Ecology, 85(11), 2004, pp. 3140–3150

POPULATION TIME SERIES: PROCESS VARIABILITY, OBSERVATION ERRORS, MISSING VALUES, LAGS, AND HIDDEN STATES

JAMES S. CLARK,[1,3] AND OTTAR N. BJØRNSTAD[2]

[1]*Department of Biology and Nicholas School of the Environment, Duke University, Durham, North Carolina 27709 USA*
[2]*Departments of Entomology and Biology, 501 ASI Buildings, Penn State University, University Park, Pennsylvania 16802 USA*

Abstract. Population sample data are complex; inference and prediction require proper accommodation of not only the nonlinear interactions that determine the expected future abundance, but also the stochasticity inherent in data and variable (often unobserved) environmental factors. Moreover, censuses may occur sporadically, and observation errors change with sample methods and effort. The state variable (usually density or abundance) may be hidden from view and known only through highly indirect observational schemes (such as public health records, hunting reports, or fossil/archeological surveys). We extend the basic state-space model for time-series analysis to accommodate these dominant sources of variability that influence population data. Using examples, we show how different types of process error and observation error, unequal sample intervals, and missing values can be accounted for within the flexible framework of Bayesian state-space models. We provide algorithms based on Gibbs sampling that can be used to obtain posterior estimates of population states and of model parameters. For models that can be linearized, results can be used for direct sampling of the posterior, including those with missing values and unequal sample intervals. For nonlinear models, we make use of Metropolis-Hastings within the Gibbs sampling framework. Examples derive from long-term census and population data. We illustrate the extension to discrete state variables with multiple stages using a Time-series Susceptible–Infected–Recovered (TSIR) model for mid 20th-century measles infection in London, where birth rates are assumed known, but susceptibles and infected individuals arise from imperfect reporting.

Key words: Bayesian analysis; fossil pollen; measles; observation error; population regulation; state-space model; stochasticity; time-series analysis.

INTRODUCTION

The combined importance of observation error and process variability in population surveys is increasingly recognized. An integrated framework for analysis is needed, because process variability is part of population growth, whereas observation errors are not (Turchin 1995, Bjørnstad et al. 2001, Dennis et al. 2001, de Valpine 2002, de Valpine and Hastings 2002, Calder et al. 2003). Both can be large, and it can be difficult to incorporate them into a coherent modeling framework.

Even in data-rich systems, models can be highly uncertain (Pascual et al. 1997), with nontrivial consequences for inference (Wood and Thomas 1999). Ecological data are often limited and indirect (Legendre 1993). Nonlinear process models are commonly needed to describe population growth (e.g., Nisbet and Gurney 1982, Turchin 1995, Bjørnstad and Grenfell 2001), and these may need to be embedded within complex ''data models'' (Wikle 2003, Clark et al. 2004). Dynamics may involve networks of interactions among different life stages (e.g., Tuljapurkar and Caswell 1997) that themselves are variable (Clark 2003, Clark et al. 2003, 2004), and interactions among different species whose abundance are not always monitored (Stenseth et al. 1997). Data often miss obscure or hidden states, as not all life stages may be observed (Bjørnstad et al. 1999). Data may be densities (continuous state) or counts (discrete state). The ''latent'' variables—unobserved abundance of key species or life stages—may require estimation (Calder et al. 2003, Wikle 2003) or accommodation using time-lagged responses (e.g., Royama 1992, Bjørnstad et al. 2001). Sample intervals may be uneven or missed entirely, and sampling methods may change over time, meaning that the error structure will be nonstationary. As a consequence, descriptive time-series analysis ''has not proven terribly useful in analysis of (population) dynamics'' (Hilborn and Walters 1992:309). Yet there are growing numbers of accumulating long-term data sets (Kendall et al. 1998, Inchausti and Halley 2001) that should hold important insights, provided that models can accommodate the principal sources of variability in the data, including population fluctuations that cannot be ascribed to measured covariates (termed ''model misspecification'' or ''process error''), nonlinear relationships, and errors in measurement (''observation error'').

Manuscript received 31 July 2003; revised 3 March 2004; accepted 29 April 2004. Corresponding Editor: J. M. Ver Hoef.
[3] E-mail: jimclark@duke.edu.

These issues are of immediate concern in population risk assessment, which depends on a full treatment of variability and uncertainty (Boyce 1992, Ludwig 1999, Ellner and Feiberg 2003, Staples et al. 2004). Inadequate data are used with models to predict ''population viability,'' extinction probability, risk of infection, or expected time to extinction. Yet predictions may be sensitive to model assumptions, and they may depend critically on missing components (J. S. Clark, G. Ferraz, and N. Oguge, *unpublished manuscript*). Simulations demonstrate that the predictions needed by ecologists and managers may be highly sensitive to model uncertainties (e.g., Ludwig 1999, Clark 2003, Ellner and Feiberg 2003; but see Akcakya and Burgman 1995 for an alternative view). Ecological inference, therefore, requires a coherent probability framework that can be used to link ''inadequate data'' with the process models that describe underlying dynamics.

Despite long recognition of these challenges (e.g., Kuno 1971, Walter and Ludwig 1981, Morris 1983), flexible statistical approaches for time-series data are recent. Ives et al. (2003) provide an approach that can be applied to multiple species when components can be linearized. Calder et al. (2003) reviewed a Bayesian state-space framework (Carlin et al. 1992, West and Harrison 1997), for accommodating process and observation errors in ecological data that provides for simultaneous inference on parameters and latent variables and readily accommodates nonlinear models. In this study, we extend this state-space approach to accommodate the sources of stochasticity that are common in population time series. We present general results and demonstrate their application to long-term sample data subject to a range of errors, including animal population surveys, fossil pollen, and underreporting in epidemic models, where states are discrete and partially observed. We focus on inference and the role of different classes of stochasticity. A companion paper treats structured populations that are partially observed through capture–recapture methods (J. S. Clark, G. Ferraz, and N. Oguge, *unpublished manuscript*).

The Basic State-space Model

Let x_t represent log density, $\ln(n_t)$. The state-space model consists of two equations (Carlin et al. 1992, Calder et al. 2003), a process equation (1a) and an observation equation (1b):

$$x_t = f(x_{t-1}) + \varepsilon_t \tag{1a}$$

$$y_t = g(x_t) + w_t \tag{1b}$$

where y_t represents the observations of log abundance. The first ''process equation'' describes how log density changes through time. Examples include a random walk

$$n_t = n_{t-1}e^{\varepsilon_t} \qquad [f(x_{t-1}) = x_{t-1}] \tag{2a}$$

exponential growth

$$n_t = n_{t-1}e^{b+\varepsilon_t} \qquad [f(x_{t-1}) = x_{t-1} + b] \tag{2b}$$

and density-dependent growth, for instance as given by the Ricker model (sometimes called the discrete logistic):

$$n_t = n_{t-1}\exp(b_0 + b_1 n_{t-1})$$

$$[f(x_{t-1}) = x_{t-1} + b_0 + b_1 e^{x_{t-1}}]. \tag{2c}$$

The process model may contain parameters to be estimated from data (the bs). The notation $f(x_{t-1}; \mathbf{b})$ is used to indicate that the growth rate function involves parameters to be estimated. For exponential growth, $\mathbf{b} = b$. For logistic growth, $\mathbf{b} = [b_0, b_1]^{\mathsf{T}}$, where T denotes the transpose, and bold font is used to signify vectors. If we assume that ε_t is a zero mean Gaussian process with variance σ^2,

$$\varepsilon_t \sim \mathcal{N}(0, \sigma^2) \tag{3}$$

then the process variance is lognormal, and the process equation (1a) represents population growth with multiplicative lognormal errors. From (Eq. 1a) and (Eq. 3) it follows that $x_t \mid x_{t-1}, \sigma^2 \sim \mathcal{N}(f(x_{t-1}), \sigma^2)$, where ''$\sim$'' means ''is distributed as,'' and ''$\mid$'' means ''conditioned on.''

The second, observational equation (1b) describes the effect of sampling, or ''observation error.'' This equation can include systematic biases that might depend on density, in terms of $g(x_t)$, and observation errors (w_t). For the simple unbiased lognormal observation error model,

$$w_t \sim \mathcal{N}(0, \tau^2) \tag{4}$$

so that $y_t \mid x_t, \tau^2 \sim \mathcal{N}(x_t, \tau^2)$, where τ^2 is the error variance. The continuous state and error structure outlined in Eqs. 1–4 assume that the population is relatively abundant and that densities are large. If abundances are low, a discrete-state model is preferable (e.g., Bjørnstad et al. 2002). If counts are low, it may be necessary to use a discrete sampling distribution for the observations (e.g., Stenseth et al. 2003, Clark et al. 2004). Both of these issues are illustrated below.

The likelihood can be written in terms of the log abundances, $\mathbf{x}$, and the time series of log observations, $\mathbf{y}$:

$$p(\mathbf{x}, \mathbf{y} \mid \mathbf{b}, \sigma^2, \tau^2)$$

$$\overbrace{\propto\ p_1(x_0 \mid \mu_0, \sigma^2) \prod_{t=1}^{T} p_1[x_t \mid f(x_{t-1}; \mathbf{b}), \sigma^2]}^{\text{process model}}$$

$$\overbrace{\times \prod_{t=1}^{T} p_2[y_t \mid g(x_t), \tau^2]}^{\text{data model}}. \tag{5}$$

There are two probability densities, one associated with the process (p_1) and one with observations (p_2). We

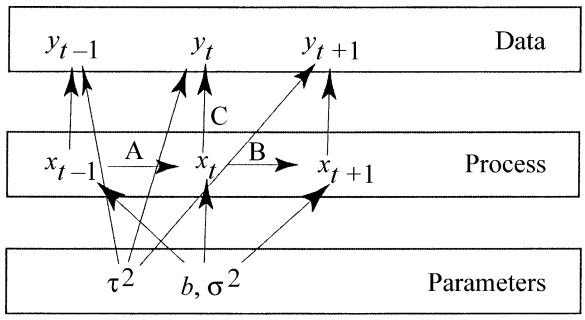

FIG. 1. The basic state-space model for exponential growth. Data constitute the upper level of the graph and consist of observations **y**. The lower two levels are not observed and, thus, must be estimated. These include the latent process **x**, the parameter describing constant per capita growth b, and variances for the process error σ^2 and observation error τ^2, which, in this example, are assumed to be constant.

refer to these as the ''process model'' and ''data model,'' respectively. Note that this likelihood is conditioned on the initial state x_0. The two sides of Eq. 5 are proportional (rather than equal) to one another, because the right hand side lacks the normalization constant that makes the integral equal to unity. This constant can be obtained by a sampling-based approach (Gelfand and Smith 1990), such as Markov Chain Monte Carlo (MCMC) integration. Previous ecological applications of MCMC for simulation of Bayesian posteriors include Bjørnstad et al. (1999), Millar and Meyer (2000), Calder et al. (2003), Clark et al. (2003, 2004), Stenseth et al. (2003), and Wikle (2003).

Fig. 1 illustrates the relationship between process and data. Change in density over time can be described by a first order Markov process $\mathbf{x} = x_1, x_2, \ldots, x_T$; the future $(x_{t-1}, x_{t+2}, \ldots)$ is conditionally independent of the past $(\ldots, x_{t-2}, x_{t-1})$. Thus we can model the process for a given time t by focusing on the arrows that are directly connected to x_t in Fig. 1. These include the observation at time t, y_t, and the state of the process immediately before and after t (i.e., x_{t-1} and x_{t+1}). The conditional distribution for a given unobserved state x_t simplifies to

$$p(x_t | x_{t-1}, x_{t+1}, y_t, \mathbf{b}, \sigma^2, \tau^2) \propto p_1[x_t | \overset{A}{f}(x_{t-1}; \mathbf{b}), \sigma^2] \times p_1[x_{t+1} | \overset{B}{f}(x_t; \mathbf{b}), \sigma^2] p_2[y_t | \overset{C}{g}(x_t), \tau^2]. \quad (6)$$

In terms of Fig. 1, the three densities on the right hand side are (A) the arrow pointing from x_{t-1} to x_t, (B) the arrow pointing from x_t to x_{t+1}, and (C) the arrow pointing from x_t to y_t. Analysis is simplified because we need only specify these ''local'' (conditional) relationships. We can use a sampling-based approach to marginalize over the full model, which contains additional densities for parameters. We illustrate with an example.

Exponential growth

The simple exponential growth model (Fig. 1) serves as a reference against which different types of complexity can be compared. The parameter b represents per capita rate of change (Eq. 2b). We employ previous assumptions for stochastic terms (Eqs. 3, 4) and a normal distribution for the unknown initial log abundance, $x_0 \sim \mathcal{N}(\mu_0, \sigma_0^2)$.

To complete the Bayesian model, we require prior distributions for unknown parameters. For convenience we use priors that are conjugate with the likelihood (Calder et al. 2003), such that prior and posterior distributions have the same form. The variance parameters σ^2 and τ^2 have inverse gamma (IG) priors, which are conjugate for the normal likelihood:

$$\sigma^2 \sim \text{IG}(\alpha_\sigma, \beta_\sigma) \qquad \tau^2 \sim \text{IG}(\alpha_\tau, \beta_\tau). \quad (7)$$

In some cases, we use ''noninformative'' priors, meaning that prior distributions are rather flat and only weakly influence parameter estimates (Hartigan 1998). A noninformative inverse gamma density (IG) prior has parameter values that are small ($\alpha, \beta \ll 1$). In other cases, we have information on variances that is incorporated by way of informative priors.

The conjugate prior distribution for the growth rate is Gaussian with prior mean B_0 and variance V_b. The prior is noninformative for large V_b. The full model is

$$p(b, \sigma^2, \tau^2, \mathbf{x} | \mathbf{y}, \ldots) \propto \mathcal{N}(x_0 | \mu_0, \sigma_0^2) \prod_{t=1}^{T} \mathcal{N}(x_t | x_{t-1} + b, \sigma^2) \prod_{t=1}^{T} \mathcal{N}(y_t | x_t, \tau^2) \times \mathcal{N}(b | B_0, V_b) \text{IG}(\sigma^2 | \alpha_\sigma, \beta_\sigma) \text{IG}(\tau^2 | \alpha_\tau, \beta_\tau) \quad (8)$$

where ''...'' represents prior parameter values. The joint posterior is normalized using a sampling-based approach. The Gibbs sampler is a MCMC technique that draws alternately from the conditional posteriors for each of the unknowns, including all of the xs and the parameters b, σ^2, and τ^2 (Gelfand and Smith 1990, Gilks et al. 1996). Due to conjugacy, each can be sampled directly from conditional posteriors (see Appendix). We immediately turn to complications that typically confront most analyses.

Black Noddy on Heron Island

Ecological data are typically not this simple. To be of much use for estimating growth rates, data sets must be ''long term.'' Long-term data accumulate stochasticity from many sources. Samples are often unevenly spaced in time or missed entirely, and observation errors vary as sampling methods and sampling effort changes. Increase in the Black Noddy (*Anous minutus*) population on Heron Island, part of Australia's Great Barrier Reef, illustrates some typical challenges. About fifty breeding pairs were observed in the early 1900s. Eighty years later the estimate exceeded 60 000 (Ogden 1993). Several population estimates are available at irregular intervals, obtained by methods that were neither standardized nor bounded by estimates of uncertainty. Only the most recent censuses in 1985 (Barnes

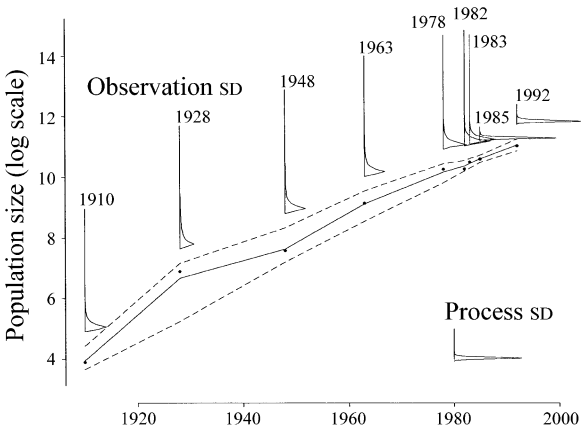

FIG. 2. Estimates of population density of the Black Noddy (1910–1992), shown as the posterior mean (solid line) and 95% credible intervals (dashed lines) for the model that admits uneven sample intervals. Posteriors for errors are taken as square roots (standard deviations) to have the same units as density (and, thus, the same scaling as the vertical axis). The base of each error posterior plot is zero. Above each sample point are posterior distributions for observation errors τ_t. At lower right is the posterior density for the square root of process error σ. All error posteriors integrate to 1.0 and have the same vertical scale, which, for clarity, is omitted.

and Hill 1989) and 1992 (Ogden 1993) report standard errors, and each is obtained by different methods. The three most obvious departures from the basic model outlined in the previous section are that (1) the distribution of observation errors varies from sample to sample, (2) our knowledge of sampling accuracy and bias varies, and (3) sampling has happened at irregular intervals.

Variable observation errors.—For sample-specific observation errors we allow that each observation y_t has a distribution with a unique variance, with the observation at time t distributed as $\mathcal{N}(y_t \mid x_t, \tau_t^2)$. In model 8, we can include a prior for each variance, depending on the available information, $\Pi_{t=1}^{T}\,\mathrm{IG}(\tau_t^2 \mid \alpha_t, \beta_t)$. For the early censuses, we have limited prior insight and use a noninformative prior IG(0.5, 0.01). For two recent censuses, we have estimates of 40 718 ± 3214 (1985; Barnes and Hill 1989) and 63 140 ± 7043 (1992; Ogden 1993). If these error estimates were known to be precisely correct, we could fix the error variances at the corresponding log values $\tau^2_{1985} = 0.00577$ and $\tau^2_{1992} = 0.0112$. Given that error estimates themselves are uncertain, we use informed priors to allow that standard errors are themselves estimated and, thus, have some degree of uncertainty. From moments of the inverse gamma we used parameter values having means equal to the reported values, and variances equal to the reported values squared, respectively, IG(3, 0.0115) and IG(3, 0.0224). These are reasonable assumptions that we use for demonstration. A full sensitivity analysis to prior values is readily accomplished using models described here. Sensitivity to the prior is judged by the effect of modifying the prior on the posterior. Implementation entails several extensions to the basic model that are given in the Appendix.

Uneven census intervals.—Although the process variance does not have an explicit time dimension, a time scale is implicit. The variance applies to a specific time increment, and variability in the process depends on time elapsed between censuses. As the interval between observations widens, the variability contributed by process error increases correspondingly. With process error, the growth that accrues over this interval is

$$n_t = n_{t-\delta_t}(e^{b+\varepsilon_t})^{\delta_t}.$$

The corresponding state-space model for log density is given by

$$x_t = x_{t-\delta_t} + \delta_t b + \delta_t \varepsilon_t \qquad y_t = x_t + w_t. \qquad (9)$$

Credible intervals on log abundance x_t (dashed lines in Fig. 2) integrate estimates of both process and observation error. Intervals are broad for early censuses, when sampling intervals are wide, and observation errors are estimated to be large (posterior densities shown above data points in Fig. 2). Intervals narrow toward recent times, when censuses are more frequent, and we have more information on observation errors. The small process error estimated here (lower right in Fig. 2) indicates that much of the scatter in data, particularly for early censuses, is assigned to observation errors.

Missing data.—The uneven sample intervals can be viewed as a special case of the general problem of missing data (Fig. 3). The representation of credible intervals for densities x_t in Fig. 2 follows standard practice—the dashed lines are interpolations between census dates. But the uncertainties only apply at times for which they were estimated. The uncertainty in population size in, say, 1935 is not accurately represented on this graph, because we have no estimate for 1935. The uncertainty in 1935 is not the interpolation between the closest estimates (1928, 1948). Rather, it depends on the full time series of observations, on the model, and on the priors. It is most strongly affected by the time elapsed since the last and until the next censuses, respectively. A probability statement for population size in, say, 1935 thus requires direct estimation.

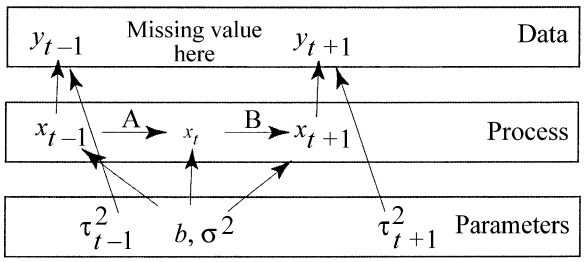

FIG. 3. Missing values in the exponential model with observation errors that differ in distribution from one sample to the next.

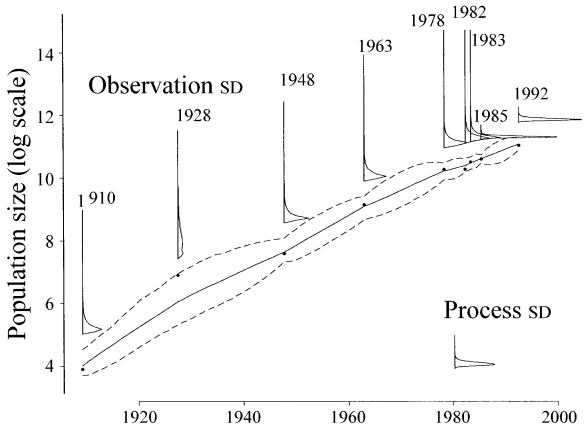

FIG. 4. Posterior densities for the Black Noddy from the model that treats noncensus years as missing values. Symbolism is as in Fig. 2.

Missing values are readily accommodated by the state-space framework. To emphasize this point, we estimate population size for every year from 1910 to 1992. In this example, every year for which there is no population estimate is treated as a ''missing value.'' Because the process itself is no longer ''unevenly spaced,'' we use the model for evenly spaced data (Eqs. 1, 2). We estimate observation errors only for years in which samples are available.

In the Gibbs sampler, we distinguish between sample times for which observations are available from those where observations are not. To sample x_t for census years, we draw values of x_t as before. For years not constrained by observations, there is no arrow connecting x_t to y_t (Fig. 3), indicating that there is no contribution from y_t (see Appendix).

Posterior estimates of population density are similar for the models that assume either uneven census intervals (Fig. 2) or missing data (Fig. 4). For the missing data model, we have predictions of x_t for each year when census data are unavailable. Predictive intervals ''balloon out'' between censuses. The widths of the balloons depend on the relative importance of process and observation error. If the process error is large, then the model is poorly specified, meaning that it fails to capture the complexity of population growth. Fat balloons reflect an uncertain process, in the sense that the growth rate ($b + \varepsilon_t$) is highly uncertain. This can result because the process model is inadequate, and thus the estimate of σ^2 is large. This contrasts with the case where most of the uncertainty is associated with observation error. When observation errors account for much of the stochasticity, credible intervals are wide at the census times themselves and predictive intervals do not increase substantially at intervening times. In such cases, large and uncertain observation errors dominate and balloons are not fat. Both situations can be observed in our analysis. We have a well-defined ''balloon'' between the last two censuses, both of which have small observation errors (Fig. 4). In other words, fat balloons mean that estimates of density could benefit from additional samples, because each census having low observation error brings large benefit.

Estimates of the population growth rates depend on model assumptions. Traditional models that assume either observation or process (but not both) show point estimates that differ from one another by almost 10% (0.081 and 0.090, respectively), and they are ''overconfident'' (Fig. 5a): credible intervals are too narrow. For the two state-space models (embracing both sources of variability, one using the irregular-sample and the other the missing-value parameterization), posterior means are in better agreement (0.080 and 0.086). The irregular-sample model steps forward with a few (nine), large intervals. The estimate for process error is especially low for this model, with greater uncertainty assigned to the growth rate parameter. The missing-value model has many growth increments (82), most of which are ''unconstrained'' by observations.

Population regulation

If $f(x_t)$ is nonlinear, we cannot sample directly from the conditional posterior for x_t. We describe the case for uneven sample intervals, with the understanding that, for evenly spaced censuses, time increments can

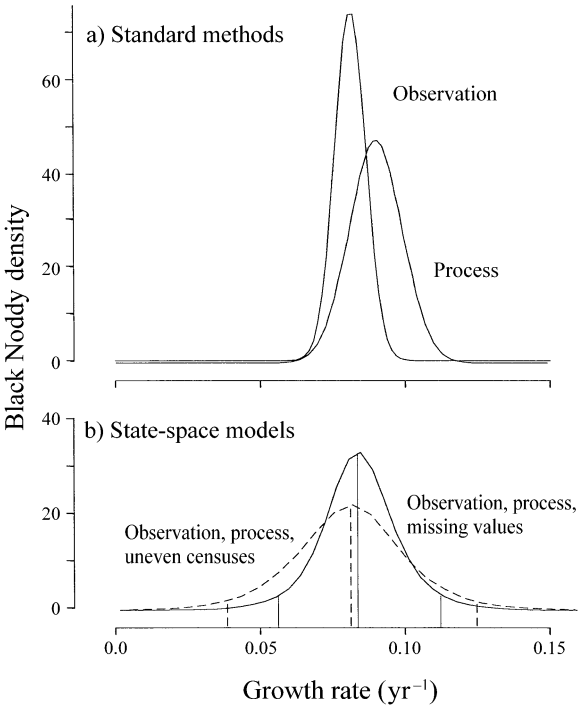

FIG. 5. Estimates of population growth rate (parameter b) for Black Noddys on Heron Island with four different assumptions. (a) Error distributions come from standard regression methods that assume either process or observation error, but not both. (b) Estimates for the model with unequal sample intervals (Fig. 2) are represented by the dashed line. Missing values (Fig. 4) are represented by the solid line.

be set equal to 1.0 (Calder et al. [2003] describe the standard model). For nonlinear models, there can be important implications of sample interval. We return to this issue in the *Discussion* section. Certain types of density dependence can be linearized, the Gompertz model being an example (e.g., Ives et al. 2003). Because many population models cannot be linearized, we use a typical example, the Ricker model (or discrete logistic). The process equation is

$$x_t = x_{t-\delta_t} + \delta_t b_0 + \delta_t b_1 e^{x_{t-1}} + \delta_t \varepsilon_t. \qquad (10)$$

This model, applied with uneven census intervals and missing values, is illustrated with two examples, characterized by different sampling concerns. The Gibbs sampler that is used for estimation is described in the Appendix.

Bialowieza moose.—The first nonlinear example consists of annual censuses of moose (*Alces alces*) in the Bialowieza Primeval Forest (BPF) that straddles of the border of Poland and Belarus (Jedrzejewska et al. 1997). A range of impacts has influenced ungulate populations, including timber removal, variable hunting pressure, and pasturing. The moose population was extirpated early in the 20th century and reintroduced. Our example concerns the post reintroduction phase. Careful censuses were conducted annually by a combination of snow tracking, hunting, drive censuses, and counts at baiting sites. The data used in this example are ''corrected'' for several known biases, and the authors of the study believe the residual observation errors to be small. Despite the careful censusing, several years are missing.

Posterior inference from the state-space model is consistent with the view that process error dominates observation errors (Fig. 6, upper panel). We used noninformative priors for all parameters. For variances, we used IG(0.01, 0.01). Credible intervals for the series closely follow observations, with lower confidence (wide posteriors) for years lacking census data (Fig. 6); credible intervals balloon out at missing censuses, because process error dominates. This contrasts with the dominant observation error estimates for the Black Noddy example (Figs. 2, 4). Estimates of parameter b_1 suggest some level of density dependence in the population growth: the posterior mean is negative, but the 95% credible interval does include zero (Fig. 6, lower right).

Gelfand and Ghosh's (1998) predictive loss provides support for the interpretation of some density dependence. This model selection index consists of a goodness-of-fit term G_m and a penalty for model complexity P_m. The more complex model results in smaller G_m. As complexity increases, P_m may initially decrease. The eventual rise in P_m with model complexity signals overfitting. We choose the model that minimizes $D_m = G_m + P_m$ (see Appendix). For the model without density dependence we obtain $D_m = 4.92 + 10.1 = 15.0$. For density dependence, we have $D_m = 4.50 + 8.38 = 12.9$.

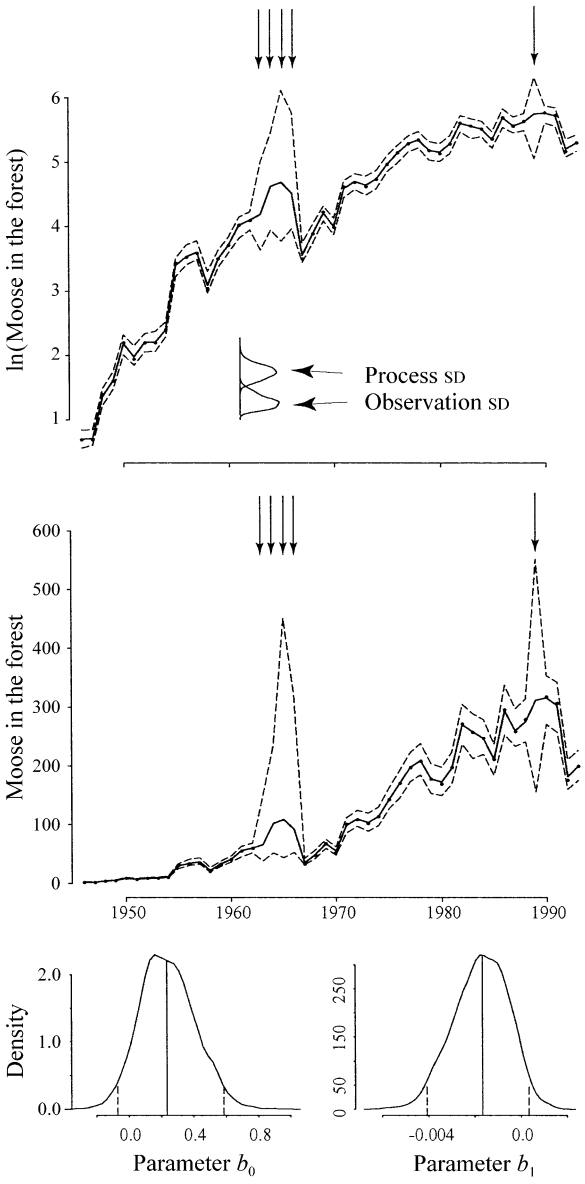

FIG. 6. Posterior distributions for parameters and latent variables for the logistic state-space model applied to the Bialowieza Primeval Forest moose population (1950–1990). Process errors are estimated to be much larger than observation errors (upper panel, inset). Posteriors for estimates of population size (dashed lines in upper and middle panels) show large uncertainty for years lacking census data (arrows). The posterior estimate for the density-dependent parameter b_1 has a 95% credible interval that includes zero, but the most probable value is negative (lower right). All error posteriors integrate to 1.0 and have the same vertical scale, which, for clarity, is omitted.

Thus, predictive loss motivates inclusion of density dependence.

Fossil pollen evidence for beech in southern Ontario, Canada.—Fossil pollen data provide evidence for growth of tree populations. Pollen extracted from sediments of lakes, together with ^{14}C estimates of sample

age, provides a rough idea of population change. Data are almost always unevenly spaced, because dates are assigned by indirect methods that are often completed or refined after samples are extracted for analysis. In addition to the fact that pollen abundance is not the same as tree abundance, there are large ''observation errors'' that are not associated with population growth. These include variability in sedimentation rates that introduce variability into pollen accumulation data that is unrelated to population growth. We demonstrate with an example of beech (*Fagus grandifolia*) from Nutt Lake, Ontario, Canada, originally analyzed by Bennett (1987). Because processes affecting observation are similar throughout, we use a single distribution for all samples. Knowing that process errors are likely to be large, we use slightly more informed priors than for the last example, with process and observation errors having priors IG(25, 5) and IG(2, 1), respectively. These have mean values of 0.2 and 1, respectively. Priors are still too weak to overwhelm the data. Time is indexed at the beginning of the series, which is ~8000 years before present.

Posterior estimates of pollen accumulation rate are broad (Fig. 7, upper, middle panels), with observation errors accounting for much of the variability in the data. The density-dependent parameter is estimated to be about zero (Fig. 7, lower right). This does not necessarily imply that density dependence is absent from the growth of the beech population. For example, observational errors may overwhelm details of the population growth process.

Lagged effects

Models with lagged effects are sometimes used to assess effects of interactions involving variables that cannot be directly observed (Royama 1992, Bjørnstad et al. 2001). Although we do not pursue an example here, we provide the approach for the lagged Ricker model (Royama 1992). If population growth rate depends on the last k time increments, we have the model

$$x_t = x_{t-1} + b_0 + \sum_{j=2}^{k} b_j e^{x_{t-j}} + \varepsilon_t \qquad y_t = x_t + w_t. \tag{11}$$

Matrix notation is more compact

$$\boldsymbol{\Delta} = \mathbf{X}\mathbf{b} + \boldsymbol{\varepsilon} \qquad \mathbf{y} = \mathbf{x} + \mathbf{w}$$

with densities $\mathbf{x}^{\mathsf{T}} = [x_1, \ldots, x_T]$, and observations $\mathbf{y}^{\mathsf{T}} = [y_1, \ldots, y_T]$. The vectors $\boldsymbol{\varepsilon}$ and $\mathbf{w}$ are sequences of zero-mean Gaussian variates. There is a response vector

$$\boldsymbol{\Delta} = \begin{bmatrix} x_1 - x_0 \\ \vdots \\ x_T - x_{T-1} \end{bmatrix}$$

a design matrix

$$\mathbf{X} = \begin{bmatrix} 1 & e^{x_{1-2}} & \cdots & e^{x_{1-k}} \\ 1 & e^{x_{2-2}} & \cdots & e^{x_{2-k}} \\ \vdots & \vdots & & \vdots \\ 1 & e^{x_{T-2}} & \cdots & e^{x_{T-k}} \end{bmatrix}$$

and a parameter vector $\mathbf{b}^{\mathsf{T}} = [b_0, \ldots, b_k]$. Note that we have negative indices in the design matrix. This sets us up to use priors for lags on the first k values, which are indexed as $(-k + 1, \ldots, 0)$. An alternative is to condition on the first k values of the series, and model only those samples beginning at $k + 1$. We have the same choice for the last k values of the series. Data modeling follows the same approach as for previous models (see Appendix).

Discrete-state Models

We have illustrated population growth with Gaussian process and observation errors on log density. In fact, Bayesian state-space models can accommodate a wide range of assumptions about the process and data. Here we use an epidemic time series to model discrete states for both the process and the data.

Epidemics of immunizing childhood diseases result from predator–prey-like interactions between infectious and susceptible individuals. To understand disease dynamics, epidemiologists construct models that include various states of the host (Anderson and May 1991, Bjørnstad et al. 2002, Grenfell et al. 2002). The TSIR (Time-series Susceptible–Infected–Recovered) model enumerates susceptible S_t and infected I_t individuals within the host population. Additional information is provided by birth rate data, B_t, which supply new susceptible individuals. The time increment in the TSIR model for measles is two weeks—the approximate duration of infection after which individuals recover. There are 26 two-week intervals in a year. Newborns become susceptible within several months. Before vaccination began in 1966, about 95–99% of all individuals in urban areas contracted the disease during their lifetimes. Our example includes the 21 years of measles case reports from London, between 1944 and 1964. Reported cases show a striking two-year cycle during this period (Fig. 8). We apply the methods in a discrete-state context to estimate how transmission rates change over the course of a year. Observables include reported infections, birth rates, and total population size for each of 546 two-week intervals in the time series. We are interested in how transmission rates vary during the course of the school year for this childhood infection. Our model accounts for the difference between reported infections I_t and actual infections $I_t^{(o)}$ due to underreporting, and the unobserved susceptible class S_t.

To estimate unobserved infection rates, we use a Bayesian framework for ''susceptible reconstruction'' as has been applied in previous efforts (Bjørnstad et al. 2002). The approach amounts to assuming that

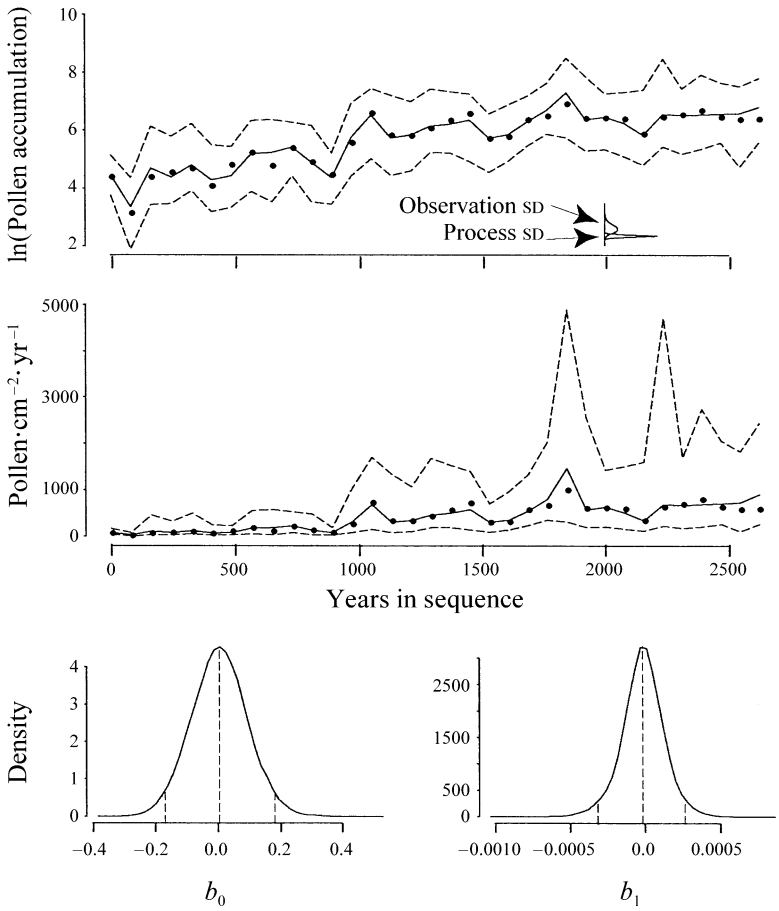

FIG. 7. Fossil pollen accumulation ([no. grains] · cm^{-2} · yr^{-1}) in Nutt Lake, Ontario, Canada (dots in upper and middle panels), and posteriors for regression parameters (lower panels). Dashed lines are 95% confidence intervals. Data are from Bennett (1987).

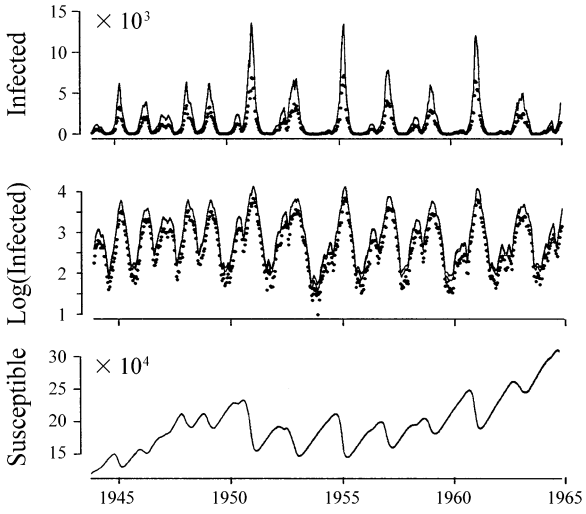

FIG. 8. Reported measles cases in London (dots) and posterior estimates (95% confidence intervals) of actual cases and susceptibles (solid lines), 1945–1965. The narrow confidence intervals result in lines that are indistinguishable from one another.

(nearly) all individuals will eventually pass through the infected state. If so, the long-term ratio of cumulative reported infections to cumulative births is the reporting fraction ρ (i.e., the fraction of cases that are reported). In previous analyses this assumption is used to estimate the mean number of susceptibles in the population (by regression), which is then used in a subsequent analysis of the TSIR model. The Bayesian approach allows us to incorporate this relationship directly in the analysis of disease transmission, together with uncertainty in the reporting rate. For clarity, our example makes use of a simplified TSIR model.

The probability that a susceptible individual will contract the disease is governed by the ''force of infection,'' φ (Bailey 1957). The time series of the unobserved susceptible class is given by

$$S_t = S_{t-1} + B_{t-1} - I_t. \tag{12}$$

Here, I_t is the loss of susceptibles due to infection, and B_t represents new susceptibles due to births. We assume that the force of infection is proportional to the fraction of infected individuals in the population. The constant

of proportionality is given by the time-varying transmission rate:

$$\varphi_{t-1} = \beta_w \times \frac{I_{t-1}}{N_{t-1}}$$

where β_w is the transmission rate for biweekly w. Note that each of the 26 parameters for β_w are identical across years. For example, β_1 is estimated from time intervals $t = [1, 27, \ldots]$, β_2 from to time intervals $t = [2, 28, \ldots]$, and so on. With these assumptions the distribution of new infections is

$$I_t \sim \text{Bin}(S_{t-1}, \varphi_{t-1}).$$

The relationship between actual and reported infections is

$$I_t^{(o)} \sim \text{Bin}(I_t, \rho).$$

The reporting rate has the expected value of

$$\frac{1}{0.95} \times \frac{\sum_{t=1}^{T} I_t^{(o)}}{\sum_{t=1}^{T} B_t} = 0.459$$

which represents the long-term average for this rate. The coefficient 0.95 represents the estimated fraction of individuals that are infected at some point during their childhood (nearly 1; Bjørnstad et al. 2002). The reporting rate of 0.495 is slightly lower than the 0.52 estimated using a more elaborate analysis (Finkenstädt and Grenfell 2000, Bjørnstad et al. 2002). We assumed uncertainty in this value. By matching moments for the Beta density, we obtain a slightly informative prior with the appropriate mean value with $\rho \sim \text{Beta}(45, 53)$. We assume flat priors on the transmission coefficients, $\text{Unif}(\beta_w \mid 0, 1000)$. The β_w cannot exceed N_{t-1}/I_{t-1}. In our analysis, values never approached this constraint. The full probability model is:

$$p[\mathbf{I}, \boldsymbol{\beta} \mid \mathbf{I}^{(o)}, \mathbf{B}, \rho, \ldots] \propto \prod_{t=1}^{T} \text{Bin}[I_t \mid S_{t-1}, \varphi_{t-1}(I_{t-1}, S_{t-1}, \beta_w)] \prod_{t=1}^{T} \text{Bin}[I_t^{(o)} \mid I_t, \rho] \times \prod_{w=1}^{26} \text{Unif}(\beta_w \mid 0, 1000)\text{Beta}(\rho \mid 45, 53)$$

where "..." refers to prior parameter values. For simplicity we condition on an initial susceptible value of $S_0 = 120\,000$. The Gibbs sampler is given in the Appendix.

Posterior estimates of transmission rates show the increase with the autumn school term and eventual decline as the school year progresses (Fig. 9, upper panel). Posterior estimates for transmission rates are well identified. The estimated actual infections show a biennial cycle, with values being roughly twice as high as those reported (Fig. 8, solid line). Given our conditioning on $S_0 = 120\,000$, the time series of susceptibles are tightly reconstructed. In reality the estimate of S_0 is uncertain, a fact that can be embraced through wider priors in a refined analysis.

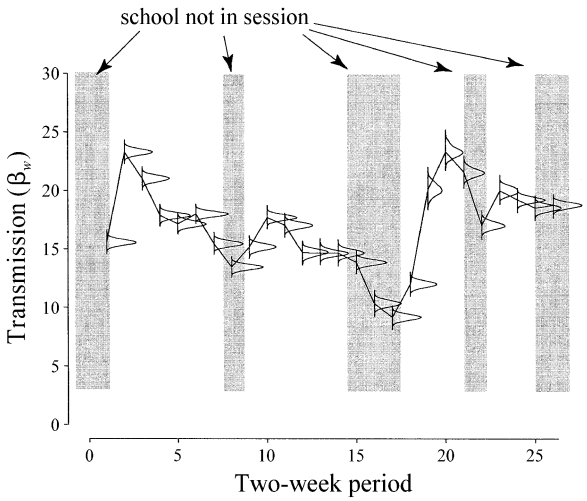

FIG. 9. Biweekly measles transmission coefficients (dimensionless), represented as the posterior means (solid line) and posterior distributions.

DISCUSSION

State-space models provide flexibility for the complex processes that generate population data. Less flexible alternatives can demand impossible data requirements and yield predictions that promote skepticism (Ludwig 1999, Clark 2003, Ellner and Feiberg 2003). When models cannot accommodate the important sources of stochasticity, statistical analysis is often piecemeal, with some things estimated first, then others from residuals. The result can be parameter estimates (e.g., demographic rates, strength of density dependence) and predictions (future population size, extinction times) that are not clearly linked to the data or the underlying process. Meanwhile, data may contain substantial information that is difficult to exploit.

State-space models admit complexity, and they yield inference on the range of unobservables that affect process and observations. Our examples address many of the common challenges that population data present. For the apparent exponentially increasing Black Noddy population on Heron Island (Ogden 1993), we allowed for process misspecification and observation errors that differed with each unevenly spaced sample (Figs. 2, 4). Observation errors were large and heterogeneous (Fig. 2). The estimate of population growth differed substantially from models assuming only observation or process errors (Fig. 5). We estimated that process error dominated the stochasticity in carefully censused moose populations from Bialowieza Primeval Forest, and we identified the extent to which uncertainty increased for missing censuses (Fig. 6). By contrast, observation error dominated the fossil pollen record of

beech expansion in southern Ontario, with inference of no identifiable density dependence (Fig. 7). The measles example illustrates how we can estimate time-varying rates and, more importantly, how we can reconstruct the dynamics of the unobserved susceptible class as an integral part of the statistical model for parameter estimation (Fig. 9).

Our treatment is not comprehensive. The approach is flexible to a variety of assumptions regarding stochasticity. Non-normal errors can be accommodated with scale mixtures (Carlin et al. 1992) or, as illustrated by our measles analysis, by incorporating specific distributions (such as binomial, Poisson). Variances that depend on population density could make use of the method of Stroud et al. (2003) involving discrete mixtures. The approach also admits hierarchically structured parameters to accommodate variability within populations. Clark et al. (2004) include random individual effects in a time-series (repeated-measures) approach to estimate the range of fecundities within tree populations. In a companion paper, J. S. Clark, G. Ferraz, and N. Oguge (*unpublished manuscript*) will apply a random effects approach to mark–recapture data, a type of longitudinal study with discrete states, where subjects are partially observed. Some related approaches in the ecological literature include McAllister et al. (1994), Gibson and Renshaw (1998), Bjørnstad et al. (1999), Millar and Meyer (2000), and Stenseth et al. (2003).

The advantages of the Bayesian state-space approach stem from a consistent probability model that applies simultaneously to the full process, the component conditional processes, the parameters, and the data. By building on low-dimensional conditional relationships, it is possible to construct a model that is fully consistent. The sampling-based approach provides the means for integrating the model.

Acknowledgments

This research was supported by Duke University's Center on Global Change and by NSF Grant DEB-9981392. For helpful comments, we thank Jay Ver Hoef, Matt Ferrari, and two anonymous reviewers.

Literature Cited

Akcakaya, H. R., and M. Burgman. 1995. PVA in theory and practice. Conservation Biology **9**:705–707.

Anderson, R. M., and R. M. May. 1991. Infectious diseases of humans: dynamics and control. Oxford University Press, Oxford, UK.

Bailey, N. T. J. 1957. The mathematical theory of epidemics. Griffin, London, UK.

Barnes, A., and G. J. E. Hill. 1989. Census and distribution of Black Noddy *Anous minutus* nests on Heron Island, November 1985. Emu **89**:129–134.

Bennett, K. D. 1987. Holocene history of forest trees in southern Ontario. Canadian Journal of Botany **65**:1792–1801.

Bjørnstad, O. N., B. F. Finkenstädt, and B. T. Grenfell. 2002. Dynamics of measles epidemics: estimating scaling of transmission rates using a time series SIR model. Ecological Monographs **72**:169–184.

Bjørnstad, O. N., J.-M. Fromentin, N. C. Stenseth, and J. Gjøsæter. 1999. Cycles and trends in cod population. Proceedings of the National Academy of Science USA **96**: 5066–5071.

Bjørnstad, O. N., and B. T. Grenfell. 2001. Noisy clockwork: time series analysis of population fluctuations in animals. Science **293**:638–643.

Bjørnstad, O. N., S. M. Sait, N. C. Stenseth, D. J. Thompson, and M. Begon. 2001. Coupling and the impact of specialised enemies on the dimensionality of prey dynamics. Nature **409**:1001–1006.

Boyce, M. 1992. Population viability analysis. Annual Reviews of Ecology and Systematics **23**:481–506.

Calder, K., M. Lavine, P. Mueller, and J. S. Clark. 2003. Incorporating multiple sources of stochasticity in population dynamic models. Ecology **84**:1395–1402.

Carlin, B. P., N. G. Polson, and D. S. Stoffer. 1992. A Monte Carlo approach to nonnormal and nonlinear state-space modeling. Journal of the American Statistical Association **87**:493–500.

Clark, J. S. 2003. Uncertainty in population growth rates calculated from demography: the hierarchical approach. Ecology **84**:1370–1381.

Clark, J. S., M. Dietze, I. Ibanez, and J. Mohan. 2003. Coexistence: how to identify trophic tradeoffs. Ecology **84**: 17–31.

Clark, J. S., S. LaDeau, and I. Ibanez. 2004. Fecundity of trees and the colonization–competition hypothesis. Ecological Monographs **74**:415–442.

Dennis, B., R. A. Desharnais, J. M. Cushing, S. M. Henson, and R. F. Costantino. 2001. Estimating chaos and complex dynamics in an insect population. Ecological Monographs **71**:277–303.

de Valpine, P. 2002. Review of methods for fitting time-series models with process and observation error and likelihood calculations for nonlinear, non-Gaussian state-space models. Bulletin of Marine Science **70**:455–471.

de Valpine, P., and A. Hastings. 2002. Fitting population models with process noise and observation error. Ecological Monographs **72**:57–76.

Ellner, S. P., and J. Fieberg. 2003. Using PVA for management despite uncertainty: effects of habitat, hatcheries, and harvest on salmon. Ecology **84**:1359–1369.

Finkenstädt, B., and B. Grenfell. 2000. Time series modeling of childhood diseases: a dynamical systems approach. Applied Statistics **49**:187–205.

Gelfand, A. E., and S. K. Ghosh. 1998. Model choice: a minimum posterior predictive loss approach. Biometrika **85**:1–11.

Gelfand, A. E., and A. F. M. Smith. 1990. Sampling-based approaches to calculating marginal densities. Journal of the American Statistical Association **85**:398–409.

Gibson, G. J., and E. Renshaw. 1998. Estimating parameters in stochastic compartmental models using Markov chain methods. IMA Journal of Mathematics Applied in Medicine and Biology **15**:19–40.

Gilks, W. R., S. Richardson, and D. J. Spiegelhalter. 1996. Markov Chain Monte Carlo in practice. Chapman and Hall, London, UK.

Grenfell, B. T., O. N. Bjørnstad, and B. Finkenstädt. 2002. Endemic and epidemic dynamics of measles. II. Scaling noise, determinism and predictability with the time series SIR model. Ecological Monographs **72**:185–202.

Hartigan, J. A. 1998. The maximum likelihood prior. Annals of Statistics **26**:2083–2103.

Hilborn, R., and C. Walters. 1992. Quantitative fisheries stock assessment. Choice, dynamics and uncertainty. Chapman and Hall, London, UK.

Inchausti, P., and J. Halley. 2001. Investigating long-term ecological variability using the global population dynamics database. Science **293**:655–657.

Ives, A. R., B. Dennis, K. L. Cottingham, and S. R. Carpenter. 2003. Estimating community stability and ecological interactions from time-series data. Ecological Monographs **73**:301–330.

Jedrzejewska, B., W. Jedrzejewski, A. N. Bunevich, L. Milkowsi, and Z. A. Krasinski. 1997. Factors shaping population densities and increase rates of ungulates in Bialowieza Primeval Forest (Poland and Belarus) in the 19th and 20th centuries. Acta Theriologica **42**:399–451.

Kendall, B. E., J. Prendergast, and O. N. Bjørnstad. 1998. The macroecology of population dynamics: taxonomic and biogeographic patterns in population cycles. Ecology Letters **1**:160–164.

Kuno, E. 1971. Sampling error as a misleading artifact in ''Key factor analysis.'' Researches on Population Ecology **13**:28–45.

Legendre, P. 1993. Real data are messy. Statistics and Computing **3**:197–199.

Ludwig, D. 1999. Is it meaningful to estimate an extinction probability? Ecology **80**:298–310.

McAllister, M. K., E. K. Pikitch, A. E. Punt, and R. Hilborn. 1994. A Bayesian approach to stock assessment and harvest decisions using the sampling-importance resampling algorithm. Canadian Journal of Fisheries and Aquatic Sciences **51**:2673–2687.

Millar, R. B., and R. Meyer. 2000. Bayesian state-space modeling of age-structured data: fitting a model is just the beginning. Canadian Journal of Fisheries and Aquatic Sciences **57**:43–50.

Morris, C. N. 1983. Parametric empirical Bayes inference: theory and applications. Journal of the American Statistical Association **78**:47–55.

Nisbet, R. M., and W. S. C. Gurney. 1982. Modelling fluctuating populations. Wiley, Chichester, UK.

Ogden, J. 1993. Population increase and nesting patterns of the Black Noddy *Anous minutus* in Pisonia forest on Heron Island: observations from 1978, 1979, and 1992. Australian Journal of Ecology **18**:395–403.

Pascual, M. A., P. Kareiva, and R. Hilborn. 1997. The influence of model structure on conclusions about the viability and harvesting of Serengeti wildebeest. Conservation Biology **11**:966–976.

Royama, T. 1992. Analytical population dynamics. Chapman and Hall, London, UK.

Staples, D. F., M. L. Taper, and B. Dennis. 2004. Estimating population trend and process variation for PVA in the presence of sampling error. Ecology **85**:923–929.

Stenseth, N. C., W. Falck, O. N. Bjørnstad, and C. J. Krebs. 1997. Population regulation in snowshoe hare and Canadian lynx: asymmetric food web configurations between hare and lynx. Proceedings of the National Academy of Sciences of the USA **94**:5147–5152.

Stenseth, N. C., H. Viljugrein, T. Saitoh, T. F. Hansen, M. O. Kittilsen, E. Bolviken, and F. Glockner. 2003. Seasonality, density dependence, and population cycles in Hokkaido voles. Proceedings of the National Academy of Sciences of the USA **100**:11478–11483.

Stroud, J. R., P. Müller, and N. G. Polson. 2003. Nonlinear state-space models with state-dependent variances. Journal of the American Statistical Association **98**:377–386.

Tuljapurkar, S., and H. Caswell. 1997. Structured-population models in marine, terrestrial, and freshwater systems. Chapman and Hall, New York, New York, USA.

Turchin, P. 1995. Population regulation: old arguments and a new synthesis. Pages 19–39 *in* N. Cappuccino and P. Price, editors. Population dynamics. Academic Press, New York, New York, USA.

Walter, C. J., and D. Ludwig. 1981. Effects of measurement errors on the assessment of stock–recruitment relationships. Canadian Journal of Fisheries and Aquatic Science **38**:704–710.

West, M., and J. Harrison. 1997. Bayesian forecasting and dynamic models. Springer Verlag, New York, New York, USA.

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

Wood, S. N., and M. B. Thomas. 1999. Super-sensitivity to structure in biological models. Proceedings of Royal Society London, B **266**:565–570.

APPENDIX

The Gibbs sampler for Bayesian state-space models is described in ESA's Electronic Data Archive: *Ecological Archives* E085-103-A1.