CONCEPTS & SYNTHESIS

EMPHASIZING NEW IDEAS TO STIMULATE RESEARCH IN ECOLOGY

WHY DO POPULATIONS CYCLE? A SYNTHESIS OF STATISTICAL AND MECHANISTIC MODELING APPROACHES

Bruce E. Kendall,[1,7] Cheryl J. Briggs,[2,8] William W. Murdoch,[1,2] Peter Turchin,[3]
Stephen P. Ellner,[4] Edward McCauley,[5] Roger M. Nisbet,[2] and Simon N. Wood[6]

[1]*National Center for Ecological Analysis and Synthesis, University of California,
Santa Barbara, California 93106 USA*
[2]*Department of Ecology, Evolution and Marine Biology, University of California,
Santa Barbara, California 93106 USA*
[3]*Department of Ecology and Evolutionary Biology, University of Connecticut,
Storrs, Connecticut 06268 USA*
[4]*Biomathematics Graduate Program, Department of Statistics, North Carolina State University,
Raleigh, North Carolina 27695-8203 USA*
[5]*Ecology Division, Department of Biological Sciences, University of Calgary, Calgary,
Alberta, Canada T2N 1N4*
[6]*Mathematical Institute, University of St. Andrews, St. Andrews, Fife KY16 9SS, UK*

Abstract. Population cycles have long fascinated ecologists. Even in the most-studied populations, however, scientists continue to dispute the relative importance of various potential causes of the cycles. Over the past three decades, theoretical ecologists have cataloged a large number of mechanisms that are capable of generating cycles in population models. At the same time, statisticians have developed new techniques both for characterizing time series and for fitting population models to time-series data. Both disciplines are now sufficiently advanced that great gains in understanding can be made by synthesizing these complementary, and heretofore mostly independent, quantitative approaches. In this paper we demonstrate how to apply this synthesis to the problem of population cycles, using both long-term population time series and the often-rich observational and experimental data on the ecology of the species in question. We quantify hypotheses by writing mathematical models that embody the interactions and forces that might cause cycles. Some hypotheses can be rejected out of hand, as being unable to generate even qualitatively appropriate dynamics. We finish quantifying the remaining hypotheses by estimating parameters, both from independent experiments and from fitting the models to the time-series data using modern statistical techniques. Finally, we compare simulated time series generated by the models to the observed time series, using a variety of statistical descriptors, which we refer to collectively as ''probes.'' The model most similar to the data, as measured by these probes, is considered to be the most likely candidate to represent the mechanism underlying the population cycles. We illustrate this approach by analyzing one of Nicholson's blowfly populations, in which we know the ''true'' governing mechanism. Our analysis, which uses only a subset of the information available about the population, uncovers the correct answer, suggesting that this synthetic approach might be successfully applied to field populations as well.

Key words: blowflies; density dependence; fitting mechanistic models to time-series data; inverse problem; Lucilia cuprina; *mechanistic population models; modeling population cycles; population cycles, determining the causes of; statistical time-series analysis; structured population models; time-series models.*

INTRODUCTION

The population densities of many species can fluctuate nearly periodically over time, with periods that cannot be explained simply by seasonal variation (Fig.

1). These regular, large-amplitude oscillations invite explanation and indeed these data sets and others like them have fascinated generations of ecologists. How

[7] Present address: Donald Bren School of Environmental Science and Management, University of California, Santa Barbara, California 93106 USA.

[8] Present address: Department of Integrative Biology, University of California, Berkeley, California 94720 USA.
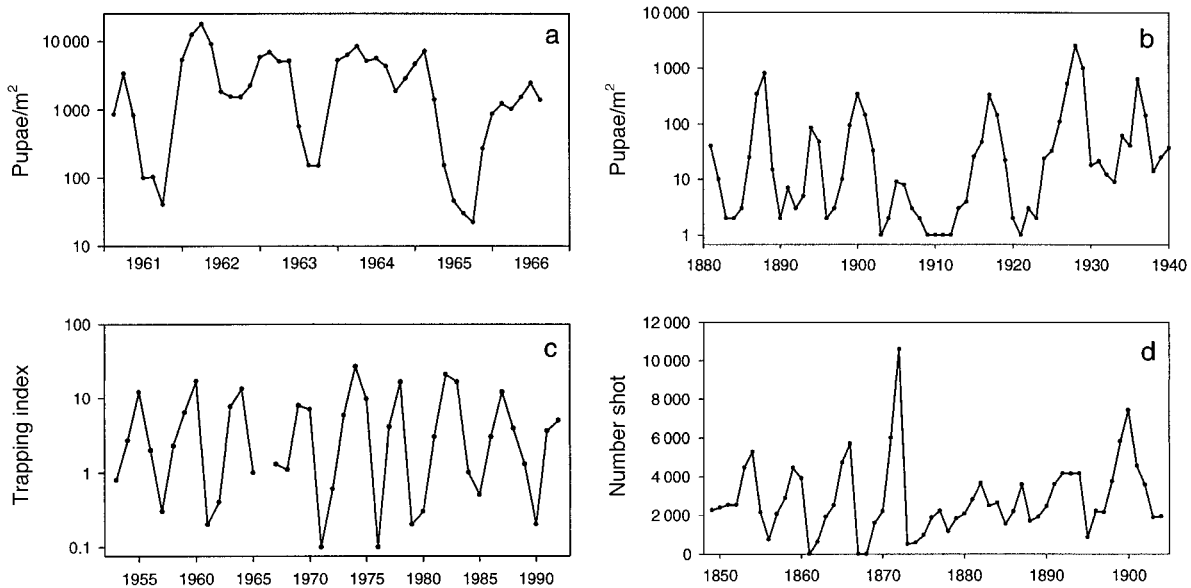
Fig. 1. Examples of cyclic population dynamics. (a) Coffee leaf-miners (*Leucoptera* spp.) at Lyamungu, Tanzania (Bigger 1973). (b) Pine looper (*Bupalus piniarius*) in Germany (Schwerdtfeger 1941). (c) Voles (*Microtus* and *Clethrionomys*) at Kilpisjärvi, northern Finland (Laine and Henttonen 1983, Hanski et al. 1993). (d) Red Grouse (*Lagopus lagopus scotius*) in Scotland (Middleton 1934).

can we uncover the mechanisms that drive cycles in population density?

A common approach to studying population cycles has been to perform short-term experimental and observational studies that look directly at population processes that might cause the cycles. This approach has been used intensively, for example, in studies of Canadian hares (Krebs 1996) and British grouse (Watson et al. 1984, Hudson et al. 1992, Moss et al. 1996). This work is invaluable for building up a list of biologically plausible causes of cycles in a population, as was summarized for many species of outbreaking forest insects in Berryman (1988). Often the empirical evidence suggests that some aspect of fecundity or mortality varies with the population's density, and each such result is put forward as a potential governing factor. Nevertheless, these hypotheses are rarely tested rigorously, and the number of competing hypotheses tends to grow through time. A decisive test can be difficult. Although some kind of density dependence is required for population cycles to arise, not all density-dependent interactions will do the job; it is almost impossible to determine the effect of the interaction from verbal models of the population. Only mathematical population models can show which factors are even capable of generating cyclic dynamics; without such models the experimental approach will not be able to solve the question of population cycles.

At the same time, ecologists have been accumulating empirical data of another sort: long time series of population abundances. There are a few dozen "classic" time series, such as those in Fig. 1, that are widely

known, but recent work has produced a deluge of new series. For example, a project collecting published time series of 10 or more years has passed the 5000 mark and is still going strong (J. Prendergast, *personal communication*). 700 of these series are longer than 25 yr, and 30% of those have periodic oscillations (Kendall et al. 1998). These data allow us to go beyond merely characterizing a population as "cyclic," and quantify aspects of the cycle such as the period, amplitude, and maximum growth rate.

How can we employ mathematical models to best use both the above-detailed information on population processes and any long-term population time series, and reach stronger conclusions about the causes of population cycles? We believe that the time is right to bring together two complementary quantitative approaches—time-series statistics and mechanistic population modeling. Four major features distinguish the two approaches: (1) the goals of the analysis; (2) the way in which the time series is treated; (3) the kinds of models that are developed; and (4) the methods used for judging the models' explanatory power. These approaches are now highly developed, and each has been applied individually to the problem of population cycles, but they have rarely been used in concert. Our purposes in this paper are to demonstrate that much is to be gained by bringing these traditions together, to outline a series of steps to achieve that end, and to show a successful case study.

The primary aims of time-series statisticians are to describe the data and extrapolate the time series into the future. This approach analyzes the time series di-

rectly to produce descriptors of the dynamics such as period, amplitude, and Lyapunov exponent. The models used are usually non-mechanistic—they either explicitly include statistical features of the data, such as the autocorrelation structure or the spectrum of periodic tendencies, or they use flexible ''nonparametric'' functions, such as splines or neural networks, to relate future population size to current and past abundances. There are well-understood quantitative assessments of model fit, such as the mean squared prediction error, which are used both for parameter estimation and to compare the fit of different models to a particular data set.

An example of a statistical model that is commonly used with population cycles is the linear autoregressive (AR) model (Royama 1992). This is used both to quantitatively characterize the period and strength of the cycle, and to look for correlations with other oscillatory variables. A well-known application of the latter is the search for the putative correlation between Canadian lynx and hare populations and the sunspot cycle (Moran 1949, Keith 1963, Sinclair et al. 1993, Ranta et al. 1997, Sinclair and Gosline 1997).

Nonlinear models of varying complexity have also been used, both to describe time series (Lindström et al. 1995) and to estimate whether the fluctuations in various populations are chaotic (Hassell et al. 1976, Turchin 1993, Ellner and Turchin 1995).

The statistical handling of time series has several strengths: the techniques are well developed (and in particular the notion of ''goodness of fit'' is quantitatively defined); the sampling properties of the statistics are well understood; and the approach makes relatively few assumptions about the data. However, the approach treats the time series simply as a string of numbers, ignoring any qualitative and quantitative information about the ecological system that generates the time series. Hence the models tend to be biologically naive. Furthermore, statistical analyses provide no information about the underlying ecological mechanisms, for there is not a unique relationship between statistical patterns and mechanisms. Indeed, a particular pattern can often be generated by a variety of different mechanisms.

The approach of theoretical population ecologists, which dates back to Lotka and Volterra, typically differs in all four regards. The goal of this approach is typically to understand the causes of a generic phenomenon, such as ''the 3–4 yr rodent cycle,'' rather than a particular time series. Thus the time series is abstracted to its period, for example, rather than being modeled directly. The models that are constructed are explicitly biological and contain the major mechanisms or processes that are believed to generate the dynamic behavior. The model's explanatory power is rarely judged on the basis of statistical goodness of fit, but rather on a grosser level of agreement, such as the

model's capacity to generate cycles of approximately the right period.

The theoretical literature on population models that can produce cycles is too vast to review here in any detail. Cycles can be caused by direct and delayed density dependence (May 1974, Gurney et al. 1983; these terms are used differently in reference to continuously and discretely breeding organisms), a variety of consumer–resource interactions, such as predator–prey, host–parasitoid, and host–pathogen interactions (Lotka 1925, Volterra 1926, Nicholson and Bailey 1935, Rosenzweig and MacArthur 1963, Anderson and May 1981), and periodic variation in the environment, which can cause population cycles with a period longer than that of the driving environmental variable (Aron and Schwartz 1984, Rinaldi et al. 1993). Furthermore, when one of the above factors is present but too weak to cause deterministic cycles, random exogenous forcing can cause cycles to appear (Royama 1992, Kaitala et al. 1996.

Godfray and Hassell (1989) illustrate the approach commonly taken by theoretical ecologists. They sought to explain cycles in continuously breeding tropical pest insects in which the period is roughly equal to the pest's generation time. They developed a model of the interaction between a host, in which only the juveniles are vulnerable to attack, and a parasitoid that displays density dependence in its attack rate. The similarities in the form and period of the cycles between the model output and the real populations were used as evidence that parasitoids may be driving the generation cycles often observed in tropical insect populations. This is the degree of interface between data and models found in most theoretical ecology studies.

Gurney et al. (1980) went one step further in their attempt to explain the oscillations in one of Nicholson's blowfly populations. They constructed a model incorporating a likely mechanism, estimated the parameters mainly from other available data, and showed that the model output ''looked like'' the data, including capturing the ''double-humped peak'' of the cycle. This study is unusual in theoretical ecology in that some parameter estimation was done from the time series itself but, as is typical, the goodness of fit was not measured statistically.

The mechanistic modeling approach has several strengths: mechanistic factors controlling the dynamics are explicitly included, enhancing our understanding and allowing predictions to be made about the consequences of environmental change; it uses information from other experiments and observations beyond the time series; and in principle the approach provides a theoretical framework for recognizing that dynamical behavior seen in different ecological systems may be manifestations of the same underlying dynamical process. However, the statistical properties of mechanistic models are not, in general, well understood—any part of the model that does not have independent supporting
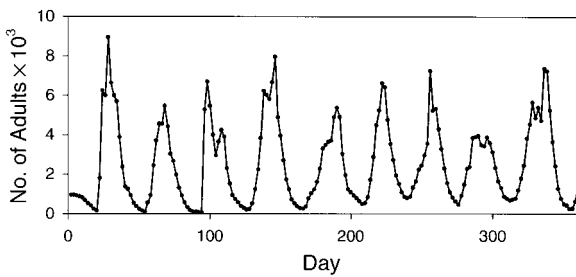
FIG. 2. A laboratory population of the sheep blowfly (*Lucilia cuprina*). This is ''population I'' from Fig. 8 of Nicholson (1957); the data are tabulated in Brillinger et al. (1980). Adults are counted every 2 d. Other data are available for this population (e.g., time series of eggs and adult mortality), but we want to illustrate this process with the sort of data one might have from a field population.

evidence requires possibly arbitrary assumptions about functional form, and there is no sharp notion of goodness of fit.

The statistical and mechanistic modeling approaches are clearly complementary, but we have found remarkably few examples of them being used in combination. As noted above, Gurney et al. (1980) did this in part when analyzing Nicholson's blowflies. Royama (1992) used a match to the autocorrelation structure of the Canadian lynx data to estimate parameters for a mechanistic predator–prey model, which matched the time-series data better than did a simple phenomenological model. Schaffer and his colleagues (Olsen and Schaffer 1990, Schaffer et al. 1990, Tidd et al. 1993) used the Lyapunov exponent and nonlinear forecast accuracy (both nonlinear time-series statistics) to show that a simple mechanistic model of measles epidemics gave a better fit to the time-series data than did a variety of phenomenological models, and did as well as a more detailed mechanistic model.

The main goal of this paper is to formulate, and illustrate through an example, how the statistical and mechanistic modeling approaches can be combined to give both deeper insight into biological mechanisms and more-rigorous testing of the models that embody these mechanisms. After giving an overview of the process, we apply our approach to the population of blowflies shown in Fig. 2. The data come from a laboratory study of the sheep blowfly (*Lucilia cuprina*) performed by A. J. Nicholson (1957), and will be familiar to most readers from introductory ecology textbooks. Why have we dredged up Nicholson's blowflies yet again? A statistician would say that we should use simulated data, in which we know the ''truth'' exactly, to check the reliability of this approach. A field ecologist would say this is all hot air unless it works with field populations. The blowflies are a compromise: being real animals, they present far more of a challenge than any simulated population could; but being a laboratory population, we know the true limiting factor.

During this analysis we pretend to know less than we actually do, limiting our knowledge to the adult time series, the general structure common to all of Nicholson's experimental populations, and various quantitative estimates of generic life-history parameters (such as maximum fecundity). This is typical of the data available for well-studied field populations.

INFERRING THE CAUSES OF POPULATION CYCLES

The biological question we wish to answer is: Which interaction (or interactions) causes the population to cycle? We rephrase this in a more quantitative form: Which interaction can generate population dynamics similar to the observed time series? Thus we must move from a qualitative hypothesis (such as ''the observed population cycles are caused by density-dependent fecundity'') to explicit predictions about population dynamics. This is not a trivial process. The population interactions in question will typically be nonlinear, and theoretical analyses have repeatedly shown that nonlinear interactions have nonintuitive dynamical consequences. The solution is to construct a mechanistic population model that embodies the hypothesis as applied to the population, use the model to generate simulated time series, and quantitatively compare the simulated time series to the data. Each of these steps is a well-established process (although new techniques are being rapidly developed), but to our knowledge they have never before been combined in this way to quantitatively evaluate hypotheses about population dynamics.

Construct mechanistic models

There are two phases to the model construction. The first phase involves writing down the mathematical equations. These must incorporate information about general population structure, such as the organism's life history and trophic interactions, that are thought to be important. Indeed, even the mathematical framework to be used (for example, discrete vs. continuous time) will depend on the population structure. Functional forms must be specified for all of the vital processes. The hypothesis for the cause of the population cycles will usually show up as density dependence in one or more of these processes. There are many subtleties to building mathematical population models; for example see Nisbet and Gurney (1982) and Tuljapurkar and Caswell (1997).

Mathematical analysis of the model may give insights into whether the hypothesis is worth pursuing further. The model may be fundamentally incapable of generating cycles, or can only do so for biologically implausible parameter values. The latter is often true for simple discrete-time density-dependent models such as the logistic map. Alternatively, the model may not be capable of generating cycles of the correct period. For example, many simple discrete-time consumer–resource models seem unable to produce cycles with

a period of less than about six generations (Lauwerier and Metz 1986, Ginzburg and Taneyhill 1994). Nevertheless, many models will be capable of producing qualitatively correct cycles with reasonable parameter values.

Parameterize model

The second phase of model construction is to narrow the parameter values to a range that is appropriate to the population at hand. There are two sources of information for this—independent experiments and observations, and the population time series itself. Each source of parameter estimates has advantages and disadvantages.

Independent parameter estimates can be based on experimental manipulations of factors such as density or on observational studies such as those leading to life-table analyses (e.g., Klomp 1966, McCauley et al. 1996). This allows the model evaluation that follows to be statistically independent of the model parameterization. However, many parameters can be difficult to estimate directly from available data, and the estimates may not be entirely appropriate: the experimental setting might not be strictly analogous to the natural population (especially if the experiments are performed in the laboratory).

An alternative approach to parameter estimation is to fit the model directly to the time series, using a variety of statistical approaches. Until recently this was difficult to do with complex models, but modern computer-intensive techniques are rapidly being developed. We describe three such approaches below; for applications to population biology see Jones and Perry (1978), Harrison (1995), Dennis et al. (1995, 1997), and Higgins et al. (1997). All parameters can be estimated simultaneously with these techniques, and they are unequivocally relevant to the population. However, this precludes a fully independent assessment of the model.

The parameterized model represents a fully specified quantitative representation of the hypothesis as applied to the population. It also represents a hypothesis about the aspects of the general population structure that are important; this constraint is unavoidable, and reflects the contingency inherent in applying ecological principles to specific systems. This process should be repeated for the various hypotheses to be evaluated, producing a suite of quantitative hypotheses, each of which can be compared to the data.

Compare model predictions to time series

Each parameterized model makes predictions about the dynamics. These predictions are revealed by simulating the model to generate synthetic time series, to which the empirical time series can be compared. We quantify this comparison with a suite of statistical time-series descriptors, such as the period, amplitude, autocorrelation function, and spectral density function,

that we refer to collectively as "time-series probes." Each of these probes quantifies some aspect of the pattern in the dynamics. We compute these probes both for the data and for the simulated time series, and ask which simulated time series has probe values most similar to the data (Olsen and Schaffer 1990, Schaffer et al. 1990, Royama 1992, Hanski et al. 1993, Tidd et al. 1993, Turchin and Hanski 1997). If the models are stochastic (as most reasonable models will be), then we can generate a large number of simulated time series for each model and look at the distribution of probes for each model (each synthetic time series should be the same length as the data). Then we can ask under which model are the data most probable.

If we have parameterized the models by fitting them to the time series, this would appear to be blatant circularity. After all, we are taught to test our models against novel data! There are three reasons why this may not be as much of a problem as it appears. First, the comparison of the different models is in some sense a second round of model selection (the first was the parameterization exercise); it is not really a "test" of the models. The outcome is not a true model, but a plausible one (probably more than one!). Second, the fitting procedures typically try to maximize fit at either very short or very long time scales; many of the features measured by the probes are of an intermediate scale and may not have been selected by the fitting procedure. Third, the models will not in general be free to fit the data arbitrarily well. Because of the constraints imposed by the population structure and the functional forms, there may be a trade-off, for example, between period and amplitude of the oscillations that cause the model to perform poorly even with the "best" parameter values.

The outcome of this exercise will be one or more dynamically plausible mechanisms for the observed population cycles. Even if there is no clear-cut winner, some hypotheses are likely to have been eliminated, and this alone will be a great contribution in many systems. Furthermore, the results can be used to guide future focussed tests of the remaining plausible hypotheses, and make predictions of future dynamics, both in unmanipulated populations and in the presence of deliberate environmental change.

AN EXAMPLE: NICHOLSON'S BLOWFLIES

Construct mechanistic models

In the blowfly population, the qualitative life-history and "environmental" information (i.e., Nicholson's [1957] experimental protocols: a single species with non-dynamic food, constant environment, continuous breeding, relatively constant juvenile development time) suggest that a simple, single-species stage-structured model with two stages (juvenile and adult) could describe the population. The basic formulation of this type of model is as follows:

CONCEPTS & SYNTHESIS

$$\frac{dJ(t)}{dt} = B(t) - R(t) - M_{\mathrm{J}}(t)$$

$$\frac{dA(t)}{dt} = R(t) - M_{\mathrm{A}}(t) \qquad (1)$$

where J and A are the densities of juveniles and adults, B is the birth rate, t is time, R is the recruitment rate of juveniles to adults, and M_{J} and M_{A} are the death rates of juveniles and adults. The constant development time (which we denote τ) implies that $R(t)$ is related to $B(t - \tau)$, leading us to the formalism of the delay-differential equation (Nisbet 1997). Models with this structure are in fact capable of displaying cycles (Gurney and Nisbet 1985), indicating that this hypothesis of the population structure is plausible.

The next step is to identify the dependencies that go into the individual processes. For example, ''density-dependent per capita fecundity'' translates into

$$B(t)/A(t) = f(A(t)).$$

We must also come up with functional forms for all of these processes. Ideally we would have information from independent experiments to suggest this. If not, we can either guess plausible forms, or use flexible functions such as splines (Wood 1994a, b, Ellner et al. 1997, Bjørnstad et al. 1998). All of the processes must be spelled out in this way, so that instead of ''cycles are caused by density-dependent fecundity,'' the hypothesis is really ''cycles are caused by density-dependent fecundity (plus a bunch of other assumptions about population structure).'' Indeed, the model may be used to explore joint hypotheses, such as the combination of predation and intrinsic density dependence (Stenseth et al. 1996).

There are a number of models fitting into the structure of Eq. 1 that might be applied to blowflies. In some of these models the effect of the density dependence is instantaneous, as when the larval death rate is a function of current larval density. These direct-feedback models only produce cycles with a period less than twice the juvenile development time (Gurney and Nisbet 1985). The blowfly development time is 12–16 d (Nicholson 1957) and the cycle period (from spectral analysis) is 40 d, so we can reject these direct-feedback mechanisms out of hand. In models where the effect of the density dependence is delayed, the cycle period is between 2 and 4 times the development time (Gurney and Nisbet 1985), consistent with the blowfly data. Thus we need to examine these delayed-feedback models in detail. Our goal in this blowfly analysis is to be illustrative rather than exhaustive, so we only present two models (other hypotheses and models can be evaluated in an analogous fashion):

(1) *Adult-competition (AC) model.*—Adult blowflies need to eat a protein meal in order to produce eggs; if the food supply is limited, then fecundity will be density dependent. We represent this mechanism with a model in which the per capita fecundity is an exponentially declining function of adult density (based on Gurney et al. 1980, 1983, Readshaw and Cuff 1980; see Appendix for details).

(2) *Larval-competition (LC) model.*—Larvae require food to grow; if the food supply is limited then they will grow more slowly at high density. Since the development time is constant, these slow-growing flies will be smaller at maturity, which in turn will reduce their per capita fecundity as adults (based on the AF model of Gurney and Nisbet 1985; see Appendix for details).

In both the AC and LC models, the remaining processes in Eq. 1 are taken in their simplest forms: the juvenile development time is constant and per capita mortality is independent of density.

Parameterize models

The above models embody the hypotheses qualitatively, but we must still adapt them quantitatively to the population at hand by estimating parameter values. Some of the parameters can be estimated from independent data, such as the life tables in Nicholson (1957) and Readshaw and van Gerwen (1983) (Table 1). However, no information is available about many parameters, especially in the LC model. Thus we must estimate the parameters by fitting the models to the time series. There is a variety of techniques, differing in their definition of ''fit'' and in their assumptions about the data; we illustrate three that we have found particularly useful in the context of strong nonlinearities and population oscillations (banishing the technical details to the Appendix), and summarize the resulting parameter estimates in Table 1.

Direct fit.—The most obvious way to proceed is to attempt to fit the models directly. The models describe a relationship between current and future values of the state variables. If we have estimates of all the relevant state variables then we can do a nonlinear least-squares regression to find the set of parameters that best fit the data. The residuals from such regressions are typically autocorrelated; this can be accounted for directly in the regression (Carpenter et al. 1994).

To apply this approach to the AC model, we need estimates of $A(t)$, $dA(t)/dt$, and $A(t - \tau)$. We have direct observations of $A(t)$; if we smooth the time series (see for example Green and Silverman 1994) then we can estimate $dA(t)/dt$ and $A(t - \tau)$ by differentiating and interpolating, respectively, the smoothed time series.

This approach is not possible, however, with the LC model, because we do not have estimates of juvenile density and adult mass. Thus we will not discuss this approach further.

Trajectory matching.—If we believe that there is little process noise, then a correct model should be able to duplicate exactly the observed time series, with deviations only due to measurement errors. Model pa-

TABLE 1. Empirical parameter estimates (''independent estimates'') and best-fit parameter estimates and goodness-of-fit values (r^2) for the adult-competition (AC) and larval-competition (LC) models, from the three fitting methods. Data are from Nicholson (1957) and Readshaw and van Gerwen (1983).

Model	τ† (d)	qs_J‡	b§	μ_A‖	A_0¶	J_0#	Γ††	r^2
Independent estimates	12–16	5.1	n.a.	0.08–0.14	n.a.	n.a.	n.a.	···
Direct fitting								
AC	12.0	4.37	···	0.079	438	···	···	0.69
LC	n.a.	···	n.a.	n.a.	···	n.a.	n.a.	n.a.
Trajectory matching								
AC	14.0	4.77	···	0.16	573	···	···	0.70
LC	13.5	···	50.0	0.125	···	4.98×10^5	0.98	0.60
Nonlinear forecasting, NLF								
AC	14.1	5.93	···	0.212	463	···	···	0.67
LC	12.0	···	24.3	0.144	···	97.9	0.097	0.61

Note: The designation n.a. indicates ''possible, but not available.''
† Development time.
‡ The product of maximum per capita fecundity and juvenile survival. For the independent estimates a 50:50 sex ratio is assumed. In the AC model q and s_J cannot be separated.
§ Fecundity per unit mass.
‖ Adult per capita death rate.
¶ Density at which total fecundity is maximum.
# Density at which the uptake rate is half its maximum.
†† Maintenance cost.

rameters (including initial values of the state variables) are then chosen by minimizing the mean squared difference between the numerical solution of the model and the data. This process finds the parameter values for which the deterministic-model time series most closely follows the data; it is usually called ''model calibration.'' This is the approach used by Jones and Perry (1978) and Harrison (1995).

We fit both the AC and LC models to the data using nonlinear least squares. The residuals from the fits of both models are strongly autocorrelated. Ordinarily, measurement errors are assumed to be independent, and such autocorrelation would be evidence of an inadequate model. However, the time series of adult blowflies was obtained from running counts of the number of pupating and dying individuals, so any measurement errors (such as a miscount of the number of dead flies) would propagate through time, generating such autocorrelations.

Nonlinear forecasting.—The direct fitting failed on the LC model because of the problem of unobserved state variables. We introduce a technique that allows us to fit a process-error model even when all state variables are not available; derived from the field of nonlinear dynamics, it is called ''nonlinear forecasting'' (NLF) (Farmer and Sidorowich 1987, Casdagli 1989). This uses the Takens embedding theorem (Takens 1981) to substitute lagged values of the observed variable for the missing state variables. In essence, given a model time series and the data time series, NLF takes each segment of the data time series, finds segments of the model time series that are ''similar'' (close in the reconstructed state space; see Appendix), and makes predictions based on the subsequent values of

the model time series. These predictions can then be compared with the observed values from the data. The fitting algorithm searches for the set of parameters that minimizes the mean forecast error from these predictions. Since both process and measurement noise enter the procedure at the step of generating the model time series, either type of noise can be added in any way that is biologically and dynamically plausible, and need not be linear. In this example, we made the per capita fecundity into a lognormally distributed, autocorrelated random variable; the variance and correlation coefficient were free parameters to be fit.

Compare model predictions to time series

We now have quantitative hypotheses of the causes of the cycles. How do we assess their absolute or relative abilities to explain the observed population dynamics? We could use the criteria that we maximized to obtain the parameter estimates (Table 1), but comparing structurally different models in this way can be difficult, and we have no independent assessment of what constitutes ''significant difference.'' Because our true interest is in the ability of the different models to reproduce the observed dynamics, we proceed by generating simulated time series from the models and comparing them directly to the observed time series with a number of time-series probes.

For the blowfly models, the deterministic time series generated by the trajectory-matching method are much more regular than the data. To generate more realistic series, we added measurement noise, based on the mean and variance of the residuals from the fits. The resulting time series have maximum values that are much greater
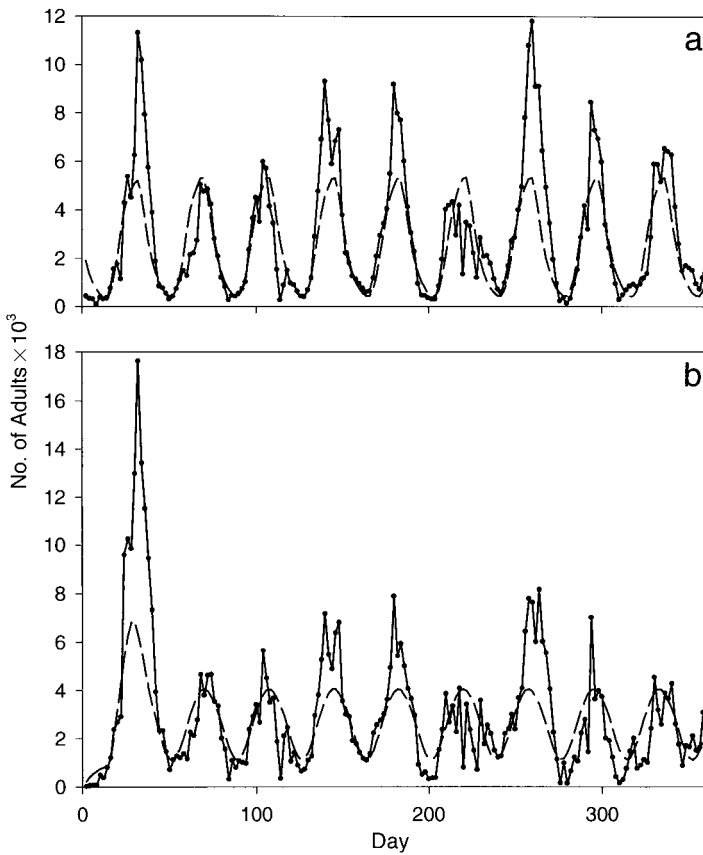
FIG. 3. Simulated time series from (a) the adult-competition (AC) model and (b) the larval-competition (LC) model using the parameter estimates and noise structure from the trajectory-matching fits (Table 1). Compare with Fig. 2. The dashed line depicts the deterministic time series; the solid line shows the stochastic realization with measurement error.

than those seen in the data, but they otherwise look qualitatively correct (Fig. 3).

The NLF fitting process incorporated dynamic noise into the model, and the variance and autocorrelation of that noise were estimated along with the other model parameters. We simply iterated the models to obtain
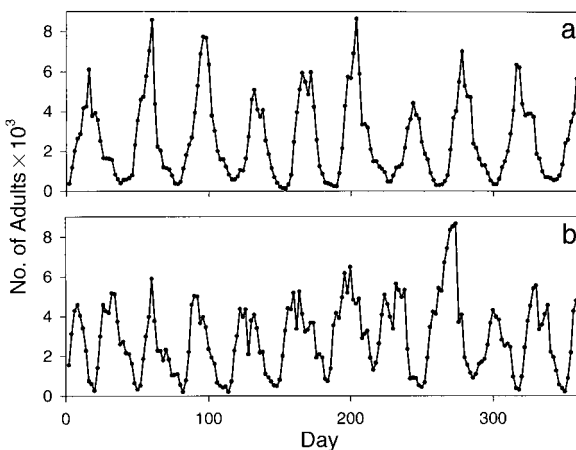

FIG. 4. Simulated time series from (a) the adult-competition (AC) model and (b) the larval-competition (LC) model using the parameter estimates and noise structure from the nonlinear-forecasting (NLF) fits (Table 1). Compare with Fig. 2.

simulated time series. The AC time series looks qualitatively correct, but the LC time series seems to have too short a period, too low an amplitude, and too little variation in peak height (Fig. 4).

Figs. 3 and 4 show single realizations of the stochastic models. These time series are short (by stochastic simulation standards), so we expect that other realizations will look somewhat different, and indeed they do. Rather than compare a single realization with the data, we compared 100 realizations from each model. The variation among these realizations allowed us to estimate confidence intervals, and ask how likely the data are under each model.

The probes fall into two groups. The first simply describes the distribution of the state variable: mean, variance, and quantiles. The other group addresses the nature and intensity of the cycles: the standard deviation of log N (a measure of variability), the period and amplitude of the cycle, the first few autocorrelation coefficients, and the Lyapunov exponent.

For each of the probes, Table 2 shows the single value for the data and the means and 90% confidence intervals (based on the 100 synthetic time series) for each of the models with each of the fitting procedures. If one of the models were ''true'' we would expect that the data value should fall within the confidence interval of that model for most or all of the probes. However,

the pattern is not so simple. For some probes (such as the 25% quantile of the distribution of densities) the AC model is correct and the LC model is incorrect, while for others (such as the Lyapunov exponent) the opposite is true. Other probes (such as period) are accurately captured by both models, but with several probes (such as the autocorrelation coefficients) both models miss entirely. In the last situation, the AC model is almost always closer to the data value than is the LC model.

We could simply count up how many probes favor each model: this gives a tie if only those probes where at least one model "fits" are considered, and favors the AC model if all probes are considered. However, the different probes are not fully independent of one another. For example, the variability and the maximum density are very strongly correlated. To make a fair comparison we need a set of independent probes of the data. We can account for the correlations among the time-series probes and create orthogonal probes using multivariate statistics (Manly 1986), such as principal-component analysis (PCA). PCA finds a set of linear combinations (called principal components) of the original measurements such that the principal components are independent of one another. In addition, the first principal component (PC1) captures as much of the variability among observations as possible (the different measurements are scaled to each have a unit variance), the second principal component captures the most variation in the residuals about PC1, and so on. The hope is that the first few principal components will capture most of the variability in the data set. This accomplishes a dual purpose of simplifying the analysis by lowering the number of variables and ensuring that the variables are independent.

We analyzed the NLF and trajectory-matching fits separately. In each case, we considered the 12 probes in Table 2 to be the measurements, and each of the 200 synthetic time series to be the observations. In both cases PC1 captured more than half of the variance in probe values among the synthetic time series; the first two principal components described 66% of the variance for trajectory-matching and 76% of the variance for NLF. In both analyses PC1 is primarily an amplitude and autocorrelation axis (Table 3). For the trajectory-matching time series, PC2 is dominated by the Lyapunov exponent and the 75% quantile, suggesting that it involves the shapes of the cycles. With NLF, the strongest contributions to PC2 are from the mean and median.

For both NLF and trajectory matching the first principal component clearly separates the AC and LC models, whereas both models have overlapping variation along PC2 (Fig. 5). We can plot the data in the same principal-component space as the models (indicated by the large "D" in Fig. 5), and see which model is most consistent with it. For the trajectory-matching analysis, the data are on the edge of the AC model values. For

the NLF analysis, the data are outside both model regions, but far closer to the AC model.

Overall, these results suggest that the AC model better describes the data than does the LC model, especially when it is fit with measurement error (trajectory matching) rather than process error (NLF).

Nicholson's blowflies revisited?

What is the right answer? Nicholson provided an unlimited supply of food to the larvae in his experiment, and while adults received an unlimited supply of sugar water (which they need for survival) they received a fixed amount of protein (which they require to produce eggs) every 2 d. It seems clear from this that the "true" mechanism is fecundity limited by adult competition for food. It is conceivable that some aspect of larval density other than food availability might have an impact, but independent experiments have shown that with limited food the per capita fecundity declines with density even when the adults came from a common larval pool (Readshaw and van Gerwen 1983). We also have some direct checks on the assumptions of the models. Nicholson recorded time series of eggs and recruits, and so we can check the development time and juvenile mortality. The cross-correlation between eggs and recruits has a maximum at 12 d; the relationship between the two series at that lag is nearly linear, confirming the assumption of density-independent larval mortality.

Nevertheless, the AC model does not predict the data perfectly. There are a number of confounding factors. Independent experiments indicate that adult mortality is not constant, but depends on both density and age (Readshaw and van Gerwen 1983). When individuals recruit into the adult stage, they take several days to produce their first batch of eggs, so there is an "immature adult" stage that competes for food but does not lay eggs. We have explored the consequences of putting both these factors into the AC model; they have no qualitative effect on the dynamics, and the quantitative effects seem small, but they could account for some of the errors.

Another problem is that the population was evolving. The population was maintained for another year following the data we analyze here, and the time series changed qualitatively; at the end of the experiment Nicholson found that the flies had an enhanced ability to reproduce at high density at the expense of a lowered maximum fecundity. However, an analysis that applied direct fits of the AC model to the data indicated that the parameters were changing systematically even during the part of the experiment we analyzed (Stokes et al. 1988). This results in "noise" with long autocorrelation times that we did not take into account.

DISCUSSION

The question we asked at the beginning of this paper was, how can we uncover the mechanisms that drive

TABLE 2. Univariate time-series ''probes'' (i.e., statistical descriptors) of the data and the simulated time series from the adult-competition (AC) and larval-competition (LC) models. For the models we report the means over 100 simulated time series, with the 5% and 95% quantiles in parentheses.

Model	Minimum	25% quantile	Median	Mean	75% quantile	Maximum
Data	60	762	1760	2480	3870	8920
Trajectory matching						
AC	60 (60; 60)	860 (620; 1090)	2260 (2010; 2530)	2560 (2430; 2670)	4050 (3790; 4320)	7620 (6750; 8650)
LC	60 (60; 60)	1600 (1360; 1830)	2830 (2620; 3060)	2910 (2760; 3060)	4080 (3850; 4330)	8330 (7210; 9680)
Nonlinear forecasting, NLF						
AC	268 (228; 314)	842 (712; 983)	2090 (1850; 2350)	2460 (2310; 2610)	3990 (3530; 4360)	6510 (5670; 7980)
LC	518 (490; 553)	1090 (988; 1190)	2240 (2070; 2460)	2320 (2200; 2510)	3570 (3410; 3960)	4360 (4110; 4650)

† Standard deviation of $\log_{10} N$.
‡ The first two coefficients of the autocorrelation function.
§ From spectral decomposition.
‖ Calculated using LENNS (Ellner et al. 1992); see Ellner and Turchin (1995). LENNS parameters: $l = 3$, $d = 7$, $k = 3$.

regular oscillations in population density? This entails inferring processes from the empirical patterns they produce. Statisticians call this the ''inverse problem'' and it is typically difficult to solve, as many processes might produce similar patterns (Wood 1997). Using Nicholson's blowfly data as an example of a cyclic population, we addressed this question with a synthetic approach that combines statistical time-series methods with mechanistic population models. We quantified hypotheses about the causes of population cycles by writing and parameterizing mathematical models embodying the putative mechanisms. We then generated simulated time series from these models, and used a variety of probes (time-series statistics) to quantitatively compare the model time series to the data. The blowfly models produced qualitatively similar dynamics, but we found that a model incorporating adult competition for food produced dynamics more like the data than did one incorporating larval competition. In particular, the data were nearly distinguishable from time series generated by the adult-competition model with observation error. The strong confirmation of this model was satisfying, for the population was in fact limited by the adult food supply and, being in a laboratory, was free from environmental variation.

Our experience with the blowfly analyses raises the hope that applying this synthetic approach to data from field populations (where the correct answer is not known beforehand) may similarly result in a rejection of at least some of the hypotheses and confirmation of others. We think that our analysis of blowfly data is a fair test of the ability of this approach to distinguish between rival hypotheses because of the constraints that we imposed on ourselves during this exercise. Specifically, we limited ourselves to the analysis of adult numbers only, since in most field case studies we will have a time series of only a single life stage. In addition, we used a model that grossly simplified the dynamical processes governing blowfly oscillations (for example, we did not incorporate the density and age dependencies that are known to affect adult mortality). In other words, our models were similar to the ones that we are likely to apply to field case studies, both being simplified caricatures of reality. Finally, our analysis was ignorant of the details of Nicholson's experimental setup—the information we used to check on whether or not we obtained the correct answer. Thus, success was by no means assured. The fact that our results did decisively point to one of the postulated mechanisms is grounds for optimism. On the other hand, the blowfly example also had several features that made our job easier: only one species was involved, the environment was relatively invariant, the population estimates were relatively precise, and there was only one major ''limiting factor'' (finite food supply).

Why are we asking simple population models to make quantitative predictions about specific populations? After all, many of these models were originally designed to answer *qualitative* questions about population dynamics, and there is no way that they can encompass all the factors affecting a population. However, population models are becoming increasingly sophisticated, to the point where we find it reasonable to challenge them quantitatively at least with model systems. Laboratory populations such as the blowflies analyzed here are clearly model systems, but so are cyclic field populations. The large amplitude and regularity of the oscillations indicate that only a few factors or interactions are dominating the population dynamics, and so populations like the ones in Fig. 1 are intermediate in complexity between the laboratory and an arbitrarily chosen field population.

TABLE 2. Extended.

Period	Variability (s)†	Autocorrelation function‡		Mean amplitude§	Lyapunov exponent‖
		AR1	AR2		
20	0.487	0.873	0.675	2530	0.169
19.7 (18; 20)	0.629 (0.573; 0.68)	0.671 (0.599; 0.724)	0.55 (0.487; 0.606)	2320 (2120; 2490)	0.0644 (−0.0193; 0.145)
18.5 (18; 20)	0.486 (0.427; 0.536)	0.466 (0.373; 0.559)	0.397 (0.303; 0.487)	1610 (1320; 1850)	0.07 (−0.0061; 0.153)
18.7 (18; 20)	0.386 (0.353; 0.413)	0.915 (0.9; 0.927)	0.733 (0.707; 0.753)	2330 (2110; 2550)	0.0697 (0.0216; 0.117)
20 (20; 20)	0.294 (0.285; 0.309)	0.934 (0.927; 0.939)	0.762 (0.744; 0.776)	1750 (1670; 1960)	0.0999 (0.0269; 0.194)

Analysis of field populations will be more difficult in several ways, even if we seek less detailed explanations for processes governing oscillations. There will be more random environmental fluctuation in the time series and probably less precise population estimates. There will likely be more candidate mechanisms to contrast against each other. Consider, for example, the spectrum of hypotheses that has been advanced to explain cycles in small mammal populations: herbivore–plant and predator–prey interactions, weather, stress, behavior, and genetics (reviews in Krebs and Myers 1974, Batzli 1992, Stenseth and Ims 1993). On the other hand, field case studies are also potentially rich in life-history and natural-history information, which as we showed in this paper can serve both to delineate likely processes and to reject other postulated mechanisms. We have been working to apply these approaches to field populations of the pine looper moth (Schwerdtfeger 1941, Broekhuizen et al. 1993). As expected, this is more challenging than the blowflies, for the populations have many more factors impinging on them and the cycles are not as regular. We will probably not be able to find a single best hypothesis, but we have succeeded in rejecting some biologically plausible hy-

potheses, such as variation in plant quality in response to herbivory (Straw 1996). Furthermore, by requiring the models to make quantitative predictions we have uncovered some novel features in the data. For example, the population growth rate appears to be depressed at extremely low densities, which might be caused by an Allee effect or a generalist predator (Turnock 1969). This factor cannot produce population cycles on its own, but it can modify the structure of cycles arising from other sources. These are valuable contributions, even for such a well-studied species, for direct analysis of the data (even detailed life tables) without reference to dynamical models yields few strong conclusions (Royama 1997).

Often the most compelling evidence for the causes of cycles comes from long-term experimental manipulations of entire populations. Recent examples of this include predator manipulations in hares (Krebs et al. 1995) and density manipulations in grouse (Moss et al. 1996). These experiments are expensive and time consuming, and so should be designed to maximize the chance of getting a decisive answer. This design is best done by using population models to make quantitative predictions about the population's response to the ma-

TABLE 3. Loadings of the probes on the first two principal components. Each value represents the correlation between the probe and the principal component; large absolute values indicate that the probe is strongly represented in the component.

	Minimum	25% quantile	Median	Mean	75% quantile	Maximum	Period	Variability (s)	Autocorrelation function		Mean amplitude	Lyapunov exponent	Variance explained (%)
									AR1	AR2			
Trajectory matching													
PC1	0†	0.37	0.35	0.36	0.04	0.2	−0.24	−0.36	−0.36	−0.34	−0.36	0.03	56
PC2	0†	−0.03	0.06	0.17	0.84	0.1	−0.23	0.13	0.07	0.10	0.19	−0.36	10
Nonlinear forecasting, NLF													
PC1	0.35	0.32	0.15	−0.27	−0.30	−0.33	0.26	−0.35	0.29	0.27	−0.35	0.12	62
PC2	−0.03	−0.16	−0.62	−0.51	−0.40	−0.02	−0.11	−0.04	−0.18	−0.27	−0.14	−0.16	14

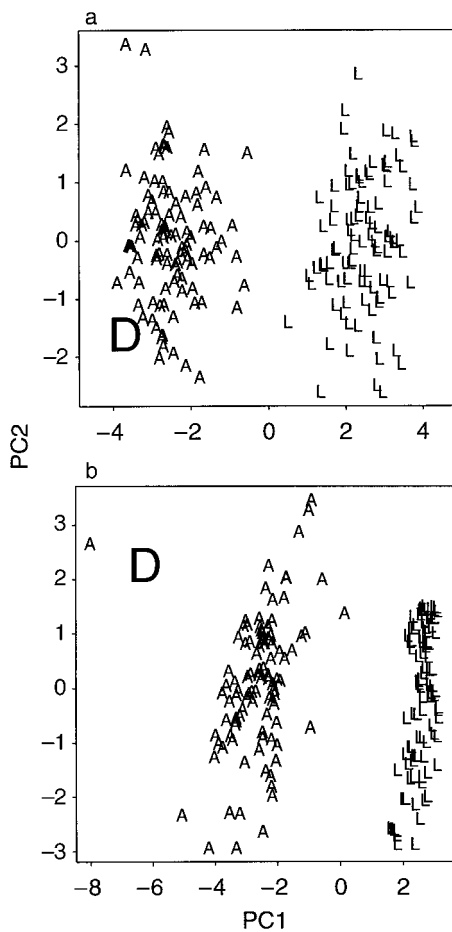† The minimum was the same for all runs of the trajectory-matching models.

FIG. 5. Principal-component analyses of the blowfly models, based on the time-series probes of the synthetic time series. A = AC model; L = LC model; D = data, plotted with the same PC scores (not included in the original analysis). (a) Models based on trajectory-matching fits; (b) models based on nonlinear-forecasting (NLF) fits.

models for such populations are not as fully developed. Often even a "failed" analysis can be useful. For example, we have been analyzing a spatially structured mite predator–prey system (Janssen et al. 1997), and found that without a detailed understanding of dispersal it is almost impossible to get a biologically sophisticated model to fit the data. This information does not exist, but our difficulties will guide future empirical research in this system.

Some of the techniques described in this paper can be applied to noncyclic populations, but we suspect that the overall approach will not work well with such populations. Certainly, the model-fitting techniques do not rely on the cyclic nature of the data (trajectory matching was originally developed for flatfish populations, and other techniques have been developed for fitting models to generic time-series data, especially in fisheries applications [e.g., Schnute 1994]). However, it is difficult to detect and accurately quantify the nonlinearities underlying the dynamics unless the population varies through a wide range of densities (Schaffer et al. 1986); this is more likely to be found in cyclic populations than those fluctuating erratically around an equilibrium. If the data are not cyclic there will be fewer probes available, so it will be more difficult to discriminate among the fitted models. Finally, the relative simplicity of the dynamical patterns in cyclic populations suggests that the dynamics might be caused by a few fairly simple processes, whereas the patterns in populations with irregular fluctuations might be the result of so many interacting processes that there is little hope of distinguishing among them.

What have we learned from the blowfly analysis that could help us in our long-term goal of uncovering mechanisms underlying cyclic dynamics in natural populations? One important lesson is that we cannot rely on a single measure of the match between model predictions and the data (such as cycle period), and should instead use a suite of probes. Multiple probes allow us to overcome the "many-to-one" problem: many different processes may cause a particular pattern observed in the data (such as a cycle with a given period). This problem is pervasive and must be remembered whenever a match between model and data is taken as evidence that the model is mechanistically valid. Indeed, we found that both the AC and LC models predict approximately the same cycle period. Additionally, both models predicted a correlation between adult density and fecundity (also observed in the data), even though in one model this was the mechanism driving the cycles, while in the other the correlation was a side effect of a completely separate process.

The use of multiple probes may be contrasted with a more conventional statistical approach in which competing models are evaluated using a single goodness-of-fit criterion (e.g., likelihood ratios). Any approach based on a single aggregate measure of how badly a model's output departs from the data is discarding a

nipulation, as is being done with laboratory populations of *Tribolium* (Constantino et al. 1995). The approaches we have advocated in this paper can play this role, as well as reducing the number of hypotheses that need to need to be considered experimentally.

If this approach is to be useful it needs to work well with systems more complex than an isolated population. Our work with pine looper suggests that it will work with multiple interacting species; an important methodological issue is how best to incorporate sporadic data about the other state variables. In the analyses described here we have treated environmental variability as "noise," but we could incorporate information about environmental variables into the models, and examine the structure of the residuals around the fits, to assess the importance of exogenous forcing in driving the cycles. Spatially structured populations will be a challenge, partly because the mechanistic

tremendous amount of information about the structure of the model's errors. Multiple probes, each getting at different features of the time series, offer the possibility of finding structural differences that a summary goodness-of-fit measure (such as r^2) glosses over. The unresolved difficulty in this approach is how to combine results from the different probes (apart from the ideal situation where all probes favor the same model). In this paper we have found that principal-component analysis is a useful tool for combining probes. We have successfully used this approach in other analyses, but do not yet know how generally applicable it will be. Another approach is to use probes that are inherently multi-valued, such as the spectral density function (Bjørnstad et al. 1996) or the distribution of ''local Lyapunov exponents'' (Ellner and Turchin 1995). Indeed, the latter proved to be the only probe that could reliably distinguish among competing models of measles epidemics (Bailey et al. 1997).

The second lesson of this case study is the value of general theory. We were able to reject immediately a large class of hypotheses that could cause cycles in blowflies only because the dynamical theory of single-species populations has matured to the point where there is a clear mapping between the dynamical mechanisms underlying cycles and some gross features of the fluctuation pattern. As far as we are aware, this sort of model synthesis has been done for only one other class of models. In populations with non-overlapping generations, feedback with a delay of one generation can produce cycles with a two-generation period, whereas feedback with a two-generation delay can generate cycles with a period of six or more generations (Royama 1977, Berryman 1996; unfortunately, these are commonly called ''direct'' and ''delayed'' density dependence, although the former is analogous to the delayed feedback discussed here). The former mechanism can arise, for example, from density-dependent fecundity, whereas the latter is typically caused by some sort of consumer–resource interactions (it can also result from maternal effects, Ginzburg and Taneyhill 1994). We would really like to see a general theory on how specific consumer–resource mechanisms map to a multidimensional space in which each axis is a specific probe (period, amplitude, stability, etc). Having such theory, and time-series data for a specific population, we would be able to quickly narrow down the spectrum of potential explanations by rejecting those that do not produce the correct quantitative measures. An additional benefit of such a theory would be a better understanding of which probes are most useful in discriminating among various explanations. This would provide guidance on the features of data that should be reported in empirical papers, or recorded in the first place. As we begin to analyze various field populations, we are frequently finding that, despite decades of field research, a critical piece of information was not recorded, or at least not reported.

Although we have learned much from the analysis of the blowfly data, many issues, especially technical ones, remain to be resolved (for the general issues of modeling population data, see Nisbet and Gurney [1982], Royama [1992], and Hilborn and Mangel [1997]). The fitting process can be improved by developing likelihood functions that incorporate our best information on the structure of the measurement and dynamic noise. We need a modeling framework that incorporates the possibility of long-term parameter drift (whether environmental or evolutionary). We also would like to know how long a time series need be to give a reasonable chance of success: How does the discriminatory power increase with series length, for example, and can many short time series substitute for one long one?

We single out one particular direction for future investigation. In the introduction we emphasized the contrast between phenomenological statistical models (no biological mechanisms specified) and mechanistic theoretical models (all relevant processes specified). These are actually two ends of a continuum. A fully mechanistic model must contain all important interactions, and must specify an explicit function for each of them; only parameter values are left to be estimated from the data. However, not all of our choices of functional forms may have a basis in data. To avoid an arbitrary choice of a functional form, we may wish to construct a model that is intermediate between the fully mechanistic and the completely phenomenological ends of the spectrum (Wood 1994*a*); we call this a ''semi-mechanistic model'' (Ellner et al. 1998). The idea is to build into the model all that we know about the biology, both quantitative and qualitative, and no more. Where a mechanistic model would require that the remaining unknown functions be guessed or chosen arbitrarily, in a semi-mechanistic model these factors would be replaced by flexible nonparametric functions, such as splines or neural networks, that would be free to find the shape that is most consistent with the time series (Ellner et al. 1997, 1998, Bjørnstad et al. 1998). Although we have not attempted such a semi-mechanistic approach in this paper, we think that this general superset of the current synthesis has great promise in helping us understand the mechanisms that cause cycles in natural populations. We hope soon to find out whether this promise is fulfilled, by applying this approach to such field case studies as cycles in forest insects and small mammals.

21535), the University of California-Santa Barbara, and the State of California.

Literature Cited

Anderson, R. M., and R. M. May. 1981. The population dynamics of microparasites and their invertebrate hosts. Philosophical Transactions of the Royal Society of London **B 291**:451–524.

Aron, J. L., and I. B. Schwartz. 1984. Seasonality and period-doubling bifurcations in an epidemic model. Journal of Theoretical Biology **110**:665–679.

Bailey, B. A., S. Ellner, and D. W. Nychka. 1997. Chaos with confidence: asymptotics and applications of local Lyapunov exponents. Pages 115–133 *in* C. D. Cutler and D. T. Kaplan, editors. Nonlinear dynamics and time series: building a bridge between the natural and statistical sciences. American Mathematical Society, Providence, Rhode Island, USA.

Batzli, G. O. 1992. Dynamics of small mammal populations: a review. Pages 831–850 *in* D. R. McCullough and R. H. Barrett, editors. Wildlife 2001: populations. Elsevier Applied Science, London, UK.

Berryman, A. A., editor. 1988. Dynamics of forest insect populations: patterns, causes, implications. Plenum, New York, New York, USA.

———. 1996. What causes population cycles of forest Lepidoptera? Trends in Ecology and Evolution **11**:28–32.

Bigger, M. 1973. An investigation by Fourier analysis into the interaction between coffee leaf-miners and their larval parasites. Journal of Animal Ecology **42**:417–434.

Bjørnstad, O. N., M. Begon, N. C. Stenseth, W. Falck, S. M. Sait, and D. J. Thompson. 1998. Population dynamics of the Indian meal moth: demographic stochasticity and delayed regulatory mechanisms. Journal of Animal Ecology **67**:110–126.

Bjørnstad, O. N., S. Champely, N. C. Stenseth, and T. Saitoh. 1996. Cyclicity and stability of grey-sided voles, *Clethrionomys rufocanus,* of Hokkaido: spectral and principal components analyses. Philosophical Transactions of the Royal Society of London **B 351**:867–875.

Brillinger, D. R., J. Guckenheimer, P. Guttorp, and G. Oster. 1980. Empirical modelling of population time series data: the case of age and density dependent vital rates. Lectures on Mathematics in the Life Sciences **13**:65–90.

Broekhuizen, N., H. F. Evans, and M. P. Hassell. 1993. Site characteristics and the population dynamics of the pine looper moth. Journal of Animal Ecology **62**:511–518.

Carpenter, S. R., K. L. Cottingham, and C. A. Stow. 1994. Fitting predator–prey models to time series with observation errors. Ecology **75**:1254–1264.

Casdagli, M. 1989. Nonlinear prediction of chaotic time series. Physica **D 35**:335–356.

Cheng, B., and H. Tong. 1992. On consistent nonparametric order determination and chaos (with discussion). Journal of the Royal Statistical Society **B 54**:427–474.

Constantino, R. F., J. M. Cushing, B. Dennis, and R. A. Desharnais. 1995. Experimentally induced transitions in the dynamic behaviour of insect populations. Nature **375**:227–230.

Dennis, B., R. A. Desharnais, J. M. Cushing, and R. F. Constantino. 1995. Nonlinear demographic dynamics: mathematical models, statistical methods, and biological experiments. Ecological Monographs **65**:261–281.

Dennis, B., R. A. Desharnais, J. M. Cushing, and R. F. Constantino. 1997. Transitions in population dynamics: equilibria to periodic cycles to aperiodic cycles. Journal of Animal Ecology **66**:704–729.

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, and C. J. Briggs. 1997. Inferring mechanism from time-series data: delay-differential equations. Physica **D 110**:182–194.

Ellner, S., D. W. Nychka, and A. R. Gallant. 1992. LENNS, a program to estimate the dominant Lyapunov exponent of noisy nonlinear systems from time series data. Institute of Statistics Mimeo Series number 2235 (BMA Series number 39). Statistics Department, North Carolina State University, Raleigh, North Carolina, USA.

Ellner, S., and P. Turchin. 1995. Chaos in a noisy world: new methods and evidence from time-series analysis. American Naturalist **145**:343–375.

Farmer, J. D., and J. J. Sidorowich. 1987. Predicting chaotic time series. Physical Review Letters **59**:845–848.

Ginzburg, L. R., and D. E. Taneyhill. 1994. Population cycles of forest Lepidoptera: a maternal effect hypothesis. Journal of Animal Ecology **63**:79–92.

Godfray, H. J. C., and M. P. Hassell. 1989. Discrete and continuous insect populations in tropical environments. Journal of Animal Ecology **58**:153–174.

Green, P. J., and B. W. Silverman. 1994. Nonparametric regression and generalized linear models: a roughness penalty approach. Chapman & Hall, London, 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., R. M. Nisbet, and J. H. Lawton. 1983. The systematic formulation of tractable single-species population models incorporating age structure. Journal of Animal Ecology **52**:479–495.

Hansen, L., and K. Singleton. 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. Econometrica **50**:1269–1286.

Hanski, I., P. Turchin, E. Korpimäki, and H. Henttonen. 1993. Population oscillations of boreal rodents: regulation by mustelid predators leads to chaos. Nature **364**:232–235.

Harrison, G. W. 1995. Comparing predator–prey models to Luckinbill's experiment with *Didinium* and *Paramecium.* Ecology **76**:357–374.

Hassell, M. P., J. H. Lawton, and R. M. May. 1976. Patterns of dynamical behaviour in single-species populations. Journal of Animal Ecology **45**:471–486.

Higgins, K., A. Hastings, J. N. Sarvela, and L. W. Botsford. 1997. Stochastic dynamics and deterministic skeletons: population behavior of Dungeness crab. Science **276**:1431–1435.

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

Hudson, P. J., D. Newborn, and A. P. Dobson. 1992. Regulation and stability of a free-living host–parasite system: *Trichostrongylus tenuis* in red grouse. I. Monitoring and parasite reduction experiments. Journal of Animal Ecology **61**:477–486.

Janssen, A., E. van Gool, R. Lingeman, J. Jacas, and G. van de Klashorst. 1997. Metapopulation dynamics of a persisting predator–prey system in the laboratory: time series analysis. Experimental & Applied Acarology **21**:415–430.

Jones, F. G. W., and J. N. Perry. 1978. Modelling populations of cyst-nematodes (Nematoda: Heteroderidae). Journal of Applied Ecology **15**:349–371.

Kaitala, V., E. Ranta, and J. Lindström. 1996. Cyclic population dynamics and random perturbations. Journal of Animal Ecology **65**:249–251.

Keith, L. B. 1963. Wildlife's ten-year cycle. University of Wisconsin Press, Madison, Wisconsin, USA.

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.

Klomp, H. 1966. The dynamics of a field population of the pine looper, *Bupalus piniarius* L. (Lep., Geom.). Advances in Ecological Research **3**:207–305.

Krebs, C. J. 1996. Population cycles revisited. Journal of Mammalogy **77**:8–24.

Krebs, C. J., S. Boutin, R. Boonstra, A. R. E. Sinclair, J. N. M. Smith, M. R. T. Dale, K. Martin, and R. Turkington. 1995. Impact of food and predation on the snowshoe hare cycle. Science **269**:1112–1115.

Krebs, C. J., and J. H. Myers. 1974. Population cycles in small mammals. Advances in Ecological Research **8**:267–399.

Laine, K., and H. Henttonen. 1983. The role of plant production in microtine cycles in northern Fennoscandia. Oikos **40**:407–418.

Lauwerier, H. A., and J. A. J. Metz. 1986. Hopf bifurcation in host–parasitoid models. IMA Journal of Mathematics Applied in Medicine & Biology **3**:191–210.

Lindström, J., E. Ranta, V. Kaitala, and H. Lindén. 1995. The clockwork of Finnish tetranoid dynamics. Oikos **74**:185–194.

Lotka, A. J. 1925. Elements of physical biology. Williams & Wilkins, Baltimore, Maryland, USA.

Manly, B. F. J. 1986. Multivariate statistical methods: a primer. Chapman & Hall, London, UK.

May, R. M. 1974. Biological populations with nonoverlapping generations: stable points, stable cycles, and chaos. Science **186**:645–647.

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

McFadden, D. 1989. A method of simulated moments for estimation of discrete response models without numerical integration. Econometrica **57**:831–860.

Middleton, A. D. 1934. Periodic fluctuations in British game populations. Journal of Animal Ecology **3**:231–249.

Monfort, A., and H. K. van Dijk. 1995. Simulation-based econometrics. Pages 1–20 *in* H. K. van Dijk, A. Monfort, and B. W. Brown, editors. Economic inference using simulation techniques. John Wiley & Sons, New York, New York, USA.

Moran, P. A. P. 1949. The statistical analysis of the sunspot and lynx cycles. Journal of Animal Ecology **18**:115–116.

Moss, R., A. Watson, and R. Parr. 1996. Experimental prevention of a population cycle in Red Grouse. Ecology **77**:1512–1530.

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

Nicholson, A. J., and V. A. Bailey. 1935. The balance of animal populations. Part I. Proceedings of the Zoological Society of London **3**:551–598.

Nisbet, R. M. 1997. Delay-differential equations for structured populations. Pages 89–118 *in* S. Tuljapurkar and H. Caswell, editors. Structured-population models in marine, terrestrial and freshwater systems. Chapman & Hall, New York, New York, USA.

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

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

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. 1992. Numerical recipes in C: the art of scientific computing. Second edition. Cambridge University Press, Cambridge, UK.

Ranta, E., J. Lindström, V. Kaitala, H. Kokko, H. Lindén, and E. Helle. 1997. Solar activity and hare dynamics: a cross-continental comparison. American Naturalist **149**:765–775.

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.

Rinaldi, S., S. Muratori, and Y. Kuznetsov. 1993. Multiple attractors, catastrophes and chaos in seasonally perturbed predator–prey communities. Bulletin of Mathematical Biology **55**:15–35.

Rosenzweig, M. L., and R. H. MacArthur. 1963. Graphical representation and stability conditions of predator–prey interactions. American Naturalist **97**:209–223.

Royama, T. 1977. Population persistence and density dependence. Ecological Monographs **47**:1–35.

———. 1992. Analytical population dynamics. Chapman & Hall, London, UK.

———. 1997. Population dynamics of forest insects: are they governed by single or multiple factors? Pages 37–48 *in* A. D. Watt, N. E. Stork, and M. D. Hunter, editors. Forests and insects. Chapman & Hall, London, UK.

Schaffer, W. M., S. Ellner, and M. Kot. 1986. Effects of noise on some dynamical models in ecology. Journal of Mathematical Biology **24**:479–523.

Schaffer, W. M., L. F. Olsen, G. L. Truty, and S. L. Fulmer. 1990. The case for chaos in childhood epidemics. Pages 138–166 *in* S. Krasner, editor. The ubiquity of chaos. American Association for the Advancement of Science, Washington, D.C., USA.

Schnute, J. T. 1994. A general framework for developing sequential fisheries models. Canadian Journal of Fisheries and Aquatic Sciences **51**:1676–1688.

Schwerdtfeger, F. 1941. Über die Ursachen des Massenwechsels der Insekten. Zeitschrift für angewandte Entomologie **28**:254–303.

Sinclair, A. R. E., and J. M. Gosline. 1997. Solar activity and mammal cycles in the Northern Hemisphere. American Naturalist **149**:776–784.

Sinclair, A. R. E., J. M. Gosline, G. Holdsworth, C. J. Krebs, S. Boutin, J. N. M. Smith, R. Boonstra, and M. Dale. 1993. Can the solar cycle and climate synchronize the snowshoe hare in Canada? Evidence from tree rings and ice cores. American Naturalist **141**:173–198.

Statistical Sciences. 1993. S-plus guide to statistical and mathematical analysis, version 3.2. MathSoft, Seattle, Washington, USA.

Stenseth, N. C., O. N. Bjørnstad, and W. Falck. 1996. Is spacing coupled with predation causing the microtine density cycle? A synthesis of current process-oriented and pattern-oriented studies. Proceedings of the Royal Society of London **B 263**:1423–1435.

Stenseth, N. C., and R. A. Ims. 1993. Population dynamics of lemmings: temporal and spatial variation—an introduction. Pages 61–96 *in* N. C. Stenseth and R. A. Ims, editors. The biology of lemmings. Linnean Society, London, UK.

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.

Straw, N. A. 1996. The impact of pine looper moth, *Bupalus piniaria* L. (Lepidoptera; Geometridae) on the growth of

Scots pine in Tentsmuir Forest, Scotland. Forest Ecology and Management **87**:209–232.

Takens, F. 1981. Detecting strange attractors in turbulence. Pages 366–381 *in* D. A. Rand and L. S. Young, editors. Dynamical systems and turbulence. Springer-Verlag, Berlin, Germany.

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 & Hall, New York, New York, USA.

Turchin, P. 1993. Chaos and stability in rodent population dynamics: evidence from nonlinear time-series analysis. Oikos **68**:167–172.

Turchin, P., and I. Hanski. 1997. An empirically based model for latitudinal gradient in vole population dynamics. American Naturalist **149**:842–874.

Turnock, W. J. 1969. Predation by larval Elateridae on pupae

of the pine looper, *Bupalus piniarius* (L.). Netherlands Journal of Zoology **19**:393–416.

Volterra, V. 1926. Fluctuations in the abundance of a species considered mathematically. Nature **118**:558–560.

Watson, A., R. Moss, P. Rothery, and R. Parr. 1984. Demographic causes and predictive models of population fluctuations in Red Grouse. Journal of Animal Ecology **53**: 639–662.

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

———. 1994b. Spline models of biological population dynamics: how to estimate mortality rates for stage structured populations with dimorphic life histories. IMA Journal of Mathematics Applied in Medicine & Biology **11**: 61–78.

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

APPENDIX

Details of the models for Nicholson's blowflies

Adult competition (AC) model.—We represent density-dependent fecundity by setting $B(t) = qA(t)e^{-A(t)/A_0}$, where q is the maximum per capita fecundity and A_0 is the density at which total fecundity is maximum. Assuming that the development time (τ) and the egg-to-adult survival (s_J) are constant, the recruitment rate is given by

$$R(t) = s_J B(t - \tau) = s_J qA(t - \tau)e^{-A(t-\tau)/A_0}$$

where t = time, $A(t)$ = adult density. Finally, assuming a constant adult mortality rate (μ_A) allows us to reduce Eq. 1 to

$$\frac{dA(t)}{dt} = s_J qA(t - \tau)e^{-A(t-\tau)/A_0} - \mu_A A(t).$$

Larval competition (LC) model.—We assume that the juveniles gain mass at a density-dependent rate

$$g(t) = \max\left\{0, \frac{1}{1 + J(t)/J_0} - \Gamma\right\},$$

where Γ is a maintenance cost and J_0 is the density at which the uptake rate is half its maximum. Again assuming a constant development time, the per capita mass at maturity is given by

$$G(t) = \int_{t-\tau}^{t} g(x)\, dx.$$

Assuming that fecundity is directly proportional to mass and mortality is density independent, we have

$$\frac{dA(t)}{dt} = bW(t - \tau)s_J - \mu_A A(t)$$

$$\frac{dJ(t)}{dt} = bW(t) - \mu_J J(t) - bW(t - \tau)s_J$$

$$\frac{dW(t)}{dt} = G(t)bW(t - \tau)s_J - \mu_A W(t)$$

where $W(t)$ is the total mass of the adult population at time t, b is the fecundity per unit mass of adults, and μ_J and μ_A are the density-independent per capita death rates of juveniles and adults, respectively; $s_J = e^{-\mu_J \tau}$. Notice that the dynamic state variables are J and W; the equation for A is included only to facilitate comparison with the observed adult densities.

Model implementation details

The deterministic implementation of the models (used for trajectory matching) was done with a second-order Runge-Kutta numerical integrator, using cubic interpolation to get the lagged values of the state variables. The stochastic implementation (for parameter estimation by nonlinear forecasting) used a simple Euler integration scheme with a short-enough time step that lagged values could be obtained without any interpolation. Stochasticity was introduced as noise in the fecundity parameter q, which was modeled by treating $\log q$ as a first-order Gaussian autoregressive process. Both the variance and the autocorrelation of $\log q$ were fitted parameters.

Fitting details

Direct fits.—The AC model was fitted by nonparametrically smoothing the observed time series, using a LOESS smooth (Green and Silverman 1994). Such a smooth yields derivative estimates for each state-variable sample, as well as a smoothed estimate of the variable itself. We now have all the relevant variables, as we can estimate $A(t - \tau)$ at times between the sample times using the smooth. For each of a closely spaced range of τ values we got the best fit to the AC model, using the nonlinear least-squares optimizer in S-plus (Statistical Sciences 1993), and chose the value of τ that gave the best overall fit. For further details see Ellner et al. (1997).

Trajectory matching.—Model parameters (including the initial values of the state variables) were chosen by minimizing the sum of squared differences between the numerical solution of the model and the data. Iteratively re-weighted least squares were used to obtain maximum-likelihood parameter estimates for non-normal error models.

Care must be taken in the numerical solution of the model for this to work reliably. The models were solved by an adaptive stepping routine in order to avoid numerical instability in the solution for any trial parameters. With delay-differential models it is important to ensure that the order of approximation of lagged variables is better than the order of the integration scheme; failure to do so can lead to a non-smooth objective function. For moderately complicated models derivatives of the model with respect to parameters are usually obtained numerically by finite differencing. This is only reliable in general if the truncation and approximation errors of the approximations are monitored and the differences adjusted accordingly. Additionally, model runs used for differencing should use the integration mesh implied by

the current parameter set, and not adjust the integration mesh by adaptive stepping, if accuracy in the derivatives is not to be thrown away.

Parameter estimation should be done by constrained methods, both to include biologically sensible parameter bounds and to ensure that impossible parameter values are not passed to the model integration method. For most ecological analyses a constrained Quasi-Newton method is an appropriate way of minimizing the least-squares objective, since the residual variability will tend to be far too large for Gauss-Newton methods to perform efficiently; at present it is unclear whether sequential quadratic programming or a more traditional approach will be faster.

Nonlinear forecasting.—Nonlinear forecasting accuracy (NLF) is closely related to what statisticians call a ''generalized method of moments'' (Hansen and Singleton 1982). A ''moment'' in this context is any quantity that describes (and can be computed directly from) the times series; all of the probes discussed in this paper are moments in this sense. The selected moments p_i are computed for the model and for the data, the differences between p_i(data) and p_i(model) are combined into an overall measure of difference, and model parameters are chosen to minimize this measure of difference. The ''method of simulated moments'' (McFadden 1989; see Monfort and van Dijk [1995] for a thorough review) refers to the situation in which model moments are estimated by simulation, as in the current analysis. The choice of moments is typically ad hoc. In NLF, moments are the conditional mean of the observed state variable $x(t)$ given past values, which we obtain by constructing a nonlinear forecasting model for the observed state variables. We used the sum of squared forecasting errors as our overall measure of difference, which is equivalent to the conventional least-squares fitting criterion for regression or time-series models.

This procedure is implemented numerically as follows (see Ellner et al. [1998] for full details). First, a time-delay state vector is created from the data:

$$\mathbf{X}_t = (x_t, x_{t-L}, \ldots, x_{t-mL})$$

where L is the lag interval and m is the number of lags, and a similar vector $\mathbf{Y}_t$ is made from the model output. In order to make a forecast prediction for x at some time in the future, $t + T_p$, we find the distance between the data point $\mathbf{X}_t$ and each model point $\mathbf{Y}_i$. Our prediction of x_{t+T_p} is the weighted mean of y_{i+T_p}, with the weighting being a function of the distance between $\mathbf{X}_t$ and $\mathbf{Y}_i$. In statistical jargon, we are fitting a kernel regression to the model output, and using the regression to predict the data. A kernel regression has only one fitted parameter, namely the bandwidth h, which controls how rapidly more distant $\mathbf{Y}_i$ are down-weighted in computing the weighted average. Here we selected h by cross-validation, which is a standard method for kernel regression (see e.g., Green and Silverman 1994).

Before applying NLF we must choose the prediction interval (T_p), the lag interval (L), and the number of lags (m). We chose these on the basis of the data, using the method of Cheng and Tong (1992), which uses kernel regression to find values for which the data are best able to predict themselves (i.e., the vectors $\mathbf{Y}_i$ described above are constructed from the data themselves rather than from model output, with suitable care that there is no overlap between the components of the $\mathbf{X}$'s and $\mathbf{Y}$'s). An exhaustive search was run over all possible combinations of these parameters, with the result that the best predictions overall were obtained with a lag of $L = 6$ d, and just one lag. After an initial drop-off the prediction r^2 was nearly constant over a range of prediction intervals up to ~20 d, so we used the average prediction r^2 at 6, 12, and 20 d as the fitting criterion.

Minimization of the nonlinear forecasting error (the fitting criterion) was carried out numerically using Powell's method (Press et al. 1992) for nonlinear minimization. To allow the minimization algorithm to converge we used the same random-number seed for all simulations of the model; this makes the estimated forecasting error a continuous function of model parameters.