Promoting the Science of Ecology

---

http://www.jstor.org

FITTING POPULATION DYNAMIC MODELS TO TIME-SERIES DATA BY GRADIENT MATCHING

STEPHEN P. ELLNER,[1,4] YODIT SEIFU,[1,2,5] AND ROBERT H. SMITH[3]

[1]Biomathematics Program, Department of Statistics, North Carolina State University,
Raleigh, North Carolina 27695-8203 USA
[2]Department of Mathematics and Statistics, University of New Brunswick, Fredericton, New Brunswick, Canada E3B 5A3
[3]Department of Biology, University of Leicester, University Road, Leicester LE1 7RH UK

Abstract. We describe and test a method for fitting noisy differential equation models to a time series of population counts, motivated by stage-structured models of insect and zooplankton populations. We consider semimechanistic models, in which the model structure is derived from knowledge of the life cycle, but the rate equations are estimated nonparametrically from the time-series data. The method involves smoothing the population time series $x(t)$ in order to estimate the gradient dx/dt, and then fitting rate equations using penalized regression splines. Computer-intensive methods are used to estimate and remove the biases that result from the data being discrete time samples with sampling errors from a continuous time process. Semimechanistic modeling makes it possible to test assumptions about the mechanisms behind population fluctuations without the results being confounded by possibly arbitrary choices of parametric forms for process-rate equations. To illustrate this application, we analyze time-series data on laboratory populations of blowflies *Lucilia cuprina* and *Lucilia sericata*. The models assume that the populations are limited by competition among adults affecting their current birth and death rates. The results correspond to the actual experimental conditions. For *L. cuprina* (where the model's structure is appropriate) a good fit can be obtained, while for *L. sericata* (where the model is inappropriate), the fitted model does not reproduce some major features of the observed cycles. A documented set of *R* functions for all steps in the model-fitting process is provided as a supplement to this article.

Key words: blowflies; gradient matching; Lucilia cuprina; Lucilia sericata; model fitting; partially specified models; population dynamics; semimechanistic models; semiparametric models; stage structured models.

INTRODUCTION

Quantitative dynamic modeling of population dynamics is an important adjunct to experimental study of the mechanisms underlying population fluctuations. Manipulative experiments are invaluable for identifying the processes operating in a population that may be responsible for observed patterns of population fluctuations. However, it has only rarely been possible to test experimentally the various hypotheses about which particular process is, for example, the underlying cause of a population cycle (Hudson et al. 1998 and Korpimäki and Norrdahl 1998 are two recent exceptions), and such experiments always involve considerable effort. Quantitative dynamic modeling thus can make an important contribution by helping to determining which processes (under which conditions) are capable of producing population fluctuations that are consistent with the data. Constructing alternative models, and comparing the ability of each to account for the observed fluctuations, can be a powerful tool for identifying likely causal mechanisms, which can then be the focus of empirical tests (e.g., Hilborn and Mangel 1997, Kendall et al. 1999, Turchin and Ellner 2000a, b).

Empirically derived mechanistic models also play an important role in applied ecology, including the management of fisheries, forests, and other natural resources. In community ecology, fitting multispecies dynamic models to data on community dynamics can be an effective way of estimating the strengths of within- and between-species interactions (Pfister 1995, Laska and Wooton 1998, Ives et al. 1999). The alternative models that would be compared in this case correspond to different assumptions about which interspecific interaction coefficients are small enough to be ignored.

Before alternative hypotheses can be tested in this way, each hypothesis has to be translated into a model. A rate equation is written for each process affecting the variables of interest, and these are combined into a state variable model such as a system of differential or difference equations (e.g., Haefner 1996, Hastings

[4] Present address: Department of Ecology and Evolutionary Biology, Cornell University, E145 Corson Hall, Ithaca, New York 14853-2701 USA. E-mail: spe2@cornell.edu

[5] Present address: Novartis Pharmaceutical Corporation, 59 Route 10, East Hanover New Jersey, 07936-1080 USA. E-mail: yodit.seifu@pharma.novartis.com.

1997, Gurney and Nisbet 1998). Rate equations are almost always built up from simple (and conventional) parametric functional forms. Familiar functional forms include

1) the Holling type-II equation for a saturating functional response, $F = Ax/(B + x)$, where F = feeding rate, x = prey abundance, and A and B are positive parameters;

2) the Nicholson-Bailey parasitism rate equation, $1 - e^{-aP}$, where P = number of parasitoids and a is a positive parameter; and

3) the bilinear interaction terms in many microparasite host–pathogen models (infection rate = βSI, where S = number of susceptible hosts, I = number of infective hosts, and β is a positive parameter) and in Lotka-Volterra competition and predation models.

Many of these conventional equations were initially derived from mechanistic assumptions, but in applications those assumptions are rarely verified empirically. A particular functional form is often used because it appears to fit the data at hand or data on a similar system, or because it was used in a previous model, without testing whether another form would fit better, or whether the data are sufficient for selecting reliably between competing functional forms.

Relying on conventional functional forms is appropriate for strategic models, whose purpose is to capture the main qualitative features of a system and increase our conceptual understanding of how the system may be operating. But when models are compared with data in order to evaluate competing hypotheses about causal processes, the choice of functional forms for each process-rate equation is an undesirable confounding factor. Model 1 may fit better than Model 2 because Model 1 makes the right mechanistic assumptions and Model 2 doesn't. But it is also possible that Model 2 simply suffers from a poor choice of functional form for a process rate that is not a part of Model 1 (type-II instead of type-III functional response, etc.). Seemingly innocuous choices between alternative functional forms can have drastic effects on model predictions (Wood and Thomas 1999). Similarly, an estimated interaction coefficient may not reflect the actual strength of an interspecific interaction, if the corresponding rate equation is not a good description of the interaction over the range of densities encountered.

Our emphasis in this paper is therefore on fitting dynamic models in which some or all of the rate equations are estimated nonparametrically. This means that parametric functional forms are replaced by a flexible family of possible functions, such as a spline, neural network, or local regression. Models of this type have been called semimechanistic (Ellner et al. 1998, Smith et al. 2000), semiparametric (Wood 1999), and partially specified (Wood 2001).

If the model is deterministic, nonparametric rate equations can be fitted by trajectory matching; Wood (1999, 2001) describes a very general method using spline rate equations. In trajectory matching, differences between model output and the data are assumed to be entirely due to measurement errors in the data. Model parameters are chosen to minimize a weighted mean-square difference between model output and the data. However, if the population dynamics are noisy (i.e., perturbed by unpredictable exogenous factors), different methods may be required (though not always: see *Discussion*).

The method presented here is suitable for noisy population dynamics, under two conditions. The first is that all state variables of the model are measured or estimated. The second is that measurements are taken frequently and accurately enough that good estimates of the instantaneous rate of change (the gradient) can be obtained for each state variable. Under these conditions, estimating a rate equation

$$dx_i/dt = F_i(x_1, x_2, \cdots, x_n)$$

becomes a problem in nonparametric regression, with (usually inexact) observations of both the independent and dependent variables. The fitting method is then efficient enough to allow resampling-based inference methods such as bootstrapping, and simulation-based methods for estimating and reducing the bias due to random sampling errors. Other fitting methods can be used on noisy dynamics if these two conditions are not satisfied, but they are computationally demanding and currently can only be used to estimate a small number of parameters (see *Discussion*).

The paper is organized as follows. In *Methods: Single-species model and fitting procedures*, we describe the fitting method in the context of a stage-structured population model, and in *Results: Simulated data*, we study the method's accuracy using simulated data. In the case studies below (*Results: Nicholson's blowflies*, Lucilia cuprina, and *Results: Lucilia sericata control populations*), we apply our procedure to data from two experiments on the dynamics of laboratory populations of blowflies, *Lucilia cuprina* (Nicholson 1957) and *Lucilia sericata* (Smith et al. 2000). In order to mimic a typical application to field data, we use only the time series of adult counts from these experiments. Finally, we draw conclusions and discuss how our method relates to other methods for fitting population dynamic models to time series data. A set of R (Ihaka and Gentlemen 1996) functions for all steps in the model-fitting process, with documentation and examples, is provided as a Supplement.

METHODS: SINGLE-SPECIES MODEL AND FITTING PROCEDURES

For specificity and looking ahead to our case studies, we present the method in the context of a simple stage structured population model:

$$dx/dt = B(x(t - \tau)) - D(x(t)). \tag{1}$$

Here $x(t)$ typically represents the number of reproductively mature adults, $D(x[t])$ is adult mortality, and $B(x(t - \tau))$ is recruitment into the adult class at age τ of individuals born τ time units previously. Models of this type have been developed for a variety of animal species with discrete life stages, including insects and zooplankton, with success at capturing both qualitative and quantitative aspects of the population dynamics (e.g., Gurney et al. 1980, McCauley et al. 1996). Models such as Eq. 1 are also coming into use for applied purposes such as biological control (Murdoch 1994, Murdoch and Briggs 1996, Briggs et al. 1999), so procedures for fitting this type of model may have some practical value. In Eq. 1, there is a strict separation in time between density effects on births (time-lagged) and deaths (instantaneous), which makes it possible to estimate B and D separately. Without a time separation, e.g., if $dx/dt = B(x(t)) - D(x(t))$, it would only be possible to estimate the net effect $B - D$.

Our goal is to estimate the rate functions B and D, and possibly the delay time τ, from a time series of observations $\{x(t_i), i = 1, 2,\cdots, n\}$. Even laboratory populations are not perfectly deterministic in their dynamics, so we need to consider stochastic models for which Eq. 1 represents the "average dynamics" in the sense that

$$E[dx/dt] = B(x(t - \tau)) - D(x(t)). \qquad (2)$$

One such model would be a differential equation with multiplicative variability in vital rates due to fluctuations in environmental conditions

$$dx/dt = \varphi(t)B(x(t - \tau)) - \psi(t)D(x(t)). \qquad (3)$$

In Eq. 3, $\varphi(t)$ and $\psi(t)$ represent random variation in vital rates due to environmental variables that are assumed to change smoothly over time, but with a random component to their pattern of change. For example, they could represent the net effects of ambient temperature, soil moisture, etc., on birth and death rates. For consistency with Eq. 2, $\varphi(t)$ and $\psi(t)$ are standardized to $E[\varphi(t)] = E[\psi(t)] = 1$, and for the statistical analysis we assume that $\varphi(t)$ and $\psi(t)$ are stationary (i.e., that they do not have any trends over the period of data collection). Eq. 3 can then be interpreted as a nonautonomous differential equation (conditional on the values of $\varphi(t)$ and $\psi(t)$), whose solutions will satisfy Eq. 2; note that this is not true of "stochastic differential equations" in the classical sense (e.g., Oksendal 1998), which have nondifferentiable solution paths because their random term is a white-noise process.

As discussed in Introduction, we are interested in nonparametric estimates of B and D. The methods presented here are conceptually similar to those in Ellner et al. (1997). However, those methods were statistically inefficient and relied on ad hoc visual criteria for selecting the complexity of the fitted rate equations. The methods here are more efficient, exploit statistical the-

ory for objectively selecting the complexity of a nonparametric model, and produce simpler models (many fewer degrees of freedom) for the same data sets without any loss of accuracy. In outline, the proposed method runs as follows:

1) Fit a smooth curve to the data $x(t_i)$ by local polynomial regression, which gives estimated values of the gradient $dx(t_i)/d(t)$.

2) Fit Eq. 2 with B and D estimated by penalized regression splines, with automatic selection of model complexity.

3) In both fitting steps, simulation-based methods are used to estimate and correct for biases that result from the data being discrete-time samples with measurement errors.

For simplicity we assume that the value of the time delay, τ, is known. In the case studies here on insect populations, τ corresponds to the egg-to-adult development time, which can be estimated with a small-scale experiment. If τ is not known, the fitting process can be repeated for a set of τ values covering the plausible range, and a best-fitting value of τ can be selected by the criterion used to select model complexity. Our approach can also be used if a time delay is absent, but the birth and death rates cannot be separated. For example, if $dN/dt = N(b(N) - d(N))$ with $b(N) = b_0 - b_1N$ and $d(N) = d_0$, the gradient is $(b_0 - d_0)N - b_1N^2$ so $b_0 - d_0$ can be estimated but not b_0 or d_0.

We now describe the fitting procedures in more detail, followed by a simulation study to assess the accuracy of the estimates.

Estimating the gradient

The first step is smoothing the data to obtain an estimate of the gradient $dx/dt(t_i)$. We used local polynomial regression (Fan and Gijbels 1996), for two reasons: (1) estimates of dx/dt at sufficiently distant times are independent because they are derived from disjoint sets of data, which is useful for fitting the model and selecting model complexity (see the Appendix) and (2) the fitted curve automatically yields a gradient estimate.

In local polynomial regression with data $\{x(t_i), i = 1, 2,\cdots, n\}$, the value of the fitted curve at any time t is obtained by fitting a low-order polynomial by weighted least-squares regression of x on t, with weights $w_i(t)$ that are large for t_i near t and small for t_i far from t. The linear term in the fitted polynomial is then the gradient estimate at time t. Large-sample properties favor polynomials of even order for estimating derivatives (Fan and Gijbels 1996). We used fourth-order polynomials, which outperformed quadratics in simulation tests because they gave better approximations to the asymmetries between increase and decrease phases at peaks and troughs.

The weights are defined as follows:

$$w_i(t) = K(|t_i - t|/h) \qquad (4)$$

where $K(d)$ is some function (called the kernel) that falls off with increasing distance d. We used the Gaussian kernel $K(d) = e^{-d^2}$; to accelerate computation we set $K(d) = 0$ for $d > 5$, which has no effect on numerical results. When the bandwidth h is very large, the weights are nearly constant ($w_i(t) \approx K(0)$ for all i,t) so the fitted curve is almost identical to a least-squares polynomial regression of x on t. As h decreases the regression becomes genuinely local (distant times are ignored because the weights are small), and the fitted curve comes closer to interpolating the data.

As we discuss in the Appendix, for our purposes a very small value of h is best, generally one to two times the sampling interval, in order to minimize bias in the gradient estimates. Even with h small, there is some bias towards zero at the extremes (very high gradients are underestimated, and vice versa). However, it is possible to estimate this bias and correct for it (see the Appendix). The bias corrections are fairly small (at most 20% on the simulated and real data series considered here, and typically smaller), but they do produce a noticeable improvement in the results reported below on simulated data.

To test this gradient-estimation method, we generated simulated data for which the exact gradient values are known and can be compared to estimates recovered from the population time series. The simulation model had birth rate $B(x) = 10xe^{-x/500}$, mortality rate $D(x) = 0.3x$, and maturation time $\tau = 18$. The death rate was deterministic ($\psi \equiv 1$), but there was environmental stochasticity in the birth rate, with $\varphi(t)$ being Gaussian distributed with mean $= 1$, standard deviation $= 0.1$, and autocorrelation $= 0.75$ at lag $t = 1$. The simulation model was actually more realistic than the fitted model in Eq. 2. Eq. 2 omits demographic stochasticity and treats population size as a continuous variable. The simulation model had discrete individuals, and incorporated demographic stochasticity by using a finite time step $dt = 0.1$ and setting the number of births in each time interval as a Poisson random variable with mean equal to the expected number of births

number of births in

$$(t,t + dt) \sim \text{Poisson } (\varphi(t)B(x(t))dt). \tag{5}$$

This mismatch between the (discrete) simulated data and the (continuous) fitted model was deliberate. A continuous-state model is always an approximation to the discrete population changes resulting from individual births and deaths, so our fitting method needs to be robust against the unavoidable discrepancy between a continuous-state model and discrete population counts. Although mean population size is high in both models, the troughs drop to roughly 20 and 40 adults, respectively, which are more extreme than the data we will be analyzing. Thus, these models provide a legitimate test of our method in the presence of demographic stochasticity.

In the absence of measurement errors (Fig. 1a), the fitted smooth curve is nearly an interpolation of the data. There is close correspondence between the estimated and true gradient values, with the estimates accounting for ~96% of the variance in the true gradient values, varying slightly from one run of the model to another ("true" gradient values being the actual finite rate of change, $[x(t + 0.1) - x(t - 0.1)]/0.2$, computed from the simulation output). In the presence of sampling errors (derived from a Poisson random sampling model depending on a capture probability p, see the Appendix) the gradient estimates are slightly less accurate, accounting for ~86% of the variance in the true values. The scatter plots also indicate that there is no systematic tendency to over- or underestimate the gradient, except very slightly at the extremes. The coefficient of variation (CV) for sampling errors in the Poisson sampling model is proportional to $1/(\text{square root of population size})$; very similar results obtain for simulated sampling errors with constant CV, up to about $CV = 0.2$.

The simulated time series in Fig. 1 is finely sampled, having ~40 data points per cycle. Many real data sets are sparser; for example the Nicholson (1957) data analyzed in the next section has ~20 points per cycle. Sampling the simulated time series every second day (instead of daily) hardly decreases the overall accuracy of gradient estimation: about 78% of the variance in true gradient values is explained by the estimate, in the presence of sampling errors. Eventually, however, too much detail is lost by sparser subsampling, and with samples taken every 6 d (7 points per cycle) the estimates account for <50% of the variance.

Estimating the rate equations

Once values of the gradient $Z_i = dx/dt(t_i)$ have been estimated, the problem of fitting the model in Eq. 2 is reduced to fitting the nonparametric additive model:

$$Z_i = B(X_i) - D(Y_i) + \varepsilon_i$$

$$X_i = x(t_i - \tau) \qquad Y_i = x(t_i). \tag{6}$$

A variety of nonparametric model families could be used to estimate B and D. We used cubic penalized-regression splines (Eilers and Marx 1996, Ruppert and Carroll 1997, 2000), because of their computational advantages. The model has the form of linear regression onto a set of nonlinear basis functions (see the Appendix). As a result, many useful constraints on the fitted functions can be imposed as linear constraints on the parameters, and the constrained model can still be fitted quickly using quadratic programming (as explained by Wood 1997 for a different family of splines, and in the Appendix). Both of those features are essential here. Quick fitting made it possible to hone and validate our methods (including the computationally intensive bias-reduction procedure described at the end of this section) by testing them on a large number of
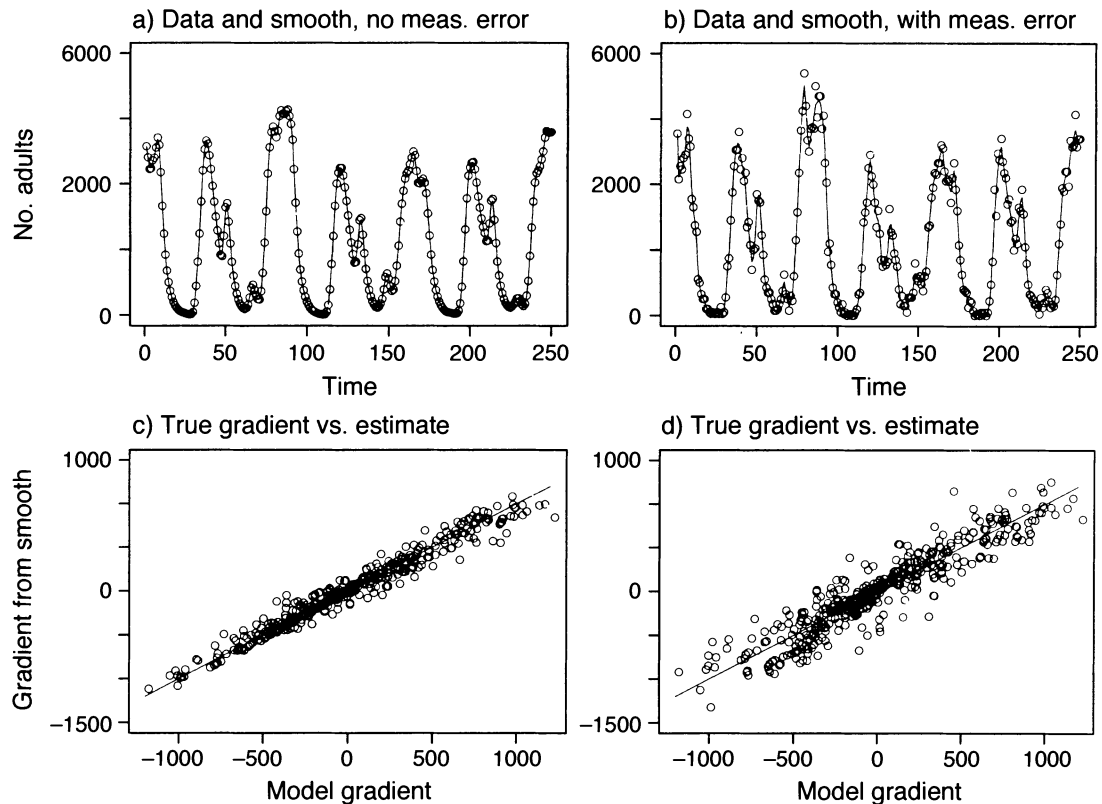
FIG. 1. Typical results from the gradient estimation step. (a) A time series generated by the model described in Eq. 3 with demographic and environmental stochasticity, and the smooth curve fitted to the data by local polynomial regression with bandwidth $h = 1$ d. The time series consisted of 250 daily counts of the number of adults $x(t)$, without any measurement error, for Eq. 3 with $B(x) = 10xe^{-x/500}$, $D(x) = 0.3x$, $\psi \equiv 1$, and $\tau = 18$, and environmental stochasticity given by Eq. 5. (b) The same population time series with sampling errors generated by the Poisson sampling model (see the Appendix) with capture probability $p = 0.02$. (c) Scatterplot of true vs. estimated gradient values for the time series in (a). (d) Scatter plot of true vs. estimated gradient values for the time series in (b).

simulated data sets. Constraining B and D to be non-negative greatly improved the accuracy of the fits, because otherwise the birth and death rates can become confounded; in the case studies described in *Results: Nicholson's blowflies* Lucilia cuprina and *Results: Lucilia sericata control populations*, the unconstrained estimates had lower death rates and birth rates and the birth rate became negative at high densities. The examples file in the software provided with this paper (see Supplement) includes a comparison of constrained and unconstrained fits to the Nicholson data.

The complexity of the fitted B and D functions depends on a pair of smoothing parameters (α_X, α_Y) (Ruppert and Carroll 1997). As either of the αs approach infinity, the fitted function in that variable reduces to a cubic polynomial. As α approaches zero, all available degrees of freedom are used without constraint, and the fitted curve becomes very wiggly. Following Ruppert and Carroll (1997), we selected the values of the smoothing parameters by minimizing the GCV (generalized cross validation) score for the fitted model (Wahba 1990). GCV is the nonparametric analog of

adjusting a regression model's sum of squared errors to correct for the model's degrees of freedom. For splines and similar models this is done using the model's effective number of parameters (as defined by Wahba 1990). Minimizing GCV as a function of (α_X, α_Y) is an objective and automatic way to select model complexity that approximates optimization of the model's ability to make predictions at (X, Y) values outside the data set used to fit the model (Wahba 1990). We used a modified GCV score, GCV2, in which the residual sum of squares is slightly overcorrected for model degrees of freedom. GCV2 creates a small bias towards simpler models but greatly reduces the chances of spuriously selecting an overly complex model (Nychka et al. 1992).

Because the rate equations are estimated by regression with adult density as the independent variable, sampling error in the population counts produces bias in the regression coefficients. For the model in Fig. 1, both the mean mortality rate and the peak birth rate are underestimated (Fig. 2a). The downward curvature in the fitted morality rate is also an artifact of mea-

surement-error bias, the corresponding distortion at low population density being ruled out by the lack of an intercept term. These biases can be greatly reduced (Fig. 2b) using the SIMEX procedure (Cook and Stefanski 1994, Carroll et al. 1995, Stefanski and Cook 1995, 1999). SIMEX is a simulation-based method (described in the Appendix) in which the independent variables are made progressively less accurate by adding additional simulated measurement errors. The effect of the increased measurement error on parameter estimates is then extrapolated back to zero error. SIMEX reduces systematic errors but it does not have any effect on the variability of estimates. In other runs of the fitting process (e.g., different runs of the model with sampling error), the post-SIMEX estimated rate equations are wrong in different ways, as illustrated in the following section.

RESULTS

Simulated data

Fig. 3 summarizes a simulation study of the accuracy of our proposed methods, for two models having the form of Eq. 3. The first model was the same as that in Fig. 1, with linear death rate D, i.e., a constant per capita death rate. The second had a quadratic death rate D, the per capita death rate increasing linearly with population density (from 0.21 d^{-1} at zero adults, to 0.39 d^{-1} at 10 000 adults). The birthrate for the second model was $B(x) = 12xe^{-x/1000}$, and the maturation time delay was $\tau = 25$ d. Environmental and demographic stochasticity were incorporated as in Fig. 1 and the first model, so again we are deliberately fitting a continuous model to "data" with discrete individuals. The standard deviation of the birthrate multiplier $\varphi(t)$ was again 0.1, but the autocorrelation was increased so as to have autocorrelation 0.5 at a time lag of 5 d. This autocorrelation was motivated by interpreting φ as random variation in the "quality" of adults maturing on a given day, having effects lasting over the lifetime of those adults and their offspring.

We generated long simulations of each model, and extracted from each 10 successive segments of 400 daily values. Rate functions were estimated for data consisting of either the first 200 daily counts from each segment, or 200 every-other-day counts. The latter case is closer to the case studies reported below (*Results: Nicholson's blowflies* Lucilia Cuprina, *Results:* Lucilia Sericata *control populations*), both of which involve every-other-day counts and a similar mean cycle period. In both cases, we applied the Poisson sampling model (with capture probability $p = 5\%$) to generate the time series of adult counts that were used to estimate the rate equation, and used SIMEX bias reduction with quadratic extrapolation (see the Appendix for details).

In general, the estimated rate functions (dashed lines in Fig. 3) recover the shape of the true functions. Some

of the estimated birth rate functions exhibit spurious small wiggles, especially for the data sets of 200 daily counts (Fig. 3b, e). The series with alternate-day sampling cover twice as many complete cycles (roughly eight, vs. roughly four with daily sampling) and this additional information appears to outweigh the small loss in the accuracy of gradient estimates due to sparser sampling. There is also a small bias towards underestimating the peak birth rate and the death rate (the mean per capita death rate, averaged over all estimates, is low by ~10% for each model). These biases are probably due to the small remaining bias in the gradient estimates at the highest and lowest gradient values. Because the bias correction applied to gradient estimates is based on a curve that is smoother than the (unknown) true population trajectory, the bias correction tends to be an undercorrection.

Nicholson's blowflies Lucilia cuprina

We analyzed the series labeled "I" in Nicholson (1957), which has been the target for several previous modeling efforts. This population was regulated by adult competition for food, which was observed to strongly limit fecundity. The single-species delay differential equation model in Eq. 1 therefore corresponds to the correct regulating mechanism, using a simplified description in which the egg-production rate reacts instantaneously to the current level of crowding. The full series consists of 350 adult counts taken every other day. Owing to the change in dynamics observed in the course of the experiment (Stokes et al. 1988), we used only the first 180 data points (Fig. 4a). The period of the cycles is roughly 38 d, giving ~19 data points per cycle.

The gradient was estimated using bandwidth $h = 2$, and the rate functions fitted by penalized regression splines as described in *Methods*. We used a time delay of $\tau = 14$ d based on Gurney's (1980) estimate that the value of τ for this experiment was in the interval $\tau = 14.8 \pm 0.4$ d. Changing τ to 12 d or 16 d had only minor effects on the fitted rate functions, and essentially no effect on the dynamics of the fitted model. It is possible to use delays other than multiples of the sampling interval by interpolating the time series (Ellner et al. 1997), but in this case the insensitivity to the precise value of τ argues for simply rounding 14.8 d to 14 d.

The fitted birth and death rate functions are plotted as the solid lines in Fig. 4b. The resulting fit has a per-capita death rate that decreases with adult density. This form of density-dependent mortality suggests that (as in the simulation studies) the downward curvature may be due to measurement error bias. In order to apply SIMEX bias reduction, we need a model of the measurement error process. Nicholson described the data as "a careful day-to-day record of the numbers of individuals" (Nicholson 1957:155), which suggests low-to-nonexistent sampling errors. However, even under
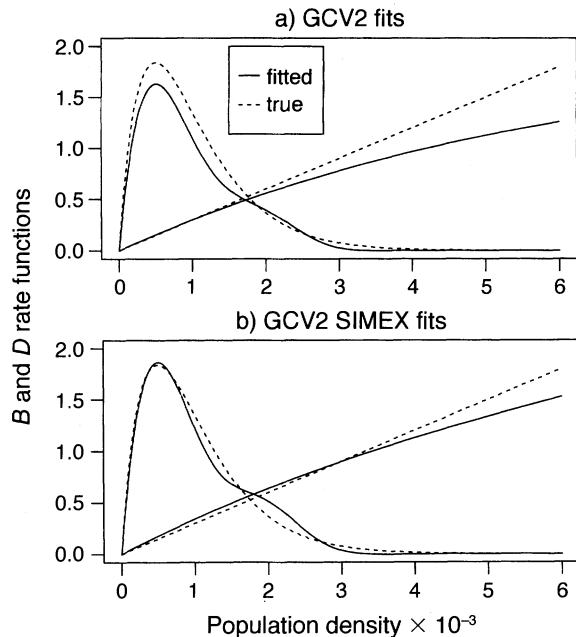
a) GCV2 fits

FIG. 2. An example of rate equation estimates from a time series generated by the model in Fig. 1. The data were 200 daily adult counts with simulated Poisson sampling errors and capture probability $p = 0.02$. The birth and death rate functions were fitted as penalized cubic regression splines, constrained to be positive at a grid of values covering the range of the data. We used 20 evenly spaced knots covering the range of the data for both the birth rate function B and the death rate function D. The one-hump curve falling to zero at high population density is the total birth rate $B(N)$, and the monotonically increasing curve is the total death rate $D(N)$. The top panel (a) and bottom panel (b), respectively, show the fits before and after SIMEX bias reduction with quadratic extrapolation. Both rate functions are plotted in units of thousands of individuals per day.

extreme assumptions about the sampling process, the curvature persists after SIMEX bias reduction. The dashed lines in Fig. 4b are the result of SIMEX bias reduction on the assumption that errors are due to sampling with capture probability $p = 0.1$; estimates assuming $p = 0.05$ are nearly identical. Very similar fits were also obtained on the assumption of lognormal sampling errors, with coefficient of variation up to $\sim 15\%$.

When the fitted model is simulated with noise, it generates dynamics very similar to those in the data. The simulation in Fig. 4c included both environmental and demographic stochasticity, as in Eq. 5. Without any noise, the model converges rapidly to a limit cycle at lower amplitude than the data, very close to the transition between single-peaked and double-peaked cycles (Fig. 4d). With increasing environmental noise (i.e., the variance of the birth rate multiplier $\varphi(t)$; see Eq. 5), the cycle amplitude can be increased to match that observed in the data. A secondary effect of environmental noise is to produce a mix of single-peak-

ed and double-peaked cycles in model trajectories, such as are observed in the data. The period of the fitted model with noise is $\sim 5\%$ longer than that of the data, producing on average ~ 8.5 cycles in the time period over which the experimental population had 9 cycles. We emphasize that the model-fitting criteria do not involve any attempt to impose a match between model trajectories and the data series, so the comparison between Fig. 4a and 4c is a demanding test of the fitted model.

It is also reassuring that the estimated rates are similar to those derived in independent analyses. Based on Nicholson's data, Kendall et al. (1999) estimated the adult mortality rate to lie between 0.08 d^{-1} and 0.14 d^{-1}; the estimated per capita mortality rate in this paper's model, averaged over the experimental time series, is 0.13 d^{-1}. Fecundity was estimated by Readshaw and Cuff (1980), using experimental data collected by Nicholson, to fall to 0 at 0.14 g·adult^{-1}·d^{-1} of protein food. For the experiment analyzed here, that rate of per-capita food supply occurs at a density of ~ 2850 adults, which is roughly where the estimated fecundity rate B drops to 0. The density at which the total egg-laying rate is maximized is less certain, with estimates (all based on Nicholson's data) ranging from 171 (Readshaw and Cuff 1980) to 600 (Gurney at al. 1980). The estimated birth rate function in Fig. 4 is maximized at an adult density of ~ 750.

The nonparametric rate functions are not quantitatively very different from the parametric forms used by Gurney et al. (1980), when fitted to the same data (Fig. 4b, dot-dash curves). However, the nonparametric birth rate function has a more precipitous decrease in egg production as density increases, which is more in line with Readshaw and Cuff's estimate of the minimum food density for egg production. The curvature in the nonparametric death rate function may be a result of systematic variation in the age structure of the adult population. Fecundity and mortality in $L.$ $cuprina$ were both found to be age-dependent in subsequent experiments (Readshaw and van Gerwen 1983). The estimated rate functions therefore could be affected by a systematic relationship between the density and age structure of the adult class, and such a relationship is likely. Nicholson (1954:21, legend to Fig. 3) observed that "adults are mostly senile as the adult minima are approached." More generally, suppose (for the sake of calculation) that the adult mortality rate is constant at some value μ. Then it can be derived (by the methods in Gurney and Nisbet 1985) that the mean age of living adults (time since the end of the pupal stage) is

$$\bar{a}(t) = \int_0^t e^{-\mu s} N(t - s) \, ds / N(t) \qquad (7)$$

where $N(t)$ is the adult population at time t. This formula applies for t large enough that no founding

a) Linear death rate

d) Quadratic death rate

b) GCV2 SIMEX fits, $T_s = 1$

e) GCV2 SIMEX fits, $T_s = 1$

c) GCV2 SIMEX fits, $T_s = 2$
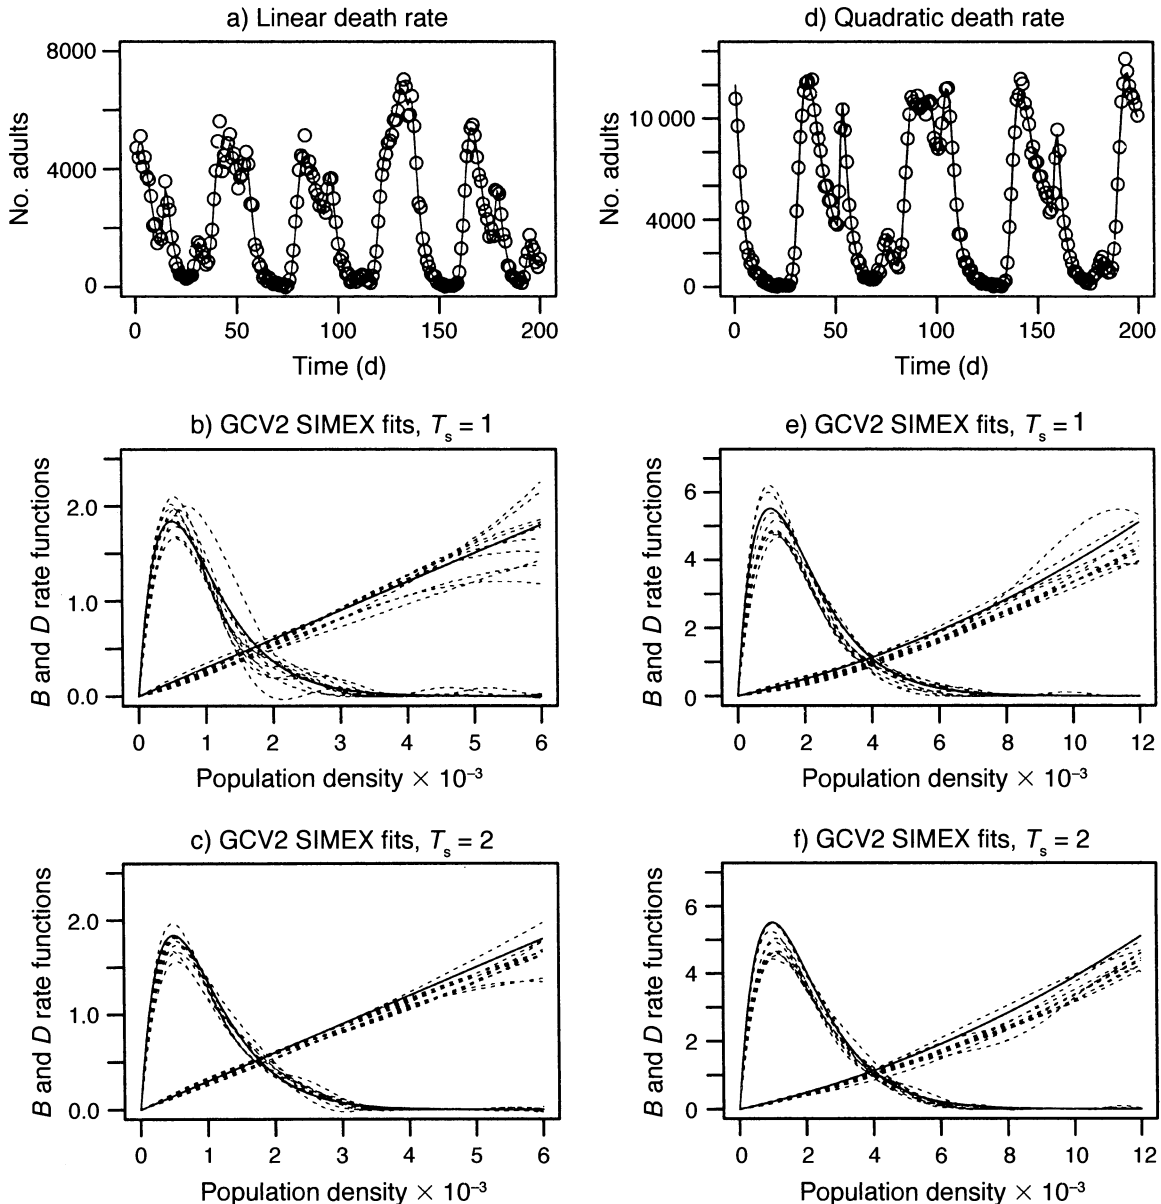
f) GCV2 SIMEX fits, $T_s = 2$

FIG. 3. Simulation study of rate equation estimation for the model described in Eq. 1 with SIMEX bias reduction. The linear death rate model (left panels) is the model in Fig. 1. The quadratic death rate model has $B(x) = 12xe^{-x/1000}$, and per capita death rate $D(x)/x$ increasing linearly from 0.21 d^{-1} at zero adults to 0.39 d^{-1} at 10 000 adults. Both models were simulated with demographic and environmental stochasticity, as described in the legend of Fig. 1. Top panels (a, d) show a typical time series, consisting of 200 daily counts of adult density, from each of the models. The middle panels (b, e) show the estimated rate equations from 10 replicate stretches of 200 daily counts (i.e., sampling time $T_s = 1$ d). The bottom panels (c, f) show the estimated rate equations from 10 replicate stretches of 200 every-other-day counts (sampling time $T_s = 2$ d).

adults are still alive. In the present case, this yields a close inverse relationship between population density and mean age of adults: the adult population is young at peaks, and old at troughs. For example, computing Eq. 7 with $\mu = 0.10$ at all sampling times $t > 40$ d after the start of Nicholson's experiment, we obtain $\ln(\bar{a}) = 8.4 - 0.78\ln(N)$ ($r^2 = 0.85$, $P < 0.001$). Thus the downward curvature in the esti-

mated mortality rate may be an accurate reflection of the fact that adults at low density are mostly older flies whose mortality rate is higher (Readshaw and van Gerwen 1983).

Lucilia sericata control populations

These experiments, described by Smith et al. (2000), involved a different blowfly species and a different
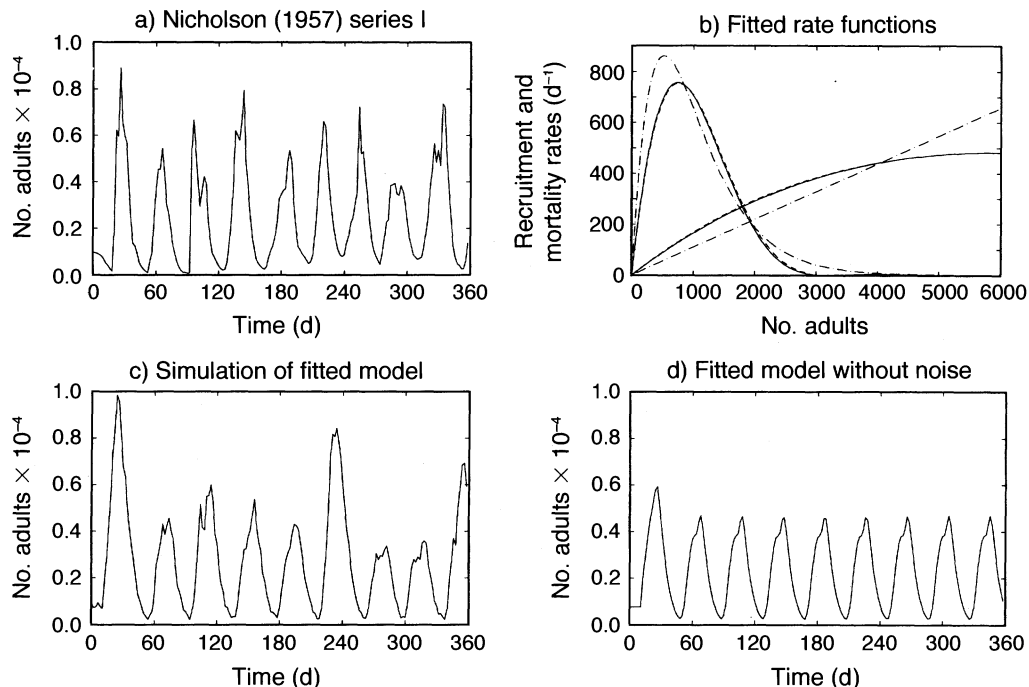
FIG. 4. Model described in Eq. 1 fitted to alternate-day *L. cuprina* adult population counts from Nicholson (1957), series "I": (a) the time series analyzed, consisting of the first 180 adult counts (360 d); (b) the estimated rate equations. The solid and dashed curves are, respectively, the nonparametric regression spline estimates before and after SIMEX bias reduction assuming the Poisson sampling model with $p = 0.1$. The dot-dash curves are the parametric model of Gurney et al. (1980) $B(x) = Rxe^{-x/A}$, $D(x) = dx$, fitted by least squares to the same estimated gradients as the regression splines. (c) Simulation of the fitted regression spline model with demographic and environmental stochasticity. The environmental stochasticity was a random birthrate multiplier as in Eq. 3, lognormally distributed with $\sigma = 0.35$. The level of environmental stochasticity was chosen to match roughly the amplitude of cycles seen in the data. Because of the large population sizes, demographic stochasticity alone has very little effect on the dynamics. (d) Simulation of the fitted regression spline model with demographic stochasticity removed, and no environmental stochasticity.

mechanism of population regulation: limited food supply for larvae. The data are alternate-day counts of larvae, pupae, and adults. The adult population was tracked by adding up the number of empty pupal cases and subtracting the number of dead or escaped flies. However, complete counts taken approximately every 3 mo when flies were transferred to clean cages showed that these indirect adult population counts were off by at most 5% (Smith et al. 2000). The data were therefore analyzed on the assumption that population counts are exact.

With larval competition as the regulating mechanism, Eq. 1 can only be justified mechanistically if larval competition is "within cohort," meaning that the egg-to-adult survival of an egg laid at time t is affected only by the density of eggs laid in a narrow time window ($t - \varepsilon$, $t + \varepsilon$), where ε is "small." The meaning of "small" here is that larval density must remain effectively constant over the time period when a given larva is affected by larval density. In *L. sericata* the ~2-d duration of larval feeding implies $\varepsilon = 2$ d, the time between successive counts. In these experiments larval density often changed markedly between successive counts: the ratio between the larger and smaller of successive larval counts (when both were nonzero) was 1.5 or higher 51% of the time, and 2.0 or higher 29% of the time. An assumption of within-cohort larval competition would be "close but not quite," so our expectation is that an attempt to fit the model described in Eq. 1 to these experiments should not be totally successful.

For all analyses, we pooled the data from six control cages that differed only in initial densities (six other cages, not analyzed here, had a mildly toxic diet). All six populations settled into a similar pattern of fluctuations (see Fig. 5a), with apparent cycles of period ~70 d (Smith et al. 2000). As in Nicholson's populations, there may have been some life history evolution during the experiments (Lingjærde et al. 2001), with some replicates appearing to change dynamics at roughly 400 d (Smith et al. 2000). We therefore used only the first 200 samples from each cage, and we again only used the adult counts to mimic the typical situation with field data where only one life stage is counted. Because the populations were established from a small number of pupae and adults, the first 20 samples (40
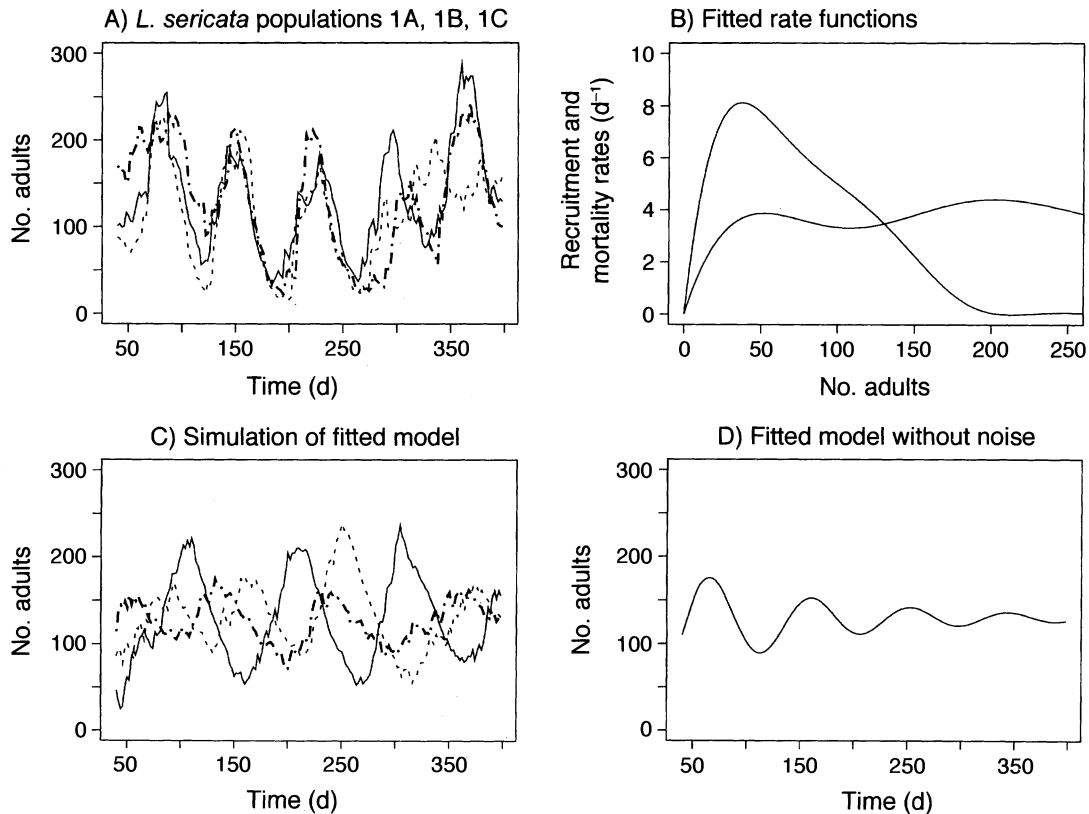
FIG. 5. Model described in Eq. 1 fitted to alternate-day adult counts from *L. sericata* control cages (Smith et al. 2000): (a) three of the six time series analyzed, each consisting of adult counts 21–200 (360 d); (b) the estimated regression spline rate equations; (c) simulations of the fitted regression spline model with demographic stochasticity, including "on/off" daily reproduction; (d) simulation of the fitted regression spline model with demographic stochasticity removed and no environmental stochasticity.

d) were discarded as possibly representing an initial transient.

The rate equations obtained from fitting the model described in Eq. 1 to the *L. sericata* data are shown in Fig. 5b. In simulating the model (Fig. 5c), we incorporated Forrest's (1996) observation that on 55% of days there was no egg laying whatsoever. Following Forrest (1996), we therefore modeled recruitment as randomly being "on" or "off" each day. On days occur with probability 0.45, and the birthrate for those days was set at (1/0.45) times the fitted birth rate function; on off days the birth rate was set to 0. This high level of demographic stochasticity caused the fitted model to exhibit periods of oscillatory dynamics (Fig. 5c), whereas in the absence of any demographic or environmental stochasticity it converges rapidly to a stable equilibrium density (Fig. 5d). The dynamics with demographic stochasticity (Fig. 5c) bear some resemblance to the data, but the model is less prone to exhibit oscillations with amplitude comparable to those in the data. When larger amplitude cycles occur (e.g., the solid curve in Fig. 5c), their shape is visibly different from the data and their duration is too long. Larger amplitude oscillations can be made more prevalent by adding some environmental stochasticity as in Eq. 3, but even then the period of oscillations remains too long by ~20%.

Another sign of trouble is the estimated mortality rate function. The adult longevity of 67 flies sampled between days 244 and 316 of the time series averaged 18.2 d (Daniels 1994), corresponding to a death rate of 0.055 d^{-1}, whereas the mean adult death rate in the fitted model is lower by ~20% (0.044 d^{-1}). In addition, an analysis of the full data set (including all life stages and age structure within stages) concluded that adult survival was density independent (Lingjærde et al. 2001), whereas the estimate here has mortality rate decreasing with density.

DISCUSSION

The option of fitting rate equations nonparametrically allows a modeler to incorporate reliably known aspects of the system (including parametric rate equations when these are justified), but to let the data "speak for themselves" about aspects that are less well known. These models have been called "semimechanistic"

(Ellner et al. 1998, Kendall et al. 1999, Smith et al. 2000) because they are intermediate between conventional mechanistic models with parametric rate equations, and phenomenological statistical models. Wood (1999) calls these "partially specified" models, and from a statistical perspective they could be described as "semiparametric" models (Wood 1997). As Perry (2000) notes, the goal of modelers has always been to exploit all available information while making up as little as possible. The semimechanistic approach carries this philosophy one step further, by embedding modern statistical tools for nonparametric function estimation into state-variable models for ecological system dynamics. This paper provides a method for fitting semimechanistic models to time-series data under certain assumptions, and has illustrated one important application: deciding whether a model structure is appropriate for a given data set, rather than testing a particular fully specified model.

A number of recent case studies illustrate the potential effectiveness of semimechanistic modeling. Wood (1994; see also Ohman and Wood 1996) used this approach to obtain more reliable estimates of time-dependent recruitment and mortality rates in structured populations. Ellner et al. (1998) showed that a semimechanistic model for measles epidemics had better forecasting accuracy than either a fully specified epidemic model or a phenomenological time-series model fitted to the same data, and produced a characterization of the dynamics as weakly stable but near the "edge of chaos." Smith et al. (2000) conducted a similar comparison for laboratory populations of blowflies *Lucilia sericata*. The semimechanistic model was stage structured, with the stages and their durations assumed a priori based on the known fly biology, but all density dependence was fitted nonparametrically. This model was far more accurate than phenomenological time-series models, which failed to replicate the observed population cycles. It also yielded the same qualitative conclusions about the dynamics as the mechanistic model, without the numerous additional experiments that had been necessary to derive rate equations for the mechanistic model. Lingjærde et al. (2001) extended the Smith et al. (2000) analyses to incorporate age structure within stages, using nonparametric generalized additive models for the rate equations. Using this approach they were able to identify stage-specific demographic changes that occurred over the course of the experiments, and to pinpoint the specific demographic differences between the control populations and those fed a diet contaminated with cadmium.

The main obstacle to semimechanistic modeling is a lack of accessible and general fitting methods. Fitting is straightforward only in the exceptional circumstance of data being available directly on each process (e.g., direct counts of births and deaths over time rather than of changes in total population size); each individual process-rate equation can then be estimated separately by nonparametric regression using commercial software. Otherwise, three other general approaches are currently available: gradient matching (described in this paper), trajectory matching, and moment matching (a method using the Kalman filter to approximate the likelihood has recently been proposed [de Valpine and Hastings 2002] but so far it has only been tested on simple one-parameter models). Trajectory matching (sometimes called total least squares, Ives et al. 1999) fits the model by minimizing a measure of total difference (such as mean-square error) between the data and a deterministic simulation of the model (e.g., Wood 1999, 2001). Moment matching refers to methods in which stochastic simulations of the model are compared with the data using either a set of "moments" (statistical measures such as the coefficients of a fitted autoregressive model) or a single summary measure such as a quasilikelihood (Gouriéroux and Montfort 1996).

Gradient matching is the most direct method, but it requires data on all state variables for which a rate function is being estimated and the quality and frequency of the data must be adequate for estimating time derivatives. When these assumptions hold, gradient matching offers the advantages of computational efficiency and relatively simple implementation (as in the software provided in the Supplement for the methods in this paper). One never really has data on each and every relevant variable, but the ones in hand may nonetheless be sufficient to represent the main processes driving system dynamics. For example, pretending that a few key variables (or even one) can capture all variation between individuals often yields informative population dynamic models from limited data (Caswell 2000). Although it has been popular recently to use time-delay embeddings as a substitute for unobserved state variables, time-delay variables have no mechanistic interpretation, and increasing the dimension of the state space also increases the amount of data needed for model fitting and selection. Thus, if a gradient-matching approach based on a candidate set of state variables fails, it would be more productive to switch to a method that can cope with unobserved variables.

Trajectory matching also offers computational efficiency, but under different assumptions. The key assumption for trajectory matching is that there are no important differences between deterministic and stochastic simulations of the model. In some cases this will be true, for example, if the deterministic dynamics are a limit cycle and random perturbations do not cause the cycles to drift in phase. The Nicholson experiment considered above appears to be one such case (Kendall et al. 1999, Wood 2001). In other cases there may be important differences, such as the *L. sericata* populations considered here. The more successful models for these experiments (the stage-structured mechanistic and semimechanistic models in Smith et al. 2000) have

the property that deterministic simulations exhibit damped oscillations to a steady state, while stochastic simulations have persistent quasicycles with a well-defined dominant period. In order to fit those data with a deterministic model, parameter estimates would have to be distorted to make the deterministic model have limit cycles. Similarly, fits of deterministic epidemic models to measles data favor parameters in the chaotic regime (Olsen and Schaffer 1990, Tidd et al. 1993) because these allow the model to mimic the highly variable severity of disease outbreaks, but nonchaotic models incorporating exogenous random variability in birth rates fitted the data better and had better forecasting accuracy (Ellner et al. 1998). Thus trajectory-matching methods should be used with caution when system dynamics are potentially affected by environmental stochasticity. However, when its assumptions are satisfied, trajectory matching is efficient enough that rate equations can be estimated nonparametrically even for complex models with unmeasured state variables (Wood 1997, 1999, 2001).

Thus, both gradient matching and trajectory matching achieve computational efficiency at the price of substantive assumptions about the data and the system being modeled. Moment-matching approaches offer statistically sound procedures for estimating parameters of any model, deterministic or stochastic, that can be simulated on the computer. Observations can be widely separated in time and subject to measurement errors, and it is not necessary for all state variables to be observed. However, statistically efficient moment-matching methods are computationally very demanding, because each evaluation of the fitting criterion involves computing the moments from a long simulation of the model. For example Turchin and Ellner (2000b) used a moment-matching approach to estimate parameters for a model of Fennoscandian vole oscillations. For their simple predator–prey model with six parameters to estimate, fitting to a single short time series (81 observations) took several hours on a then-current UNIX workstation, and bootstrapping to obtain confidence intervals took several days. Moment matching is therefore limited to parametric models with few enough parameters that iterative optimization of the fitting criterion is feasible.

Nicholson's (1957) experiment data have been analyzed extensively in order to understand quantitatively the cause of the steady, large amplitude oscillations. In 1980, two different parametric versions of the model described in Eq. 1 were proposed independently for these data (Gurney et al. 1980, Readshaw and Cuff 1980). Subsequently, one of these was used as the basis for examining the hypotheses that evolution of life history traits produced the changes in dynamics between the first and second half of the data series (Stokes et al. 1988). Elaborations of this basic model are now coming into regular use for stage-structured insects and

other animal populations (Tuljapurkar and Caswell 1997, Gurney and Nisbet 1998).

For Nicholson's (1957) experiment "I" on *Lucilia cuprina*, our results support the validity of the model described in Eq. 1 for the mean dynamics. The estimated rate functions are qualitatively similar to the parametric model of Gurney et al. (1980). In contrast, Eq. 1 is far less successful for the *L. sericata* experiments (Daniels 1994, Forrest 1996, Smith et al. 2000). The best-fit model produces cycle periods not too dissimilar to those observed (20% too long), but it did not match well the shape or range of the cycles in adult density. These contrasting findings are consistent with the expectations described above based on the mechanisms of population regulation in the two experiments: competition between adults in *L. cuprina*, and competition between feeding larvae in *L. sericata*.

Of course, even for Nicholson's experiments, the model described in Eq. 1 is an oversimplification. In particular, it neglects the age dependence of vital rates within the adult class, which is known to occur (Readshaw and van Gerwen 1983). Our successful fit to the data supports the hypothesis that the adult competition mechanism embodied in the model is the primary factor in the observed oscillations. Apparently, it is not too much of an error either to ignore age dependence in adults or (as the fitted model may be doing) to account for age dependence implicitly via rate equations that reflect the average relationship between adult density and age structure.

Our topic in this paper has been fitting models to data, and issues of model testing were mostly treated informally using qualitative comparisons with the data. Given that all models require simplifying assumptions and cannot be exactly valid, "testing" a model's mechanistic validity typically means comparing it against alternatives based on different mechanistic assumptions (Hilborn and Mangel 1997). One aspect of testing is to quantify the prediction error variance of the fitted model, and our procedures are quick enough that this can be done by cross-validation. However, goodness-of-fit comparisons of this sort are often less informative than simulating the model and comparing model output with the data. Considered strictly as regression models for predicting the expected rate of population change, the models for *L. cuprina* and *L. sericata* are more or less equally successful: the correlation coefficients between observed and predicted gradients are not too different ($\rho = 0.55$ for *L. cuprina*, 0.37 for *L. sericata*), and there are no systematic patterns in the residuals between observed and predicted rates of change in either case (Fig. 6a, b). But simulating both models and comparing predicted trajectories (Figs. 4c and 5c) and autocorrelation functions (Fig. 6c, d) with the corresponding data, the *L. sericata* model is clearly less successful. Other methods for comparing population dynamics models are available (see Tidd et al. 1993, Dennis and Taper 1994, Kendall et al. 1999, and Tur-

a) *L. cuprina*: residual plot

b) *L. sericata*: residual plot

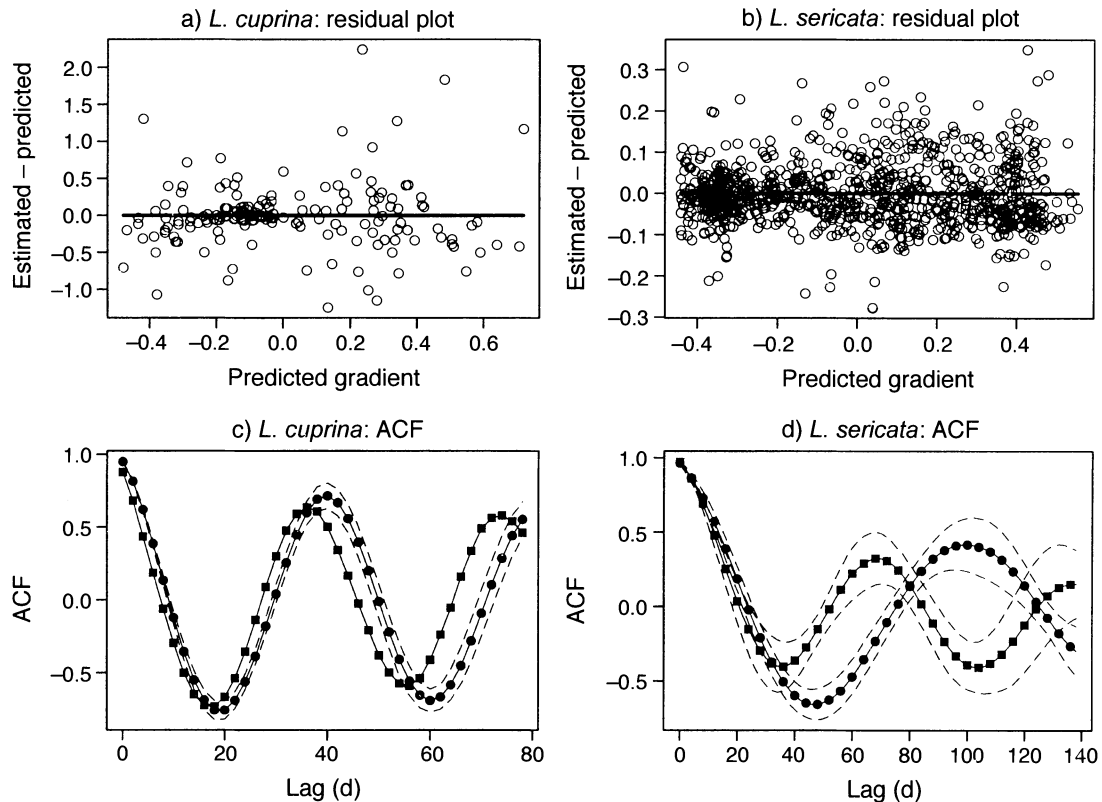### c) *L. cuprina*: ACF

d) *L. sericata*: ACF

FIG. 6.   Graphical tests of the fitted models. (a, b) Residual plots for the *L. cuprina* and *L. sericata* models, respectively, showing the differences between estimated gradients $dx(t_i)/dt$ and the predicted values from the models, $B(x(t_i - \tau)) - D(x(t_i))$, where B and D are the fitted birth and death rate functions. (c, d) Autocorrelation functions (ACF's) for the fitted models, compared with those of the data. Each ACF is plotted out to approximately twice the estimated cycle period. Solid circles are the mean ACF from 10 replicate model simulations (corresponding to Figs. 4c and 5c) equal in length to the actual data series, with dashed lines showing ± 2 SE over the 10 replicates. The solid squares are the ACF of the data. For *L. sericata* the ACF of the data is the mean ACF over the six replicates analyzed, with dashed lines showing ± 2 pointwise standard errors over the replicates. Nicholson's experiment did not include replication, so no standard errors are shown.

chin and Ellner 2000*b* for applications), which make it possible to test quantitatively which of two mechanistic models is better able to account for observed population trajectories.

LITERATURE CITED

Briggs, C. J., W. W. Murdoch, and R. M. Nisbet. 1999. Recent developments in theory for biological control of insect pests by parasitoids. Pages 22–42 *in* B. A. Hawkins and H. V. Cornell, editors. Theoretical approaches to biological control. Cambridge University Press, Cambridge, UK.

Carroll, R. J., J. D. Maca, and D. Ruppert. 1999. Nonparametric estimation in the presence of measurement errors. Biometrika 86:541–554.

Carroll, R. J., D. Ruppert, and L. W. Stefanski. 1995. Measurement error in nonlinear models. Chapman and Hall, New York, New York, USA.

Caswell, H. 2000. Matrix population models. Second edition. Sinauer, Sunderland, Massachusetts, USA.

Cook, J. R., and L. A. Stefanski. 1994. Simulation–extrapolation estimation in parametric measurement error models. Journal of the American Statistical Association 89:1314–1328.

Daniels, S. 1994. Effects of cadmium toxicity on population dynamics of the blowfly *Lucilia sericata*. Dissertation. University of Reading, Reading, UK.

Dennis, B., and M. Taper. 1994. Density dependence in time series observations of natural populations: estimation and testing. Ecological Monographs **64**:205–224.

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

Eilers, P. H. C., and B. D. Marx. 1996. Flexible smoothing with B-splines and penalties (with discussion). Statistical Sciences **11**:89–121.

Ellner, S. P., B. A. Bailey, G. V. Bobashev, A. R. Gallant, B. T. Grenfell, and D. W. Nychka. 1998. Noise and nonlinearity in measles epidemics: combining mechanistic and statistical approaches to population modeling. American Naturalist **151**:425–440.

Ellner, S. P., B. E. Kendall, S. N. Wood, E. McCauley, C. J. Briggs. 1997. Inferring mechanism from time-series data: delay-differential equations. Physica D **100**:182–194.

Fan, J., and I. Gijbels. 1996. Local polynomial modeling and its applications. Chapman and Hall, New York, New York, USA.

Forrest, B. 1996. Toxins and blowfly population dynamics. Dissertation. University of Leicester, Leicester, UK.

Gouriéroux, C., and A. Montfort. 1996. Simulation-based econometric inference. Oxford University Press, Oxford, UK.

Gurney, W. S. C., S. P. Blythe, and R. M. Nisbet. 1980. Nicholson's blowflies revisited. Nature **287**:17–21.

Gurney, W. S. C., and R. M. Nisbet. 1985. Fluctuation periodicity, generation separation, and the expression of larval competition. Theoretical Population Biology **28**:150–180.

Gurney, W. S. C., and R. M. Nisbet. 1998. Ecological dynamics. Oxford University Press, Oxford, UK.

Haefner, J. W. 1996. Modeling biological systems: principles and applications. Chapman and Hall, New York, New York, USA.

Hastings, A. M. 1997. Population biology: concepts and models. Springer-Verlag, New York, New York, USA.

Hilborn, R., and M. Mangel. 1997. The ecological detective: confronting models with data. Princeton University Press, Princeton, New Jersey, USA.

Hudson, P. J., A. P. Dobson, and D. Newborn. 1998. Prevention of population cycles by parasite removal. Science **282**:2256–2258.

Ihaka, R., and R. Gentleman. 1996. R: a language for data analysis and graphics. Journal of Computational and Graphical Statistics **5**:299–314.

Ives, A. R., S. R. Carpenter, and B. Dennis. 1999. Community interaction webs and zooplankton responses to planktivory manipulations. Ecology **80**:1405–1421.

Kendall, B. E., C. J. Briggs, W. W. Murdoch, P. Turchin, S. P. Ellner, E. McCauley, R. Nisbet, and S. N. Wood. 1999. Why do populations cycle? A synthesis of statistical and mechanistic modeling approaches. Ecology **80**:1789–1805.

Korpimäki, E., and K. Norrdahl. 1998. Experimental reduction of predators reverses the crash phase of small-rodent cycles. Ecology **79**:2448–2455.

Laska, M. S., and J. T. Wootton. 1998. Theoretical concepts and empirical approaches to measuring interaction strength. Ecology **79**:461–476.

Lingjærde, O. C., N. C. Stenseth, A. B. Kristoffersen, R. H. Smith, S. J. Moe, J. M. Read, S. Daniels, and K. Simkiss. 2001. Exploring non-linearities in the stage-specific density-dependent structure of experimental blowfly populations using non-parametric additive modeling. Ecology **82**:2645–2658.

McCauley, E., R. M. Nisbet, A. M. DeRoos, W. W. Murdoch, and W. S. C. Gurney. 1996. Structured population models of herbivorous zooplankton. Ecological Monographs **66**:479–501.

Murdoch, W. W. 1994. Population regulation in theory and practice. Ecology **75**:271–287.

Murdoch, W. W., and C. J. Briggs. 1996. Theory for biological control: recent developments. Ecology **77**:2001–2013.

Nicholson, A. J. 1954. An outline of the dynamics of animal populations. Australian Journal of Zoology **2**:9–65.

Nicholson, A. J. 1957. The self-adjustment of populations to change. Cold Spring Harbor Symposia on Quantitative Biology **22**:153–173.

Nychka, D. W., S. Ellner, A. R. Gallant, and D. McCaffrey. 1992. Finding chaos in noisy systems (with discussion). Journal of the Royal Statistical Society Series B **54**:399–426.

Ohman, M. D., and S. N. Wood. 1996. Mortality estimates for planktonic copepods: *Pseudocalanus newmani* in a temperate fjord. Limnology and Oceanography **41**:126–135.

Oksendal, B. 1998. Stochastic differential equations: an introduction with applications. Fifth edition. Springer-Verlag, New York, New York, USA.

Olsen, L. F., and W. M. Schaffer. 1990. Chaos versus noisy periodicity: alternative hypotheses for childhood epidemics. Science **249**:499–504.

Perry, J. N. 2000. Overview. Pages 173–190 *in* J. N. Perry, R. H. Smith, I. P. Woiwod, and D. Morse, editors. Chaos in real data: the analysis of non-linear dynamics from short ecological time series. Kluwer Academic, Dordrecht, The Netherlands.

Pfister, C. 1995. Estimating competition coefficients from census data: a test with field manipulations of tidepool fishes. American Naturalist **146**:271–291.

Readshaw, J. L., and W. R. Cuff. 1980. A model of Nicholson's blowfly cycles and its relevance to predation theory. Journal of Animal Ecology **49**:1005–1010.

Readshaw, J. L., and A. C. M. van Gerwen. 1983. Age-specific survival, fecundity and fertility of the adult blowfly, *Lucilia cuprina*, in relation to crowding, protein food, and population cycles. Journal of Animal Ecology **52**:879–887.

Ruppert, D., and R. J. Carroll. 1997. Penalized regression splines. Technical Report TR1249. Department of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York, USA.

Ruppert, D., and R. J. Carroll. 2000. Spatially adaptive penalties for spline fitting. Australian and New Zealand Journal of Statistics **42**:205–223.

Smith, R. H., S. Daniels, K. Simkiss, E. D. Bell, S. P. Ellner, and B. Forrest. 2000. Blowflies as a case study in nonlinear population dynamics. Pages 137–172 *in* J. N. Perry, R. H. Smith, I. P. Woiwod, and D. Morse, editors. Chaos in real data: the analysis of non-linear dynamics from short ecological time series. Kluwer Academic, Dordrecht, The Netherlands.

Stefanski, L. A., and J. R. Cook. 1995. Simulation-extrapolation: the measurement error jackknife. Journal of the American Statistical Association **90**:1247–1256.

Stokes, T. K., W. S. C. Gurney, R. M. Nisbet, and S. P. Blythe. 1988. Parameter evolution in a laboratory insect population. Theoretical Population Biology **34**:248–265.

Tidd, C. W., L. F. Olsen, and W. M. Schaffer. 1993. The case for chaos in childhood epidemics. II. Predicting historical epidemics from mathematical models. Proceedings of the Royal Society of London B **254**:257–273.

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

Turchin, P., and S. P. Ellner. 2000*a*. Modelling time-series data. Pages 33–48 *in* J. N. Perry, R. H. Smith, I. P. Woiwod, and D. Morse, editors. Chaos in real data: the analysis of non-linear dynamics from short ecological time series. Kluwer Academic, Dordrecht, The Netherlands.

Turchin, P., and S. P. Ellner. 2000*b*. Living on the edge of chaos: population dynamics of Fennoscandian voles. Ecology **81**:3099–3116.

Wahba, G. 1990. Spline models for observational data. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, USA.

Wood, S. N. 1994. Obtaining birth and mortality patterns from structured population trajectories. Ecological Monographs **64**:23–44.

Wood, S. N. 1997. Inverse problems and structured-population dynamics. Pages 555–586 *in* S. Tuljapurkar and H. Caswell, editors. 1997. Structured-population models in marine, terrestrial, and freshwater systems. Chapman and Hall, New York, New York, USA.

Wood, S. N. 1999. Semi-parametric population models. Pages 41–50 *in* Challenges in applied population biology. Aspects of applied biology. Volume 53. Association of Applied Biologists, Warwick, UK.

Wood, S. N. 2001. Partially specified ecological models. Ecological Monographs **71**:1–25.

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

APPENDIX

Details of the methods used to estimate models are available in ESA's Electronic Data Archive: *Ecological Archives* E083-042-A1.

SUPPLEMENT

A documented set of R functions for all steps in the model-fitting process is available in ESA's Electronic Data Archive: *Ecological Archives* E083-042-S1.