Chaos in microtine populations

PETER TURCHIN

Department of Ecology and Evolutionary Biology, University of Connecticut, Storrs, Connecticut 06269-3042, U.S.A.

SUMMARY

On the basis of nonlinear time-series analysis of vole data, Turchin (1993) proposed that there is a latitudinal shift from stability to chaos in the dynamics of vole populations. Falck *et al.* (1995) question several aspects of the approach I used. In particular, they point out that no measure of uncertainty has been calculated for each point estimate of the dominant Lyapunov exponent (a positive Lyapunov exponent implies chaos, while a negative one implies stability). In this note I respond to the criticisms of Falck *et al.* (1995). As part of my response, I analyse fluctuations of pooled numbers of microtine rodents at all European locations for which I currently have data. Treating each geographic location as a replicate, I show that the mean Lyapunov exponent is significantly greater than zero in northern populations (latitude greater than 60 °N).

1. INTRODUCTION

Early attempts to apply ideas from nonlinear dynamics to ecological systems employed noise-free deterministic models (Schaffer 1985; Schaffer & Kot 1986). These approaches have not succeeded in convincing ecologists that chaos may be found among natural populations, because populations are affected by both endogenous (density-dependent) and exogenous (density-independent) processes (Hastings *et al.* 1993). Recently, Ellner and co-workers (Ellner *et al.* 1991; Nychka *et al.* 1992; Ellner & Turchin 1995) have developed a theoretical framework for detecting chaos in dynamical systems affected by noise. I have used an approach based on this framework (implementation details are given in Turchin & Millstein 1993) to examine the latitudinal gradient in microtine population dynamics (Turchin 1993). My tentative conclusion was that microtine populations in Europe exhibit a dynamical shift from stability in the South to chaos in the North.

In their note, Falck *et al.* (1995) question this conclusion. Their main criticism is that the approach I used lacks a measure of uncertainty associated with each point estimate of the dominant Lyapunov exponent, and they propose a method for calculating confidence intervals. They also question several other aspects of the logical foundations of the approach. I answer these criticisms in turn.

2. BIAS AND UNCERTAINTY IN LE1 ESTIMATES

When estimating any quantity, we should be concerned about two important issues: presence and the degree of bias, and the degree of uncertainty associated with each point estimate. Falck *et al.* do not address the first issue, although, in my opinion, it is potentially a more serious problem. I have investigated whether the response surface method (RSM) for esti-

mating the dominant Lyapunov exponent (LE1) is biased by generating 'data' series of realistic length using ecological models incorporating such realistic features as multispecies interactions, dynamical noise, and measurement errors (Turchin & Millstein 1993). I have also generated such test data with the best-supported mechanistic model that we have for the specific system, microtine populations in Europe (Hanski *et al.* 1993; Turchin & Hanski unpublished results; the model assumes that oscillations are driven by a predator–prey interaction between *Microtus* voles and least weasels). The results suggest that there is a bias in the LE1 estimate, but the general tendency is to misclassify chaotic dynamics as stable, rather than vice versa. In other words, the RSM approach is actually conservative with respect to detecting chaos. It appears that any mechanism that increases data scatter around the estimated response surface – an insufficient embedding dimension, lack-of-fit errors, dynamical and observational noise – increases the degree of this bias (Turchin & Millstein 1993).

The tests with computer-generated data also indicated that the degree of uncertainty associated with each LE1 estimate based on 20–25 years of data is quite high: the standard deviation of the LE1 estimate is either similar or somewhat greater than the estimate's absolute value (tables 1–3 in Turchin & Millstein 1993). This result suggests that we will not be able to decisively support or reject the hypothesis of chaos, if we have only a single typical microtine data set to work with (a hardly unexpected conclusion!). For this reason, from the very beginning I have focused on an overall analysis of microtine data sets collected at a number of geographical localities (Turchin 1991, 1993).

Table 1 shows the results of analysing all the European data that I have found up to date (I did not include North American and Asian data sets because we are primarily interested in whether there is a shift

Proc. R. Soc. Lond. B (1995) **262**, 357–361
Printed in Great Britain

357

Table 1. *Summary of time-series analysis of European vole data*

(Abbreviations: lat = latitude in °N, d = the estimated embedding dimension, q = the estimated polynomial order, detr = whether the data were detrended (Y) or not (N), S = the standard deviation of log-transformed vole numbers. Sources of data: Finse, Framstad *et al.* (1993); Kilpisjärvi, Henttonen *et al.* (1984); Pallasjärvi, Henttonen *et al.* (1987); Kola, Koshkina (1966); Umeå, Hörnfeldt (1994); Ruotsala, Loppi and Sotkamo, Hanski *et al.* (1993) (data collected by A. Kaikusalo); Alajoki, Korpimäki & Norrdahl (1991); Karelia, Tataria, Serpukhov and Tula, Ivanter (1981); Grimsö, Lindström *et al.* (1994); Zvenigorod, Ivankina (1987); Bialowieza, Pucek *et al.* (1993); and Wytham Wood, Southern (1979))

locality	years	lat	d	q	detr	S	LE1
1. Finse[a]	1970–94	60	3	2	N	1.23	1.36
2. Kilpisjärvi 1[b]	1949–70	69	2	2	N	0.58	0.12
3. Kilpisjärvi 2[b]	1971–92	69	3	1	N	0.68	0.12
4. Pallasjarvi	1970–92	68	2	2	N	0.73	1.36
5. Kola	1946–64	67	2	2	N	0.79	0.70
6. Umeå	1971–91	64	2	2	Y	0.54	0.16
7. Sotkamo	1966–92	64	2	2	Y	0.49	0.21
8. Ruotsala	1973–92	63	3	2	N	0.65	0.87
9. Alajoki	1977–92	63	3	2	Y	0.77	1.63
10. Loppi	1972–92	61	1	2	N	0.27	−2.06
11. Karelia	1957–77	61	2	1	Y	0.35	−0.62
12. Grimso	1973–91	59	2	2	Y	0.50	0.38
13. Zvenigorod	1956–86	57	2	2	N	0.34	−0.18
14. Tataria	1936–58	56	1	1	Y	0.34	−2.69
15. Serpukhov	1936–58	55	2	2	N	0.26	−0.52
16. Tula	1936–58	54	2	1	Y	0.32	−0.46
17. Bialowieza	1971–91	52	1	1	Y	0.42	−1.02
18. Wytham Wood	1948–70	51	1	2	N	0.27	−0.79

[a] Because Finse is situated at the elevation of 1200 m above sea level, it is grouped with the northern data sets (for example, it has snow cover lasting from mid-September to the end of June (Framstad *et al.* 1993)).

[b] Kilpisjarvi data were analysed in two separate pieces to compensate for nonstationarity in this long data set.

from stability to chaos in European populations). Only one time-series per location is analysed: pooled numbers of microtine rodents trapped in the autumn. Population fluctuations from different locations are not significantly correlated, and thus we can use each location as a true replicate. Such a lack of spatial autocorrelation may be the result of, among other factors, a systematic shift in the statistical period of oscillations from almost five years in the far north to three years in the transition zone (Hanski *et al.* 1993). Another potential mechanism is the sensitive dependence on initial conditions that characterizes chaotic oscillations, which also promotes spatial asynchrony (Allen *et al.* 1993).

An informal meta-analysis of the expanded database (table 1) supports my previous conclusions (Turchin 1993): positive LE1 estimates in the North, and negative in the South. The mean (±s.e.) RSM estimate for the northern series (numbers 1–9) is 0.73 (±0.21), whereas for the southern series (numbers 10–18) the mean is −0.88 (±0.32). Assuming normality, the

respective 95 % confidence intervals are (0.27, 1.19) and (−1.61, −0.15). I hasten to add that this is not a formal meta-analysis, because I have implicitly assigned the same weight to each time series. However, I believe that in this case such an assumption is not unwarranted, since data from each location are of comparable length, and were collected using similar methods (this is especially true for the Fennoscandian locations).

Another relevant point is that there is a high degree of variation between estimates within each group (northern versus southern). The standard deviations for the northern and southern LE1 estimates are 0.61 and 0.87, respectively: on the same order of magnitude as LE1 estimates themselves. This result agrees with the patterns observed in the analysis of model-generated data. (Incidentally, it appears that obtaining positive LE1 estimates for all northern populations was likely a matter of chance; given such a high variability, we should expect a few negative estimates even when true LE1 is positive.) Again, my conclusion is that any single location cannot provide a definitive answer as to the sign of LE1. Thus, one would expect that estimating confidence limits for any of these data sets separately is unlikely to resolve the issue of whether vole oscillations in the North are chaotic or not. The only hope for obtaining any definitive answers rests with an approach that will combine numerous data sets.

3. SOME PROBLEMS WITH THE APPROACH OF FALCK *et al.* (1995)

Despite doubts expressed above, I welcome the attempt by Falck *et al.* to a develop a general method for calculating confidence limits for a single time series. Such a method may prove useful in the analysis of longer time series, for example, the famous Canadian lynx data. Additionally, a measure of uncertainty associated with each LE1 estimate would allow one to conduct a formal meta-analysis of all the data available.

Although I support the general direction taken by Falck *et al.* (1995), there are a number of problems in their method that need to be worked out before we can use it with confidence. The chief problem with their approach is that it requires a quantification of the noise component of the system that produced the time-series data. Estimating noise is probably the thorniest problem in statistical analysis. Simply calculating the variance of residuals is not a satisfactory procedure, because residuals are composed of the combined effects of dynamical noise, observation errors, and lack-of-fit errors. Therefore, the Falck *et al.* method probably inflates the magnitude of the dynamical noise. At the same time, their parametric bootstrap procedure assumes that the observation noise is absent. Yet, these two sources of noise can have very different implications for the estimation bias and precision. It is unknown how the method of Falck *et al.* is affected by these assumptions about the noise structure. By contrast, the statistical approach developed by Ellner *et al.* does not require an estimate of the dynamic noise (Ellner & Turchin 1995). Although our approach does

assume that the measurement noise is absent, the effect of making this assumption is known (a negative bias in the LE1 estimate).

Another problem is the logically inappropriate use of the confidence intervals calculated by Falck *et al.* (1995). They are interested in distinguishing between two hypotheses: H0 that LE1 is zero (or less) and H1 that LE1 > 0. The standard way to approach this problem is to estimate the distribution of the statistic (LE1) assuming the null hypothesis and to determine if the actual value falls into the rejection region. Instead, Falck *et al.* estimated the distribution of the statistic assuming H1. The problem with this approach is that the estimate of LE1 depends greatly on selecting the appropriate model complexity. Models with too few parameters, or too low embedding dimension, may not be able to approximate dynamics; whereas a model with too many parameters will fit the noise instead of only fitting the endogenous feedbacks (Ellner & Turchin 1995). Thus, fitting the data with an overly simple model will bias the LE1 estimate negatively, while using an overly complex model will bias the estimate positively. If the response surface model estimated for a data set was overly complex, then it would not only yield a supurious positive LE1, but also an equally spurious confidence interval.

A third problem is the censored nature of many sets of pseudodata on which confidence limits are based. In some cases, more than 90 % of generated time series had to be thrown out. While Falck *et al.* acknowledge this problem, they offer no potential solution for it, nor any indication as to how it might affect their results. This problem arises for the following reason. When the estimated response surface model with noise is iterated, the trajectory occasionally jumps outside the range of observed N_t values, where the fitted response surface does not provide meaningful predictions. A frequent consequence of leaving the data-circumscribed region is that the iterated trajectory will increase (or decrease) without limit. This problem was noted by Turchin & Taylor (1992) and Perry *et al.* (1993), who did not, however, offer a satisfactory solution to it. The Jacobian method used in Turchin (1993) does not suffer from this problem, because it does not require generating trajectories (see point 4 below).

Finally, I cannot agree with how Falck *et al.* interpret their results. As expected, their method yielded large confidence intervals, most of which overlapped zero (although, surprisingly, as many as six confidence intervals from the northern data were strictly positive). Just because the confidence intervals for the majority of data contain zero, however, is not a valid reason to conclude that there is no evidence for chaos in the northern populations. Any test of H1 versus H0 can be rendered 'non-significant' by dividing the data set into such small subsets, that H0 is not rejected on any of them due to lack of power. Thus, valid results may be obscured when all data available is not analysed as a whole. This point can be illustrated with the following example. Suppose one performs 20 studies measuring effects of an experimental drug. Each of the 20 studies indicates that the drug has a positive effect, but because the group of patients (data

points) in each study is always small, the effect is never statistically significant. Should we conclude from this that the drug really has no effect? On the contrary! Any kind of a meta-analysis combining all 20 studies together would show that the drug's effect is highly significant. Similarly, before we can conclude anything about the dynamics of northern populations in general, we need to combine together the results of the separate analyses of each data set for an overall test of significance.

4. OTHER ISSUES

There are several other statements made by Falck *et al.* (1995) that I would like to address:

1. In *Introduction*, Falck *et al.* imply that I have used a deterministic-based definition of chaos. In actuality, my approach is based on recent theoretical developments that specifically address how to analyse chaos in stochastic ecological systems (Ellner *et al.* 1991; Nychka *et al.* 1992). I cannot agree more with Falck *et al.* that natural populations are affected by both endogenous feedbacks and by exogenous influences. Therefore, an approach to the analysis of time-series data should incorporate both deterministic and stochastic components, as the response surface method does. These points are further discussed in Ellner & Turchin (1995).

2. As Falck *et al.* point out, the noise term in the response surface model is assumed to be state-independent and additive. This is certainly a weakness of the approach. My defence is that including a noise term at all is a great improvement over previous methods that assumed complete determinism. Further, the RSM approach performed adequately in tests with model-generated data, even though dynamical noise was included in models in a non-additive, state-dependent way. Nevertheless, I believe that more sophisticated ways of modeling noise (dynamical as well as measurement) would be a very fruitful direction to pursue in the future.

3. Falck *et al.* state that estimating the sign of LE1 requires 200–500 or more data points. This statement should be qualified as follows. First, this may be true for applications (e.g. in physics) that require a quantitative LE1 estimate. With the vole data, our goal is less ambitious: simply to determine whether LE1 is positive. Secondly, the relevant number is not really data points, but the number of excursions within the phase space (or oscillations). For example, 1000 points delineating a single oscillation would not be terribly useful, because 99 % of the data points will contain redundant information. Thirdly, long data series are not a panacea, especially in ecological applications where long-term data are almost always nonstationary.

The nine northern populations in table 1 have a combined total of 195 data points, corresponding to about 50 oscillations (at approximately four data points per oscillation). I believe this is a respectable data set on which to base one's conclusions. In addition, I view multiple localities as a virtue, rather than a problem: I would rather have nine series each 20 or so

years long, than a single series of 195 points. A broad agreement between results from different localities found in table 1 suggests that the finding of chaos is not limited to a freak combination of conditions peculiar to some location, but is rather a result of some general mechanism relevant to all microtine populations in northern Fennoscandia.

4. I am puzzled by the claim of Falck *et al.* that in estimating LE1 I make a restrictive assumption that two nearby trajectories are exposed to exactly the same sequence of noise. The confusion may have arisen because the approach for estimating LE1 from the fitted response surface has evolved with time. Originally, Turchin & Taylor (1992) characterized the dynamics of noise-free 'skeleton' by iterating the fitted response surface. A modification employed in Turchin (1991) was to estimate noise-dependent LE1 by iterating two nearby trajectories perturbed by the same sequence of dynamical noise. These approaches have been superseded by the Jacobian-based approach (Turchin 1993; Turchin & Millstein 1993). In this approach, a response surface function is fitted as before, but its Jacobians are evaluated at each data point. The LE1 is then calculated from the product of these Jacobians. Subjecting two nearby trajectories to the same sequence of noise is not an assumption of the approach, but rather a heuristic device employed by us to illustrate what chaos in a stochastic system means (see Ellner & Turchin 1995, pp. 348–350).

5. Falck *et al.* propose that instead of focusing on sensitive dependence, we should test for nonlinearity. Their test finds greater evidence for nonlinearity in northern, compared with southern time series. However, we already know that some dynamical shift occurs with latitude, as evidenced, for example, by a change in amplitude of fluctuations (Hansson & Henttonen 1985) and the strength of periodicity (Hanski *et al.* 1993). The point in question is whether the shift is extreme enough, so that northern populations are actually chaotic. The shift is not from linear to nonlinear dynamics: all population dynamics are inherently nonlinear (Royama 1981; Turchin & Millstein 1993).

In general, it is not clear to me whether a search for nonlinearity would really advance our understanding of population dynamics. Dynamics of all population systems are inherently nonlinear. Thus, if a test finds nonlinearity, then nothing new is learned. If the test cannot reject the hypothesis of linearity, then it simply means that either dynamics are approximately linear, or that the data are too noisy to make any conclusion. Again, nothing new is learned. By contrast, an approach that fits models to time-series data can tell us much more than whether the system is stable or chaotic (Turchin 1995). The estimated dimension (or order) of dynamics can be used to test hypotheses about biological mechanisms that may be driving oscillations. For example, presence of delayed density dependence is often an indication that trophic interactions may be important. Furthermore, nonlinear modeling of data raises the possibility that we eventually may be able to successfully forecast population dynamics, although this potential for prediction has not yet been realized.

5. CONCLUSION

The paper of Falck *et al.* is clearly an honest effort to tackle a very difficult problem. However, several problems with their approach remain to be resolved. It is particularly important to extensively test any proposed approach with simulated data to see whether it really works as expected. Furthermore, whatever method is proposed, it should be capable of combining confidence intervals estimated from different data sets into an overall measure of significance. Although it is underused by ecologists, meta-analysis is both tractable and powerful (Arnqvist & Wooster 1995).

Here is a suggestion for an alternative (and I believe more viable) approach for calculating confidence limits, based on a direct resampling of data. A simple resampling of the time series will, of course, not work, because it would destroy the time-dependent structures in the data, that is, the dependences of N_t on N_{t-1}, N_{t-2}, and so on. This difficulty can be overcome if we realize that our actual data are not a collection of N_t points, but rather a collection of $\{N_t, N_{t-1}, ..., N_{t-d}\}$ 'multiplets', where d is the maximum embedding dimension we are willing to consider (typically, no more than 3–4). The number of these multiplets is the time-series length minus d. The proposed bootstrap procedure samples with replacement among these multiplets, and then uses each pseudo-dataset to do a complete round of RSM fitting (however, the ordinary cross-validation should probably be avoided in favour of some other approach such as the general cross-validation, or the Bayesian information criterion). The resampling and fitting is repeated many times, and a bootstrapped distribution of LE1 estimates is constructed in a usual fashion. A meta-analysis can be accomplished with a hierarchical resampling scheme: first one samples among time series from different localities in the usual manner, and then within each time series as suggested above. I am currently investigating this approach by testing it with model-generated data in an approach similar to that in Turchin & Millstein (1993). A resampling approach, such as the one I sketched here, is possible only because fitting models with response surface methodology is quite fast on modern personal computers. For large data sets and more complex models, where multiple re-fittings are computationally infeasible, B. Bailey (Department of Statistics, North Carolina State University) has developed an alternative approach that calculates approximate confidence intervals based on likelihood-ratio tests.

In closing, any method may yield spurious results, and therefore, my conclusions that northern vole populations are chaotic should be considered as tentative, and subjected to further testing. Yet, the support for this hypothesis is not limited to evidence from time-series analyses. Parameterizations of the best-supported mechanism explaining vole oscillations – interactions with specialist mammalian predators – result in chaotic models (Hanski *et al.* 1993; Hanski & Korpimaki 1995, Turchin & Hanski, unpublished results). Furthermore, there is a widespread agreement that some sort of a bifurcation occurs in European vole population dynamics around 60° N. Southern popula-

tions are stable, characterized by strongly negative LE1s. The controversy is whether Lyapunov exponents in northern populations are significantly greater than zero, implying chaos, or near zero, implying *quasi-chaos* (Turchin 1995). Even if it turns out that the RSM approach is for some reason biased, and northern populations are actually quasi-chaotic, the implications of this result would not greatly differ from a finding of chaos. This is because LE1 measures the long-term average of trajectory divergence. When the long-term Lyapunov exponent is near zero, the dynamical system alternates between periods characterized by trajectory divergence with periods characterized by trajectory convergence. To quantify such transient effects, the long-term Lyapunov exponent may be replaced by state-dependent local Lyapunov exponents (Ellner & Turchin 1995). The point here is that there is no hard boundary between chaos and stability, so that a quasi-chaotic system with LE1 ≈ 0 will behave in many ways like a weakly chaotic system with LE1 > 0.

I thank Steve Ellner and Doug Nychka for suggestions, two anonymous referees for their comments on the manuscript, and the editors of the *Proceedings of the Royal Society of London* for providing an opportunity for me to publish this response.

REFERENCES

Allen, J. C., Schaffer, W. M. & Roscoe, D. 1993 Chaos reduces species extinction by amplifying local population noise. *Nature, Lond.* **364**, 229–232.

Arnqvist, G. & Wooster, D. 1995 Meta-analysis: synthesizing research findings in ecology and evolution. *Trends Ecol. Evol.* **10**, 236–240.

Ellner, S., Gallant, A. R., McCaffrey, D. & Nychka, D. 1991 Convergence rate and data requirements for Jacobian-based estimates of Lyapunov exponents from data. *Phys. Lett.* A **153**, 357–363.

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

Falck, W., Bjornstad, O. N. & Stenseth, N. C. 1995 Bootstrap estimated uncertainty of the dominant Lyapunov exponent for Holarctic microtine rodents. *Proc. R. Soc. Lond.* B (In the press.)

Framstad, E., Stenseth, N. C. & Ostbye, E. 1993 Time series analysis of population fluctuations of *Lemmus lemmus*. In *The biology of lemmings* (ed. N. C. Stenseth & R. A. Ims), pp. 97–115. London: Linnean Society.

Hanski, I. & Korpimäki, E. 1995 Microtine rodent dynamics in northern Europe: parameterized models for the predator-prey interaction. *Ecology* **76**, 840–850.

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

Hansson, L. & Henttonen, H. 1985 Gradients in density variations of small rodents: the importance of latitude and snow cover. *Oecologia* **67**, 394–402.

Hastings, A., Hom, C. L., Ellner, S., Turchin, P. & Godfray,

H. C. J. 1993 Chaos in ecology: is mother nature a strange attractor? *A. Rev. Ecol. Syst.* **24**, 1–33.

Henttonen, H., Tast, J., Viitala, J. & Kaikusalo, A. 1984 Ecology of cyclic rodents in northern Finland. *Mem. Soc. Fauna Flora Fennica* **60**, 84–92.

Henttonen, H., Oksanen, T., Jortikka, A. & Haukisalmi, V. 1987 How much do weasels shape microtine cycles in the northern Fennoscandian taiga? *Oikos* **50**, 353–365.

Hörnfeldt, B. 1994 Delayed density dependence as a determinant of vole cycles. *Ecology* **75**, 791–806.

Ivankina, E. V. 1987 Numerical dynamics and population structure of bank vole near Moscow. Ph.D. Thesis. Moscow University, Faculty of Biology. Moscow, Russia.

Ivanter, E. V. 1981 Population dynamics. In *European red vole* (ed. N. V. Bashenina), pp. 245–267. Moscow: Nauka.

Korpimäki, E. & Norrdahl, K. 1991 Do breeding avian predators dampen population fluctuations of small mammals? *Oikos* **62**, 195–208.

Lindström, E. R., Andrén, H., Angelstam, P. *et al.* 1994 Disease reveals the predator: sarcoptic mange, red fox predation, and prey populations. *Ecology* **75**, 1042–1049.

Nychka, D., Ellner, S., McCaffrey, D. & Gallant, A. R. 1992 Finding chaos in noisy systems. *J. R. statist. Soc.* B **54**, 399–426.

Perry, J. N., Woiwod, I. P. & Hanski, I. 1993 Using response-surface methodology to detect chaos in ecological time series. *Oikos* **68**, 329–339.

Pucek, Z., Jedrzejewski, W., Jedrzejewski, B. & Pucek, M. 1993 Rodent population dynamics in a primeveal deciduous forest (Bialowieza National Park) in relation to weather, seed crop, and predation. *Acta Theriol.* **38**, 199–232.

Royama, T. 1981 Fundamental concepts and methodology for the analysis of animal population dynamics, with particular reference to univoltine species. *Ecol. Monogr.* **51**, 473–493.

Schaffer, W. M. 1985 Order and chaos in ecological systems. *Ecology* **66**, 93–106.

Schaffer, W. M. & Kot, M. 1986 Chaos in ecological systems: the coals that Newcastle forgot. *Trends Ecol. Evol.* **1**, 58–63.

Southern, H. N. 1979 The stability and instability of small mammal populations. In *Ecology of small mammals* (ed. D. M. Stoddard), pp. 103–134. London: Chapman and Hall.

Turchin, P. 1991 Nonlinear modeling of time series data: limit cycles and chaos in forest insects, voles, and epidemics. In *Chaos and insect ecology* (ed. J. A. Logan & F. P. Hain), pp. 39–62. Virginia Experiment Station Information Series 91–3. Virginia: Blacksburg.

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

Turchin, P. 1995 Population regulation: old arguments and a new synthesis. In *Population dynamics* (ed. N. Cappucino & P. Price), pp. 19–40. New York: Academic Press.

Turchin, P. & Millstein, J. A. 1993 EcoDyn/RSM: Response surface modeling of nonlinear ecological dynamics. I. Theoretical background. Applied biomathematics, Setauket. New York. Available on request from the senior author.

Turchin, P. & Taylor, A. D. 1992 Complex dynamics in ecological time series. *Ecology* **73**, 289–305.