👤 **Peer Reviewed**

Title:
Population cycles in the pine looper moth: Dynamical tests of mechanistic hypotheses

Author:
Kendall, Bruce E.; Ellner, S P; McCauley, E; Wood, S N; Briggs, C J; Murdoch, W M et al.

Abstract:
The forest insect pest Bupalus piniarius (pine looper moth) is a classic example of a natural population cycle. As is typical for Populations that exhibit regular oscillations in density, there are several biological mechanisms that are hypothesized to be responsible for the cycles; but despite several decades of detailed study there has been no definite conclusion as to which mechanism is most important. We evaluated three hypotheses for which there was direct experimental evidence: (1) food quality (nutritional value of pine needles affected by defoliation); (2) parasitoids (trophic interactions with specialist parasitoids), and (3) maternal effects (maternal body size affects the performance of offspring). We reviewed the empirical evidence for each of these hypotheses and expressed each hypothesis in the form of a mechanistic dynamic model. We used a nonlinear forecasting approach to fit each model to three long-term Population time series in Britain that exhibit some degree of regular cycling, and we used parametric bootstrap to evaluate the significance of differences between models in their goodness of fit to the data. The results differed among the three forests: at Culbin, the parasitoid and maternal effects models fit equally well; at Roseisle, the food quality and maternal effects models fit equally well; and at Tentsmuir, the parasitoid model fit best. However, the best-fit parasitism models required that the parasitism rate vary between nearly 0 and nearly 1 during a cycle, greatly exceeding the range of parasitism rates that have been observed in the field. In contrast, the required variation in the observable maternal quality variable (pupal mass) was within the range of empirical observations. Under mild

constraints on the parasitism rate (though allowing a much wider range than has been measured in B. piniarius at any location), the fit of the parasitism model fell off dramatically. The maternal effects model then had uniformly strong support, outperforming the constrained parasitism model at all three sites and the food quality model at two; it performed slightly better than the food quality model at the remaining site. This represents the first system in which the maternal effects hypothesis for population cycles has been supported by both strong biological and dynamical evidence.

POPULATION CYCLES IN THE PINE LOOPER MOTH: DYNAMICAL TESTS OF MECHANISTIC HYPOTHESES

Bruce E. Kendall,[1,2,3,9] Stephen P. Ellner,[3] Edward McCauley,[4] Simon N. Wood,[5,10] Cheryl J. Briggs,[6] William W. Murdoch,[7] and Peter Turchin[8]

[1]*Donald Bren School of Environmental Science and Management, University of California,
Santa Barbara, California 93106-5131 USA*
[2]*National Center for Ecological Analysis and Synthesis, University of California, Santa Barbara, California 93106 USA*
[3]*Department of Ecology and Evolutionary Biology, Cornell University, Ithaca, New York 14853 USA*
[4]*Ecology Division, Department of Biology, University of Calgary, Calgary, Alberta, T2N 1N4, Canada*
[5]*The Mathematical Institute, North Haugh, St. Andrews, Fife KY16 9SS, UK*
[6]*Department of Integrative Biology, University of California, Berkeley, California 94720 USA*
[7]*Department of Ecology, Evolution, and Marine Biology, University of California,
Santa Barbara, California 93106-9610 USA*
[8]*Department of Ecology and Evolutionary Biology, University of Connecticut, Storrs, Connecticut 06269 USA*

Abstract. The forest insect pest *Bupalus piniarius* (pine looper moth) is a classic example of a natural population cycle. As is typical for populations that exhibit regular oscillations in density, there are several biological mechanisms that are hypothesized to be responsible for the cycles; but despite several decades of detailed study there has been no definite conclusion as to which mechanism is most important. We evaluated three hypotheses for which there was direct experimental evidence: (1) food quality (nutritional value of pine needles affected by defoliation); (2) parasitoids (trophic interactions with specialist parasitoids); and (3) maternal effects (maternal body size affects the performance of off-spring). We reviewed the empirical evidence for each of these hypotheses and expressed each hypothesis in the form of a mechanistic dynamic model. We used a nonlinear fore-casting approach to fit each model to three long-term population time series in Britain that exhibit some degree of regular cycling, and we used parametric bootstrap to evaluate the significance of differences between models in their goodness of fit to the data. The results differed among the three forests: at Culbin, the parasitoid and maternal effects models fit equally well; at Roseisle, the food quality and maternal effects models fit equally well; and at Tentsmuir, the parasitoid model fit best. However, the best-fit parasitism models required that the parasitism rate vary between nearly 0 and nearly 1 during a cycle, greatly exceeding the range of parasitism rates that have been observed in the field. In contrast, the required variation in the observable maternal quality variable (pupal mass) was within the range of empirical observations. Under mild constraints on the parasitism rate (though allowing a much wider range than has been measured in *B. piniarius* at any location), the fit of the parasitism model fell off dramatically. The maternal effects model then had uniformly strong support, outperforming the constrained parasitism model at all three sites and the food quality model at two; it performed slightly better than the food quality model at the remaining site. This represents the first system in which the maternal effects hypothesis for population cycles has been supported by both strong biological and dynamical evidence.

Key words: food quality; host–parasitoid dynamics; maternal effects; nonlinear dynamics; population cycles.

Introduction

Population cycles in forest defoliating insects typically have periods of 6–12 years, whereas the insects themselves live for one year or less. Models attempting to understand these cycles have incorporated parasitoids (Hassell and Varley 1969, Beddington et al. 1976), predators (Turchin 2003), pathogens (Anderson and

May 1980, 1981, Dwyer 1994), variation in food quality (Edelstein-Keshet and Rausher 1989) and quantity, maternal effects (Myers 1988, Ginzburg and Taneyhill 1994), and rapid evolution (Chitty 1967, Baltensweiler 1984, Myers 1988), each of which can generate long-period cycles in models. Weather fluctuations are not sufficiently periodic to explain the regular oscillations. In contrast, direct density dependence cannot produce long-period cycles by itself (Morris 1990), but can stabilize otherwise divergent oscillations produced by the other mechanisms.

Careful examination of a particular population usually reveals empirical evidence that several of these

processes are occurring, so it is difficult to identify which is driving the dynamics. For example, in Switzerland some larch budmoth (*Zeiraphera diniana*) outbreaks were attacked by a viral pathogen, which caused up to 50% annual mortality during the subsequent population collapse, so the host–pathogen interaction was for a time considered to be a likely explanation for the budmoth cycle (Auer 1968, Anderson and May 1980). However subsequent cycles proceeded in the absence of pathogen attack, without any evident change in period or amplitude (Baltensweiler et al. 1977). So although pathogen interactions were demonstrably present and strong, a natural experiment (extinction of the pathogen) demonstrated they were simply ''along for the ride'' while other factors actually caused the large-scale changes in population density.

Consequently, the driving mechanisms of cycles have been identified with some degree of confidence for only two species. Larch budmoth cycles in the Swiss Alps are driven primarily by parasitoids, with food quality dynamics modifying, but not driving, the cycles (Turchin et al. 2003). This conclusion was reached by analyzing time series of not only the insect, but also of parasitoids and needle quality. In contrast, the cycles of southern pine beetle (*Dendroctonus frontalis*) in the southeastern United States appear to be driven by interactions with a specialist beetle predator (Turchin 2003). This conclusion was reached by comparing the results from multiyear experimental manipulations to predictions from models (Turchin et al. 1999). We are not aware of any other strong conclusions about forest insect cycles.

The pine looper moth (*Bupalus piniarius*) often exhibits regular oscillations in pine plantations; this has been extensively documented in Germany and Britain (e.g., Schwerdtfeger 1942, Barbour 1988). At peak densities, the insect can completely defoliate the trees; while the trees survive, their growth in reduced, and *Bupalus* outbreaks are of concern to foresters. The life history and ecology of the insect has been studied extensively and in great detail, including two long-term life table studies (Klomp 1966, Broekhuizen et al. 1994) and a number of short-term behavioral experiments (Gruys 1970, 1971, Šmits and Larsson 1999, Šmits et al. 2001). These studies have documented strong evidence (summarized in *The data: Possible causal mechanisms*) for three of the above-mentioned processes: interactions with parasitoids, maternal effects on offspring quality, and reduced food quality in the wake of defoliation. But despite this concerted empirical effort (which resulted in at least three Ph.D. dissertations [Gruys 1970, Barbour 1980, Broekhuizen 1991] and over 880 pages of peer-reviewed publications), no definite conclusion has emerged as to which of the possible mechanisms causes the cycles.

The ideal experimental test would manipulate conditions in the field so that all but one putative mechanism was eliminated. Such experiments are expensive and time-consuming (e.g., Krebs et al. 1995), and in many populations are impossible or prohibitively difficult. Instead, information must be extracted from observational data and limited experimental manipulations. We have argued (Kendall et al. 1999) that modern statistical methods and computational power greatly increase the amount of information that can be obtained by this approach. By constructing models representing (or constrained by) experimental studies of potential mechanisms, and using the observed dynamics to estimate unknown parameters, we come as close as possible to simulating the ideal experiment by asking how well each possible mechanism can account for all available data.

Here we follow the philosophy of Kendall et al. (1999) to determine which of the empirically supported mechanisms is most likely to be driving the dynamics of three British *Bupalus* populations. First, we develop mathematical models that encapsulate each mechanism. These models are not ''off-the-shelf'' conventional models (such as the parasitoid models in Beddington et al. [1976] and the maternal effects models in Ginzburg and Taneyhill [1994]); rather, we developed models from the ground up to specifically match the currently known biology of *Bupalus*. Then we fit these models to the empirical time series of *Bupalus* abundance: we use nonlinear minimization to find the model parameter values that allow the models to best match the observed dynamics. We use these parameterized models as representatives of the respective mechanisms, and compare them, based on their goodness of fit and the nature of the time series they produce under simulation. We also examine whether the required fluctuations in the unobserved dynamic variables, such as parasitism rate and maternal quality, are biologically plausible and consistent with the available data.

THE DATA

Population dynamics

Bupalus is univoltine, overwintering as pupae with adults emerging in late May and June in the United Kingdom (Broekhuizen et al. 1994). Adults live for 10–14 days without feeding, and their eggs hatch after ~20 days. There are generally five larval instars lasting a total of 4–5 months. In late autumn the final larval stage descends from the canopy into the forest litter and passes through a brief prepupal stage before pupating.

Since 1953 the Forestry Commission of Great Britain has been keeping records of *Bupalus* abundance in 44 forests. The primary data are density estimates from an annual winter survey of pupal abundance. This is a stratified sample, with transects of small plots being taken across a number of forest ''compartments.'' For a full description of the sampling protocol, see Broekhuizen et al. (1993). We analyzed time series of pupal

TABLE 1.   Nonlinear predictability of *Bupalus* pupal density time series.

Forest	1 year	2 years	3 years
Cannock	0.11	†	†
Culbin	0.31	0.27	0.22
Laughton	†	†	†
Lossie	†	†	†
Monaughty	0.02	0.05	†
Roseisle	0.42	0.16	0.08
Sherwood 2	0.20	†	†
Sherwood 3	0.01	†	†
Sherwood 4	†	†	†
Sherwood 5	†	†	†
Speymouth	0.15	0.04	0.08
Tentsmuir	0.51	0.43	0.43
TunstaII	0.16	†	†

Notes: Predictability was estimated by the method of Cheng and Tong (1992). Bjørnstad et al. (2001: Box 1) give a detailed explanation of the method. It involves fitting a nonparametric kernel autoregression model and scanning over possible combinations of lags to optimize the cross-validation estimate of prediction error variance. As per Nychka et al. (1992) we used ''leave out k'' cross validation with $k = 7$, so that a block of data roughly equal to the duration of an outbreak cycle, centered on the time at which a prediction is made, was omitted in making each prediction of a future *Bupalus* density. The results in this table are the r^2 (proportion of variance explained) for the optimized combination of lags to predict *Bupalus* larval density 1, 2, or 3 years into the future, on a 0.2-power transformed scale. Daggers (†) indicate an estimate of complete unpredictability, such that out-of-sample prediction is estimated to be no more accurate than ignoring the recent past and using the overall mean density as the ''prediction.'' Characteristics of the various forests are described by Broekhuizen et al. (1993).

density averaged across all compartments in each forest, as presented in Broekhuizen et al. (1993), and extended with more recent data from the Forestry Commission. The data were analyzed following 0.2-power transformation, for reasons explained below in the section *Fitting models by nonlinear forecasting*.

Many of the forests are not suitable for analysis due to obvious nonstationarities in the data, such as trends in mean density. The remainder were tested for the presence of regular oscillations, based on their nonlinear predictability (Table 1). A mechanistic population model specifies the ''rules'' relating past, present, and future population densities. In order for a data set to be useful for distinguishing among models, the signal-to-noise ratio of those rules must be sufficiently high in the data. That is, the data must exhibit repeatable patterns rather than each *Bupalus* outbreak being completely idiosyncratic, so that population changes during one period can be predicted based on the patterns of change at other times in the same data set. The degree of predictability for *Bupalus* density dynamics varied greatly among forests (Table 1). On the basis of these estimates, we restricted our analysis to the three forests with the greatest degree of regularity in their dynamics, all in Scotland: Culbin, Roseisle, and Tentsmuir. The remaining time series are either fluctuating unpredictably around an equilibrium or display multiyear out-

breaks that are each idiosyncratic. The three time series that we analyze are 34–38 years long, and are displayed at the bottom of Figs. 2–4.

Possible causal mechanisms

Parasitoids.—A wide variety of parasitoids attack *Bupalus* but only two species are both abundant and fairly specialized on *Bupalus*. *Cratichneumon viator* attacks *Bupalus* pupae, having two generations per year; parasitism rates as high as 75% have been recorded for the two generations combined (Davies 1961). *C. viator* may have a third generation on an alternate host; this is a question that has not been satisfactorily resolved. Intensive studies in one British forest (Cannock) revealed parasitism rates of 8–41% in the first generation and 0–36% in the second (Broekhuizen et al. 1994). *C. viator* is a generalist, but *Bupalus* appears to be its primary host in Britain (Broekhuizen 1991).

Dusona oxyacanthae is a specialist parasitoid of *Bupalus* larvae; there is a strong negative relationship between *Dusona* density and *Bupalus* population growth rate (Barbour 1988). Parasitism rates of 2–74% were recorded in the study at Cannock (Broekhuizen et al. 1994). The *Dusona* attack rate appears to have a type I (linear) functional response, and declines with increasing parasitoid density (Broekhuizen 1991). The two parasitoids do not seem to be strongly synchronized.

For comparing our models with data on parasitism rates it is useful to have a summary measure of the overall intensity of parasitism. Following Broekhuizen et al. (1994) let $\pi_{t,i}$, $i = 1, 2, 3$ denote the stage-specific fraction of *Bupalus* in year t parasitized by *Dusona* and the two generations of *C. viator*, respectively. In the absence of any other sources of mortality, the net ''survival'' (the fraction of individuals that escape parasitism) would be $(1 - \pi_{t,1})(1 - \pi_{t,2})(1 - \pi_{t,3})$. As a measure of the total parasitism rate we therefore use the quantity

$$\prod_t = 1 - (1 - \pi_{t,1})(1 - \pi_{t,2})(1 - \pi_{t,3}). \qquad (1)$$

Over the 11 years studied at Cannock the total parasitism rate, calculated from the values of $\pi_{t,i}$ in Table 6 of Broekhuizen et al. (1994), ranged from 0.16 to 0.81.

Maternal effects.—Mean pupal mass varies substantially among years, from 130 to 190 mg in The Netherlands and from 85 to 140 mg in Cannock (Klomp 1968, Broekhuizen et al. 1994). This affects the fecundity of the adults that emerge: adults that emerge from large pupae have higher fecundity (Klomp 1966, 1968). In turn, pupal mass is a declining function of the density of young larvae earlier that year (Klomp 1966). This collapses into a form of direct density dependence, which would not generate long-period cycles. However, our analysis of Klomp's (1968) data (see

Models: Maternal effects) reveals that pupal size also positively affects the larval survival of the offspring generation. This is a true maternal effect, and provides the delayed feedback necessary for long-period cycles. Studies in Cannock did not reveal a relationship between pupal mass and larval density (Broekhuizen et al. 1994), and pupal masses are generally low compared with the Dutch studies. However, Barbour (1988) found that pupal mass fluctuates through the *Bupalus* cycle, with mass peaking in advance of *Bupalus* density; Barbour (1980) found a negative relationship between pupal mass and density two years previously.

Food quality.—*Bupalus* larvae prefer to feed on older foliage (Barbour 1988). Recent studies in Sweden (Šmits and Larsson 1999) have quantified that preference, and showed that larvae forced to feed on young needles have lower survival rates than those allowed to feed on mature needles. Needle age also affects fecundity: adults forced to lay eggs on new needles had lower fecundity than those laying on mature needles; furthermore, eggs have a substantial chance of falling off new needles (Šmits et al. 2001). Not surprisingly, ovipositing adults display a strong preference for mature needles (Šmits et al. 2001). The relationship between the needle age distribution and *Bupalus* densities have not been quantified, but at outbreak densities *Bupalus* can completely defoliate the pine trees. This would mean that in the following year all needles would be new, and thus less favorable for the *Bupalus* population. This provides a second-order feedback that could lead to population cycles.

MODELS

We describe here the assumptions of our models, and the biological evidence on which they are based. Some additional details are in Appendix A. All of the models assume that direct density-dependent population regulation acts during *Bupalus* larval survival, as this is the stage where Klomp (1966) found most of the important density dependence in his *k*-factor analysis. Each model adds to this a single form of delayed density dependence.

Interaction with parasitoids

Neither of the two major parasitoids that attack *Bupalus* in Britain complies with standard models for host–parasitoid dynamics.

C. viator attacks *Bupalus* pupae, and is multivoltine. It goes through two generations on a single generation of *Bupalus* pupae, and there is controversy over whether there is a third generation on an alternate host (Broekhuizen 1991). Here we assume that there is no alternate host, or that the alternate host is not dynamically coupled to the *Bupalus–C. viator* system. The first generation of *C. viator* attacks before the pupal survey in our data sets, and attacked pupae are not included in the *Bupalus* density estimates. The model assumes that the first generation of *C. viator* attacks

young pupae after *Dusona* parasitism and density-dependent mortality, but before the survey. The second generation attacks after the survey, before reproduction occurs.

Dusona attacks *Bupalus* larvae, but does not kill them immediately. Parasitized larvae continue feeding and growing, and therefore contribute to density-dependent mortality of *Bupalus* larvae. They die in the process of pupating, and are converted into *Dusona* pupae. Consequently, the attack rate by *Dusona* is a function of the density of young *Bupalus* larvae, before density-dependent larval mortality (assumed to be proportional to *Bupalus* egg density). However, the number of *Dusona* produced is proportional to *Bupalus* pupal density after density-dependent mortality.

For both species, the general parasitoid model uses an attack rate equation incorporating both parasitoid interference and parasitoid functional response. That is, the fraction of *Bupalus* attacked per parasitoid is a decreasing function of both *Bupalus* and parasitoid density. In order to limit the number of parameters in the model, we assume that the *C. viator* attack parameters are the same in the two generations.

The resulting model is as follows. Let N_t be the *Bupalus* pupal density at the time of the survey. *Bupalus* adults (A_t) result from the survivors of attack by the second *C. viator* generation, while attacked *Bupalus* pupae give rise to the first *C. viator* generation the following year:

$$A_t = s_A N_t \exp\left(-\frac{a_c C_{2,t}}{1 + a_c h_c N_t + a_c b_c C_{2,t}}\right) \qquad (2)$$

$$C_{1,t+1} = s_{c_1} N_t \left[1 - \exp\left(-\frac{a_c C_{2,t}}{1 + a_c h_c N_t + a_c b_c C_{2,t}}\right)\right] \quad (3)$$

where s_A is the survival of late pupae, s_{c_1} is the fraction of *C. viator* eggs that survive to maturity, a_c is the attack rate, h_c is the handling time, and b_c is the parasitoid interference coefficient.

Density-independent fecundity which results in young larvae

$$L_{t+1} = A_t e^r \qquad (4)$$

is followed by density-dependent survival of all larvae, including those parasitized by *Dusona*:

$$\tilde{L}_{t+1} = L_{t+1}[\exp(-\mu_0 - cL_{t+1})] \qquad (5)$$

where $\tilde{L}$ is the density of old larvae. The attack rate by *Dusona* is a function of L, but the resulting mortality acts on $\tilde{L}$:

$$P_{t+1} = \tilde{L}_{t+1} \exp\left(-\frac{a_d D_{t+1}}{1 + a_d h_d L_{t+1} + a_d b_d D_{t+1}}\right) \qquad (6)$$

$$D_{t+2} = s_d \tilde{L}_{t+1}\left[1 - \exp\left(-\frac{a_d D_{t+1}}{1 + a_d h_d L_{t+1} + a_d b_d D_{t+1}}\right)\right] \quad (7)$$

where s_d, a_d, h_d, and b_d are the *Dusona* equivalents of s_{c_1}, a_c, h_c and b_c and P is the density of new pupae.

The first generation of *C. viator* attacks *Bupalus* before the next year's survey:

$$N_{t+1} = s_p P_{t+1} \exp\left(-\frac{a_c C_{1,t+1}}{1 + a_c h_c P_{t+1} + a_c b_c C_{1,t+1}}\right) \quad (8)$$

$$C_{2,t+1} = s_{c_2} P_{t+1}\left[1 - \exp\left(-\frac{a_c C_{1,t+1}}{1 + a_c h_c P_{t+1} + a_c b_c C_{1,t+1}}\right)\right] \quad (9)$$

where s_p is the survival of early pupae.

The number of parameters can be reduced by applying standard steps: rescaling the (unmeasured) parasitoid density variables (e.g., $D \leftarrow D/s_d$) and associated attack parameters, grouping parameters (e.g., absorbing s_A, s_p, and $e^{-\mu_0}$ into the value of r), and eliminating L using Eq. 4. The resulting model is

$$A_t = N_t \exp\left(-\frac{a_c C_{2,t}}{1 + a_c h_c N_t + a_c b_c C_{2,t}}\right)$$

$$C_{1,t+1} = s_c N_t\left[1 - \exp\left(-\frac{a_c C_{2,t}}{1 + a_c h_c N_t + a_c b_c C_{2,t}}\right)\right]$$

$$L_{t+1} = A_t \exp[r(1 - A_t/K)]$$

$$P_{t+1} = L_{t+1} \exp\left(-\frac{a_d D_{t+1}}{1 + a_d h_d A_t + a_d b_d D_{t+1}}\right)$$

$$D_{t+2} = L_{t+1}\left[1 - \exp\left(-\frac{a_d D_{t+1}}{1 + a_d h_d A_t + a_d b_d D_{t+1}}\right)\right]$$

$$N_{t+1} = P_{t+1} \exp\left(-\frac{a_c C_{1,t+1}}{1 + a_c h_c P_{t+1} + a_c b_c C_{1,t+1}}\right)$$

$$C_{2,t+1} = P_{t+1}\left[1 - \exp\left(-\frac{a_c C_{1,t+1}}{1 + a_c h_c P_{t+1} + a_c b_c C_{1,t+1}}\right)\right] \quad (10)$$

where L_{t+1} now represents old larvae (parasitized and unparasitized). Note that unless $s_{c_1} = s_{c_2}$, only one of these can be eliminated by rescaling (equivalently, we could set both of these to one by allowing different values of a_c in the two generations; in terms of *Bupalus* dynamics these two rescalings are equivalent). A, L, and P could be eliminated from the model, but the resulting equations would be terrifying to behold.

Model 10 has nine dynamic parameters. To reduce this somewhat (so that all models being compared have a similar number of parameters), we mainly considered the model with type-I parasitoid functional response (i.e., $h_c = h_d = 0$), which proved to fit our data almost as well as the model with a type-II functional response (recall that Broekhuizen 1991 found a type-I functional response by *Dusona* at Cannock). We tried eliminating parasitoid interference but this gave much poorer fits for Tentsmuir, where the evidence for parasitoid effects is strongest.

We also examined models with one or the other of the parasitoids acting by itself, to represent the hypothesis that only one of the two parasitoids is strongly coupled to *Bupalus*. In these models we restored the parasitoid functional response parameter. The equations for these reduced parasitoid models are given in Appendix A.

In fitting these models to data we imposed the constraint that all parasitoid species included in the model had to persist (with a density of 10^{-8} as the extinction threshold). Without this constraint, the optimized parameter fits to the *Bupalus* data sometimes had the unrealistic property that one or both parasitoids would go extinct within a few generations.

Maternal effects

Our model for maternal effects is based on the work of Klomp (1966, 1968). Because these studies are from the Netherlands while our population time series are from the United Kingdom, we use Klomp's results to derive the form of the model but not to estimate or constrain the numerical values of any parameters.

The measurable index of individual quality is pupal mass (W). Klomp (1966) demonstrated that a female's mass as a pupa had positive effects on her subsequent fecundity as an adult, and on the egg-to-adult survival of her offspring. In contrast to the original theoretical model for population cycles driven by maternal effects (Ginzburg and Taneyhill 1994), pupal mass is not a dynamic variable but rather is determined by *Bupalus* density in the current generation. Pupal mass is modeled as decreasing exponentially to a minimum as a function of the density of *Bupalus* eggs (Fig. 1a):

$$W_t = P_{min} + P_0 e^{-\beta E_t} \quad (11)$$

where β is positive.

Based on Klomp's data we assume that fecundity increases linearly with the individual's mass during the pupal stage (Fig. 1b). Thus, the number of eggs is related to the previous generation's adult density by

$$E_{t+1} = A_t(-a + bW_t) \quad (12)$$

where W_t is pupal mass and a, b are positive parameters.

Egg-to-adult survival is modeled as declining exponentially with egg density (recall that this is due to direct density dependence during the larval stage), and increasing exponentially with maternal pupal mass (Fig. 1c,d). Thus adult density is related to egg density as

$$A_{t+1} = rE_{t+1} \exp(-sE_{t+1} + uW_t) \quad (13)$$

where r, s, and u are positive parameters, constrained so that $A_{t+1} < E_{t+1}$.

Assuming that adult density is proportional to pupal density, we have

$$E_{t+1} = N_t(-a + bW_t) \quad (14)$$

$$N_{t+1} = rE_{t+1} \exp(-sE_{t+1} + uW_t) \quad (15)$$

$$W_{t+1} = P_{min} + P_0 \exp(-\beta E_{t+1}) \quad (16)$$

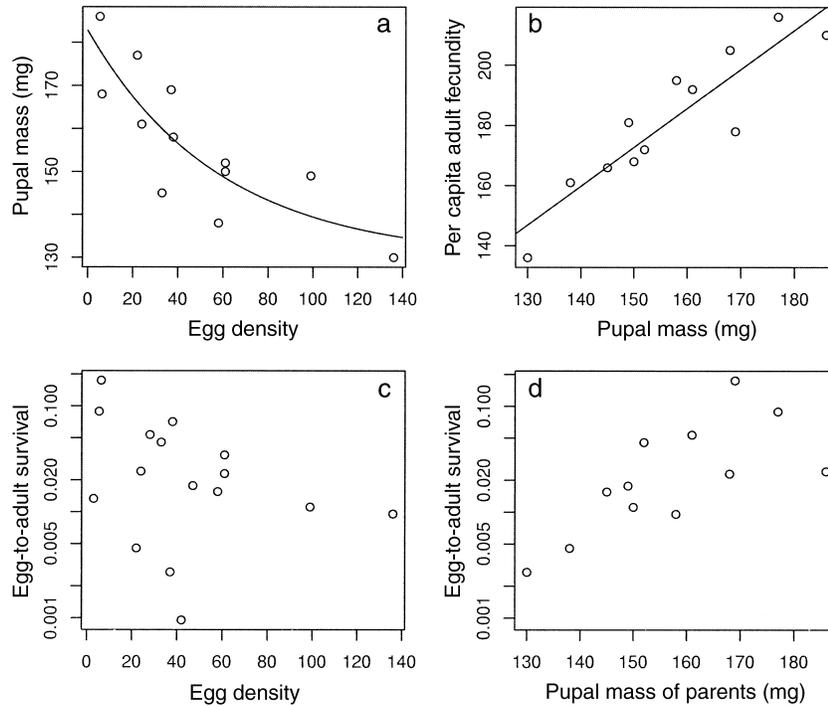where N_t is pupal density from the pupal survey. We

FIG. 1. Evidence suggesting a maternal effects mechanism in Dutch *Bupalus* populations (data from Klomp [1966, 1968]). (a) Effect of egg density on pupal mass. The line is the fitted nonlinear regression (Eq. 11). (b) Effect of pupal mass on per capita adult fecundity. The line is the fitted linear regression equation ($r^2 = 0.82$, $F_{1,12} = 54.6$, $P < 0.001$). (c, d) Relationship between egg-to-adult survival (plotted on a log scale) and (c) egg density and (d) maternal pupal mass. Multiple $r^2 = 0.49$, $F_{2,11} = 5.37$, $P = 0.02$ for linear regression of log survival on egg density and maternal pupal mass.

can then use Eq. 14 to eliminate E from the model, and then rescale variables and combine parameters to simplify the model (see Appendix A). The resulting model is

$$N_{t+1} = rN_tX_t\exp(-sN_tX_t + uX_t)$$

$$X_{t+1} = x_{\min} + \exp(-\beta N_tX_t). \qquad (17)$$

Here N_t is pupal density at the time of the annual survey, and the individual quality variable X_t is directly proportional to the effect of pupal mass on adult fecundity. In Klomp's data X_t varies by a factor of roughly two over the course of outbreak cycles.

Food quality

We used evidence from two recent studies in Sweden (Šmits and Larsson 1999, Šmits et al. 2001) to derive a mechanistic model for effects of food quality. Because this model is totally de novo and specifically represents *Bupalus* on Scots pine, we present it in more detail than the others. As with the model for maternal effects, we use the results to define the form of the model but not to estimate parameter values.

Šmits and Larsson (1999) examined the relative preference of *Bupalus* larvae for new pine needles produced that spring vs. old needles, and relative performance (larval survival) on the two types of needles. They found that the larvae preferred to feed on mature needles, causing 2.5–3.4 times greater damage on mature needles than on new needles. When forced to feed only on new needles, survival of very young larvae was reduced relative to those allowed to feed on mature needles. Šmits et al. (2001) examined adult fecundity and oviposition behavior on new and mature needles. They had three main findings. First, there was a strong preference by adults to lay eggs on mature needles whenever there was a choice (Šmits et al. 2001: Fig. 1). Second, effective fecundity was lower when adults were forced to lay on new needles (Šmits et al. 2001: Fig. 2). *Bupalus* fecundity was affected by female quality (pupal mass), but the slopes for the regressions of fecundity on pupal mass did not differ significantly between the two treatments. Finally, the lower effective fecundity resulted from a lower egg laying rate (119.4 $\pm$ 8.8 eggs on twigs with only current-year needles, vs. 141.1 $\pm$ 7.5 eggs on normal twigs), and a higher fraction of eggs falling off new needles (Šmits et al. 2001: Fig. 5).

The model begins with pupae at the time of the annual survey (N_t) who then mature into two groups of adults: those ovipositing on mature needles ($A_{(m)t}$) and those ovipositing onto new needles ($A_{(n)t}$). The proportion of adults in the two groups depends on F_t, the fraction of mature needles when the adults are ovipositing (since this is at the start of the growing season, F_t is also the fraction of mature needles encountered

by their offspring as larvae). The data on relative preference for mature needles (Šmits et al. 2001: Fig. 1) are fitted well by a power function (see Appendix A), leading to

$$A_{(m)t} = N_t[1 - (1 - F_t)^{\alpha}]^{1/\alpha} \tag{18}$$

$$A_{(n)t} = N_t - A_{(m)t} \tag{19}$$

where $\alpha > 1$.

The two groups of adults have different fecundities, R and Rb, where b reflects both the difference in egg laying and the eggs dropping from new needles. Thus,

$$L_{t+1} = R(A_{(m)t} + bA_{(n)t}). \tag{20}$$

Rather than explicitly model larval feeding choice, we simply assume that the new needles effectively provide less food than the mature needles. This encapsulates both the fact that larvae spend less time feeding on new needles, and that the new needles are poorer food. Thus the total resource available to the larvae is proportional to $\bar{F}_t = F_t + q(1 - F_t)$, $0 < q < 1$. We assumed that larval survival was a saturating (type-II) function of the total food resource per larva:

$$P_{t+1} = L_{t+1} \frac{K\bar{F}_t/L_{t+1}}{1 + K\bar{F}_t/L_{t+1}} \tag{21}$$

$$= L_{t+1} \frac{K\bar{F}_t}{L_{t+1} + K\bar{F}_t}. \tag{22}$$

This yields an undercompensating relationship between larval density and pupal density, which is consistent with Klomp's (1966: Table 23) limited data on larval and subsequent pupal densities. A Ricker-type model with overcompensation in larval survival was also tried, but the resulting model was less successful at fitting the population time series data.

Note that we have eliminated both pupal mortality and density-independent larval mortality by rescaling unobserved variables and absorbing these losses into the fecundity parameters.

The total amount of needles consumed (and hence replaced by new needles next year) is assumed to be a function of *Bupalus* larval density. Lacking any direct data on needle dynamics, we posit that a higher density of *Bupalus* larvae in the current year results in higher needle consumption and therefore a higher fraction of new needles next year. There is also some fraction of new needles appearing each year independent of defoliation by *Bupalus*, resulting from needle turnover and shoot growth. Our model is

$$F_{t+1} = (1 - f_0)e^{-\gamma L_{t+1}} \tag{23}$$

where f_0 is the fraction of new needles in the absence of *Bupalus* defoliation. The time indices in Eq. 23 are the same because F_{t+1} characterizes the foliage subsequently encountered by the L_{t+1} larvae when they mature (as adults A_{t+1}) and by the larval offspring of those adults (L_{t+2}).

The needle proportions in Eq. 23 respond immediately to the current level of defoliation. We also tried a lagged response model, $F_{t+1} = (1 - \delta)[(1 - f_0)e^{-\gamma L_{t+1}}] + \delta F_t$. However in two of the three sites the best-fit parameters for this model collapsed to immediate response ($\delta = 0$) and in the third it made only a very slight improvement, so we used the simpler model (Eq. 23).

Environmental stochasticity

The dynamic models derived previously are only the deterministic ''skeleton'' of the population dynamics, omitting effects of exogenous variables (notably climate) which influence *Bupalus* density variation. Lacking data on these variables—indeed, not knowing with any certainty which variables have strong impacts—we therefore model their aggregate effects descriptively as process noise. We adopted the most commonly used model for effects of environmental stochasticity, multiplicative lognormal perturbations to the population densities generated by the deterministic skeleton, representing environmentally driven variation in birth and survival rates (which act multiplicatively on population density). So for example, when simulating the maternal effects model (Eq. 17) the population was actually updated each year by

$$N_{t+1} = rN_tX_t\exp(-sN_tX_t + uX_t)\exp[\sigma Z_N(t)] \tag{24}$$

where $\{Z_N(1), Z_N(2), \ldots\}$ are independent Gaussian random variables with zero mean and unit variance. The same kind of lognormal perturbations were applied to each stage transition (e.g., pupae to adults, egg laying by adults) in the parasitoid and food quality models. We assumed a common value of σ for each transition, because the perturbations are unobserved (given only annual pupal counts) so we could not hope to identify separately the components at different stages of the life cycle.

Measurement noise

Our method for fitting models to the *Bupalus* time series (described in the next section) requires that model output incorporate simulated measurement errors in *Bupalus* density, corresponding to the sampling variability in the actual data. For the three forests analyzed, we estimated the magnitude of sampling variability in *Bupalus* pupal density using data in the Appendix of Barbour (1980), which included average density and the among-compartment standard error for all forests until 1978. There is an allometric (power) relation between the standard error and the mean density μ:

$$(SE)^2 = C\mu^b. \tag{25}$$

We estimated the parameters C and b of this relationship separately for each forest: at Culbin, $C = 0.024$ and $b = 1.342$; at Roseisle, $C = 0.046$ and $b = 1.345$; at Tentsmuir, $C = 0.071$ and $b = 1.641$.

To simulate the sampling variability in the data, after a time series of "true" *Bupalus* density values had been produced by the model, each value was replaced by a gamma-distributed random variable with mean given by the true density, and variance determined by the allometric parameters corresponding to the data set being fitted.

Fitting Models by Nonlinear Forecasting

To fit and compare the models we use the method of nonlinear forecasting (NLF) described and used by Kendall et al. (1999) and Turchin and Ellner (2000). NLF is a simulation-based solution to the problem of fitting complicated dynamic models whose likelihood function is intractable: impossible to compute analytically and computationally too demanding for current computers. Our models' likelihood functions are intractable due to the combination of measurement noise, process noise, nonlinear dynamics, and the fact that we have data on only one of several state variables. When determinism is strong such models can be fitted efficiently by "trajectory matching" (forcing the model to closely approximate that observed time series; Wood 2001), but in this case there is too much variability among outbreaks for trajectory matching to be suitable. Similarly, methods for deterministic models based on expressing the deterministic skeleton of the model as a delay difference equation for the observed state variable yielded poor parameter estimates in trials on simulated data.

NLF is an example of simulated quasi-maximum likelihood (SQML; Smith 1993), an "indirect inference" method (Gourieroux and Montfort 1996) in which a mechanistic model is compared to data via an intermediate statistical model chosen to have two properties: (1) it provides a good statistical description of time series for the observed state variable produced by the mechanistic model; and (2) it has a likelihood function that is easy to compute and optimize. Indirect inference methods have two general advantages. First, the mechanistic models can be arbitrarily complicated, so long as it is computationally feasible to produce long simulations. Second, a uniform method and goodness-of-fit criterion can be applied to all of the alternative models using exactly the same fitting code, regardless of how they differ in assumptions. Indirect inference allows nonlinear state space models to be fitted within the frequentist paradigm, so that we can (for example) attach a conventional *P* value to the relative performance of alternative models.

The SQML goodness-of-fit criterion is to first fit the intermediate statistical model to output from the mechanistic model and then compute the likelihood of the data under the fitted statistical model, the so-called quasilikelihood function (QLF) of the data. Under some mild technical assumptions (Smith 1993), the QLF functions to a large extent like the true likelihood function for the mechanistic model. Parameter estimates based on maximizing the QLF are consistent and asymptotically normal, and the optimized QLF provides a criterion for comparing alternative models. The QLF cannot be used for likelihood ratio tests, but even with the true likelihood these would be questionable here due to the small sample size.

For the intermediate statistical model we used a nonlinear Gaussian autoregressive model for the observed state variable (pupae/per square meter). In past applications we have used a nonparametric kernel autoregression model (Kendall et al. 1999, Turchin and Ellner 2000). However, for this study a kernel model was computationally infeasible, because a high-dimensional statistical model (3–4 past values) was needed to obtain good forecasts of output from the simulation models. This increases greatly the length of model simulations needed to fit the kernel, and the time required to compute the kernel regression scales as the square of the simulation length. We therefore used instead a Generalized Additive Model:

$$F(x_t, x_{t-1}, x_{t-2}, x_{t-3}) = \sum_{i=1}^{4} \sum_{j=0}^{3} a_{ij} f_i(x_{t-j})$$

where the f_i are radial basis functions with fixed location and scale parameters (see Appendix B for details). Because the f_i are specified in advance, the model is linear in parameters and can be fitted quickly by linear least squares. The data were 0.2-power transformed so that they would conform better to the autoregression model's assumption of Gaussian process noise, without exaggerating measurement errors at low densities (as would occur with a log transformation). Forecasting two years ahead was necessary to capture nonlinear dynamic aspects of the data and models; the data exhibit strong enough linear autocorrelations at a lag of one year that accurate one-year-ahead forecasting is possible based strictly on the linear autocorrelation between successive values.

The QLF thus has the form

$$\text{QLF}(\theta)$$
$$= \prod_{t} \frac{1}{\sqrt{2\pi}\hat{\sigma}} \exp\left\{ -\frac{[x_{t+2} - F_\theta(x_t, x_{t-1}, x_{t-2}, x_{t-3})]^2}{\hat{\sigma}^2} \right\}$$
$$(26)$$

where $x_t = N(t)^{0.2}$, θ is the vector of parameters for the mechanistic model, F_θ is the generalized additive model fit to the time series generated by the mechanistic model with parameters θ, and $\hat{\sigma}$ is the residual standard deviation of the generalized additive model. We fitted the model by minimizing the negative log-quasilikelihood. Numerical methods for model fitting by NLF are described in Appendix B.

The quasilikelihood is a relative measure of the ability of the model to capture the dynamics in the data. For an absolute measure of model quality, we also present values of the prediction r^2, defined as $1 - \text{NMSE}$, where NMSE is the normalized mean square error, de-

TABLE 2. Fits of the models in three forests.

Model	Constraints		Culbin		Roseisle		Tentsmuir	
	Min	Max	LL	r^2	LL	r^2	LL	r^2
Maternal effects			1.442	0.639	0.796	0.488	0.826	0.604
Parasitoids (both)	0.00	1.00	1.370	0.585	0.676	0.351	1.004	0.722
	0.01	0.99	1.347	0.564	0.608	0.282	0.910	0.665
	0.03	0.97	1.251	0.471	0.533	0.227	0.770	0.605
	0.05	0.95	1.220	0.444	0.466	0.093	0.608	0.404
Parasitoids (*Dusona* only)			1.195	0.414	0.490	0.209	0.861	0.634
Parasitoids (*C. viator* only)	0.00	1.00	1.435	0.634	0.586	0.220	1.035	0.739
	0.01	0.99	1.193	0.416	0.441	0.010	0.776	0.581
	0.03	0.97	1.145	0.441	0.437	0.133	0.726	0.589
	0.05	0.95	1.018	0.335	0.149	−0.855	0.300	0.216
Food quality			1.354	0.569	0.778	0.469	0.572	0.343

Notes: LL is proportional to the optimized log-quasilikelihood function, and r^2 is the prediction r^2 of the best-fitting model. ''Constraints'' refers to the minimum and maximum allowable parasitism levels.

fined as the mean squared deviation between the predictions and the data divided by the variance of the data. In the nonlinear dynamics literature comparison of prediction methods based on minimization of NMSE is common. For example entries submitted to the Santa Fe Institute Forecasting Competition (Weigend and Gershenfeld 1994) were ranked based on NMSE. The popularity of NMSE probably traces to an influential paper by Farmer and Sidorowich (1987), who used NMSE to compare different choices of time-delay embedding parameters. The switch to prediction r^2, because ecologists are familiar with its intuitive meaning, was introduced to our knowledge by Tidd et al. (1993). Positive values of this measure can be interpreted in the same way as the R^2 from a linear regression: the fraction of the variance in the data that is explained by the model. When $r^2 = 0$, the model is not capturing any of the dynamics, simply predicting the unconditional mean of the data at each time step. In contrast to linear regression, a nonlinear forecasting model can actually perform worse than this, producing a negative correlation between predictions and observations. This is represented by a negative value of the prediction r^2.

Comparing models by parametric bootstrap

We used a parametric bootstrap procedure to test whether the differences among models in optimized quasilikelihood are statistically significant. We used the worse-fitting model, at its best-fitting parameter values, to generate 25 independent simulations each of *Bupalus* pupal density. Each simulation was 350 years long, with the first 250 years dropped to eliminate any possible transients. The first 35 years of each remaining 100 years of data were fitted with the two models to be compared, using the same procedures as were used on the actual data. For each site, this yields 25 simulated values for the difference in log-quasilikelihood (LL) between the two models, which we compare against the actual difference on the real data.

This comparison shows whether or not the improvement in LL achieved by the better fitting model on the

actual data would be likely to occur by chance if the poorer fitting model were actually the correct description of the mechanism generating the data. In principle it is identical to a standard statistical comparison between alternative models for a data set, e.g., comparing models with and without a treatment effect: we take the poorer fitting model as our null hypothesis, and compare the observed improvement in goodness-of-fit against its distribution under the null. The choice of null models is post hoc, but conventional hypothesis testing always uses the worse-fitting model as the null. In our case, because the models are not nested, we simply do not know a priori which one will have the worse fit. In addition, for comparisons involving parasitoid models, it is reasonable to treat the parasitoid model as the null, for that is the default explanation that most biologists (including us!) use to explain population cycles in forest insects.

RESULTS

The goodness of fit of the five models (with parameter values optimized independently for each population) is summarized in Table 2. The *C. viator* model fit just about as well as the two-parasitoid model, but the *Dusona* model did quite poorly. Dropping the latter model from consideration, we are left with four contending models representing the three possible mechanisms. There is essentially a tie in Culbin between parasitoids and maternal effects, and best fits by maternal effects in Roseisle and parasitoids in Tentsmuir. The food quality model does not provide the best fit in any of the three sites, although it is a ''close second'' in Roseisle. Bootstrap significance tests (see *Results: Parametric bootstrap*) confirm that the maternal effects model fits the data significantly better than the food quality model in Culbin and Tentsmuir, but the difference in Roseisle is not significant.

The best-fit parameter values are listed in Appendix C, and simulations of the fitted models are shown in Figs. 2–4. Since these are stochastic models, we show five replicate simulations of each model. The two-par-
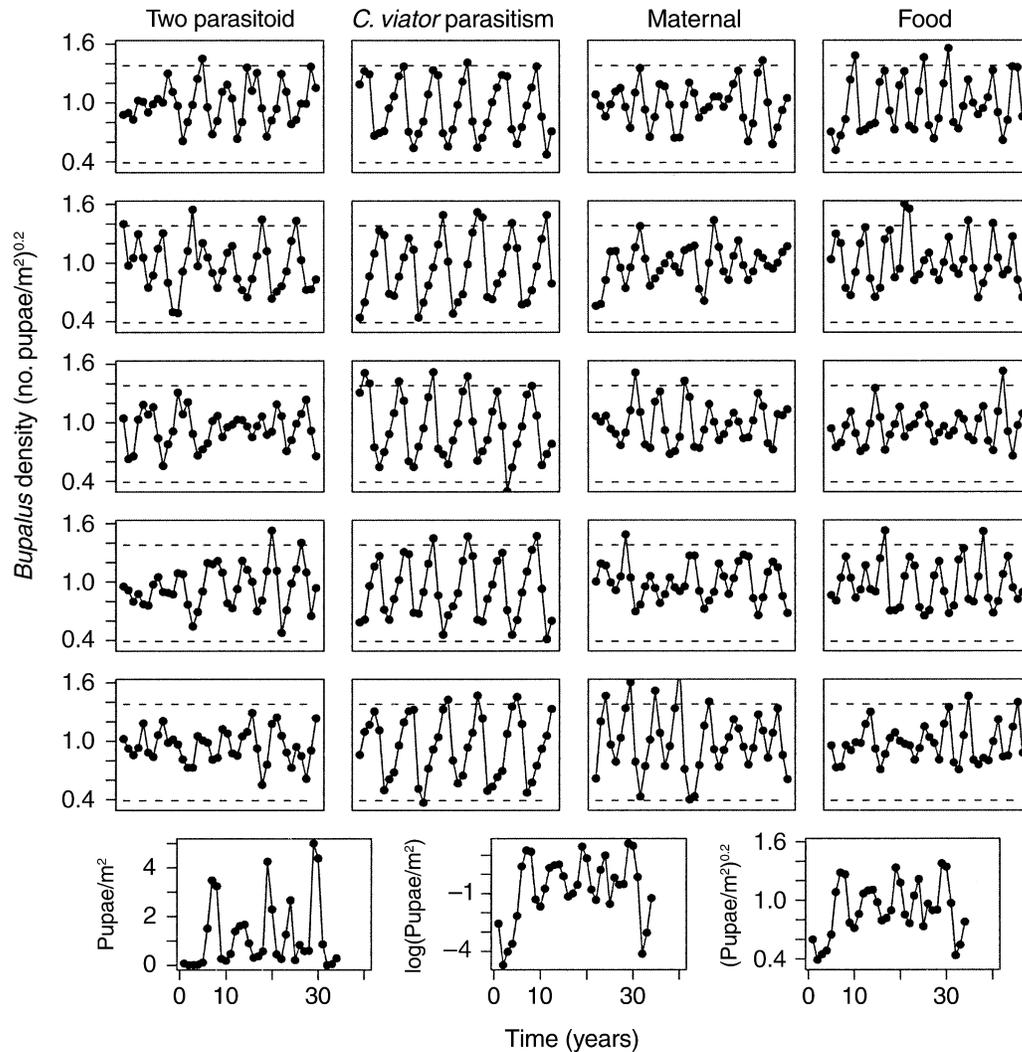
FIG. 2. Five replicate simulations of each contending model, using parameters found by fitting the models to Culbin. Dashed lines are drawn at the minimum and maximum observed pupal density values in the empirical time series. The actual data are shown in the bottom row, in untransformed, log-transformed, and fifth-root-transformed scales.

asitoid model shows a great deal of variation among replicates. For example, in Tentsmuir, where it has a much better fit than the maternal model, some replicates appear to the eye to be very similar to the data (in terms of the period, amplitude, and variation among outbreaks) while others barely depart from equilibrium. In contrast, the *C. viator* parasitism model in Culbin and Tentsmuir is nearly deterministic, for the fitting algorithm chose an unreasonably low level of noise. This results in cycles that are much more regular than observed in the data. However, if the process noise variance is set to a level comparable with the other fits, a great deal of variability among replicates also occurs in this model (simulations not shown). The maternal effects model is quite variable among replicates, but in Culbin and Tentsmuir, most replicates appear to the

eye to be similar to the data. The food quality model generally seems to have a shorter period than the data.

In simulations of the fitted maternal effects model, maternal quality (X_t) varied by a factor of roughly 1.8 in Culbin, 3 in Roseisle, and 2.6 in Tentsmuir. Recall that X_t is proportional to the effect of pupal mass on fecundity; this variation is less than the variation in fecundity observed in Cannock, so it is biologically plausible. In contrast, simulations of the fitted parasitoid models all showed annual total parasitism rates Π_t (see Eq. 1) coming very close to zero and one (Table 3). Because these are outside the range of observed parasitism rates (reviewed previously), we examined the effects of refitting the models with modest constraints on the parasitism rates, by imposing severe penalties on the quasilikelihood if the parasitism rate exceeded the given threshold values.
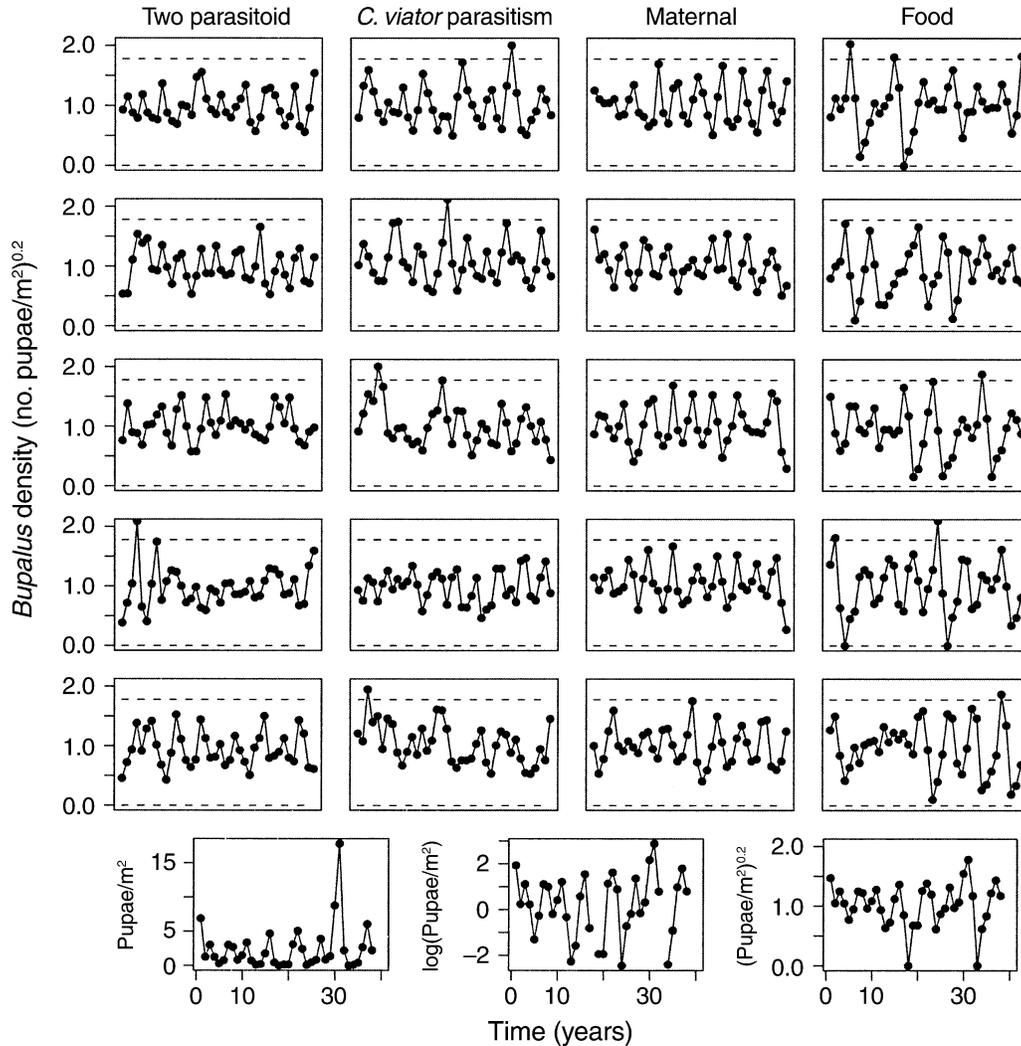
FIG. 3. Five replicate simulations of each contending model, using parameters found by fitting the models to Roseisle. Dashed lines are drawn at the minimum and maximum observed pupal density values in the empirical time series. The actual data are shown in the bottom row, in untransformed, log-transformed, and fifth-root-transformed scales.

Constrained parasitoid model fits

The fit of the two-parasitoid model declines gradually as increasingly tighter constraints are imposed on the possible range of variation in annual total parasitism rates (Table 2). The decline is most rapid in Tentsmuir, suggesting that the extreme parasitism rates are critical for the good fit of the model at that site. In contrast, even the most modest of constraints on parasitism rate (between 0.01 and 0.99) causes the fit of the *C. viator* model to fall off dramatically.

Simulations of the constrained two-parasitoid model (using the intermediate constraints [0.03, 0.97]) continue to somewhat resemble the data, but the occurrence of simulations that (by eye) match the data well is less common, especially for Tentsmuir (Fig. 5).

Parametric bootstrap

Let $\hat{D}$ be the observed difference in log-quasilikelihood between the maternal effects model and the contesting model, and D_i^*, $i \in \{1, 2, \ldots, 25\}$ be the values from the parametric bootstrap. For the parametric bootstrap comparing the maternal effects model to the constrained two-parasitoid model, all 25 bootstrap replicates (assessed under the hypothesis that the data actually were generated by the parasitoid model) had values substantially below the observed value at all three sites (Culbin, $\hat{D} = 0.233$, $-0.194 < D_i^* < 0.046$; Roseisle, $\hat{D} = 0.263$, $-0.059 < D_i^* < 0.156$; Tentsmuir, $\hat{D} = 0.218$, $-0.224 < D_i^* < 0.094$). For the bootstrap comparison of maternal effects with the food quality model, a similar result was found at Tentsmuir ($\hat{D} = 0.254$, $-0.143 < D_i^* < 0.040$). However, at Culbin,
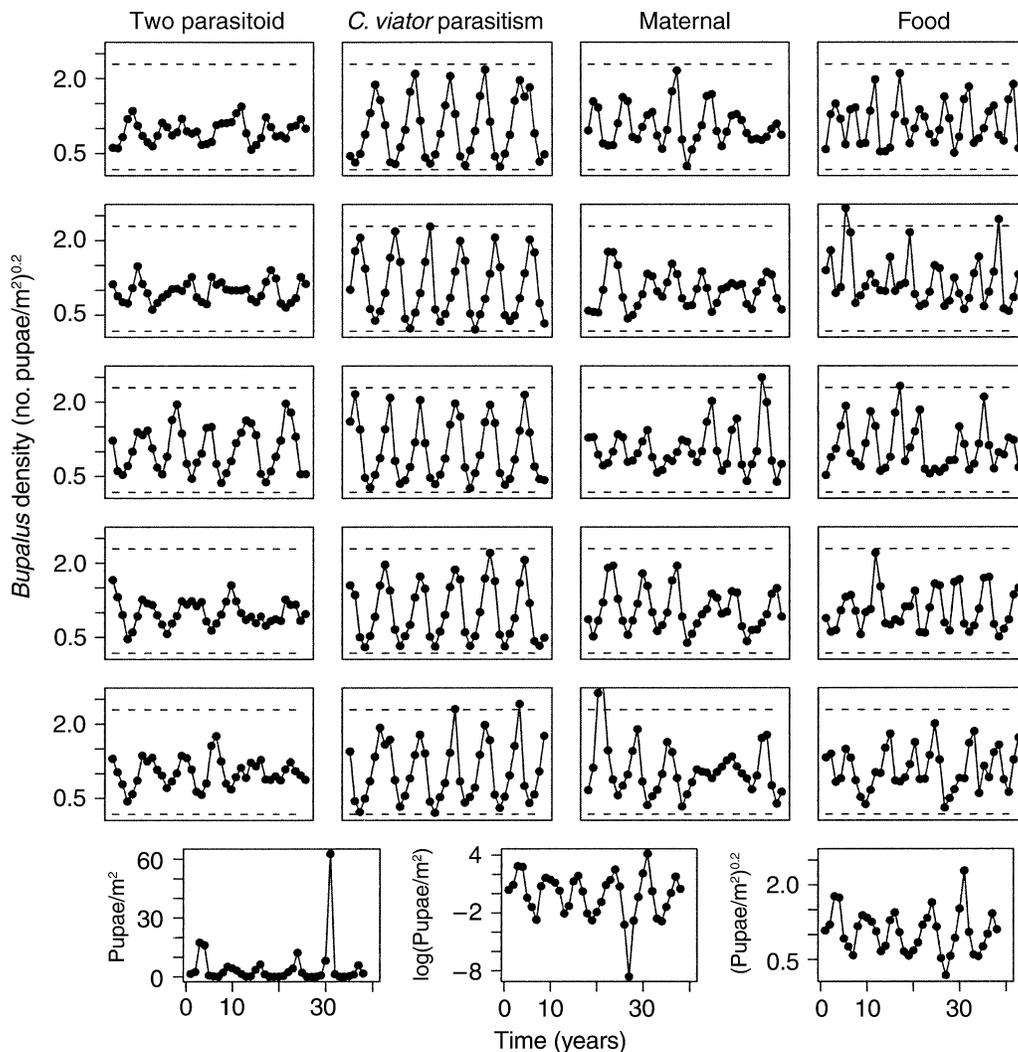
FIG. 4. Five replicate simulations of each contending model, using parameters found by fitting the models to Tentsmuir. Dashed lines are drawn at the minimum and maximum observed pupal density values in the empirical time series. The actual data are shown in the bottom row, in untransformed, log-transformed, and fifth-root-transformed scales.

one of the 25 bootstrap replicates approached the data value ($\hat{D} = 0.088$, $-0.089 < D_i^* < 0.088$), and in Roseisle the bootstrap distribution was centered on the observed value ($\hat{D} = 0.018$, $-0.151 < D_i^* < 0.174$). We ran an additional 25 bootstrap replicates at Culbin; none of them approached the data value, so $P \approx 0.02$ for the hypothesis that the food quality model generated the data at Culbin. Figures showing the distributions of bootstrap replicates may be found in Appendix D.

Thus under fairly mild constraints on the parasitism rate, the maternal effects model provides a significantly better fit to the data than the parasitoid model in all three data sets. The maternal effects model also provides a significantly better fit than the food quality model at two sites (Tentsmuir, Culbin) while at the third site the difference in fit is small enough that it could easily occur by chance if the food quality model described the mechanism driving the cycles.

DISCUSSION

We found that the parasitoid models best fit the *Bupalus* time series at Tentsmuir and performed as well as the maternal effects model at Culbin if parasitism rates were allowed to come arbitrarily close to zero and one. However, if parasitism rates were constrained to be between the observed rates, the goodness of fit of the parasitism models dropped dramatically, and the maternal effects model fit significantly better than the parasitoid model across all three sites. The food quality model performed well only at Roseisle, where it fit as well as the maternal effects model.

We were initially very skeptical of the maternal effects mechanism; indeed, this study would have been completed several years previously had we not struggled so hard to find some grounds for rejecting the hypothesis. However, the weight of evidence led us to

TABLE 3. Range of total parasitism rates Π, observed in replicate simulations of parasitoid models fitted without constraints on parasitism rate.

Model	Culbin	Roseisle	Tentsmuir
Two parasitoid			
Minimum	0.0666 (0.0076, 0.214)	0.292 (0.105, 0.426)	0.0637 (0.0018, 0.441)
Maximum	0.979 (0.964, 0.987)	0.995 (0.991, 0.997)	0.994 (0.980, 0.997)
C. viator only			
Minimum	0.0001 (0.0000, 0.0005)	0.0138 (0.0001, 0.170)	0.0016 (0.0002, 0.078)
Maximum	0.994 (0.989, 0.997)	0.985 (0.982, 0.986)	1.000 (1.000, 1.000)

Notes: Table entries are the median and (5th percentile, 95th percentile) of the minimum and maximum annual value of Π over 500 simulations of length 35 years each.
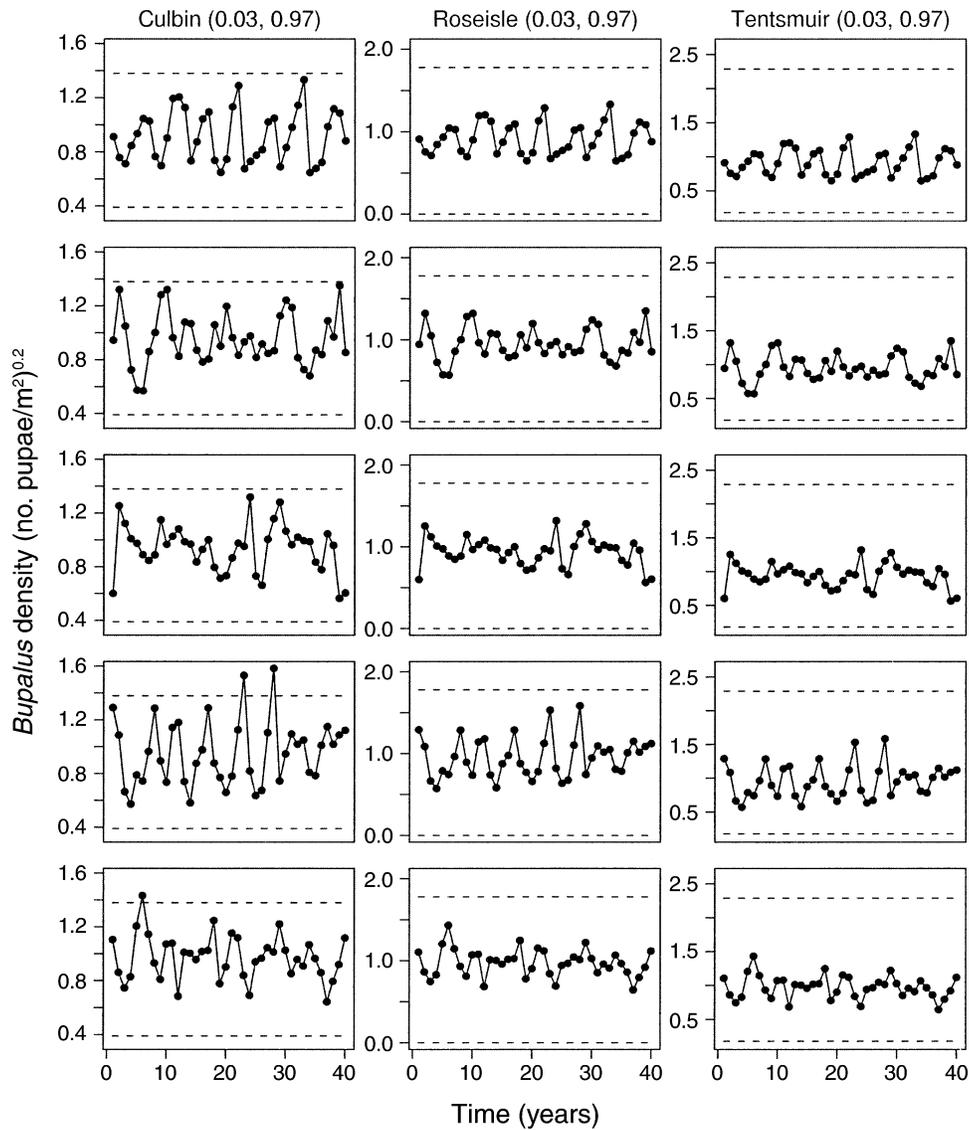

FIG. 5. Five replicate simulations each of constrained two-parasitoid models for each of the sites. Above each column, values in parentheses are the minimum and maximum allowed values of the total parasitism rate that were imposed when fitting the model.

TABLE 4. Fits of the off-the-shelf maternal effects and parasitoid models.

Model	Culbin		Roseisle		Tentsmuir	
	LL	r^2	LL	r^2	LL	r^2
Maternal effects	1.406	0.615	0.685	0.379	1.158	0.790
Parasitoids (unconstrained)	1.405	0.612	0.608	0.354	0.949	0.693

conclude that maternal effects are probably the dominant mechanism driving the cycles in British *Bupalus* populations. Although maternal effects have been hypothesized as drivers of population cycles (Myers 1988, Ginzburg and Taneyhill 1994), this is the first convincing evidence that observed maternal effects are sufficient to account quantitatively for the observed population variation in a particular species.

Our skepticism about the maternal effects mechanism stemmed from the fact that parasitoids have been widely considered to be the most likely explanation for long-period insect population cycles (e.g., Berryman 1996). We found that parasitoid models could only explain the British *Bupalus* cycles if parasitism rates are allowed to reach values much more extreme than have been observed in any empirical studies. However, we caution that it would be premature to totally dismiss the parasitoid mechanism. Despite the many years of empirical study it is not impossible that the full range of parasitism rates goes beyond what is represented in the existing data. Parasitism rates are difficult to estimate, were estimated at sites other than the study sites, and are often based on limited sample sizes. Furthermore, the range of total parasitism rates will be affected by the degree of synchrony among parasitoid species' variations in density, which could be stronger in the more cyclic populations where the total parasitism rates have not been monitored.

Even without constraints on the parasitism rates, the fitted parasitoid models consistently had the property that qualitative dynamics similar to the data series "come and go" during long simulations of the model. In one sense this could be viewed as a fitting failure: we do nothing to penalize the model if it spends time in a part of the state space that is far away from the data. On the other hand, the data are but a single realization of a stochastic process, and in this sense are entirely consistent with a stochastic model that sometimes behaves like the data, and sometimes behaves differently. Indeed, given that *Bupalus* at many other sites appears to have been fluctuating around an equilibrium for 35 years, a mechanism that can exhibit these alternate dynamical behaviors has some appeal.

The food quality model was generally the least successful at accounting for the observed dynamics, and was rejected statistically at two of the three sites analyzed: Culbin and Tentsmuir. A caveat to these findings is that we had no direct empirical data that could be used to model the defoliation dynamics, so it is conceivable that a more realistic defoliation model

could lead to improved fits. One possible approach to this would be to develop a "partially specified" model (Wood 1997, 2001, Ellner et al. 1998) using a flexible or nonparametric model for defoliation dynamics. However, once the food quality hypothesis emerged as the across-the-board "loser," we did our best to avoid biasing the analysis against this mechanism through poor model choice, by trying several plausible model variants at the stages for which we had the least data—larval density dependence and defoliation dynamics. For the other mechanisms, where the picture is more complete, we used a single model that was developed from the empirical studies of that mechanism in *Bupalus*. So if anything our analysis was unfairly biased in favor of the food quality hypothesis, and it therefore seems unlikely that additional information relevant to food quality dynamics (which would constrain our search for the "best possible" model) would change the conclusions.

Similarly, the parameters of the food quality model were not constrained by experimental information because the relevant experiments were done in Sweden, but it seems likely that such constraints would further reduce this mechanism's credibility. Fitted values of q are always near 0, meaning that new needles are totally valueless to larvae. Experiments say otherwise: larvae fed on only new needles still had survival of ~80%, compared to ~90% when mature needles were available (Šmits and Larsson 1999). If the model is constrained so that larval survival can never be <20% of what it would have been on a diet of all mature needles, its ability to match the data is significantly reduced (LL = 1.15, r^2 = 0.36 for Culbin, LL = 0.52, r^2 = 0.12 for Roseisle, LL = 0.47, r^2 = 0.22 for Tentsmuir).

The failure of constrained parasitoid models can almost be inferred directly from the data, by estimating the quantitative impact of the limiting factor whose identity we are trying to determine. When each *Bupalus* population is in the middle of its density distribution (33rd–67th percentiles), the population is little affected by density dependence, and so the difference in population growth rates between the ascending and descending phases of the cycle must be driven by parasitism rate. Even discounting extreme values (which might be driven by measurement error) the ratios between the smallest and largest growth rates in these phases are ~0.05 for Culbin and Roseisle, and 0.01 for Tentsmuir. Thus, for parasitoids to drive the cycles, maximum parasitism rates would need to exceed 0.95

in the first two sites, and 0.99 in Tentsmuir, far beyond the levels that have been observed in the field.

We can also understand the dynamical role of parasitoids by looking at trajectories of the two-parasitoid model. In the unconstrained model, episodes of 2–3 years when the total parasitism rate (Π_t) stays very high cause a drastic decrease in *Bupalus* density population and then in the parasitoids. These high-parasitism periods are generated in the model by offset outbreaks of the two parasitoids, with the peak of *C. viator* preceding that of *Dusona* by one year. Afterward there is gradual regrowth of *Bupalus* and parallel increase in the parasitism rate, until the next outbreak. With realistic constraints on the parasitism rate, this scenario becomes impossible. Instead the fitted models rely on the exogenous noise to generate occasional large *Bupalus* outbreaks (without noise the fitted models all converge to equilibrium). When such outbreaks occur they resemble the data, but as the constraints on Π become more strict the model is increasingly at the mercy of the noise and loses the longer term outbreak pattern.

In the maternal effects model, the initial crash is caused by high density, which reduces both density and quality the subsequent year. The resulting low individual quality in the following year causes a second year of population decline, even if population density is already low. After the second decline individual quality has recovered, so the population can begin to increase, but density has fallen so much in the two decrease years that several years are required before another outbreak occurs. The interaction of quality and density also explains the fact that outbreaks may last either one or two years. Depending on the population density in the year preceding the first outbreak year, individual quality in the first outbreak year may be high or low. In the former case, even though density is high in the first outbreak year, the population still can increase or remain nearly constant. The second outbreak year then has both high density and low quality, so the outbreak collapses the next year. In the latter case, because density is high and quality is low in the first outbreak year, the outbreak terminates immediately.

In contrast to the parasitism model, however, the variation in maternal quality required to cause these dynamics is consistent with independent data. Across 15 years in the Netherlands, pupal mass varied from 130 to 190 mg, and the associated fecundity varied by a factor of 1.6 (Klomp 1966). Across six years in the British forest of Cannock Chase, fecundity varied by a factor of four while mass varied from 85 to 140 mg (Broekhuizen et al. 1994). Across a single outbreak each, mass varied from 72 to 108 mg in Culbin and from 89 to 112 mg in Tentsmuir (Barbour 1980). The fitted maternal effects model generated two- to three-fold fecundity variation (less than that observed in Cannock Chase) without any constraints needing to be imposed.

This difference between the maternal and parasitoid models is not an artefact of the particular functional forms used to represent the interactions. Although we do not show the results here, we have examined a variety of functional forms within the basic structure of the models, with no qualitative effects on the results. In addition, we have examined established off-the-shelf models: a Nicholson–Bailey parasitoid model with host density dependence and parasitoid interference, and a maternal effects model from Ginzburg and Taneyhill (1994). With one exception, the off-the-shelf models fit less well than the corresponding models that we developed specially for *Bupalus*. At Tentsmuir the off-the-shelf maternal effects model actually fits better than any of our models. Comparing just the two off-the-shelf models, they fit equally well at Culbin, but the maternal effects model fits better at Roseisle and Tentsmuir (Table 4). Unfortunately, we have no way to check the plausibility of this maternal effects model, as the quality variable does not map directly onto any clearly identified biological quantity.

Although we conclude that the likely cause of the cycles at Culbin and Tentsmuir is maternal effects, parasitoids are known to be present and may play a role in modulating the cycles. As one check on their potential importance, we developed a model that included parasitoids as well as maternal effects. This model had the advantage of many more free parameters (11, vs. 6 for maternal effects alone), but with parasitism rate Π limited to (0.05, 0.95) it could not improve on the fit of the maternal effects model (LL = 1.437 for Culbin, 0.833 for Tentsmuir). These results are consistent with an interpretation where parasitism rates are along for the ride, driven by the *Bupalus* cycles rather than driving them.

Some readers may be skeptical of our strong inferential claims based ''merely'' on time series analysis. However, information is not just coming from the time series. Here and previously (Turchin et al. 2003) success depends on the availability and use of additional data that constrain the model or its parameters. It is notable that the third site (Roseisle), for which we could not uniquely identify a best-fitting mechanism, exhibited by far the lowest degree of nonlinear predictability in the observed dynamics (Table 1). The difference between this site and the others, in terms of the ''signal to noise ratio'' in the dynamics, indicates a practical limit to the methods used here for field data. In order to get beyond the short-term dynamics where linear effects dominate (and that any sufficiently complex nonlinear model could match), we had to fit models based on two-year-ahead predictions. At that range, only 16% of the variability is predictable for Roseisle. It is perhaps not surprising in retrospect that the predictable component of the variability did not contain enough information to uniquely identify a driving mechanism that accounted best for the dynamics. On the other hand, it is very encouraging for future studies

that 27% predictability (at Culbin) was sufficient for some clear conclusions to emerge.

Our conclusion—that cycles in British *Bupalus* are driven predominantly by maternal effects and not parasitoids—runs counter to established wisdom. Thus it would be desirable to challenge this conclusion with focused empirical study. In particular, it would be valuable to quantify both the parasitism rates and the pupal size fluctuations in the strongly fluctuating populations. Some of this information is collected as part of the annual pupal survey, although not entirely consistently (H. Evans, *personal communication*). Together with intensive studies on the cycling populations, such data may help resolve whether either parasitism rates or maternal quality are fluctuating enough to drive the population cycles. Experimental manipulations of maternal size and parasitoid densities may also provide valuable information. Information on the relationship between needle defoliation and larval density would be valuable to better specify needle dynamics for the food quality model. Our results also indicate that a simple experiment to quantify larval survival as a function of needle age and larval density might be sufficient to decisively eliminate the food quality hypothesis. If the results from Sweden are confirmed in UK populations, a realistically constrained food quality model would drop to a distant third place finish at all three sites.

The models might be made more sensitive to the dynamical mechanisms if explicit information on exogenous factors (such as climate and needle biomass) were incorporated into the models, to provide a mechanistic and observable replacement for the ''process noise'' that is currently treated as pure white noise. Another important exogenous factor is management: forests are logged and thinned at various times, and in some instances forests were sprayed with insecticide in response to severe *Bupalus* outbreaks (Crooke 1959, Scott and Brown 1973, Straw 1996). We did develop a variant of the model that included spraying, based on the inferred criteria for spraying, but it did not improve the fit of the models (see Appendix E). Finally, the spatial averaging of *Bupalus* density across an entire forest might obscure local interaction patterns. However, it is unclear whether the potential increase in determinism at smaller spatial scales would outweigh the problems associated with the smaller sample size and associated uncertainties in local population density.

Literature Cited

Anderson, R. M., and R. M. May. 1980. Infectious diseases and population cycles of forest insects. Science **210**:658–661.

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.

Auer, C. 1968. Erste Ergebnisse einfacher stochastischer Modelluntersuchungen uber die Ursachen der Populationsbewegung des grauen Larchenwicklers *Zeiraphera diniana*, Gn. (= *Z. griseana Hb.*) im Oberengadin, 1949/66. Zeitschrift für angewandte Entomologie **62**:202–235.

Baltensweiler, W. 1984. The role of environment and reproduction in the population dynamics of the larch bud moth, *Zeiraphera diniana* Gn. (Lepidoptera, Tortricidae). Advances in Invertebrate Reproduction **3**:291–302.

Baltensweiler, W., G. Benz, P. Bovey, and V. Delucchi. 1977. Dynamics of larch bud moth populations. Annual Review of Entomology **22**:79–100.

Barbour, D. A. 1980. Population dynamics of the pine looper moth *Bupalus piniaria* (L.) (Lepidoptera, Geometridae) in British pine forests. Ph.D. Thesis. University of Edinburgh, Edinburgh, UK.

Barbour, D. A. 1988. The pine looper in Britain and Europe. Pages 291–308 *in* A. A. Berryman, editor. Dynamics of forest insect populations: patterns, causes, implications. Plenum Press, New York, New York, USA.

Beddington, J. R., C. A. Free, and J. H. Lawton. 1976. Dynamical complexity in predator–prey model framed in simple difference equations. Nature **225**:58–60.

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

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

Broekhuizen, N. 1991. The population dynamics of the pine looper moth, *Bupalus piniaria* L. (Lepidoptera: Geometridae). Ph.D. Thesis. University of London, London, UK.

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.

Broekhuizen, N., M. P. Hassell, and H. F. Evans. 1994. Common mechanisms underlying contrasting dynamics in two populations of the pine looper moth. Journal of Animal Ecology **63**:245–255.

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.

Chitty, D. 1967. The natural selection of self-regulatory behaviour in animal populations. Proceedings of the Ecological Society of Australia **2**:51–78.

Crooke, M. 1959. Insecticidal control of the pine looper moth in Great Britain. I. Aerial spraying. Forestry **32**:166–196.

Davies, J. M. 1961. The pine looper moth, Bupalus piniarius, at Cannock Chase in 1960. Report on Forest Research **1961**:176–182.

Dwyer, G. 1994. Density dependence and spatial structure in the dynamics of insect pathogens. American Naturalist **143**: 533–562.

Edelstein-Keshet, L., and M. D. Rausher. 1989. The effects of inducible plant defenses on herbivore populations. 1. Mobile herbivores in continuous time. American Naturalist **133**:787–810.

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.

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.

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

Gruys, P. 1970. Growth in Bupalus piniarius (Lepidoptera: Geometridae) in relation to larval population density. Centre for Agricultural Publishing and Documentation, Wageningen, the Netherlands.

Gruys, P. 1971. Mutual interference in Bupalus piniarius (Lepidoptera, Geometridae). Pages 199–207 *in* P. J. den Boer and G. R. Gradwell, editors. Proceedings of the Advanced Study Institute on Dynamics of Numbers in Populations (Oosterbeek, 1970). Pudoc, Wageningen, the Netherlands.

Hassell, M. P., and G. C. Varley. 1969. New inductive population model for insect parasites and its bearing on biological control. Nature **223**:1133–1137.

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

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.

Klomp, H. 1968. A seventeen-year study of the abundance of the pine looper, *Bupalis piniarius* L. (Lepidoptera: geometridae). Symposia of the Royal Entomological Society of London **4**:98–108.

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.

Morris, W. F. 1990. Problems in detecting chaotic behavior in natural populations by fitting simple discrete models. Ecology **71**:1849–1862.

Myers, J. H. 1988. Can a general hypothesis explain population cycles of forest Lepidoptera? Advances in Ecological Research **18**:179–242.

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

Schwerdtfeger, F. 1942. Über die Ursachen des Massenwechsels der Insekten. Zeitschrift für angewandte Entomologie **28**:254–303.

Scott, T. M., and R. M. Brown. 1973. Control of the pine looper moth in Great Britain. III. Aerial spraying with tetrachlorinvos. Forestry **46**:81–93.

Smith, A. A. 1993. Estimating nonlinear time-series models using simulated vector autoregression. Journal of Applied Econometrics **8**:S63–S84.

Šmits, A., and S. Larsson. 1999. Effects of previous defoliation on pine looper larval performance. Agricultural and Forest Entomology **1**:19–26.

Šmits, A., S. Larsson, and R. Hopkins. 2001. Reduced realised fecundity in the pine looper *Bupalus piniarius* caused by host plant defoliation. Ecological Entomology **26**:417–424.

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.

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 Biological Sciences **254**:257–273.

Turchin, P. 2003. Complex population dynamics: a theoretical empirical synthesis. Princeton University Press, Princeton, New Jersey, USA.

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

Turchin, P., A. D. Taylor, and J. D. Reeve. 1999. Dynamical role of predators in population cycles of a forest insect: an experimental test. Science **285**:1068–1071.

Turchin, P., S. N. Wood, S. P. Ellner, B. E. Kendall, W. W. Murdoch, A. Fischlin, J. Casas, E. McCauley, and C. J. Briggs. 2003. Dynamical effects of plant quality and parasitism on population cycles of larch budmoth. Ecology **84**: 1207–1214.

Weigend, A. S., and N. A. Gershenfeld. 1994. Time series prediction: forecasting the future and understanding the past. Addison-Wesley, Reading, Massachusetts, USA.

Wood, S. N. 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 and Hall, New York, New York, USA.

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

APPENDIX A

Details of model development can be found in ESA's Electronic Data Archive: *Ecological Archives* M075-009-A1.

APPENDIX B

Details of fitting methods can be found in ESA's Electronic Data Archive: *Ecological Archives* M075-009-A2.

APPENDIX C

Parameter estimates for unconstrained models can be found in ESA's Electronic Data Archive: *Ecological Archives* M075-009-A3.

APPENDIX D

Parametric bootstrap distributions can be found in ESA's Electronic Data Archive: *Ecological Archives* M075-009-A4.

APPENDIX E

Analysis of a model that incorporates the effects of insecticide spraying can be found in ESA's Electronic Data Archive: *Ecological Archives* M075-009-A5.