

Ecological Applications, 16(1), 2006, pp. 74–86
© 2006 by the Ecological Society of America

HIDDEN PROCESS MODELS FOR ANIMAL POPULATION DYNAMICS

K. B. NEWMAN,^{1,4} S. T. BUCKLAND,¹ S. T. LINDLEY,² L. THOMAS,¹ AND C. FERNÁNDEZ³

¹*School of Mathematics and Statistics, University of St. Andrews, St. Andrews, Scotland*

²*National Marine Fisheries Service, Santa Cruz Laboratory, 110 Shaffer Road, Santa Cruz, California 95060 USA*

³*Department of Mathematics and Statistics, Fylde College, Lancaster University, Lancaster, UK*

Abstract. Hidden process models are a conceptually useful and practical way to simultaneously account for process variation in animal population dynamics and measurement errors in observations and estimates made on the population. Process variation, which can be both demographic and environmental, is modeled by linking a series of stochastic and deterministic subprocesses that characterize processes such as birth, survival, maturation, and movement. Observations of the population can be modeled as functions of true abundance with realistic probability distributions to describe observation or estimation error. Computer-intensive procedures, such as sequential Monte Carlo methods or Markov chain Monte Carlo, condition on the observed data to yield estimates of both the underlying true population abundances and the unknown population dynamics parameters. Formulation and fitting of a hidden process model are demonstrated for Sacramento River winter-run chinook salmon (*Oncorhynchus tshawytscha*).

Key words: *endangered species; kernel smoothing; sequential importance sampling; state-space models.*

INTRODUCTION

Fish and wildlife scientists and managers collect data on an animal population and build mathematical models based on those data to answer questions about the population. The questions include inquiries about the current state of the population, the future state of the population, the nature of the dynamics of the population and the influence of environmental factors, and the likely consequences of various management actions. Some examples are: How many juvenile winter-run chinook salmon (*Oncorhynchus tshawytscha*) are currently present in the Sacramento River system? Given estimates of the adult returns of these salmon this year, how many juveniles will be produced next year? How does precipitation during the spawning season influence egg survival? If water export pumps near the mouth of the Sacramento River are shut down during the period of smolt outmigration, how much will the juvenile survival rate increase?

Models for population dynamics necessarily are time indexed and take the form either of update equations for discrete time models or differential equations for continuous time models (Gurney and Nisbet 1998, Gotelli 2001). Similarly, measurements made on the population are often collected over time and are thus a time series of observations. To estimate the parameters of

the population models, three approaches have been taken (Hilborn and Walters 1992: Section 6.13, Schnute 1994, Quinn and Deriso 1999: Section 5.4). Process error models allow for stochasticity in the evolution over time of the animal population, but observations are treated as if they were made without error, and possible temporal and/or spatial dependence of the observations is ignored. Measurement error models treat the dynamics of the animal population as deterministic but recognize errors in the observations. Models that only account for one source of error frequently underestimate the real uncertainty about population sizes or processes, hence the subject of this paper is the third approach, where both process and measurement errors are included.

Models for animal population dynamics, particularly of fish populations, that simultaneously allow for process and measurement errors have for computational convenience assumed normal (Gaussian) distributions for both errors and a linear formulation for the linkage of past and present abundances and the linkage of observations to abundances (Mendelssohn 1988, Sullivan 1992, Speed 1993, Schnute 1994, Newman 1998). Parameters were estimated by the maximum likelihood method with the Kalman filter (Kalman 1960) used to calculate the likelihood. Such models have been labeled normal dynamic linear models (West and Harrison 1989) and are a special case of state-space models.

The assumptions of normality and linearity are quite restrictive and often fail to accurately describe the stochastic nature of the population dynamics and the ob-

Manuscript received 31 March 2004; revised 30 June 2004; accepted 26 August 2004. Corresponding Editor: D. S. Schimel. For reprints of this Invited Feature, see footnote 1, p. 3.

⁴ E-mail: ken@mcs.st-and.ac.uk

servations. Approximations (e.g., the extended Kalman filter) and transformations of more realistic stochastic models have been done to attempt to shoehorn models into normality and linearity (Newman 1998, Besbeas et al. 2002, Lindley 2003). For example, Newman (1998) started with underlying multinomial and binomial processes, and then used the corresponding expected values and variances in a normal linear model.

Unfortunately, it is often not possible to linearize or normalize realistic population dynamics models without substantially altering the behavior of the models. Developments in computer-intensive statistical estimation procedures, such as sequential Monte Carlo methods (Liu and Chen 1998, Liu and West 2001) and Markov chain Monte Carlo (Gilks et al. 1996), hereafter MCMC, combined with powerful desktop computers, are making feasible the fitting of more realistic models for the animal population dynamics and associated observations. The freely available WinBUGS software (*available online*)⁵ is designed to make implementation of MCMC simpler for applied statisticians and has been used to fit simple one-dimensional animal population dynamics models with observation errors (Meyer and Millar 1999). Recent applications to animal populations with process and measurement errors that assume neither normality nor linearity include models for South Atlantic albacore (*Thunnus alalunga*) (Meyer and Millar 1999), red deer (*Cervus elaphus*) (Trenkel et al. 2000), and grey seals (*Halichoerus grypus*) (Thomas et al. 2005).

The purpose of this paper is to introduce ecologists to a structure for simultaneously modeling the dynamics of animal populations while allowing for both process and measurement errors, models that we label hidden process models (HPMs). The style of this paper is intended to be tutorial and expository. We present a particular procedure for fitting such models, a variation on sequential importance sampling, and demonstrate it with an application to winter-run chinook salmon.

HIDDEN PROCESS MODELS

A HPM is a model for two time series running in parallel. One time series describes the underlying (but unknown) true situation and the other time series consists of measurements on or estimates of this truth. We will here define HPMs in the context of an animal population.

At a given point in time, an animal population can be described by the number of animals falling into particular categories. We denote the vector of abundances by category at time t by \mathbf{n}_t , $t = 1, 2, \dots, T$. For example, with a red deer population, \mathbf{n}_t could be the number of young, the number of breeding adult females, the number of nonbreeding adult females, and

⁵ <http://www.mrc-bsu.cam.ac.uk/bugs/winbugs/contents.shtml>

the number of adult males at time t . We refer to this vector as the state vector and denote the entire time series by $\mathbf{n}_{1:T} = \mathbf{n}_1, \dots, \mathbf{n}_T$, and more generally, $\mathbf{n}_{1:t} = \mathbf{n}_1, \dots, \mathbf{n}_t$. Note that the equally spaced indexing is just for convenience and to reduce notation, i.e., $\mathbf{n}_{t_1}, \mathbf{n}_{t_2}, \dots, \mathbf{n}_{t_r}$, where $t_1 < t_2 < \dots < t_r$, is possible.

Assuming measurements or estimates of the state vector are made at time t , there is an associated vector that we call the observation vector and denote by \mathbf{y}_t . Again, denote the time series from 1 to T by $\mathbf{y}_{1:T}$. Note that the state and observation vectors need not have the same number of components, i.e., not all states need be measured or separately observed. For example, in a census, one may not observe the sex of the animals even though, in the state vector, these are distinguished by sex. In this case, the dimension of the observation vector is lower than that of the state vector.

A hidden process model describes both the evolution of the state vector over time and the link between the state and observation vectors. The evolution of \mathbf{n}_t is the model for population dynamics and can include both demographic and environmental stochasticity (“process variation”). This evolution we call the state process and it can consist of several subprocesses.

For example, with salmon, suppose that abundance is measured immediately following a period of harvest. The sequence of subprocesses is assumed to be natural mortality, movement, and harvest. A presentation of this evolution that is familiar to ecologists is the Leslie matrix (Gotelli 2001: Chapter 3), but here we assume that such a matrix is only an approximation of the expected changes with time:

$$E[\mathbf{n}_t] \approx \mathbf{A}_t \mathbf{n}_{t-1} = (\mathbf{I} - \mathbf{H}_t) \mathbf{M}_t \mathbf{S}_t \mathbf{n}_{t-1}$$

where \mathbf{S}_t , \mathbf{M}_t , and \mathbf{H}_t are matrices representing subprocesses of survival, movement, and harvest, respectively. The matrix representation will be an approximation in many cases, such as when subprocesses are density dependent, as this will result in $E[\mathbf{n}_t]$ not being a linear function of \mathbf{n}_{t-1} . The matrix representation is largely for conceptual convenience and does not necessarily play a role in the actual fitting of models to data, thus the quality of the approximation is relatively unimportant. We call \mathbf{A}_t a generalized Leslie matrix (Buckland et al. 2004). The matrix representation of $E[\mathbf{n}_t]$ is made explicit (three regions are shown just as a simplifying example):

$$E \begin{pmatrix} n_{1t} \\ n_{2t} \\ n_{3t} \end{pmatrix} \approx \begin{pmatrix} 1 - h_{1t} & 0 & 0 \\ 0 & 1 - h_{2t} & 0 \\ 0 & 0 & 1 - h_{3t} \end{pmatrix} \begin{pmatrix} p_{1 \rightarrow 1,t} & p_{2 \rightarrow 1,t} & p_{3 \rightarrow 1,t} \\ p_{1 \rightarrow 2,t} & p_{2 \rightarrow 2,t} & p_{3 \rightarrow 2,t} \\ p_{1 \rightarrow 3,t} & p_{2 \rightarrow 3,t} & p_{3 \rightarrow 3,t} \end{pmatrix} \times \begin{pmatrix} \varphi_{1t} & 0 & 0 \\ 0 & \varphi_{2t} & 0 \\ 0 & 0 & \varphi_{3t} \end{pmatrix} \begin{pmatrix} n_{1,t-1} \\ n_{2,t-1} \\ n_{3,t-1} \end{pmatrix}$$

where, from time period $t - 1$ to period t , φ_{it} is the

expected survival rate for area i , $p_{i \rightarrow j,t}$ is the probability of a fish in area i moving to area j , and h_{ii} is the expected harvest rate in area i . Thus, the next-to-last matrix of φ_{it} s is the survival matrix, S_t ; the immediately preceding matrix of $p_{i \rightarrow j,t}$ is the movement matrix that redistributes the survivors; and the first matrix is the identity matrix minus a diagonal matrix of harvest rates, i.e., it generates the expected remaining fish in each region following harvest.

We believe that the generalized Leslie matrix perspective on the state process evolution is conceptually useful because (1) it is familiar to ecologists, (2) one can easily create variations by simply changing the ordering of the matrices for the subprocesses, and (3) such modularization into subprocesses facilitates the formulation of alternative hypotheses; e.g., density-independent vs. density-dependent survival. At the same time, we emphasize that such a matrix perspective need only be an approximation for more realistic descriptions, i.e., nonlinear functions for state processes. For implementation, one should think of the evolution of the state process through its subprocesses in terms of a linked set of probability density or mass functions, which we generically label pdfs.

Probability models that link the observations, \mathbf{y}_t , to the state vector, \mathbf{n}_t , can be extremely flexible, too. They can reflect a wide variety of sampling and estimation procedures. For example, \mathbf{y}_t can include (1) catch per unit effort data, (2) mark-recapture or band-recovery data, or (3) attempted censuses, line, or point transect surveys. Alternatively, the components of the observation vector can be direct estimates of the state vector components rather than “raw” data. For example, raw data in a mark-recapture study can include numbers of marks, sample sizes, and numbers of recaptures over different sampling periods; rather than work directly with these numbers, point estimates of abundance could be viewed as the observations with some associated estimation error. The application to chinook salmon given later works with estimates of fish abundance, for example, instead of the original sample data.

We now formally define a hidden process model. There are three classes of pdfs to specify: (1) one that generates the initial state vector, (2) one that describes the evolution of the state vector from one time period to the next, and (3) one that links the observation vector to the state vector:

$$\text{Initial state pdf: } g_0(\mathbf{n}_0 | \theta) \quad (1)$$

$$\begin{aligned} \text{State } t \text{ pdf: } & g_t(\mathbf{n}_t | \mathbf{n}_{t-1}, \mathbf{n}_{t-2}, \dots, \mathbf{n}_0, \theta) \\ & = g_t(\mathbf{n}_t | \mathbf{n}_{0:t-1}, \theta) \end{aligned} \quad (2)$$

$$\text{Observation } t \text{ pdf: } f_t(\mathbf{y}_t | \mathbf{n}_t, \theta). \quad (3)$$

Note that the special case where \mathbf{n}_t only depends upon \mathbf{n}_{t-1} is called a state-space model. The vector of the parameters of the state and observation processes, θ ,

can include survival probabilities, harvest rates, maturation probabilities, and measures of observation noise. We will consider a Bayesian inference setting and, hence, we will also specify a prior distribution on θ :

$$\text{Prior pdf for } \theta: g_0(\theta). \quad (4)$$

More discussion on Bayesian inference will follow later.

The pdf f_t represents the observation process, stochastic or deterministic. We have implicitly assumed that \mathbf{y}_t given \mathbf{n}_t is independent of all other states and observations (Eq. 3), but this is not strictly necessary. An example would be where \mathbf{y}_t is a weighted combination of current survey data, thus a function of \mathbf{n}_t , and historical estimates, \mathbf{y}_{t-1} , for example.

The pdf g_t generically represents “process variation” due to the mortality, movement, birth, maturation, etc. subprocesses. Assuming that these processes occur sequentially, rather than simultaneously, the pdf g_t can be modeled by a set of linked probability distributions characterizing each of the subprocesses. Let $\mathbf{u}_{r,t}$ denote the state of the population following the occurrence of the r th subprocess. If there are k subprocesses occurring between time periods $t-1$ and t , the evolution of the state vector can be described as

$$\mathbf{u}_{1,t} \sim G_{1,t}(\mathbf{n}_{0:t-1}, \theta)$$

$$\mathbf{u}_{2,t} \sim G_{2,t}(\mathbf{n}_{0:t-1}, \mathbf{u}_{1,t}, \theta)$$

$$\vdots$$

$$\mathbf{n}_t = \mathbf{u}_{k,t} \sim G_{k,t}(\mathbf{n}_{0:t-1}, \mathbf{u}_{1,t}, \dots, \mathbf{u}_{k-1,t}, \theta)$$

where each $G_{r,t}$ is a distribution corresponding to the appropriate pdf; e.g., $\mathbf{u}_{2,t}$ has pdf $g_{2,t}(\mathbf{u}_{2,t} | \mathbf{n}_{0:t-1}, \mathbf{u}_{1,t}, \theta)$. We note that directly evaluating the resulting pdf $g_t(\mathbf{n}_t | \mathbf{n}_{0:t-1}, \theta)$ can be quite complicated because it involves integrating over the intermediate \mathbf{u} 's; i.e.,

$$\begin{aligned} & g_t(\mathbf{n}_t | \mathbf{n}_{0:t-1}, \theta) \\ & = \int_{\mathbf{u}_{1:k-1,t}} \prod_{r=1}^k g_{r,t}(\mathbf{u}_{r,t} | \mathbf{n}_{0:t-1}, \mathbf{u}_{1:r-1,t}, \theta) d\mathbf{u}_{1:k-1,t} \end{aligned}$$

where $\mathbf{u}_{1:r,t} = \mathbf{u}_{1,t}, \dots, \mathbf{u}_{r,t}$. A schematic picture of the linked subprocesses and the observations is given in Fig. 1.

A complete specification of the probability distribution for states, observations, and parameter vector, which includes the intermediate states, is the following:

$$\begin{aligned} & \Pr(\mathbf{n}_0, \mathbf{u}_{1:k,1}, \dots, \mathbf{u}_{1:k,T}, \mathbf{y}_{1:T}, \theta) \\ & = g_0(\mathbf{n}_0 | \theta) g_0(\theta) \\ & \quad \times \prod_{t=1}^T \left[f_t(\mathbf{y}_t | \mathbf{n}_t, \theta) \prod_{r=1}^k g_{r,t}(\mathbf{u}_{r,t} | \mathbf{n}_{0:t-1}, \mathbf{u}_{1:r-1,t}, \theta) \right] \end{aligned}$$

where $\mathbf{n}_t = \mathbf{u}_{k,t}$.

The generality of HPMs for a variety of population dynamics models and observation or sampling proce-

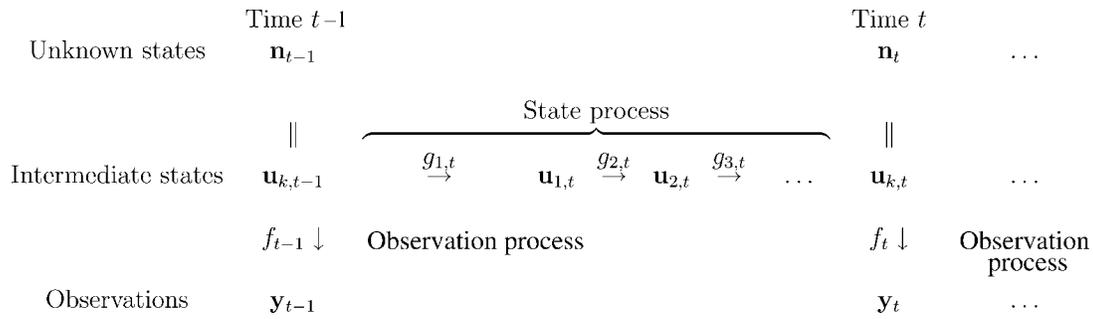


FIG. 1. Schematic picture of state and observation processes with the linked subprocesses for the state’s evolution. State \mathbf{n}_{t-1} evolves to \mathbf{n}_t by intermediate subprocesses, the $\mathbf{u}_{r,t}$ s. The observations \mathbf{y}_{t-1} and \mathbf{y}_t are connected to the corresponding states \mathbf{n}_{t-1} and \mathbf{n}_t by the observation process pdfs f_{t-1} and f_t .

dures was described by Buckland et al. (2004) and some examples are summarized in Table 1.

STATISTICAL INFERENCE

With HPMs, there are two primary inference objectives: to estimate $\mathbf{n}_{1:T}$ and to estimate θ given $\mathbf{y}_{1:T}$. In the case of multiple intermediate subprocesses, one may wish to estimate the $\mathbf{u}_{r,t}$ as well. For realistic and often complex HPMs, computer-intensive procedures are the only ones feasible. In a Bayesian context, we explore the posterior distribution of the states $\mathbf{n}_{1:T}$ and the parameter vector θ via simulation. Two general computational approaches are sequential Monte Carlo procedures and MCMC. Likelihood estimation via MCMC is possible (Geyer 1996), but, pragmatically, the Bayesian solution is simpler in that maximum likelihood involves finding the maxima of what will often be a high-dimensional surface. Conversely, a potential difficulty with, and criticism of, Bayesian methods is the selection of suitable prior distributions for θ and the states.

In this paper, we choose to describe just one of the two general computational approaches to Bayesian inference, sequential Monte Carlo methods. The particular approach that we describe is called sequential importance sampling with resampling and kernel smoothing (SISR/KS). Before describing the SISR/KS algorithm, and in keeping with the tutorial intent of this paper, the next few sections are brief expositions of Bayesian inference and Monte Carlo methods. What may be helpful for the reader to keep in mind is that

the final output from SISR/KS are computer generated samples of the states and the parameters conditional on the observations.

To keep the length of this paper down, we do not explain the other general computational approach to Bayesian inference, MCMC, other than to say that one simulates states and parameters in a Markovian fashion. Given the current values in the chain, a new value (for a state or parameter component) is proposed using a relatively arbitrary probability distribution, called the proposal distribution. The chain moves to the proposed value with a probability that depends on the ratio of the posterior and proposal distributions, where the ratios are evaluated at the proposed value and the current value of the chain. In related current research we are comparing the use of SISR/KS and MCMC for complex state-space models and are in agreement with the following statement from Godsill et al. (2004) on the point of using MCMC in these cases: “. . . constructing an effective MCMC sampler in models with significant degrees of nonlinearity and non-Gaussianity is not always straightforward. Specifically, in these cases it can be hard to construct effective proposal distributions either over collections of states simultaneously or even for single states conditional on all others.”

However, we point out that MCMC has been used successfully for state-space models and also note that SISR/KS has its own drawbacks (particle depletion and choice of the smoothing parameter, both to be discussed). See Meyer and Millar (1999) for an application of MCMC to a relatively simple state-space model

TABLE 1. Examples of HPM components for animal populations (from Buckland et al. 2004).

States (\mathbf{n}_t)	Subprocesses	Observations (\mathbf{y}_t)
Two age classes	survival, aging, birth	census with errors
Five age/stage/sex classes	winter survival, maturation, aging, sexing, summer survival, removal	removal experiment
Two age classes	survival, aging, birth, marking	mark-recapture survey
Four age classes	density-dependent survival and birth, harvest	line transect survey
Two metapopulations	survival, movement, aging, birth	adult female census
Predator and prey abundances	birth, survival	strip transect survey

(with one dimensional state and observation vectors), and Brooks et al. (2004) for multidimensional models.

Finally, we recommend Liu (2001) as an excellent source for general Monte Carlo methodologies including sequential Monte Carlo procedures, MCMC, and inference for state-space models.

Bayesian inference

Here we describe the three general steps in Bayesian inference in the context of HPMs.

Step 1.—The inference procedure begins with a statement of prior knowledge about θ , \mathbf{n}_0 , $\mathbf{n}_{1:T}$, that is quantified by a prior probability distribution, $g(\mathbf{n}_0, \mathbf{n}_{1:T}, \theta)$. This is the product of the state pdfs in Eqs. 1 and 2 and the prior pdf (Eq. 4), as follows:

$$g(\mathbf{n}_0, \mathbf{n}_{1:T}, \theta) = g_0(\mathbf{n}_0 | \theta)g_0(\theta) \times \prod_{t=1}^T g_t(n_t | \mathbf{n}_{1:t-1}, \mathbf{n}_0, \theta).$$

Step 2.—The data $\mathbf{y}_{1:T}$ are viewed as samples from probability distributions that are functions of the states and the parameters. This is simply the observation process model in Eq. 3, which, when viewed as a function of the states and the parameters, is called the likelihood function:

$$L(\mathbf{n}_0, \mathbf{n}_{1:T}, \theta | \mathbf{y}_{1:T}) \equiv \prod_{i=1}^T f_i(\mathbf{y}_i | \mathbf{n}_i, \theta).$$

Step 3.—The prior distribution or knowledge is updated on the basis of the data to yield a posterior distribution for θ , \mathbf{n}_0 , and $\mathbf{n}_{1:T}$, using Bayes' theorem:

$$g(\mathbf{n}_0, \mathbf{n}_{1:T}, \theta | \mathbf{y}_{1:T}) = \frac{g(\mathbf{n}_0, \mathbf{n}_{1:T}, \theta) \times L(\mathbf{n}_0, \mathbf{n}_{1:T}, \theta | \mathbf{y}_{1:T})}{f(\mathbf{y}_{1:T})}. \quad (5)$$

Samples from the posterior distribution provide complete information about the states and parameters. Various summary measures of nearly arbitrary complexity can be readily calculated given such a sample; e.g., means, medians, and other percentiles, and correlations between states and/or parameters.

Historically, the difficulty with Bayesian methods has been that calculation of the posterior distribution is usually a high-dimension integration problem (e.g., calculating $f(\mathbf{y}_{1:T})$ or evaluating appropriate summary measures of this distribution). It has been the relatively recent use of Monte Carlo integration techniques by statisticians that has made Bayesian inference feasible. These methods explore the posterior distribution via simulation and avoid explicit calculation of $f(\mathbf{y}_{1:T})$.

Monte Carlo methods

Here, we sketch some of the general principles of Monte Carlo methods for doing integration and for generating samples from distributions. In the next section, we give the algorithm for the particular Monte Carlo method that we have used for HPMs. Let X be a random

variable with pdf $p(x)$. To begin, note that the expected value of a function of X , $E[\varphi(X)]$, is by definition an integral:

$$E[\varphi(X)] = \int \varphi(x)p(x) dx. \quad (6)$$

The methods we now describe are procedures for calculating such integrals by computer simulation and for generating samples from the pdf $p(x)$. We refer to $p(x)$ as the target pdf. Note that, in the case of Bayesian inference, $p(x)$ is the posterior distribution; e.g., $p(x) \equiv g(\mathbf{n}_0, \mathbf{n}_{1:T}, \theta | \mathbf{y}_{1:T})$.

1. *Simple Monte Carlo integration.*—When one can directly generate samples from the target pdf $p(x)$, one can do simple Monte Carlo integration. Denote such a sample by $x^{*1}, x^{*2}, \dots, x^{*N}$, where N is relatively large; e.g., $N = 10\,000$. The integral (expectation) in Eq. 6 is estimated as follows:

$$\hat{E}[\varphi(X)] = \frac{1}{N} \sum_{i=1}^N \varphi(x^{*i}).$$

This is known as simple or exact Monte Carlo integration.

2. *Probabilities \equiv integrals.*—Determining the probability of X falling within a specific range of values is an integration problem and can thus be estimated by Monte Carlo integration. For example, given the sample $x^{*1}, x^{*2}, \dots, x^{*N}$, $\Pr(a \leq X \leq b)$ is estimated by

$$\hat{\Pr}(a \leq X \leq b) = \frac{1}{N} \sum_{i=1}^N I(a \leq x^{*i} \leq b).$$

$I()$ is an indicator function equaling 1 when $a \leq x^{*i} \leq b$ is true and 0 otherwise. Thus the estimated probability is simply the fraction of times the simulated value falls between a and b .

3. *Importance sampling.*—Direct sampling from the target pdf $p(x)$ is often not feasible. Instead one generates the sample $x^{*1}, x^{*2}, \dots, x^{*N}$ from an alternative, feasible pdf $q(x)$, sometimes called the trial pdf. The generated sample values must then be adjusted or weighted properly to account for the x^* s coming from $q(x)$ instead of $p(x)$. The weight $w_1(x^*)$ for a particular simulated value x^* is the ratio $p(x^*)/q(x^*)$. Estimates of the integral in Eq. 6 are then

$$\hat{E}[\varphi(X)] = \frac{1}{N} \sum_{i=1}^N \varphi(x^{*i}) \frac{p(x^{*i})}{q(x^{*i})} = \frac{1}{N} \sum_{i=1}^N \varphi(x^{*i}) w_1(x^{*i}).$$

4. *Importance sampling for Bayesian problems.*—Besides dealing with situations where direct sampling from $p(x)$ is not feasible, a particular version of importance sampling can be used when evaluating $p(x)$, and thus the weight $w_1(x)$, is difficult. The difficulty usually arises when $p(x) = h(x)c$, where $h(x)$ is tractable, but the constant c is not. This is often the case with Bayesian inference where $p(x)$ is the posterior distribution and the intractable constant is the denom-

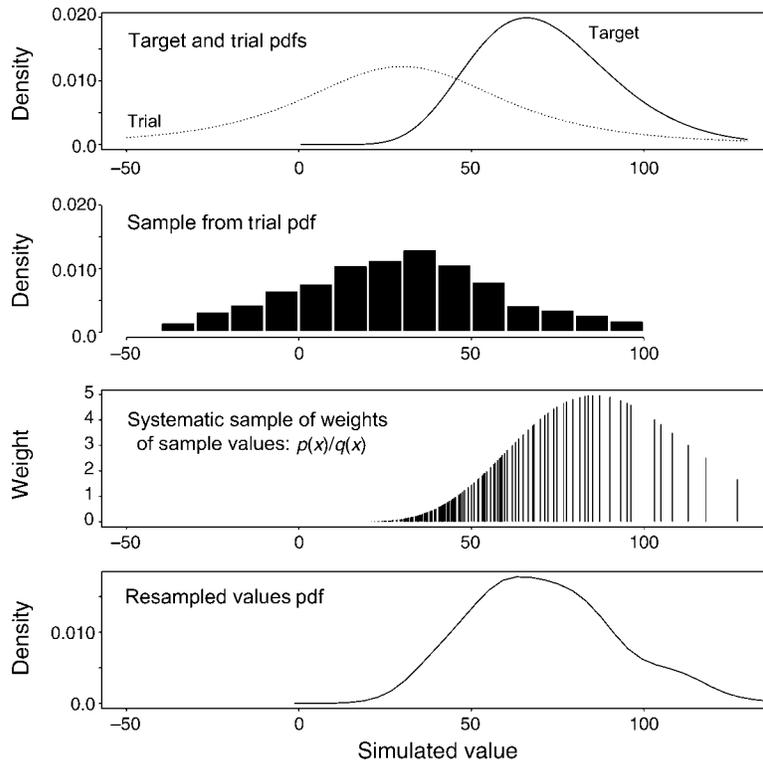


FIG. 2. Demonstration of importance sampling. The top panel shows the target pdf (solid line) and the trial pdf (dotted line). The second panel shows a histogram of a sample of $N = 1000$ from the trial pdf. The third panel graphically displays a systematic sample of the weights attached to each simulated value. The bottom panel is estimated pdf based on the importance sampler.

inator in Bayes theorem, $f(\mathbf{y}_{1:T})$ (see Eq. 5). In this case, the weights are calculated in a different manner so that the constant c cancels out and need not be calculated at all:

$$\hat{E}[\varphi(X)] = \sum_{i=1}^N \varphi(x^{*i}) w_2(x^{*i})$$

where

$$w_2(x^{*i}) = \frac{c \frac{h(x^{*i})}{q(x^{*i})}}{\sum_{j=1}^N c \frac{h(x^{*j})}{q(x^{*j})}} = \frac{h(x^{*i})}{\sum_{j=1}^N \frac{h(x^{*j})}{q(x^{*j})}}$$

5. *Bootstrap resampling.*—When importance sampling is done using a trial pdf $q(x)$, the resulting sample can be resampled so that it becomes a sample from the target pdf $p(x)$. The procedure is called bootstrap resampling or sampling importance resampling (SIR) and simply involves resampling the x^* s according to the weights $w_2(x^*)$ s. There is an analogy here with the concept of natural selection in biology—the “fittest” x^* s, those with the highest weights, tend to survive and multiply, while the “unfit” x^* s are not resampled and become “extinct.” Fig. 2 is a picture of the process of importance sampling combined with resampling.

6. *For HPMs: Sequential importance sampling with resampling (SISR).*—Now we consider the particular problem of making inferences about the unknown states, \mathbf{n}_t , of an HPM, and to simplify discussion, we assume the parameters, θ , are known. In the next section, we show what we do when θ and \mathbf{n}_t are unknown. For an HPM, generation of the unknown states is carried out sequentially by combining importance sampling with bootstrap resampling for time $t = 1$, then time $t = 2$, up to time T . In the simplest form of SISR, at each point in time, the trial pdf is the state process distribution for \mathbf{n}_t conditional on the states in the previous time periods. In other words, $q(\mathbf{n}_t) = g_t(\mathbf{n}_t | \mathbf{n}_{1:t-1}, \theta)$, with g_t as specified in Eq. 2, so that the generated value is a projected value based on the time series state process. The weight $w_2(\mathbf{n}_t)$ is proportional to (state pdf \times observation pdf)/state pdf, and by cancellation of terms in the numerator and denominator, $w_2(\mathbf{n}_t)$ is simply proportional to the observation pdf, $f_t(\mathbf{y}_t | \mathbf{n}_t, \theta)$. Thus, the resampling results in selecting \mathbf{n}_t s that are the “fittest” values according to the observations. For each value of t , this procedure provides a sample from the posterior distribution of \mathbf{n}_t given $\mathbf{y}_{1:t}$. Inference on \mathbf{n}_t based on data just up to time t , and not beyond, is called filtering.

Algorithm: SISR/KS

Sequential importance sampling with resampling is technically feasible with unknown parameters, but it is generally impractical due to a problem known as particle depletion of parameter values. In contrast to simulated states that are resampled just once (when doing filtering), the initial set of N simulated parameter values gets progressively resampled and subsequently thinned and the number of unique remaining values can be quite small. Here, we describe a modification of SISR, sequential importance sampling-resampling with kernel smoothing of the parameters (SISR/KS), which can be used when both the parameters and states are unknown. The kernel smoothing component to the algorithm aims to mitigate the problem of particle depletion of simulated parameter values.

To initialize the algorithm, generate N sets of \mathbf{n}_0 and θ from $g_0(\mathbf{n}_0 | \theta)$ and $g_0(\theta)$; call these draws \mathbf{n}_0^{*i} and θ^{*i} , where $i = 1, \dots, N$. Set $t = 1$.

Step 1.—Generate a sample of size N of \mathbf{n}_t from a trial pdf, $q(\mathbf{n}_t)$, with the sample values denoted \mathbf{n}_t^{*i} , $i = 1, 2, \dots, N$.

Step 2.—Calculate N weights that are the ratio of the product of the likelihood of the data and the state pdf to the trial pdf:

$$w_t^{*i} \propto \frac{f_t(\mathbf{y}_t | \mathbf{n}_t^{*i}, \theta^{*i}) g_t(\mathbf{n}_t^{*i} | \mathbf{n}_{1:t-1}^{*i}, \mathbf{n}_0^{*i}, \theta^{*i})}{q(\mathbf{n}_t^{*i})}$$

where, again, $f_t(\cdot)$ is the likelihood, $g_t(\cdot)$ is the state pdf, and $q(\cdot)$ is the trial pdf.

Step 3.—Resample both \mathbf{n}_t^{*i} and θ^{*i} according to w_t^{*i} to yield a sample denoted \mathbf{n}_t^{**i} , θ^{**i} , $i = 1, \dots, N$.

Step 4.—For each $i = 1, \dots, N$, kernel smooth the parameters. The resulting smoothed parameters, denoted θ^{***} , are weighted combinations of the original values (θ^{**}) and simulated values from a multivariate normal distribution:

$$\theta^{***i} = \lambda \theta^{**i} + (1 - \lambda) Z_i$$

where $0 \leq \lambda \leq 1$ and $Z_i \sim \text{Multivariate Normal}(\bar{\theta}^{**}, \Sigma_{\theta^{**}})$. Z_i is a vector of parameter values simulated from a multivariate normal distribution with a mean vector $\bar{\theta}^{**}$ where the j th component of the vector is the sample average of the N values of the j th parameter, i.e., the average over i of all the θ_j^{**i} . Similarly, the covariance matrix for the multivariate normal, $\Sigma_{\theta^{**}}$, is a matrix of variances and covariances of the vectors θ^{**i} . The constant λ is a tuning parameter that controls the degree of smoothing. When $\lambda = 1$, there is no smoothing and $\theta^{***i} = \theta^{**i}$. When $\lambda = 0$, there is maximal smoothing in the sense that $\theta^{***i} = Z_i$ and the influence of θ^{**i} appears only in its contribution to $\bar{\theta}^{**}$ and $\Sigma_{\theta^{**}}$.

Step 5.—For $i = 1, \dots, N$, redefine $\mathbf{n}_t^{*i} = \mathbf{n}_t^{**i}$ and $\theta^{*i} = \theta^{***i}$.

Step 6.—If $t < T$, increment t to $t + 1$ and go to Step 1 and repeat.

After t time steps, the set of draws \mathbf{n}_t^{*i} , θ^{*i} , $i = 1, \dots, N$ is an approximate sample from the posterior distribution $g(\mathbf{n}_t, \theta | \mathbf{y}_{1:t})$.

The magnitude of the initial sample size N is critical to the consistency of the resulting posterior samples. The number of unique values (called particles by SIS practitioners) is successively reduced at each point in time as simulated values which are less likely than others, given the observations, are successively removed in the resampling process. As mentioned previously, this phenomenon is known as particle depletion, and, if severe enough, can lead to quite different results for different simulation runs, i.e., high Monte Carlo variation. The purpose of the kernel smoothing is to lessen this Monte Carlo variation by introducing new parameter values close to those already present in the set of particles (think “mutation” in a Darwinian analogy). Kernel smoothing of the states is not required here because these are generated anew at each time period by the stochastic state process. One must tune the degree of smoothing (λ in Step 4): lots of smoothing reduces the Monte Carlo variation, but too much introduces substantial bias in some estimates. The optimal amount of smoothing is an area of research.

There are many refinements to SISR/KS. The solution that we actually use for the winter-run chinook salmon is one such variation called auxiliary particle filtering with kernel smoothing of parameters (Liu and West 2001). The auxiliary particle filter (Pitt and Shephard 1999) is a refinement of SIS that further reduces particle depletion, and thus lessens Monte Carlo variation. There are additional refinements (residual resampling, tempering the likelihood) that are described in papers in Doucet et al. (2001). Example code is sketched in Thomas et al. (2005).

APPLICATION TO WINTER-RUN CHINOOK SALMON

Sacramento River winter-run chinook salmon have suffered drastic population declines since the 1960s (Fig. 3) and was declared an endangered species in 1994. Declines were due to many factors, including the building of dams that prevented salmon from reaching spawning areas, habitat degradation, and overfishing.

A schematic drawing of the salmon life history is shown in Fig. 4. The sequence of subprocesses for a given cohort of juveniles J_t begins with sexing (where p_m is the probability of a fish being male) and survival, with probability $\varphi_{2,t+1}$, to age 2 (when they are in the ocean), then maturation (with sex-specific probabilities ρ_{2m} and ρ_{2f}). Maturing fish return to spawn and then die. For age 2 fish that do not mature, the subprocesses are again survival (with probability φ_3) and then maturation (with probabilities ρ_{3f} and ρ_{3m}). Finally the age 3 fish that do not mature experience ocean mortality (surviving with probability φ_4) and then maturation (this time with certainty).

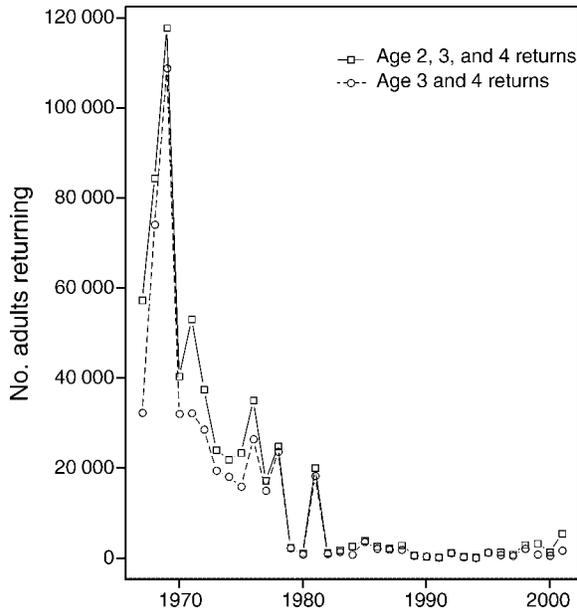


FIG. 3. Estimated adult salmon returns to Red Bluff Diversion Dam for 1967–2001. The solid line and open squares represent age 2, 3, and 4 returns combined; the dotted line and open circles represent age 3 and 4 returns only. Data are from the California Department of Fish and Game.

Two different HPMs were formulated, one called the life history (LH) model and the other the Hallock and Fisher (HF) model. The LH model is based closely on the salmon’s life history. The state vector has 11 components

$$\mathbf{n}_t = (J_t, O_{2mt}, S_{2mt}, O_{2ft}, S_{2ft}, O_{3mt}, S_{3mt}, O_{3ft}, S_{3ft}, S_{4mt}, S_{4ft})$$

where J_t are the number of juveniles, S_{ast} are the number of age a , sex s spawners, and O_{ast} are the number of ocean dwelling, immature fish of age a and sex s . The observation vector has only three components, estimates of juvenile outmigrants, age 2 spawners, and

ages 3 and 4 spawners combined, without sex distinctions.

The HF model is based on work by Hallock and Fisher (1985) and the state vector is more tailored to the currently available data. Hallock and Fisher (1985) analyzed returns of fish that had been tagged as juveniles during brood years 1969–1971. They calculated average return-by-age probabilities conditioning on fish that did return and estimated that 25%, 67%, and 8% returned at ages 2, 3, and 4. A fundamental distinction from the LH model is that maturation probabilities are the same for both sexes; hence, $\rho_{2m} = \rho_{2f} = \rho_2$, and $\rho_{3m} = \rho_{3f} = \rho_3$. The state vector for the HF model lacks sex distinctions and ignores abundances of immature fish that remain in the ocean and only has the following four components: $\mathbf{n}_t = (J_t, S_{2t}, S_{3t}, S_{4t})$. We denote the return-by-age probabilities by r_2, r_3, r_4 , and note that they can be shown to be lower-dimensional functions of maturation probabilities and ages 3 and 4 survival probabilities. In particular,

$$\begin{aligned} r_2 &\propto \rho_2 \\ r_3 &\propto (1 - \rho_2)\varphi_3\rho_3 \\ r_4 &\propto (1 - \rho_2)\varphi_3(1 - \rho_3)\varphi_4 \end{aligned}$$

where the proportionality constant is determined by the restriction that $r_2 + r_3 + r_4 = 1$.

The observation processes are identical for both models. The observation vector contains spawning escapement estimates and juvenile outmigrant estimates with lognormal observation errors. The equations for the state processes for both models, along with the uncommon observation equations, are given in Table 2. Note that the LH model is a state-space model in that the current state pdf depends only on the values in the previous time period. The HF model is not a state-space model, it is third-order Markov; the pdf for the current state depends on states in the previous three time periods. If desired, a third-order Markov model could also

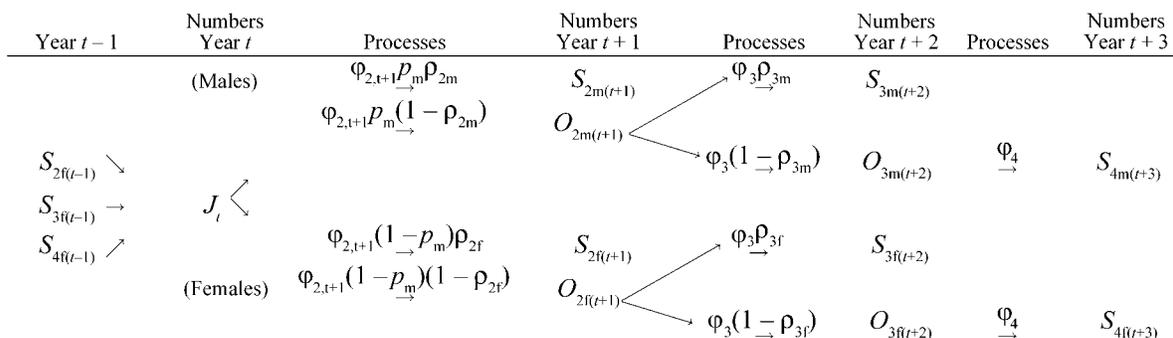


FIG. 4. Schematic of winter-run chinook salmon life history. S denotes returning spawners with subscripting for age, sex, and year of spawning; J are juveniles, while O are immature fish remaining in the ocean. The parameter φ denotes survival (age-specific; survival to age 2 also year-specific), ρ denotes age- and sex-specific maturation probabilities, and p_m is the probability of a fish being male (fixed at 0.5).

TABLE 2. Hidden process model formulations, state, and observation processes, for winter-run chinook salmon.

Model	Equations
Life history (LH) model state process	$J_t \sim [\text{Lognormal}] \left\{ \log \left[\frac{S_{2f+3f+4f,t-1}\alpha}{1 + \beta(S_{2f+3f+4f,t-1})} \right], \log(\text{CV}_{\text{BH}}^2 + 1) \right\}$ $(O_{2mt}, S_{2mt}, O_{2ft}, S_{2ft}) \sim \text{Multinomial}[J_{t-1}, \varphi_{2t}p_m(1 - \rho_{2m}), \varphi_{2t}p_m\rho_{2m}, \varphi_{2t}(1 - p_m)(1 - \rho_{2f}), \varphi_{2t}(1 - p_m)\rho_{2f}]$ $(O_{3mt}, S_{3mt}) \sim \text{Trinomial}[O_{2m,t-1}, \varphi_3(1 - \rho_{3m}), \varphi_3\rho_{3m}]$ $(O_{3ft}, S_{3ft}) \sim \text{Trinomial}[O_{2f,t-1}, \varphi_3(1 - \rho_{3f}), \varphi_3\rho_{3f}]$ $S_{4mt} \sim \text{Binomial}(O_{3m,t-1}, \varphi_4)$ $S_{4ft} \sim \text{Binomial}(O_{3f,t-1}, \varphi_4)$
Hallock and Fisher (HF) model state process	$J_t \sim [\text{Lognormal}] \left[\log \left\{ \frac{p_m S_{2+3+4,t-1} \alpha}{1 + \beta(1 - p_m)(S_{2+3+4,t-1})} \right\}, \log(\text{CV}_{\text{BH}}^2 + 1) \right]$ $S_{2t} \sim \text{Binomial}(J_{t-1}, \varphi_{2t}\rho_2)$ $S_{3t} \sim \text{Binomial} \left(J_{t-2} - S_{2,t-1}, \frac{\varphi_{2,t-1}\rho_2 r_3 / r_2}{1 - \varphi_{2,t-2}\rho_2} \right)$ $S_{4t} \sim \text{Binomial} \left(J_{t-3} - S_{2,t-2} - S_{3,t-1}, \frac{\varphi_{2,t-2}\rho_2 r_4 / r_2}{1 - \varphi_{2,t-2}\rho_2 - \varphi_{2,t-2}\rho_2 r_3 / r_2} \right)$
Observation process	$y_{Jt} \sim \text{Lognormal}[\log(J_t), \log(\text{CV}_J^2 + 1)]$ $y_{S_{2t}} \sim \text{Lognormal}[\log(S_{2t}), \log(\text{CV}_A^2 + 1)]$ $y_{S_{3+4t}} \sim \text{Lognormal}[\log(S_{3t} + S_{4t}), \log(\text{CV}_A^2 + 1)]$

Notes: The notation “[Lognormal]” means a discretized version of a lognormal random variable. For the LH model, $S_{2f+3f+4f,t-1} = S_{2f,t-1} + S_{3f,t-1} + S_{4f,t-1}$. Also, $S_{at} = S_{aft} + S_{omt}$ for $a = 2, 3, 4$ and $S_{2+3+4,t-1} = S_{2,t-1} + S_{3,t-1} + S_{4,t-1}$. For definitions of parameters, see *Application to winter-run chinook salmon*. See Fig. 4 for a description of other notation.

be derived from the LH model by integrating out the O_{ast} analytically.

Pseudocode for fitting the LH model

Before giving some of the actual implementation details for fitting both HPMS, we first present pseudocode, in the spirit of Hilborn and Mangel’s (1997) *Ecological Detective*, for the SISR/KS algorithm as it applies to fitting the LH model.

Generate N sets of \mathbf{n}_0 , which for the LH model is

$$\mathbf{n}_0 = (J_0, O_{2m0}, S_{2m0}, O_{2f0}, S_{2f0}, O_{3m0}, S_{3m0}, O_{3f0}, S_{3f0}, S_{4m0}, S_{4f0}) \quad (7)$$

and N sets of θ (using prior distributions $g_0(\mathbf{n}_0 | \theta)$ and $g_0(\theta)$), where

$$\theta = (\alpha, \beta, \text{CV}_{\text{BH}}, \varphi_{2,1}, \dots, \varphi_{2,T}, \varphi_3, \varphi_4, \rho_{2f}, \rho_{2m}, \rho_{3f}, \rho_{3m}, \text{CV}_J, \text{CV}_A). \quad (8)$$

Call the generated sample $\mathbf{n}_0^{*i}, \theta^{*i}, i = 1, \dots, N$. Set $t = 1$.

Step 1.—Generate a sample of size N of $\mathbf{n}_t = (J_t, \dots, S_{4ft})$, from a trial pdf, $q(\mathbf{n}_t)$, which was chosen to be the state pdf (see the LH model state process equations in Table 2 for the state pdf); e.g., generating J_t from a lognormal form of the Beverton-Holt recruitment model, generating $(O_{2mt}, O_{2ft}, S_{2mt}, S_{2ft})$ from a

multinomial distribution and so on. Call the sample $\mathbf{n}_t^{*i}, i = 1, \dots, N$.

Step 2.—Calculate N weights that are the ratio of the product of the likelihood of the data at time t and the state pdf to the trial pdf. Because the trial pdf was the state pdf, the weights are proportional to the likelihood, which is the product of lognormal pdfs for $y_{Jt}, y_{S_{2t}},$ and $y_{S_{3+4t}}$ (see the observation process equations in Table 2):

$$w_t^{*i} \propto f_t[y_{Jt} | \log(J_t^{*i}), \log(\text{CV}_J^{*i,2} + 1)] \\ \times f_t[y_{S_{2t}} | \log(S_{2t}^{*i}), \log(\text{CV}_A^{*i,2} + 1)] \\ \times f_t[y_{S_{3+4t}} | \log(S_{3t}^{*i} + S_{4t}^{*i}), \log(\text{CV}_A^{*i,2} + 1)]$$

where the $f_t(\cdot)$ s are lognormal density functions.

Step 3.—Resample both \mathbf{n}_t^* and θ^* according to w_t^* to yield a sample denoted $\mathbf{n}_t^{**}, \theta^{**}$.

Step 4.—For each $i = 1, \dots, N$, kernel smooth θ^{**i} to yield θ^{***i} with $\theta^{***i} = \lambda\theta^{**i} + (1 - \lambda)Z_i$, where $Z_i \sim \text{Multivariate Normal}(\theta^{**}, \Sigma_{\theta^{**}})$ with $\theta^{**} = \alpha^{**}, \dots, \text{CV}_A^{**}$ and

$$\Sigma_{\theta^{**}} = \begin{bmatrix} \text{Var}(\alpha^{**}) & \text{Cov}(\alpha^{**}, \beta^{**}) & \dots & \text{Cov}(\alpha^{**}, \text{CV}_A^{**}) \\ \vdots & \vdots & \vdots & \vdots \\ \text{Cov}(\text{CV}_A^{**}, \alpha^{**}) & \text{Cov}(\text{CV}_A^{**}, \beta^{**}) & \dots & \text{Var}(\text{CV}_A^{**}) \end{bmatrix}$$

TABLE 3. Estimates of outmigrating juveniles and adult returns to Redd Bluff Diversion Dam.

Observation year	No. juvenile outmigrants (fry)	No. young returns (age 2)	No. adult returns (age 3 + 4)
1992	NA	80	1160
1993	NA	137	250
1994	NA	124	62
1995	NA	31	1266
1996	1 816 984	629	708
1997	469 183	352	528
1998	2 205 162	923	2079
1999	5 000 416	2466	822
2000	1 366 162	789	563

Note: The data sources are Martin et al. (2001) for juveniles (shown in fry equivalents) and D. Kilham (*personal communication*) for young and adult returns.

Step 5.—For $i = 1, \dots, N$, redefine $\mathbf{n}_i^{*i} = \mathbf{n}_i^{**i}$ and $\theta^{*i} = \theta^{**i}$.

Step 6.—If $t < T$, increment t to $t + 1$ and go to Step 1 and repeat.

Fitting the HPMs

The observation time series was relatively short, five years of juvenile outmigration estimates (1996–2000) and nine years of adult returns estimates (1992–2000; Table 3). Given such a short time series and the number of parameters (20 for the LH formulation and 17 for

the HF formulation; see Table 2), we expect confounding between some of the parameters as well as non-negligible influence of prior distributions on the posterior distributions. Therefore, different prior distributions were used to evaluate the sensitivity to priors.

The actual fitting procedure was an extension of SISR/KS, the auxiliary particle filter with kernel smoothing. The algorithm was written in S-Plus (Insightful Corporation, Seattle, Washington, USA). The input data and S-Plus code are available in the Supplement.

The number of simulated values was set at $N = 400\,000$ and the smoothing parameter λ was 0.9 (relatively moderate smoothing). With $N = 400\,000$ the Monte Carlo variation, as measured by the coefficient of variation for the mean parameter estimates, averaged 1.4% for the LH model and 1.5% for the HF model.

RESULTS

The results of fitting the LH and HF HPMs are samples from the posterior distributions of the states and of the parameters. Fig. 5 plots the filtered means of the state variables from both models against years along with the observed values. When it came to estimating states, both models gave quite similar results. The filtered state values were smoother in time than the observations.

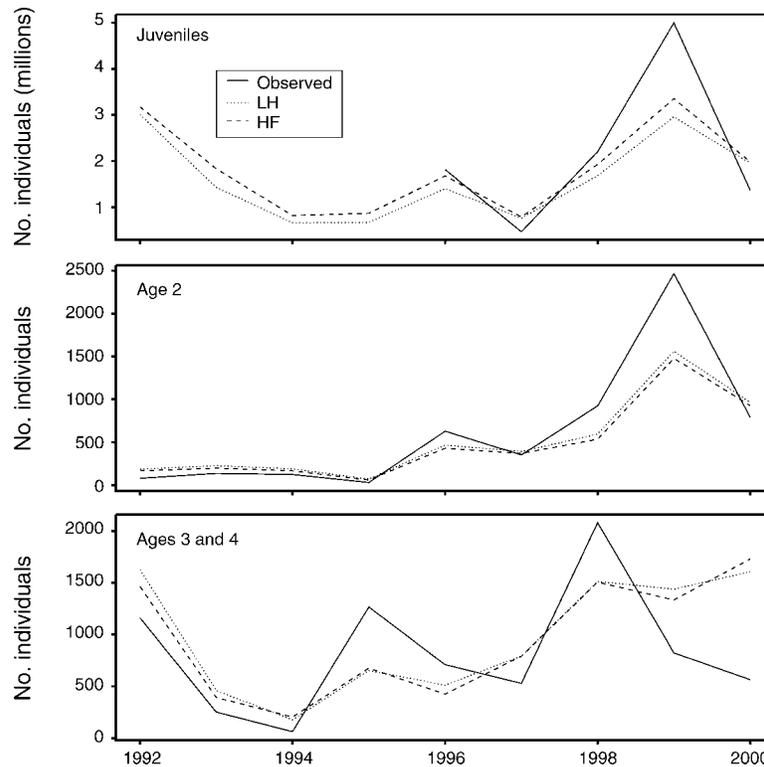


FIG. 5. Comparison of the LH and HF models' mean filtered juveniles, young adults, and older adults from LH and HF models with observed (estimate) values for the years 1992 through 2000. There were no juvenile observations prior to 1996.

TABLE 4. Summary of predictions from LH and HF models. Values are numbers of fish.

Class	Observation	LH prediction	HF prediction
Juveniles	1 366 162	3 021 317	3 180 551
Age 2	789	991	950
Ages 3 + 4	563	3368	2656

To compare the LH and HF models' predictive performance, we made predictions of the numbers of juveniles, age 2 returns, and ages 3 and 4 combined returns for the year 2000 and compared the predictions to the observations. For a Bayesian, prediction of the states in year 2000 (index 00) corresponds to computing the posterior distribution of \mathbf{n}_{00} given all data available up to the previous year, i.e., $g(\mathbf{n}_{00} | \mathbf{y}_{92:99})$. Once this distribution has been computed, one has at one's disposal not only a point estimate of \mathbf{n}_{00} , but also measures of uncertainty and correlation between the components of the state vector, \mathbf{n}_{00} , which accurately reflect one's knowledge and uncertainty at the time the prediction is formulated.

To make the predictions, the first eight years of data were used to generate samples from the posterior distributions for parameters and states. These sample values were then plugged into the state process equations for both the LH and the HF models, and juvenile and spawner abundances were generated from the appropriate lognormal, binomial, and/or trinomial distributions (see Table 2). For the HF model we used the filtered state values for years 6, 7, and 8, whereas, strictly speaking, one should use smoothed state values for years 6, 7, and 8, because the latter reflects all the

information as of 1999. Both models yielded quite similar predictions which we summarize in Table 4. Both had predictions of juveniles and age 3 and age 4 combined returns considerably higher than estimated values.

Table 5 gives the means and standard deviations of the posterior distributions for the parameters. Estimates of in-common parameters were relatively similar for the LH and HF models; e.g., the juvenile recruitment function parameters, α and β , the year-specific age 2 survival probabilities, and the coefficients of variation for juvenile and adult estimates. For both models there was a strong positive correlation between α and β . Both models indicated that estimates of adult returns were more variable than estimates of juvenile outmigrants (as measured by CV_A and CV_J). Using estimates of φ_3 , φ_4 , and the ρ 's, r_2 and r_3 , the HF model's age 2 and age 3 conditional return rates, can also be calculated under the LH model; the posterior means were 0.35 and 0.52 for r_2 and r_3 , somewhat similar to the HF estimates of 0.30 and 0.62.

Relatively strong influence by the priors, or relatedly, relative weak information in the data, is suggested for some of the parameters. The maturation parameters, in particular, had prior and posterior means that were quite similar. This is perhaps not surprising given the lack of sex-specific information about adult returns. The HF model's posterior mean return rates, r_2 and r_3 were relatively similar to the prior means; on the other hand, the priors were data-based using Hallock and Fisher's (1985) tag recovery information.

Sensitivity of the posterior distributions to the priors was also examined by using alternative prior distri-

TABLE 5. Comparison of means for the posterior distributions of some of the parameters for the life history (LH) and Hallock and Fisher (HF) models.

Parameters	Prior distribution	Prior mean	Posterior means	
			LH	HF
α	Gamma(4.68,1068)	5000	4132 (1598)	4224 (1604)
β	Beta(0.988,997)	0.001	0.0010 (0.00091)	0.0009 (0.0024)
CV_{BH}	Beta(4.44,3.30)	0.57	0.54 (0.17)	0.54 (0.18)
$\varphi_{2,97}$	Beta(1.55,308)	0.005	0.0065 (0.0046)	0.0069 (0.0047)
$\varphi_{2,98}$	Beta(1.55,308)	0.005	0.0072 (0.0049)	0.0082 (0.0056)
$\varphi_{2,99}$	Beta(1.55,308)	0.005	0.0063 (0.0044)	0.0063 (0.0055)
$\varphi_{2,00}$	Beta(1.55,308)	0.005	0.0049 (0.0039)	0.0051 (0.0040)
φ_3	Beta(1.00,1.05)	0.49	0.28 (0.15)	NA
φ_4	Beta(1.00,1.05)	0.49	0.40 (0.26)	NA
ρ_{2f}	Beta(1.60,14.5)	0.10	0.087 (0.066)	NA
ρ_{2m}	Beta(1.60,14.5)	0.10	0.091 (0.068)	NA
ρ_{3f}	Beta(3.00,2.00)	0.60	0.62 (0.20)	NA
ρ_{3m}	Beta(3.00,2.00)	0.60	0.55 (0.21)	NA
ρ_2	Beta(1.60,14.5)	0.10	NA	0.070 (0.042)
r_2	Dirichlet(5.0,13.4,1.6)	0.25	NA	0.30 (0.074)
r_3	Dirichlet(5.0,13.4,1.6)	0.67	NA	0.62 (0.084)
CV_J	Gamma(3.68,0.286)	1.05	0.93 (0.45)	0.86 (0.43)
CV_A	Gamma(3.68,0.286)	1.05	1.23 (0.38)	1.23 (0.37)
Unique particles			6910	5804

Notes: Values reported are averages of means over five runs with $N = 400\,000$. Values in parentheses are the standard deviations from a sample run and represent parameter uncertainty, not Monte Carlo variation. For definitions of parameters, see Table 2 and *Application to winter-run chinook salmon*.

butions. Uniform distributions were used instead of beta distributions for β , survival, and maturation rates, and instead of gamma distributions for α , CV_{BH} , CV_J , and CV_A . The means of the posterior distributions for α and β were higher than for the first set of priors and the same positive correlation between the two parameters was evident. There was less variability between years in age 2 survival estimates (values ranged from 0.49 to 0.58). Estimates of CV_{BH} were lower. Posterior means for sex- and age-specific maturation rates were as before quite similar to the prior means, again reflecting the lack of sex-specific information. The means of the posterior distributions of the states were quite similar for both the LH and HF models; the goodness of fit of mean state values to observations, especially for juvenile observations, was generally worse than for the original set of priors.

Thus with this relatively short time series and an observation vector with only three components, the choice of prior distributions did have a noticeable effect on the parameters. We believe, though we have not analytically justified this, that the influence of prior distributions increases as the dimension of the state vector increases relative to the dimension of the observation vector, and as the length of the time series decreases. Tentative advice for making inferences about the state vector is to keep the dimension of the state vector of roughly the same order of magnitude as the observation vector. This general problem of non-identifiability (or weak identifiability) as it applies to HPMs is an area needing research; potentially the ideas presented in several papers of Catchpole and Morgan (e.g., Catchpole and Morgan 1997, 2001) may provide useful diagnostics for determining the degree of influence of priors on particular parameters and/or states.

A PROCEDURE FOR ECOLOGISTS

HPMs are a useful framework for discrete time models of animal population dynamics in the presence of process and observation errors. To make the formulation and fitting of HPMs to animal population dynamics more widely accessible, user-friendly software is needed. The software needs to include automated fitting algorithms, such as SISR/KS or MCMC. It also needs to provide measures of the degree to which a HPM suffers from parameter non-identifiability problems. For example, if one formulates a state vector with 15 components and an observation vector with one component, it will, in most circumstances, be impossible to separately estimate each component of the state vector, let alone estimate all the parameters. Finally, software should include a suite of diagnostic tools, as exists for the fitting of linear regression models. The diagnostics would detect outliers and influential observations, as well as provide measures of goodness of fit and predictive performance.

Until such off-the-shelf software is developed for ecologists to use, we see the development, fitting, and use of HPMs as necessarily a collaborative effort between ecologists and statisticians. Ecologists would take the lead with model formulation and statisticians would focus primarily on model fitting. Both would work together to criticize and interpret the results. With this collaboration in mind we offer general suggestions for steps to take in the formulation, fitting, and use of HPMs. Implicit to these steps we assume that the ecologist has a clear picture of the intended use of the HPM and understands what the available data are, how data were collected, and limitations of the data.

- 1) Formulate the state process. Identify the subprocesses (e.g., birth, mortality, movement); formulate theories about the nature of the subprocesses (e.g., density dependence or independence of survival); develop reasonable probability distributions for the subprocesses (e.g., multinomial, Poisson); begin an initial list of the components of \mathbf{n}_t (e.g., young distinguished by sex or not).

- 2) Formulate the observation process. Identify the available, relevant data and tentatively determine the components of the state vector; describe the measurement or estimation processes; link the observation to the state process; develop reasonable probability distributions for the linkage.

- 3) Rethink the components and dimensions of the state and observation vectors. Rethink them in light of the potential for nonidentifiability. The state vector may need to be reduced in dimension, collapsing several components into a single component.

- 4) Write computer code that implements the Monte Carlo estimation procedure. When writing code from scratch, sequential importance sampling and its variations are perhaps easier than MCMC to program, but MCMC may be potentially more computationally efficient and accurate once the code is written.

- 5) Post-fitting analyses. Study the effect of prior distributions on posterior distributions by varying the priors and explore to the degree possible correlations between parameters and between states.

We close with mention of entry points to the literature that will generally be accessible to the quantitatively minded ecologist. The papers by Buckland et al. (2004) and Thomas et al. (2005) describe the general framework of HPMs for animal population dynamics, with the Thomas et al. paper providing a real application. The first chapters of the books edited by Gilks et al. (1996), and Doucet et al. (2001) are excellent introductions to MCMC and sequential Monte Carlo methods, respectively.

ACKNOWLEDGMENTS

K. B. Newman thanks the U.S. National Marine Fisheries Service for financial support under NMFS Contract Order #40ABNF101005. This work was supported in part by the National Science Foundation, grant DEB-0091961.#8221. We

also thank the editor, two anonymous reviewers, and Charles Paxton for helpful suggestions aimed at improving readability.

LITERATURE CITED

- Besbeas, P., S. N. Freeman, B. J. T. Morgan, and E. A. Catchpole. 2002. Integrating mark–recapture–recovery and census data to estimate animal abundance and demographic parameters. *Biometrics* **58**:540–547.
- Brooks, S. P., R. King, and B. J. T. Morgan. 2004. A Bayesian approach to combining animal abundance and demographic data. *Animal Biodiversity and Conservation* **27**:515–529.
- Buckland, S. T., K. B. Newman, L. Thomas, and N. B. Kesters. 2004. State-space models for the dynamics of wild animal populations. *Ecological Modelling* **171**:157–175.
- Catchpole, E. A., and B. J. T. Morgan. 1997. Detecting parameter redundancy. *Biometrika* **84**:187–196.
- Catchpole, E. A., and B. J. T. Morgan. 2001. Deficiency of parameter-redundant models. *Biometrika* **88**:593–598.
- Doucet, A., N. Freitas, and N. Gordon. 2001. *Sequential Monte Carlo methods in practice*. Springer-Verlag, New York, New York, USA.
- Geyer, C. 1996. Estimation and optimization of functions. Pages 241–258 in W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. *Markov chain Monte Carlo in practice*. Chapman and Hall, New York, New York, USA.
- Gilks, W. R., S. Richardson, and D. J. Spiegelhalter. 1996. *Markov chain Monte Carlo in practice*. Chapman and Hall, New York, New York, USA.
- Godsill, S. J., A. Doucet, and M. West. 2004. Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association* **99**:156–168.
- Gotelli, N. 2001. *A primer of ecology*. Third edition. Sinauer Associates, Sunderland, Massachusetts, USA.
- Gurney, W. S. C., and R. M. Nisbet. 1998. *Ecological dynamics*. Oxford University Press, Oxford, UK.
- Hallock, R. J., and F. W. Fisher. 1985. Status of winter-run chinook salmon, *Oncorhynchus tshawytscha*, in the Sacramento River. Anadromous Fisheries Branch Office Report. California Department of Fish and Game, Sacramento, California, USA.
- Hilborn, R., and M. Mangel. 1997. *The ecological detective*. Princeton University Press, Princeton, New Jersey, USA.
- Hilborn, R., and C. J. Walters. 1992. *Quantitative fisheries stock assessment: choice, dynamics, and uncertainty*. Chapman and Hall, New York, New York, USA.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Transactions of ASME Journal of Basic Engineering* **82**:35–45.
- Lindley, S. T. 2003. Estimation of population growth and extinction parameters from noisy data. *Ecological Applications* **13**:806–813.
- Liu, J. S. 2001. *Monte Carlo strategies in scientific computing*. Springer-Verlag, Berlin, Germany.
- Liu, J. S., and R. Chen. 1998. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association* **93**:1032–1044.
- Liu, J., and M. West. 2001. Combined parameter and state estimation in simulation-based filtering. Pages 197–223 in A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo methods in practice*. Springer-Verlag, Berlin, Germany.
- Martin, C. D., P. D. Gaines, and R. R. Johnson. 2001. Estimating the abundance of Sacramento River juvenile winter chinook salmon with comparisons to adult escapement. Final Report, Red Bluff Research Pumping Plant Report Series Volume 5. U.S. Fish and Wildlife Office, Red Bluff, California, USA.
- Mendelssohn, R. 1988. Some problems in estimating sizes from catch-at-age data. *Fishery Bulletin* **86**:617–630.
- Meyer, R., and R. B. Millar. 1999. Bayesian stock assessment using a state-space implementation of the delay difference model. *Canadian Journal of Fisheries and Aquatic Sciences* **56**:37–52.
- Newman, K. B. 1998. State-space modeling of animal movement and mortality with application to salmon. *Biometrics* **54**:1290–1314.
- Pitt, M. K., and N. Shephard. 1999. Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association* **94**:590–599.
- Quinn, T. J., and R. B. Deriso. 1999. *Quantitative fish dynamics*. Oxford University Press, Oxford, UK.
- Schnute, J. 1994. A general framework for developing sequential fisheries models. *Canadian Journal of Fisheries and Aquatic Sciences* **51**:1676–1688.
- Speed, T. 1993. Modelling and managing a salmon population. Pages 267–292 in V. Barnett and K. F. Turkman, editors. *Statistics for the environment*. John Wiley and Sons, New York, New York, USA.
- Sullivan, P. 1992. A Kalman filter approach to catch-at-length-analysis. *Biometrics* **48**:237–257.
- Thomas, L., S. T. Buckland, K. B. Newman, and J. Harwood. 2005. A unified framework for modelling wildlife population dynamics. *Australian and New Zealand Journal of Statistics* **47**(1):19–34.
- Trenkel, V. M., D. A. Elston, and S. T. Buckland. 2000. Calibrating population dynamics models to count and cull data using sequential importance sampling. *Journal of American Statistical Association* **95**:363–374.
- West, M., and J. Harrison. 1989. *Bayesian forecasting and dynamic models*. Springer-Verlag, New York, New York, USA.

SUPPLEMENT

S-Plus code used to fit life history (LH) and Hallock and Fisher (HF) hidden process models, along with data files (*Ecological Archives* A016-006-S1).