed constant) are used to project fall abundance in each breeding area prior to migration (F^n_j); (5) historical data involving band releases stratified by breeding and harvest areas, and recoveries and recaptures stratified by harvest and breeding areas, are used to estimate rates of movement from breeding to harvest areas (ψ_{ij}), and of return (fidelity) to breeding areas (ϕ_{ji}; Zimpfer, 2004); (6) historical data involving band recoveries, harvest regulations, and hunter numbers stratified by harvest regions, are used to develop predictions of harvest rates conditional on harvest regulations and hunter numbers.

For the purposes of this study, we assumed fixed values for the parameter estimates of non–harvest survival (S^a, S^b) and movement (ψ_{ij}, ϕ_{ji}) rates, and perfect ability to control harvest rates at specified values via harvest regulations. We further conditioned on the observed system state $\mathbf{X}_t = [N_1, N_2, N_3, M_1, M_2, M_3]_t$. Thus, focus is on uncertainty in stratum–specific estimates of age ratio, which, in turn, induce uncertainty in the projection model, from two sources: statistical uncertainty in the parameter values, conditioned on an assumed model structure $\Pr(\theta|M)$, and uncertainty in model structure $\Pr(M)$. Predictions under these alternative, estimated models are then compared to observations of the consequent states of system $\mathbf{X}_t$. Stratum–specific age ratios are estimated under a joint likelihood of wing survey and band recovery data (fig. 6). This consists of the following components:

$$
\begin{aligned}
W_y &\sim \text{Bin}(\alpha, W) \\
m_y &\sim \text{Bin}(h_y, R_y) \\
m_a &\sim \text{Bin}(h_a, R_a) \\
\alpha &= A/(A + \tau) \\
\tau &= h_a/h_y
\end{aligned}
\tag{6}
$$

where W_y is the number of juvenile wings in the harvest survey for a specific reproduction area, W is total (adult + juvenile) wings, $\{m_i, R_i; i = y, a\}$ are number of recoveries and bands, respectively, for juveniles and adults, $\{h_i; i = y, a\}$ are band recovery rates, τ is relative vulnerability of adult to juvenile harvest, and α is proportion of young in the harvest (adjusted from population age ratio A by τ. These likelihoods are, in turn, embedded in a model describing the relationship between age ratio and black duck and mallard abundance:
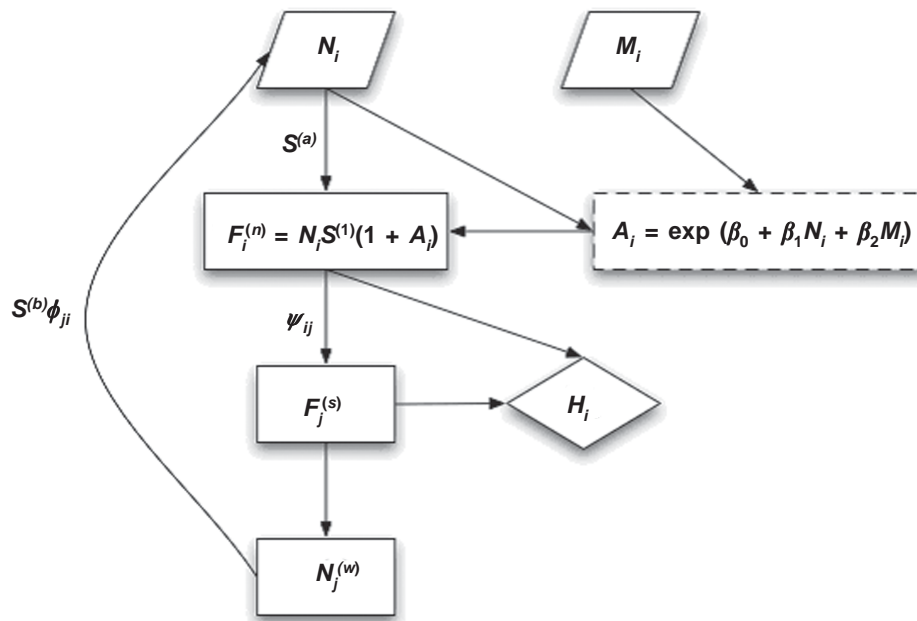
Fig. 5. Flow diagram of black duck state dynamics model. Stratum–specific numbers of black ducks (N_i, i = 1,2,3) survive from spring to fall at the constant rate S^a. Stratum–specific black duck and mallard (M_i) abundance influence fall age ratios A_i according to a production model (see (6) for details; age ratios are applied to the surviving adult population to determine fall abundance prior to migration and harvest ($F^{(s)}_i$). Surviving birds migrate to southern harvest/wintering areas j = 1,2,3 according to fixed rates ψ_{ij}, with post–harvest (H_j) abundance leaving N_j, j = 1,2,3 in each wintering area. Survival over winter (S^b_j) is determined by a density–dependent model as a function of area–specific abundance $N^{(w)}_j$, with fidelity to breeding areas at a fixed rate ϕ_{ij}. In present study movement and survival parameters are fixed, with only the parameters of the reproduction estimation and prediction model (6) estimated from data.

Fig. 5. Diagrama de flujo del modelo de la dinámica del estado del ánade sombrío. Los números de ánades sombríos específicos al estrato (N_i, i =1,2,3) sobreviven de primavera a otoño a una razón constante de S^a. La abundancia de ánades sombríos y ánades azulones específica al estrato (M_i) influye en las tasas de edad durante el otoño A_i según un modelo de producción (para detalles, ver (6); las tasas de edad se aplican a la población adulta superviviente para determinar la abundancia otoñal antes de la migración y la recolección ($F^{(s)}_i$). Las aves supervivientes migran hacia áreas de recolección/hibernación meridionales j = 1,2,3 según tasas fijas ψ_{ij}, de manera que la abundancia posterior a la recolección (H_j) deja N_j, j = 1,2,3 en cada área de hibernación. La supervivencia durante el invierno (S^b_j) viene determinada por un modelo dependiente de la densidad como una función de la abundancia específica a un área $N^{(w)}_j$, con fidelidad a las áreas de reproducción a una tasa fija ϕ_{ij}. En el presente estudio, los parámetros de movimiento y de supervivencia son fijos, de manera que los parámetros del modelo de predicción y de estimación de reproducción (6) son los únicos que se estiman a partir de los datos.

$$A_t(N_t, M_t) = \exp(\beta_0 + \beta_1 N_t + \beta_2 M_t) \qquad (7)$$

where $\{\beta_j\}$ are coefficients to be estimated. Harvest rates could not be directly estimated during the period of our study, because of problems induced by the conversion to toll–free solicitation of bands (see Conroy et al., 2002). We instead used estimates from hunter surveys in the U.S. and Canada to obtain an estimate of total annual harvest per harvest region $\mathbf{H}_t$, which was assigned to age and sex categories in our model

according to differential vulnerability estimates τ. These estimates were used together with the projected pre–harvest population in each region (according to our process model) to estimate harvest rates and harvest mortality, with the constraint that the latter could not exceed unity. Conditioned on observed $\mathbf{Y}_t = (\mathbf{N}_t, \mathbf{M}_t, \mathbf{m}_{gt}, \mathbf{m}_{st}, \mathbf{W}_{gt})$, we used our process model to generate posterior predictions of stratum–specific abundance for black ducks at the next survey period by the relationship
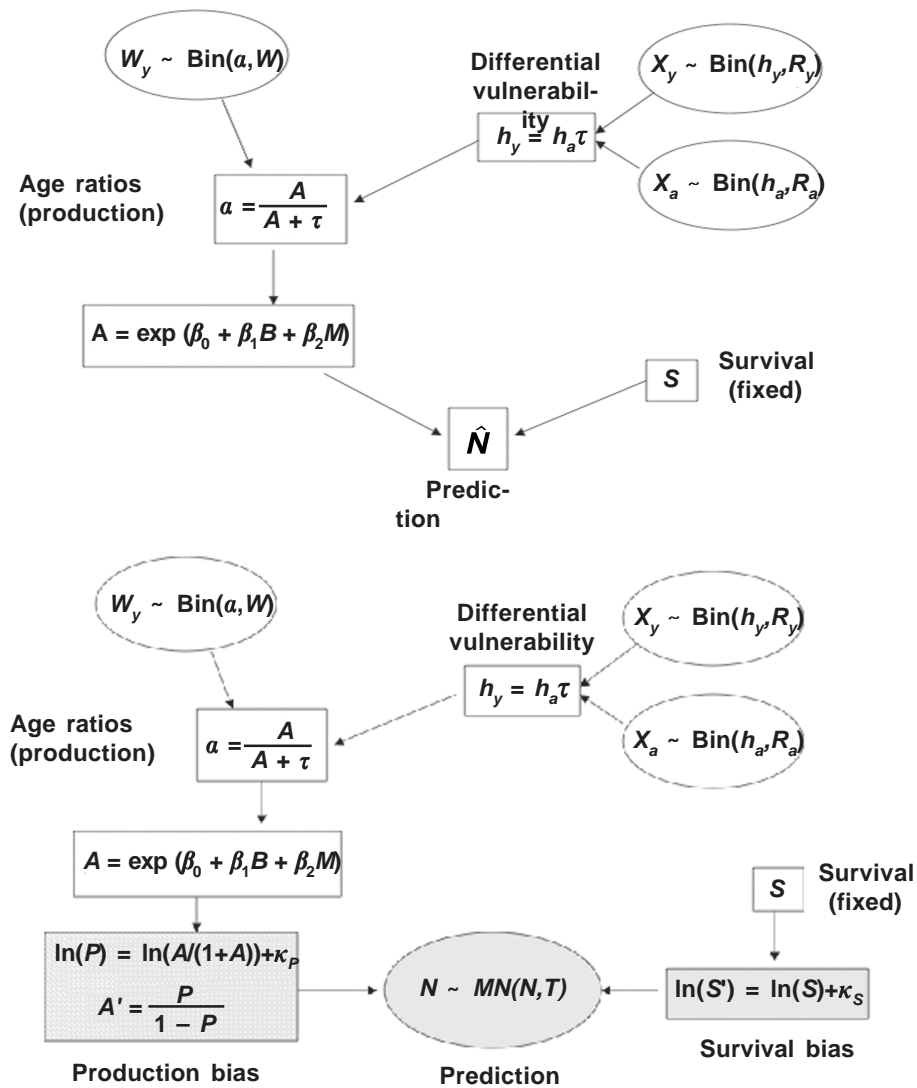
Fig. 6. Model for black duck reproduction estimation and prediction using non–integrated (top) and integrated (bottom) frameworks. Binomial likelihoods (ellipses) are calculated for proportion of juveniles in the harvest (a) from wing recovery data (W), and for stage–specific harvest rates (h) based on bandings (R) and dead recoveries (X). These are, in turn, used to calculate age ratios (A) that form the basis for the reproduction model, which includes density–dependent terms for black ducks (B) and mallards (M). This, along with a fixed survival function, is then used to predict the next year's population size. The integrated form is identical to the non–integrated, except that the predicted population is now part of a multivariate log–normal likelihood, as the expected value of the observed population size. Each integrated model incorporates zero or more bias parameters (κ).

Fig. 6. Modelo para la predicción y estimación de reproducción del ánade sombrío utilizando marcos no integrados (parte superior) y marcos integrados (parte inferior). Las probabilidades binomiales (elipses) se calculan para la proporción de individuos jóvenes en la recolección (a) a partir de los datos de recuperación de alas (W), y para las tasas de recolección específicas a una etapa (h) basándose en anillamientos (R) y en la recuperación de aves muertas (X). Tales probabilidades se utilizan, a su vez, para calcular las tasas de edad (A) que forman la base para el modelo de reproducción, incluyendo términos dependientes de la densidad para el ánade sombrío (B) y el ánade real (M). Esto, junto con una función de supervivencia fija, se utiliza despues para predecir el tamaño que tendrá la población el próximo año. La forma integrada es idéntica a la no integrada, salvo que ahora la población prevista forma parte de una probabilidad logarítmica normal multivariante, expresada como el valor previsto del tamaño de la población observada. Cada modelo integrado incorpora cero o más parámetros de sesgo (κ).

$$\tilde{\mathbf{N}}_{t+1} = f\left(\mathbf{Y}_t, \hat{\theta} \,\middle|\, \gamma\right)$$

where $\qquad \hat{\theta} = \left[\hat{h}_a, \hat{h}_y, \hat{\tau}, \hat{a}, \{\hat{\beta}_i\}\right] \qquad$ (8)

Here, γ are parameters fixed as constants for this analysis (*e.g.*, movement and survival rates), and f(.) is the functional form specified under our process model. Finally, we modeled observed abundance $\mathbf{N}_{t+1}$ via a multivariate log–normal distribution, centered at the predicted abundance $\tilde{\mathbf{N}}_{t+1}$:

$$\mathbf{N} \sim MLN\left(\tilde{\mathbf{N}}, \mathbf{T}\right) \alpha \left(\prod_{i=1}^{3} \mathbf{N}_i\right)^{-1} |\mathbf{T}|^{\frac{1}{2}}$$

$$\exp\left[-\frac{1}{2}\left(\log\mathbf{N} - \log\tilde{\mathbf{N}}\right)'\mathbf{T}\left(\log\mathbf{N} - \log\tilde{\mathbf{N}}\right)\right] \quad (9)$$

where subscripts are suppressed for notational simplicity. Previous approaches have involved independent estimation of θ, then using these estimates to predict future states $\mathbf{N}_{t+1}$. Experience with the black duck and other duck harvest models has uncovered apparent, systematic over–prediction from this approach. An integrated, hierarchical approach, outlined above, endeavors to remediate this problem by establishing feedback between prediction and estimation (fig. 6). In our case, this would be achieved via adjustments of the coefficients of the age ratio process models A_t = (N_t, M_t). However, it is not known whether over–prediction is due to the reproduction model, or the survival process model (which, for the purposes of this study, we have assumed has fixed coefficients) (Conroy et al., 2002). Indeed, it is possible that neither component induces the bias, but rather, some aspect of the sampling process itself is flawed. Therefore, in addition to the (implicit) adjustment contained in the integrated age ratio process model, we explicitly modeled systematic biases in predicted survival and age ratio:

$$\log(S'_t) = \log S_t + \kappa_s$$
$$\log(A'_t) = \log A_t + \kappa_a \qquad (10)$$

where S_t, A_t are the values of survival and age ratio, respectively, predicted from the model conditioned on current states and data, and κ_s, κ_a are log–scale biasing factors which are estimated; the values S'_t, A'_t are then used in prediction.

MCMC implementation

We implemented the black duck parameter estimation and prediction models using Python (http://python.org), an open source, object–oriented programming language. Python is a modular development environment, with a wide selection of third–party scientific and numerical tools suitable for biometric applications. We developed a Python module, PyMC (http://pymc.sourceforge.net), that implements an adaptive random walk Metropolis-Hastings algorithm for MCMC sampling. At each iteration of the algorithm, new parameter values are proposed according to a random walk. The

increment $\in_t$ is generated using a N(0,σ) density. The advantage of the random walk approach is that no problem–specific restrictions regarding the form of the proposal distribution need to be considered. Proposal distributions for Metropolis–Hastings sampling must be enveloping, such that $q(x) \geq \pi(x)\ \forall x$. Therefore, most proposal distributions must be chosen manually for each variable in each problem. In contrast, the random walk algorithm functions independently of the form of the target $\pi(\theta)$. The disadvantage of the random walk is that when steps are relatively large (i.e., large σ), proposed values may frequently fall in the tails of the target distribution, resulting in unacceptably low acceptance rates; similarly, if increments are relatively small, acceptance rates may be very high, but the rate of mixing (exploration of the support of θ) will be correspondingly low. Either extreme is inefficient, and therefore undesirable. The adaptive random walk implemented in PyMC addresses this inefficiency by adapting the scale parameter of the proposal distribution according to the recent acceptance probability for each parameter during the simulation. Every *k* iterations (*k* = 100 is default), the variance is decreased for acceptance rates below 20%, and increased for those above 50% (arbitrarily chosen), thereby balancing proposal acceptance and mixing. Adaptation occurs in the burn–in phase of the algorithm, and continues until all parameter acceptance rates fall within the aforementioned interval.

The joint likelihood for the integrated model included three components related to estimation and prediction. Binomial likelihoods were calculated for the adult and juvenile harvest rates, based on banding and recovery data for these groups. Additionally, the probability of the proposed proportion of young in the harvest a was calculated using wing data in a binomial likelihood. As illustrated in (6), a is also related to differential vulnerability and age ratio. The final component is the likelihood of the observed population given the predicted value of the model. In each year, the likelihood of the observed population in each area was calculated based on a multivariate log–normal density centered at the array of predicted values. The sum of logarithms for these likelihoods were passed to the Metropolis–Hastings sampler in PyMC after every proposal of a new parameter; this joint log–likelihood was used to either accept or reject the proposed value. Because all parameters are assumed to have equal priors, these cancel out from the ratio of posterior densities used to calculate a. The proposal distributions $q(\theta^{(t)}_i | \phi)$ and $q(\phi | \theta^{(t)}_i)$]similarly drop out, since the probability of jumping from $\theta^{(t)}_i$ to ϕ is equal to the reverse jump, under the random walk strategy.

Model scenarios

A suite of 10 distinct model scenarios was specified. The first was a *null* model which separately estimated parameters of the age–ratio reproduction function, then used this function to predict area–specific

black duck populations for 11 years (1991–2001). Four additional scenarios each used the integrated approach described above to simultaneously estimate reproduction model parameters and predict subsequent population size. Each of the integrated models estimated some combination of vital rate bias parameters described in (10): Model 10 assumed reproduction bias, but no survival bias; Model 01 assumed survival bias, but no reproduction bias; Model 11 estimated both bias terms; while Model 00 estimated neither. These five models were replicated under two alternative biological models for black duck reproduction, one incorporating a Mallard competitive effect, the other excluding this effect, for a total of 10 scenarios.

The Metropolis–Hastings sampler in PyMC produced predictions and estimates for the model set. A total of 100,000 sampling iterations for each model were executed, with the first 50,000 conservatively discarded as "burn–in" samples, assuming that convergence had been achieved by that stage. Manual inspection of sample traces suggested convergence and adequate mixing of each chain. We compared reproduction model parameter estimates β among model scenarios, as well as bias factor estimates, where relevant, using 95% Bayesian credible intervals derived from the posterior distribution of the final 50,000 samples. We also calculated the log–bias of each population $i = 1,2,3$ in each of $t = 1,...,10$ years:

$$B = \log(\tilde{N}_{it}) - log(N_{it}) \qquad (11)$$

Model selection was performed using Akaike's Information Criterion (Burnham & Anderson, 2002; Akaike, 1973, AIC), calculated for each model at each iteration. Burnham & Anderson (2002) illustrate the equivalence of AIC model weights and Bayesian posterior model weights, provided that model priors are equivalent (as we have specified). Use of AIC greatly simplifies model selection in a Bayesian framework relative to other approaches, such as reversible jump MCMC (King & Brooks, 2002). The lack of random effects and the relatively small set of models in this study eliminated the need for procedures that are far more complex to implement. The calculation of AIC at each MCMC iteration yielded a distribution of values, rather than the typical scalar value, which explicitly characterizes parametric uncertainty and its interaction with model selection uncertainty.

Results

Figure 7 illustrates the systematic over–prediction resulting from separately estimating vital rates, then using those rates in a predictive model. This effect is most severe in the Western and Central populations; moreover, over–prediction is higher in earlier years relative to later. Predictions derived from an integrated framework (fig. 8) are only subtly less positively biased overall. The ad-
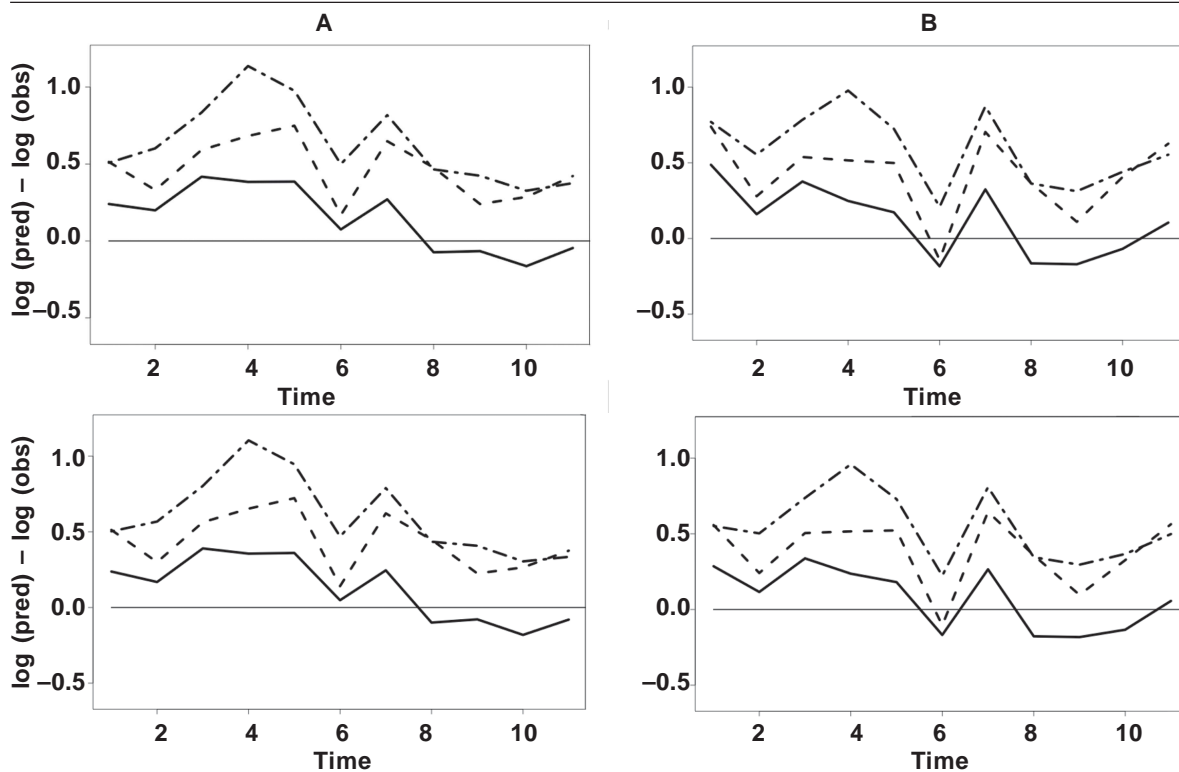
dition of a reproduction bias term (fig. 9) or both bias terms (fig. 11) produces a more dramatic reduction in prediction bias, particularly earlier in the time series. Here, the relative over–prediction of the West and Central populations is balanced by under–prediction in the East to achieve relative unbiasedness overall. The addition of the survival bias parameter did not appreciably improve prediction (fig. 10). Calculated AIC–based model weights reinforce the influence of reproduction bias on prediction, as these models account for over 90% of total weight (table 1). Much biological uncertainty also remains in the form of the reproduction function, with neither model dominating the other consistently, with respect to AIC weight.

Estimates of reproduction and survival rate bias factors are summarized in table 2. A positive reproduction bias was estimated when survival bias was assumed absent, under Model 10, while a negative survival bias is discovered in the absence of reproduction bias using Model 01 (though 95% credible intervals include zero). Specifying dual bias results in a positive reproduction bias under the no–competition reproduction model (Model 11), virtually no bias when competition is assumed (Model 11c), and a positive survival bias estimate in either case. Complementary to these estimates are those of the reproduction model parameters (table 3). The age ratio reproduction model parameter estimates are strikingly similar among statistical bias models, and between integrated and non–integrated models. All have approximately equivalent intercepts and a pattern of increasing negative density dependence west to east. Stronger differences are evident between biological models, where mallard competition effects are balanced by generally larger intercept values relative to those of the non–competitive models.

Discussion

The pattern of over–prediction that pervades black duck population models may well be independent of the quality of the model parameter estimators, or the data used by them. A potential explanation for these systematic and unidentified biases in observed population size is a flawed breeding survey (or wintering survey, depending on which model is employed (Conroy et al., 2002)). Some have suggested that the current survey design is inadequate for reliably estimating the breeding population (Bordage, 2000); a non–standard survey over the past decade may in fact be responsible for the particularly acute over–prediction in the early part of the time series (D. Bordage, pers. comm.). This type of bias cannot be accounted for by our integrated models.

Assuming, however, that an important component of the existing bias is due to the models or data for estimating vital rates, an integrated framework such as that presented here may prove beneficial.
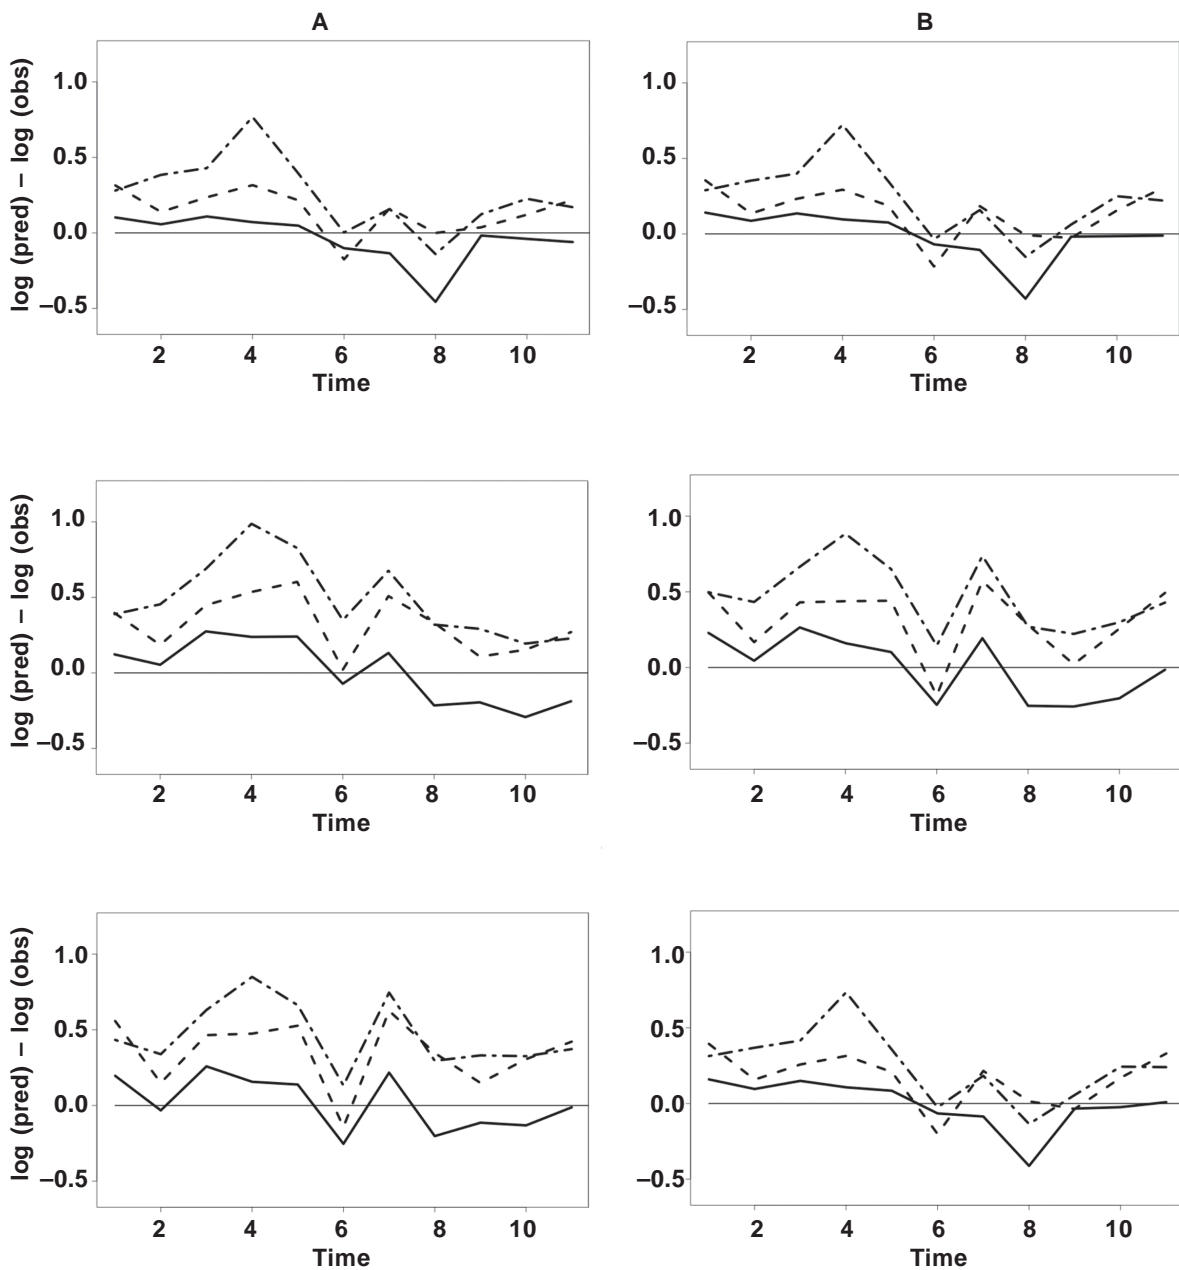
Figs. 7, 8. Bias in population predictions from non–integrated model (7) and from integrated model with no bias factors (*Model 00*) (8), relative to actual population estimates for production with mallard competition (A) and without (B) over 11 years (1991–2001). Plots indicate differences log(predicted) – log(observed) for Western (dashed line), Central (dot–dashed) and Eastern (solid) populations of black ducks.

*Figs. 7, 8. Sesgos en las predicciones de población a partir del modelo no integrado (7) y a partir del modelo integrado sin factores de sesgo (*modelo 00*), con relación a las estimaciones de población actuales para la producción con competencia de ánades reales (A) y sin competencia (B) a lo largo de 11 años (1991–2001). Las representaciones gráficas indican diferencias logarítmicas (previstas) – logarítmicas (observadas) para las poblaciones de ánades sombríos del estrato geográfico occidental (línea discontinua), central (línea discontinua punteada) y oriental (línea continua).*

Model parameter estimates that were informed by the consequent population prediction resulted in less biased population estimates relative to their non–integrated counterparts, particularly when reproduction bias parameters were specified. Remaining bias showed spatial and temporal patterns of heterogeneity. In particular, estimates were more positively biased in the first half of the time series; where negative biases occurred, they tended to be in the second half. Again, this trend may be due to inconsistencies in survey methodologies over this time period. Spatially, more over–prediction occurred in the Western and Central populations, balanced by relative unbiasedness or even under–prediction in the East. This may also be generally related to survey problems in these areas, because the West and Central regions are characterized by large, unsurveyed areas in Ontario and Quebec, in contrast to the relatively smaller, well–surveyed Atlantic provinces in the East. The quality of predictions among integrated models were not obviously different according to which vital rate bias parameters were estimated.

Each of the integrated models not only represents a different explanation for the source of over–prediction, but also for the factors influencing production in general. Having quantified structural and biological uncertainty, future development could link weighted predictions across models to a dynamic decision optimization procedure, thereby providing a complete decision analysis system to inform management. Though not currently incorporated into the modeling framework presented here, methods of dynamic optimization exist which complement the integrated, stochastic simulation approach outlined thus far (e.g. reinforcement learning).

Figs. 9–11. Bias in population predictions from integrated model with reproduction bias (*Model 10*) (9), from integrated model with survival bias (*Model 01*) (10), and from integrated model with both reproduction and survival bias (*Model 11*) (11), relative to actual population estimates for production with mallard competition (A) and without (B) over 11 years (1991–2001). Plots indicate differences log (predicted) – log (observed) for Western (dashed line), Central (dot–dashed) and Eastern (solid) populations of black ducks.

*Figs. 9–11. Sesgos en las predicciones de población a partir del modelo integrado con sesgos de reproducción (*Modelo 10) (9), a partir del modelo integrado con sesgos de supervivencia (*Modelo 01) (10) y a partir del modelo integrado con sesgos de reproducción y sesgos de supervivencia (*Modelo 11) (11), con relación a las estimaciones de población actuales para la producción con competencia de ánades reales (A) y sin competencia (B) a lo largo de 11 años (1991–2001). Las representaciones gráficas indican diferencias logarítmicas (previstas) – logarítmicas (observadas) para las poblaciones de ánades sombríos del estrato geográfico occidental (línea discontinua), central (línea discontínua punteada) y oriental (línea continua).*

Table 1. Mean AIC values, along with AIC values and associated model weights of 10 competing models for black duck population dynamics, based on the final 50,000 of 100,000 total MCMC iterations.

Tabla 1. Valores medios de AIC junto con valores de AIC y pesos de modelos asociados de 10 modelos alternativos para la dinámica poblacional del ánade sombrío, a partir de las 50.000 de un total de 100.000 iteraciones MCMC finales.

Model	AIC	AIC	Weight
10c	25827.94	0.00	0.519
10	25828.59	0.65	0.375
11	25831.61	3.67	0.082
00	25835.88	7.94	0.010
00c	25836.61	8.67	0.007
01c	25838.70	10.76	0.002
01	25838.77	10.83	0.002
11c	25839.86	11.92	0.001
Null	25877.29	49.35	0.000
Nullc	25884.37	56.43	0.000

Table 2. Production and survival bias parameter estimates (log scale) for each of eight integrated model scenario combinations (95% Bayesian credible intervals in parentheses). Zero values indicate no estimate for given model scenario.

Tabla 2. Estimaciones de parámetros de producción y sesgos de supervivencia (escala logarítmica) para cada una de las ocho combinaciones de modelos integrados (ente paréntesis los intervalos bayesianos creíbles al 95%). Los valores cero no indican ninguna estimación para el modelo dado.

Model	Production Bias	Survival Bias
00,00c	0	0
10	0.701(0.641, 0.762)	0
10c	0.696(0.638, 0.750)	0
01	0	−0.085(−0.478, 0.250)
01c	0	−0.121(−0.473, 0.186)
11	0.679(0.573, 0.754)	0.051(−0.128, 0.251)
11c	−0.100(−0.211, 0.027)	0.241(−0.298, 0.760)

Table 3. Production model parameter estimates for each model scenario (95% Bayesian credible intervals in parentheses). Spatially–explicit parameters listed on multiple lines for each model, where appropriate: West (top), Central (middle), East (bottom).

Tabla 3. Estimaciones de parámetros del modelo de producción para cada modelo (entre paréntesis los intervalos bayesianos creíbles al 95%). Los parámetros espacialmente explícitos se detallan en líneas múltiples para cada modelo: estrato geográfico occidental (línea superior), estrato geográfico central (linea media), estrato geográfico oriental (línea inferior).

Model	Intercept	Black Duck effect	Mallard effect
00	1.872 (1.402,2.316)	−0.501 (−0.724,−0.256)	0
	1.473 (1.060,1.870)	−0.617 (−0.957,−0.275)	
	1.627 (1.376,1.871)	−0.877 (−0.993,−0.759)	
01	1.843 (1.379,2.299)	−0.478 (−0.713,−0.238)	0
	1.478 (1.050,1.888)	−0.608 (−0.952,−0.248)	
	1.640 (1.395,1.886)	−0.879 (−0.987,−0.766)	
10	1.914 (1.367,2.462)	−0.494 (−0.794,−0.185)	0
	1.284 (0.819,1.828)	−0.700 (−1.074,−0.360)	
	1.893 (1.679,2.095)	−0.928 (−1.022,−0.837)	
11	1.843 (1.272,2.379)	−0.457 (−0.752,−0.142)	0
	1.269 (0.761,1.775)	−0.671 (−1.032,−0.294)	
	1.805 (1.545,2.056)	−0.883 (−0.985,−0.763)	
Null	1.495 (0.955,2.019)	−0.255 (−0.531, 0.038)	0
	1.655 (1.220,2.058)	−0.738 (−1.070,−0.386)	
	1.679 (1.421,1.938)	−0.902 (−1.013,−0.792)	

Table 3. (Cont.)

Model	Intercept	Black Duck effect	Mallard effect
00c	2.727 (2.391,3.063)	−0.794 (−0.887,−0.692)	−0.265 (−0.482, −0.048)
	1.705 (1.471,1.934)		
	1.493 (1.260,1.719)		
01c	2.710 (2.351,3.072)	−0.787 (−0.884,−0.693)	−0.253 (−0.478, −0.033)
	1.704 (1.470,1.921)		
	1.482 (1.260,1.704)		
10c	2.820 (2.452,3.176)	−0.837 (−0.935,−0.721)	−0.283 (−0.503, −0.036)
	1.541 (1.338,1.748)		
	1.693 (1.466,1.911)		
11c	2.757 (2.358,3.141)	−0.800 (−0.898,−0.701)	−0.275 (−0.519, −0.025)
	1.716 (1.477,1.947)		
	1.523 (1.292,1.759)		
Nullc	2.729 (2.325,3.116)	−0.805 (−0.904,−0.707)	−0.220 (−0.503,0.065)
	1.736 (1.511,1.969)		
	1.507 (1.277,1.731)		

Though incomplete, we have presented an integrated framework for modeling population dynamics. The feedback between predictions and parameter estimates achieved by sampling from a full joint posterior via Markov chain Monte Carlo results in vital rate estimates that are better predictors of population change. This holistic approach is a more efficient use of all available information, relative to standard modeling procedures that estimate parameters and project population states in serial. The availability of complementary procedures for dynamic decision analysis hold promise for the development of a truly integrated natural resource decision–making tool.

References

Akaike, H., 1973. Information theory as an extension of the maximum likelihood principle. In: *Second International Symposium on Information Theory:* 267–281 (B. N. Petrov & F. Csaki, Eds.). Akademiai Kiado, Budapest.

Bordage, D., 2000. Black duck joint venture helicopter survey – Québec. Technical report, Canadian Wildlife Service, Québec Region, Environment Canada, Sainte–Foy, QC.

Burnham, K. P. & Anderson, D. R.. 2002. Model Selection and Multi–Model Inference: A Practical, Information–theoretic Approach. Springer, New York.

Conroy, M. J., Miller, M. W. & Hines, J. E., 2002. Identification and synthetic modeling of factors affecting American black duck populations. *Journal of Wildlife Management* 6 (Wildlife Monograph No.150).

Gamerman, D., 1997. *Markov Chain Monte Carlo: statistical simulation for Bayesian inference*. Chapman and Hall, London, first edition.

Hastings, W. K., 1970. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57: 97–109.

King, R. & Brooks, S. P., 2002. Model Selection for Integrated Recovery/Recapture Data. *Biometrics,* 58: 841–851.

Lubow, B. C., 1995. SDP: Generalized software for solving stochastic dynamic optimization problems. *Wildlife Society Bulletin*, 23: 738–742.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E., 1953. Equations of state calculations by fast computing machine. *J. Chem. Phys.* 21: 1087–1091.

Robert, C. P. & Casella, G., 1999. *Monte Carlo statistical methods*. Springer–Verlag, New York.

Wikle, C. K., 2003. Hierarchical Bayesian models for predicting the spread of ecological processes. *Ecology,* 84: 1382–1394.

Williams, B. K., Nichols, J. D., & Conroy, M. J., 2002. *The Analysis and Management of Animal Populations*. Academic Press, San Diego, CA.

Zimpfer, N. L., 2004. *Estimating movement and production rates in American black ducks*. Master's thesis, Univ. of Georgia, Athens.