



HAL
open science

Detecting dynamical changes in nonlinear time series using locally linear state-space models

Anthony Ives, Vasilis Dakos

► **To cite this version:**

Anthony Ives, Vasilis Dakos. Detecting dynamical changes in nonlinear time series using locally linear state-space models. *Ecosphere*, 2012, 3 (6), pp.58. 10.1890/ES11-00347.1 . hal-02195160

HAL Id: hal-02195160

<https://hal.science/hal-02195160v1>

Submitted on 7 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Detecting dynamical changes in nonlinear time series using locally linear state-space models

ANTHONY R. IVES^{1,†} AND VASILIS DAKOS^{2,3}

¹*Department of Zoology, University of Wisconsin, Madison, Wisconsin 53706 USA*

²*Department of Aquatic Ecology & Water Quality Management, Wageningen University, P.O. Box 47, 6700AA Wageningen, The Netherlands*

Citation: Ives, A. R., and V. Dakos. 2012. Detecting dynamical changes in nonlinear time series using locally linear state-space models. *Ecosphere* 3(6):58. <http://dx.doi.org/10.1890/ES11-00347.1>

Abstract. Interest is growing in methods for predicting and detecting regime shifts—changes in the structure of dynamical processes that cause shifts among alternative stable states. Here, we use locally linear, autoregressive state-space models to statistically identify nonlinear processes that govern the dynamics of time series. We develop both time-varying and threshold models. In time-varying autoregressive models with p time lags, $AR(p)$, and vector autoregressive models for n -dimensional systems of order $p = 1$, $VAR(1)$, we assume that coefficients vary with time. We can infer an approaching regime shift if the coefficients indicate critical slowing down of the local dynamics of the system. In self-excited threshold models, we assume that the time series is governed by two autoregressive processes; the state variable switches between them when the time series crosses a threshold value. We use the existence of a statistically significant threshold as an indicator of alternative stable states. All models are fit to data using a state-space form that incorporates measurement error, and maximum likelihood estimation allows for statistically testing alternative hypotheses about the processes governing dynamics. Our model-based approach for forecasting regime shifts and identifying alternative stable states overcomes limitations of other common metric-based approaches and is a useful addition to the toolbox of methods for analyzing nonlinear time series.

Key words: alternative stable states; critical slowing down; critical transition; early warning signals; Kalman filter; regime shift; SETAR; SETARSS; threshold autoregressive models; time-varying autoregressive models; TVAR; TVARSS.

Received 7 December 2011; revised 14 May 2012; accepted 16 May 2012; final version received 5 June 2012; **published** 29 June 2012. Corresponding Editor: J. Drake.

Copyright: © 2012 Ives and Dakos. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits restricted use, distribution, and reproduction in any medium, provided the original author and sources are credited.

³ Present address: Integrative Ecology Group, Estación Biológica de Doñana, C/ Américo Vespucio S/N, E-41092 Sevilla, Spain.

† **E-mail:** arives@wisc.edu

INTRODUCTION

Ecological dynamics may experience abrupt transitions between alternative stable states, often termed regime shifts (Scheffer et al. 2001, Carpenter 2003). Although regime shifts are difficult to predict, recent work has proposed a series of statistical properties that change in predictable ways before a system shifts to an

alternate state, and these properties can be used as generic early warning signals (Scheffer et al. 2009, Clarke and Signorino 2010). Theoretically, most of these signals are manifestations of critical slowing down; the return rate of a system towards one of the stable states slows down as a bifurcation is approached at which the system shifts to the alternative state (Wissel 1984, Strogatz 1994). Generic early warning signals

include a rise in variance (Carpenter and Brock 2011), a decline in recovery rate after a perturbation (van Nes and Scheffer 2007), an increase in autocorrelation (Held and Kleinen 2004), and a peak in skewness (Guttal and Jayaprakash 2008).

An advantage of using these types of generic metrics as early warning signals is that they are relatively simple to measure. A disadvantage, however, is that they are descriptive; they only describe the increasingly slow dynamics around a stable state that is gradually forced towards a regime shift (Scheffer et al. 2009). This leads to at least four handicaps. First, statistical inference is problematic; it is difficult, for example, to statistically test hypotheses about changes in dynamics, such as whether the dynamics are in fact slowing down. Second, it is difficult to interpret changes in the metrics in cases where there are multiple simultaneous changes (Carpenter and Brock 2011). For example, as a regime shift is approached, there could be changes not only in critical slowing down, but also in the location (mean value) of a stable state, and these two patterns need to be separated to obtain a true measure of the return rate to the stable state. Third, simple metrics provide no way to account for measurement error, and variation due to measurement error can muffle early warning signals (Carpenter and Brock 2010). Finally, identifying regime shifts using metrics often requires long records that are not commonly available for ecological systems. Thus, it is no surprise that so far the best empirical examples of successful early warning metrics come from long paleoclimate records (Dakos et al. 2008) or from experimental work where the signals are compared to a control treatment (Drake and Griffen 2010, Carpenter et al. 2011).

Some of these handicaps have been addressed in part by alternative approaches. For example, null models have been used as a way to estimate significance levels for some metrics, such as autocorrelation measured in paleoclimate time series (Dakos et al. 2008). Conditional heteroscedasticity has been proposed as an alternative metric to the variance and can be associated with probability values (Seekell et al. 2011). Generalized modeling approaches that incorporate multiple sources of data are suggested as ways to directly estimate return rates in cases where data are scarce (Lade and Gross 2012). Nonparametric regression

has been used to fit Drift-Diffusion-Jump models to time series where the underlying dynamics are affected by multiple processes (Carpenter and Brock 2011). Still, these approaches require a predefined model, do not explicitly account for measurement error, or need long records.

To try to overcome some of these handicaps, we developed locally linear state-space models that can detect changes through time in the stochastic dynamics of time series. Although this approach can be used to detect many types of changes in dynamics, we specifically use it to detect either critical slowing down of dynamics approaching a regime shift or “flickering” of a state variable between alternative states. Conceptually, the approach is a stochastic analogue to performing linear approximations of nonlinear, deterministic systems to infer the local dynamics around stationary points (May 1974). Even for nonlinear deterministic systems, linear stability analysis is informative about qualitative dynamics and bifurcations. In a similar way, the locally linear stochastic approach that we take can be used to infer qualitative changes in nonlinear stochastic systems. Linear stochastic models can give good approximations to the variance and lagged autocorrelations of nonlinear time series and thereby are useful when insufficient information is available to build more-realistic nonlinear models of a system.

We present two specific formulations of locally linear autoregressive state-space models: (1) time-varying autoregressive state-space models of lag order p , TVARSS(p), and (2) self-excited threshold autoregressive state-space models, SETARSS(p). Time-varying AR(p) models approximate the local dynamics of a time series by estimating autoregressive coefficients that are allowed to vary through time. This gives information about the “local stability” of the system. In addition to univariate TVARSS(p) models, we also develop time-varying vector autoregressive models of order $p = 1$, TVVARSS(1), that contain more than one state variable. TVVARSS(1) models can be fit to multiple time series simultaneously, assuming time series interact such that the values of one time series directly affect the values of another. SETARSS(p) models are autoregressive models in which the coefficients switch between different parameter sets, with the switch occurring when the state variable

crosses a threshold value; they are the state-space extension of self-excited threshold autoregressive models, SETAR(p) (Tong 1990). SETARSS(p) models could be used, for example, to explore the existence of alternative stable states, in which the dynamics within each state are approximated by a separate AR(p) model and the time series flickers between the domains of attraction to the different alternative states.

METHODS

Autoregressive models

To explain the analyses of the time-varying and threshold AR models, it is useful to begin with a brief exposition of the properties of standard AR models. A standard AR(p) model (Box et al. 1994) is

$$x(t) = b_0 + \sum_{i=1}^p b_i (x(t-i) - b_0) + \varepsilon(t) \quad (1)$$

where $x(t)$ is the possibly transformed state variable (e.g., species density), b_0 sets the mean (location) of the time series, b_i ($i \geq 1$) are the autoregression coefficients, and $\varepsilon(t)$ is the environmental variability associated with changes in the state variable; $\varepsilon(t)$ is assumed to be a Gaussian random variable with mean zero and variance σ_ε^2 . Although linear, the AR(p) model is flexible enough to describe many types of complex dynamics. Even for small values of p , an AR(p) can well-approximate the variance and autocorrelations of even highly nonlinear time series (e.g., Ives and Jansen 1998).

Whereas the coefficient b_0 determines the mean of the time series, the coefficients b_i ($i \geq 1$) determine the dynamics around the mean. A useful quantity in describing these dynamics is the inverse of the characteristic root (Box et al. 1994), which we denote λ . This quantity is similar to the magnitude of the dominant eigenvalue of the Jacobian matrix computed at a stationary point of a deterministic discrete-time model (Strogatz 1994). Values of λ near zero imply that the state variable is drawn rapidly towards the mean b_0 , and as values approach one, the central tendency of the time series diminishes (Ives et al. 2003). When λ exceeds one, the process is nonstationary; the variance of the process increases continuously through time. This is analogous to the magnitude of the

dominant eigenvalue of the Jacobian matrix exceeding one, in which case the stationary point becomes unstable.

The standard AR(p) assumes that the underlying process generating a time series does not itself change. Specifically, the location of the mean b_0 , and the inverse of the characteristic root λ (which is determined by the coefficients b_i , $i \geq 1$) are time invariant. Nonetheless, the properties of time-invariant AR(p) models have counterparts in time-varying and threshold autoregression models, and can be used to give insights into the dynamics of time-varying and nonlinear processes.

TVARSS(p) models

The general form of the suite of TVARSS(p) models we consider is

$$x(t) = b_0(t-1) + \sum_{i=1}^p b_i(t-1) \times (x(t-i) - b_0(t-1)) + \varepsilon(t) \quad (2)$$

$$b_i(t) = b_i(t-1) + \phi_i(t).$$

The top equation is a standard AR(p) model but with coefficients expressed as functions of time, and the bottom equation allows the autoregression coefficients $b_i(t)$ ($i \geq 0$) to vary as random walks, with the rates of the random walks dictated by the variances σ_i^2 of $\phi_i(t)$. Although in principle the values of the autoregressive coefficients are unbounded (since they follow random walks), the values are constrained by fitting to the data. In addition to the process Eq. 2, we assume that there is a measurement equation

$$x^*(t) = x(t) + \alpha(t) \quad (3)$$

in which $x^*(t)$ is the observed value of the state variable, and $\alpha(t)$ is a Gaussian random variable with mean zero and variance σ_α^2 depicting measurement error.

The parameters in this model that must be estimated are the initial ($t = 0$) values of the coefficients b_i , the variances in the state variables σ_ε^2 and σ_i^2 , and the measurement error variance σ_α^2 . This time-varying model is conditionally Gaussian; conditional on the Gaussian distribution of state variables at time $t - 1$, the distributions of $x(t)$ and $b_i(t)$ are Gaussian at time t . Therefore, the likelihood of parameter

estimates can be computed exactly using a Kalman filter (Harvey 1989). This also permits likelihood ratio tests for statistical inference about the parameter estimates and likelihood-based model selection such as Akaike's Information Criterion, AIC (Akaike 1973). In practice, for realistically short time series (<200 points) and higher-order models ($p > 1$), the likelihoods are not unimodal (Ives et al. 2010); therefore, caution is required when numerically computing the maximum likelihood parameter estimates.

In likelihood ratio tests, because the possible values of the variances σ^2 are constrained to be positive, the asymptotic values of the log-likelihood ratios computed for the full model and the model reduced by q parameters are distributed by a mixture of χ^2 distributions (Self and Liang 1987, Stram and Lee 1994). For the case in which the variances σ_i^2 of $\phi_i(t)$ are assumed to be independent, the distribution of the likelihood ratio is asymptotically given by

$$\sum_{k=0}^q \binom{q}{k} 2^{-q} \chi_k^2$$

(case 9 in Self and Liang 1987). For example, for testing a single parameter ($q = 1$) as in the case of a TVARSS(1), the test distribution is a 50:50 mix of χ^2_0 and χ^2_1 distributions; thus, the P -values given by the constrained likelihood ratio test are one-half the values that would be calculated from a standard likelihood ratio test using χ^2_1 .

The Kalman filter propagates estimates of both the mean and the variance of the state variables ($x(t)$ and $b_i(t)$) (Harvey 1989). This makes it possible to calculate the time-dependent confidence intervals of λ . For the case of TVARSS(1) models, $\lambda(t) = b_1(t)$, and therefore the distribution of the estimates of $\lambda(t)$ is Gaussian. For TVARSS($p > 1$) models, $\lambda(t)$ depends in a nonlinear way on $b_i(t)$ ($i \geq 1$), and therefore the distribution of its estimate is not Gaussian. Nonetheless, we can approximate the distribution of the estimate numerically to give confidence intervals; here, we report the 67% confidence intervals that coincide with standard errors (which give the 67% confidence intervals for Gaussian distributions).

Numerous modifications and constraints can be added to the model. For example, as formulated above, the variance parameters σ_i^2 are unbounded above, so that large estimates that

allow rapid changes to the coefficients b_i are possible. Instead, an upper limit on σ_i^2 could be imposed to force the coefficients to vary smoothly through time. Also, the model assumes that there is measurement error ($\sigma_\alpha^2 > 0$) and estimates this simultaneously with the other parameters. If the magnitude of measurement error is known from independent information, then this estimate could be used, thereby alleviating the need to estimate σ_α^2 . This will be particularly advantageous if the measurement error changes through time or according to the values in the time series.

Two related decisions need to be made: selecting the time step of the time series if it is very long, and whether to detrend the data. The state-space structure of the model allows data points to be unequally spaced by iterating the process equations (e.g., Eq. 2) without updating the Kalman filter with the measurement equations (e.g., Eq. 3) (Harvey 1989). Therefore, for short time series with unequally spaced samples, it is best to formulate a model with unequally spaced sample points rather than select a longer iteration interval to create equally spaced points. For very long time series, it might be advantageous to subsample the data, especially if the measurement error is known to be high, because in this case variation between sequential points will largely be caused by measurement error. As for detrending, this will often be unnecessary, because the mean location of the time series, $b_0(t)$, is allowed to vary through time. Thus, detrending is performed as part of the estimation of the TVARSS model, with the degree of detrending determined by the maximum likelihood estimate of σ^2_0 . Nonetheless, if there are strong trends in the data, then simple linear or quadratic detrending can be used (e.g., fitting a line $x(t) = c_0 + c_1t$ and analyzing the residuals). We recommend against using a smoothing algorithm, because it is unclear how this might affect the estimates of the parameters of the TVARSS model. Decisions like these about how to formulate the TVARSS, and how to transform and initially process the data, will generally depend on the data set (Clarke and Signorino 2010).

TVVARSS(1) models

The time-varying vector autoregressive state-space model, TVVARSS(1), is closely related to

TVARSS(p) models, although multiple time series can be analyzed simultaneously. If $\mathbf{X}(t)$ denotes the n -dimensional vector of state variables $x_i(t)$ ($i = 1, \dots, n$), then

$$\begin{aligned} \mathbf{X}(t) &= \mathbf{B}_0(t-1) + \mathbf{B}(t-1) \\ &\quad \times (\mathbf{X}(t-1) - \mathbf{B}_0(t-1)) + \mathbf{E}(t) \\ b_{ij}(t) &= b_{ij}(t-1) + \phi_{ij}(t) \end{aligned} \tag{4}$$

where $\mathbf{B}_0(t)$ is the n -dimensional vector containing the coefficients $b_{i,0}(t)$, $\mathbf{B}(t)$ is the $n \times n$ matrix containing coefficients $b_{i,j}(t)$ ($i, j = 1, \dots, n$), and $\mathbf{E}(t)$ is the n -dimensional Gaussian random variable with mean zero and covariance matrix Σ_ϵ . The coefficients $b_{i,j}(t)$ are assumed to follow random walks with rates given by the variances of $\phi_{i,j}$, $\sigma^2_{i,j}$ ($i = 1, \dots, n; j = 0, \dots, n$). The state variables $\mathbf{X}(t)$ are assumed to be measured with error

$$\mathbf{X}^*(t) = \mathbf{X}(t) + \mathbf{A}(t) \tag{5}$$

where $\mathbf{A}(t)$ is an n -dimensional Gaussian random variable with covariance matrix Σ_α .

We only consider TVVARSS(1) processes, rather than TVVARSS($p > 1$). Although there are no technical restrictions limiting the approach to $p = 1$, for most applications $p = 1$ is likely to be sufficient. Numerous constraints can further be imposed on the model, such as assuming that environmental variability and/or measurement error are independent between state variables (i.e., Σ_ϵ and/or Σ_α are diagonal). As with the TVARSS(p) model, the Kalman filter can be used to compute the likelihood, allowing maximum likelihood estimation and statistical inference.

SETARSS(p) models

The self-exciting threshold autoregressive state-space model, SETARSS(p), assumes that there are two possible AR(p) models governing dynamics, with the possibility that the state variable switches between them when it crosses a threshold. Specifically,

$$\begin{aligned} x(t) &= b_0 + \sum_{i=1}^p b_i (x(t-i) - b_0) + \epsilon(t) \\ &\quad \text{when } x(t-1) > c \\ x(t) &= b'_0 + \sum_{i=1}^p b'_i (x(t-i) - b'_0) + \epsilon(t) \end{aligned}$$

$$\text{when } x(t-1) \leq c \tag{6}$$

where the coefficients b_i and b'_i ($i = 0, \dots, p$) denote separate sets of coefficients; this is the SETAR model (Tong 1990) that assumes no measurement error. Here, we assume that there is measurement error that does not differ between the two AR(p) systems and is therefore given by Eq. 3. As with the time-varying models, the SETARSS(p) is conditionally Gaussian, and therefore the Kalman filter can be used to compute its exact likelihood. In addition to the two sets of autoregression parameters b_i and b'_i , parameters to be estimated are the threshold c , the variance of the process error σ^2_ϵ , and the variance of the measurement error σ^2_α .

Data simulations

To illustrate the methods, we used two stochastic simulation models that produce different types of dynamics. The first, the single-species harvesting model of May (1977), produces two alternative stable points over a range of values of the parameter that governs the harvesting rate, a . We use this model to produce two data sets. In the first, we increased the value of a so that one of the two stable points collapses at a fold bifurcation; this provides a test for whether the TVARSS(p) models can detect the impending collapse before it occurs. In the second, we set a at a value giving alternative stable points and added sufficiently high environmental variability that the population jumps between the domains of attraction to the two stable points; this provides a test of the SETARSS(p) models to identify flickering between alternative stable states. The second model we used is the Nicholson-Bailey predator-prey model (Nicholson and Bailey 1935) with prey density dependence that undergoes a Neimark-Sacker bifurcation (discrete-time version of a Hopf bifurcation). Reducing the strength of prey density dependence leads to a deterministic stable limit cycle. The simulated data set consisting of only the density of prey provides a test for TVARSS(p) models to identify a regime shift to greater population cyclicity, and the data set consisting of both prey and predator densities provides a similar test for the TVVARSS(1) model.

May harvesting model.—The harvesting model

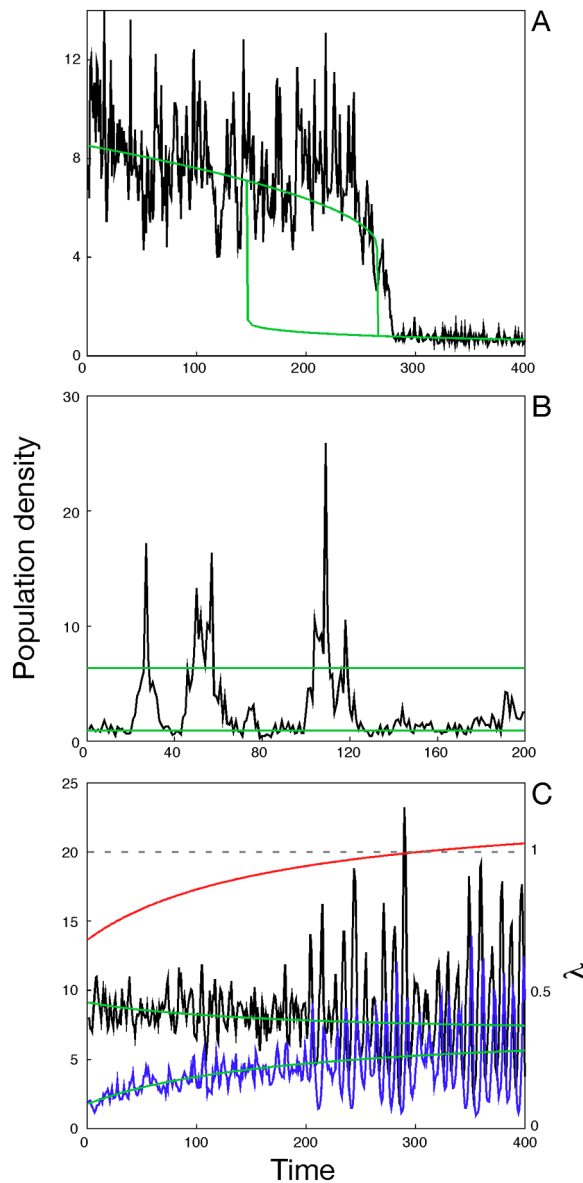


Fig. 1. Two simulation models illustrating different scenarios that we used to test the locally linear autoregressive state-space models. (A) The harvesting model of May (1977) shows alternative stable states (green lines) over a range of values of the harvesting rate, a , which in the simulation is increased through time from $a = 1$ ($t = 0$) to $a = 3$ ($t = 400$). The stochastic time series initially stays within the domain of attraction to the upper stable point, but as time (and a) increases, the upper stable point collapses, and the time series drops to the lower, persisting stable state through a fold bifurcation. Other parameters are $r = 0.75$, $K = 10$, $P = 3$, $H = 1$, $\sigma_N = 0.4$, and $\sigma_\alpha = 0.2$. (B) The

of May (1977) describes the discrete-time changes in density of a population $N(t)$ as

$$N(t) = N(t-1) \exp \left[r \left(1 - \frac{N(t-1)}{K} \right) - \frac{aN(t-1)^Q}{N(t-1)^Q + H^Q} + \varepsilon_N(t) \right] \quad (7)$$

where a gives the harvesting effort, the carrying capacity K sets the strength of density dependence, r is the intrinsic rate of increase, and H and Q determine how the harvesting rate depends on $N(t)$. We assume that there is environmental variability given by the Gaussian random variable $\varepsilon_N(t)$ that has mean zero and variance σ_N^2 . In addition, we assume that there is lognormal measurement error, so that the observed density $N^*(t)$ is

$$N^*(t) = N(t) \exp[\alpha_N(t)] \quad (8)$$

where $\alpha_N(t)$ is a Gaussian random variable with mean zero and variance σ_α^2 .

Over a range of values of a , the May harvesting model has alternative stable points. If the population is initially in the domain of attraction to the upper stable point, increasing a will eventually lead to the annihilation of this point and a rapid drop in density to the lower stable point through a fold bifurcation (Fig. 1A). To test

harvesting model of May (1977) shows flickering between the domains of attraction to two stable points. In contrast to (A), in (B) no parameter changes through time, and the switches between domains of attraction are caused by the stochasticity of the system. Green lines give the alternative stable states. Parameter values are $r = 0.75$, $K = 10$, $a = 2$, $P = 3$, $H = 1$, $\sigma_N = 0.2$, and $\sigma_\alpha = 0.1$. (C) The Nicholson-Bailey predator-prey model shows cyclic dynamics, and as the carrying capacity of the prey, K , decreases through time from 12 ($t = 0$) to 30 ($t = 400$), the cycles become more pronounced. The dominant eigenvalue of the Jacobian matrix around the stable point (red line) increases through time, reaching 1 at about $t = 300$ corresponding to a Neimark-Sacker bifurcation. With increasing K , the prey density decreases (black line) while that of the predator increases (blue line), with the deterministic stationary point given by the green lines. Other parameter values are $r = 0.75$, $k = 0.1$, $\sigma_N = 0.1$, $\sigma_P = 0.1$, and $\sigma_\alpha = 0.05$.

the ability of TVARSS(p) models to detect this regime shift, we truncated the data when it reached the midpoint between the upper and lower stable points immediately before the collapse of the upper stable point.

When there is a fixed value of a giving two stable points, and when the environmental variance σ^2_N is large, the population jumps between the domains of the two stable points (Fig. 1B). The underlying existence of the stable points is suggested because densities linger within each of the domains of attraction for successive time steps, generating flickering between the two stable points. We used this simulated data set to test whether the SETARSS(p) model could detect the existence of these alternative stable points.

Nicholson-Bailey predator-prey model.—The Nicholson-Bailey model (1935) is a simple model of predator-prey interactions in which predators search for prey randomly, and those prey that are attacked develop into predators in the following generation. Because there is a one-to-one match between prey killed and predators produced, this model is suited for parasitoid-host systems in which parasitized hosts give rise to the next generation of parasitoids. The population dynamics of prey $N(t)$ and predators $P(t)$ are given by

$$\begin{aligned} N(t) &= N(t-1) \exp \left[r \left(1 - \frac{N(t-1)}{K} \right) \right. \\ &\quad \left. - kP(t-1) + \varepsilon_N(t) \right] \\ P(t) &= N(t-1) \exp \left[r \left(1 - \frac{N(t-1)}{K} \right) \right] \\ &\quad \times \{ 1 - \exp[-kP(t-1)] \} \exp[\varepsilon_P(t)] \end{aligned} \quad (9)$$

where prey have intrinsic rate of increase r and experience density dependence with carrying capacity K . The proportion of prey that escape predation is $\exp[-kP(t)]$, the zero term in a Poisson distribution with mean $kP(t)$, and so the proportion killed is $1 - \exp[-kP(t)]$. Both prey and predator experience environmental variability incorporated as independent Gaussian random variables $\varepsilon_N(t)$ and $\varepsilon_P(t)$ that have means zero and variances σ^2_N and σ^2_P . Furthermore,

both prey and predator densities are assumed to be measured with error, so that their observed densities $N^*(t)$ and $P^*(t)$ are given by

$$\begin{aligned} N^*(t) &= N(t) \exp[\alpha_N(t)] \\ P^*(t) &= P(t) \exp[\alpha_P(t)]. \end{aligned} \quad (10)$$

As K increases, prey density dependence weakens, and the population dynamics show greater cyclicity. Increasing K eventually leads to a Neimark-Sacker bifurcation in which the stable point of the deterministic skeleton of Eqs. 9 (i.e., the equations stripped of environmental stochasticity) converts into a stable limit cycle. In the fully stochastic model, however, there is no marked transition in dynamics at this bifurcation (Fig. 1C), as the environmental variability makes cycles persist before the bifurcation is reached; nonetheless, increasing K leads to greater intrinsically driven cycles. For analyses, we truncated the data at the bifurcation point. We tested TVARSS(p) models using only the prey density and the TVARSS(1) models with both prey and predator densities.

RESULTS

TVARSS(p) models

We fit TVARSS(p) models with $p = 1, 2,$ and 3 to the data simulated from the May harvesting model (Fig. 1A); the data were log-transformed, normalized, and truncated immediately before the collapse of the stable point. The TVARSS(1) model (Fig. 2) fit the data best according to the AIC ($\Delta\text{AIC} = 0.04$ and 0.47 for $p = 2$ and 3 , respectively). We then fit TVARSS(p) models in which only the estimate of the mean of the AR process, $b_0(t)$, was allowed to vary through time ($\sigma^2_i = 0$ for $i \geq 1$). The log-likelihood ratio between the two models with $p = 1$ was 3.59 , implying that the model allowing the coefficient b_1 to vary was statistically significantly better ($0.5\chi^2_0 + 0.5\chi^2_1 = 7.18$, $P < 0.004$). A similar test for the difference between the models with and without time variation in b_i ($i \geq 1$) was significant for both $p = 2$ ($0.25\chi^2_0 + 0.5\chi^2_1 + 0.25\chi^2_2 = 6.56$, $P < 0.02$) and 3 ($0.125\chi^2_0 + 0.375\chi^2_1 + 0.375\chi^2_2 + 0.125\chi^2_3 = 8.40$, $P < 0.02$). We also estimated the inverse of the characteristic root, λ , of the TVARSS(1) model and its 67% confidence interval (standard errors) at each point in the time series from the estimates of $b_1(t)$. As the time

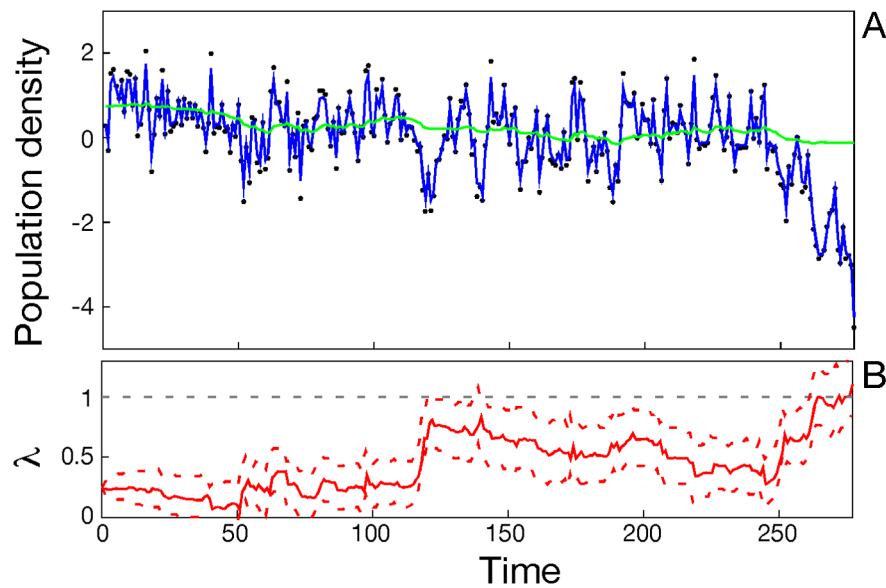


Fig. 2. Illustration of a TVARSS(1) model fit to simulated data from the harvesting model of May (1977) (Eqs. 7 and 8). The harvesting rate a increases through time to cause the collapse of the upper stable point through a fold bifurcation (Fig. 1A). (A) Differences between the fitted trajectory (blue line) and the simulated data (black dots) are attributed in the model to measurement error. The green line gives the time-varying estimate of $b_0(t)$ from the TVARSS(1). (B) The inverse of the characteristic root for the TVARSS(1) is simply $|b_1(t)|$, which is plotted through time with its 67% confidence interval (dashed lines). Parameter estimates are: initial $b_0 = 0.737$, initial $b_1 = 0.241$, $\sigma_\varepsilon = 0.558$, $\sigma_\alpha = 0.328$, $\sigma_1 = 0.039$, and $\sigma_2 = 0.016$, and the log likelihood is -288.85 .

series approached the collapse of the upper stationary point, the inverse of the characteristic root approached and exceeded 1 (Fig. 2). Values of λ approaching 1 imply critical slowing down, since the tendency of $x(t)$ to return towards the mean is lowered. Values of $\lambda > 1$ imply loss of stationarity, with $x(t)$ being pushed away from the mean value given by $b_0(t)$.

In this analysis, the confidence interval of the estimate of λ includes 1 as the regime shift is closely approached at the end of the time series. Therefore, if we were to ask whether the prediction of an impending regime shift is statistically significant, we would conclude no. This is not surprising, however, because the regime shift has not yet occurred, so the statistical power to identify a loss of stationarity at the final point in the time series is low. At the same time, however, there is strong statistical support for time-varying autoregression coefficients, and the estimate of λ from these coefficients increases towards nonstationarity. In a real data set, we would use this as evidence that an

imminent regime shift is possible, although without detailed knowledge of the nonlinearities of the system, we could not ascribe a degree of certainty to this prediction. Because the confidence bounds on λ are fairly wide, any conclusion about whether the time series has actually become nonstationary must be made cautiously; while the best (i.e., maximum likelihood) estimates from the time series imply nonstationarity, considerable uncertainty in the nonstationarity still exists. Unfortunately, the system will not stay in a nonstationary state very long and, if there is a regime shift, it will occur quickly. Therefore, from a management perspective, if there is a way to mitigate a regime shift, a manager will not have the leisure to wait for statistical significance.

These results are given only for a single, example time series. A full exploration of the method requires detailed analyses of simulated data. We are strong advocates of simulation studies, yet simulation studies are most informative when they are tailored for specific data sets that a researcher has in hand; the researcher can

then produce simulations that mimic the real data to explore the power and possible bias of a given statistical method. To illustrate the type of analyses we think should routinely be performed, we simulated the May harvesting model to generate 500 simulated data sets under the same parameter values as Fig. 2. For each simulated data set, we estimated parameters of a TVARSS(1) under the null hypothesis H_0 that coefficient b_1 did not change ($\sigma^2_1 = 0$) and under the hypothesis H_1 that $b_1(t)$ varied through time (Fig. 3). In only 212 of the 500 simulated data sets was H_0 rejected at the $\alpha = 0.05$ level. Thus, the statistical power of the test for this particular model is not high. In no case was the estimate of b_1 under H_0 greater than 1 (Fig. 3A); for this model, $|b_1(t)|$ equals the inverse of the characteristic root λ , so no data set was identified as nonstationary when analyzed under the null hypothesis. In contrast, when fit under H_1 , 36% of the estimates of λ were greater than 1 at the last time point in the data set. For those 212/500 data sets for which H_0 was rejected, 68% had estimates of $\lambda > 1$ at the last time point, suggesting that the time series was at a regime shift. Even for data sets that did not have the final values of $\lambda > 1$, for the data sets for which H_0 was rejected, 98% had a value of λ at the last time point greater than the value at the first time point, indicating that the time series was becoming less stationary and suggesting an approaching regime shift. Our intent in presenting these simulations is to illustrate the value of simulations when assessing the results of these or any methods of analyzing time series; however, because the power and performance of our methods will depend on the length of data sets and the strength of signal within them, we would not expect the results from our simulations to be representative of other applications.

We also performed simulations to test whether the TVARSS(p) models might give false positives (type I errors), rejecting the null hypothesis (no change through time) when it is in fact true. We simulated 500 data sets under the null hypothesis by fixing $a = 1$ in Eq. 7, and then fit each data set to a TVARSS(1) model and performed a likelihood ratio test against the null hypothesis that b_1 is time-invariant. The null hypothesis was rejected (the time-varying model was significantly better) in 10/500 data sets, giving a incorrectly

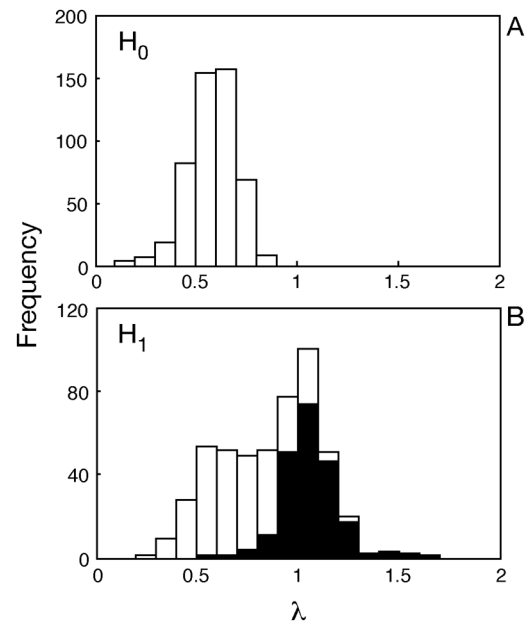


Fig. 3. Simulation study of the TVARSS(1) model. Five hundred data sets were simulated from the May harvesting model under parameter values as in Fig. 1A that produce a regime shift, with the data sets truncated before the regime shifts occurred (when the population dropped to the midpoint between the upper and lower stationary points immediately before the alternative states collapse). (A) The distribution of the resulting 500 estimates of the inverse of the characteristic root, λ , when fitting the null model assuming the autoregressive coefficient b_1 does not change. (B) The distribution of λ estimated at the last time point in the time series for the full model with time-varying estimates of $b_1(t)$. In (B), the 212/500 simulations for which the null hypothesis was rejected are shown in black. For simulations, the parameter $a = 1$ ($t = 0$) to $a = 3$ ($t = 400$), and the other parameter values are $r = 0.75$, $k = 0.1$, $\sigma_N = 0.1$, $\sigma_P = 0.1$, and $\sigma_\alpha = 0.05$.

low rejection rate of 0.02 at the α -level of 0.05. Thus, in this scenario the TVARSS(1) model did not reject as many data sets as it should have, suggesting low power. In many situations, it is statistically preferable to err on the side of false negatives than false positives, so that we do not infer patterns that in fact are not real. However, in the case of regime shifts, erring on the side of false positives is probably more often preferable; if the consequences of a regime shift are

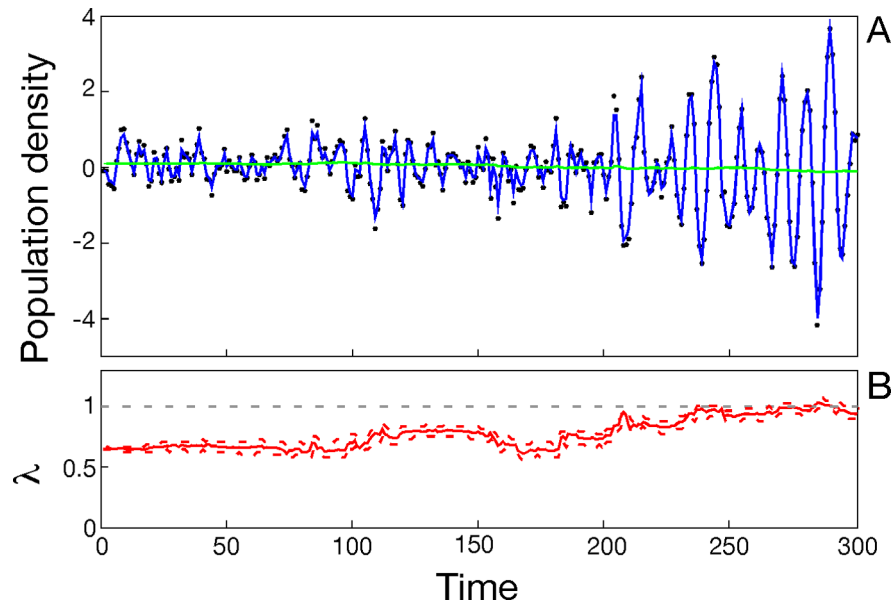


Fig. 4. Illustration of a TVARSS(2) model fit to simulated data for prey from the Nicholson-Bailey predator-prey model (Eqs. 9 and 10). The carrying capacity of the host, K , increases through time to cause a Neimark-Sacker bifurcation at about $t = 300$ (Fig. 1C). (A) Differences between the fitted trajectory (blue line) and the simulated data (black dots) are attributed in the model to measurement error. The green line gives the time-varying estimate of $b_0(t)$ from TVARSS(2). (B) The inverse of the characteristic root for the AR(2) is computed from $b_1(t)$ and $b_2(t)$, which is plotted through time with its 67% confidence interval (dashed lines). Parameter estimates are: initial $b_0 = 0.096$, initial $b_1 = 0.961$, initial $b_2 = -0.429$, $\sigma_\varepsilon = 0.302$, $\sigma_\alpha = 0.253$, $\sigma_0 = 0.0094$, $\sigma_1 = 0.027$, and $\sigma_2 = 0.026$, and the log likelihood is -206.92 .

sufficiently great, missing a warning should be avoided. Therefore, we suggest that researchers do not take as dogma an α -level of 0.05 but instead select a (higher) α -level consistent with the risks of a regime shift, and then perform simulations using a model tailored to their specific system to assess possible type I and II errors.

In addition to simulations of the May harvesting model, we investigated whether TVARSS(p) models could identify changes in the dynamics leading up to the Neimark-Sacker bifurcation in the predator-prey model (Fig. 1C) using only the time series for the prey. The TVARSS(2) model which allowed all coefficients b_i to vary through time was selected over the models with $p = 1$ ($\Delta\text{AIC} = 212$) and $p = 3$ ($\Delta\text{AIC} = 0.97$). The TVARSS(2) allowing the coefficients $b_i(t)$ ($i \geq 1$) to vary provided a statistically better fit than the TVARSS(2) model that treated these coefficients as fixed ($0.25\chi^2_0 + 0.5\chi^2_1 + 0.25\chi^2_2 = 12.98$, $P < 0.001$), implying that the dynamics of the time

series changed through time. Performing the same comparison for the pair of TVARSS(3) models similarly rejected the hypothesis that the coefficients b_i ($i \geq 1$) were constant ($0.125\chi^2_0 + 0.375\chi^2_1 + 0.375\chi^2_2 + 0.125\chi^2_3 = 13.97$, $P < 0.001$); we did not perform this comparison for the TVARSS(1) models, because their fits to the data were poor, presumably because they could not capture the cyclicity in the data. The time-varying estimate of the inverse of the characteristic root, λ , exceeded one as the Neimark-Sacker bifurcation was approached, suggesting loss of stationarity of the AR(2) process (Fig. 4). Although the 67% confidence intervals of the estimates of λ were fairly narrow, never did the lower bound exceed 1, and therefore never was it significantly greater than 1 even at the $\alpha = 0.33$ level.

In addition to the TVARSS(p) models, we also applied the TVVARSS(1) model using both prey and predator time series (Fig. 5). For the fitting, we assumed that both Σ_ε and Σ_α were diagonal.

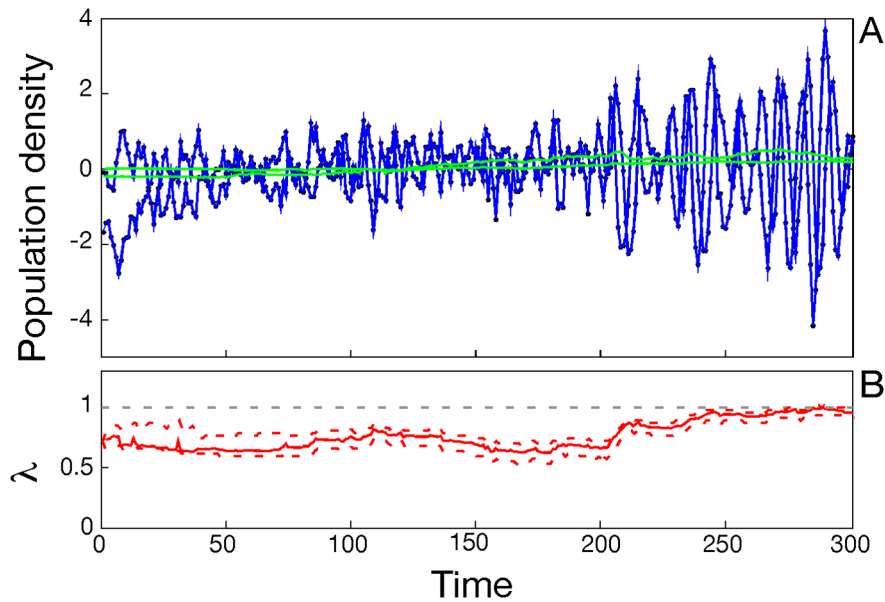


Fig. 5. Illustration of a TVVARSS(1) model fit to the two time series simulated from the Nicholson-Bailey predator-prey model (Eqs. 9 and 10). The carrying capacity of the host, K , increases through time to cause a Neimark-Sacker bifurcation (Fig. 1C). (A) Differences between the fitted trajectory (blue line) and the simulated data (black dots) are attributed in the model to measurement error. The green lines give the time-varying estimates of $b_{1,0}(t)$ and $b_{2,0}(t)$. (B) The inverse of the characteristic root for the TVVARSS(1) is computed from $b_{1,1}(t)$, $b_{1,2}(t)$, $b_{2,1}(t)$ and $b_{2,2}(t)$ (Box et al. 1994), which is plotted through time with its 67% confidence interval (dashed lines). Parameter estimates are: initial $b_{1,0} = 0.0132$, initial $b_{2,0} = -0.198$, initial $b_{1,1} = 0.467$, initial $b_{1,2} = 0.472$, initial $b_{2,1} = -0.0827$, initial $b_{2,2} = 0.880$, $\sigma_{\epsilon,1} = 0.388$, $\sigma_{\epsilon,2} = 0.314$, $\sigma_{\alpha} = 0.0862$, $\sigma_{1,0} = 0.0080$, $\sigma_{2,0} = 0.0080$, $\sigma_{1,1} = 0.0202$, $\sigma_{1,2} = 0.0160$, $\sigma_{2,1} = 0.0000$, and $\sigma_{2,2} = 0.0153$, and the log likelihood is -255.52 .

The model in which the autoregression coefficients $b_{i,j}(t)$ ($i, j \geq 1$) were allowed to vary fit the two time series much better than the model keeping these coefficients fixed ($\sum_{k=0}^4 \binom{4}{k} 2^{-4} \chi_k^2 = 57.7$, $P < 0.0001$). The statistical strength of this result was much greater than the case in which only the single time series for the prey is analyzed (Fig. 4); this is not surprising, because the predator data set added more information. As found for the analysis of only the prey time series, the inverse of the characteristic root λ increased through time, approaching 1 close to the Neimark-Sacker bifurcation, although the lower bound of the 67% confidence interval never exceeded 1. The estimate of λ was very similar between the TVARSS(2) and TVVARSS(1) models (compare Figs. 4B and 5B).

Finally, we fit the SETARSS(p) model to the data set simulated with the harvesting model of May (1977) in which the time series flickered between domains of attraction to the two

alternative states (Fig. 1B). The best-fitting model was SETARSS(1), with less-good fits for $p = 2$ ($\Delta\text{AIC} = 0.4$) and $p = 3$ ($\Delta\text{AIC} = 2.84$). Comparing the SETARSS(1) to a simple AR(1) fit in the same manner, the fit of the SETARSS(1) was significantly better ($\chi^2_3 = 24.2$, $P < 0.0001$) (Fig. 6). The tests of the same comparison were similarly significant for the SETARSS(2) ($\chi^2_4 = 26.0$, $P < 0.0001$) and SETARSS(3) ($\chi^2_5 = 20.2$, $P < 0.002$).

Potential advantages of locally linear state-space models

The approaches that we have presented using locally linear state-space models to anticipate regime shifts and to identify alternative stable states have potential advantages over existing metric-based methods. The existence and magnitude of these advantages will vary according to which methods are compared and what types of data are analyzed. We have not tested whether these advantages are realized in specific situa-

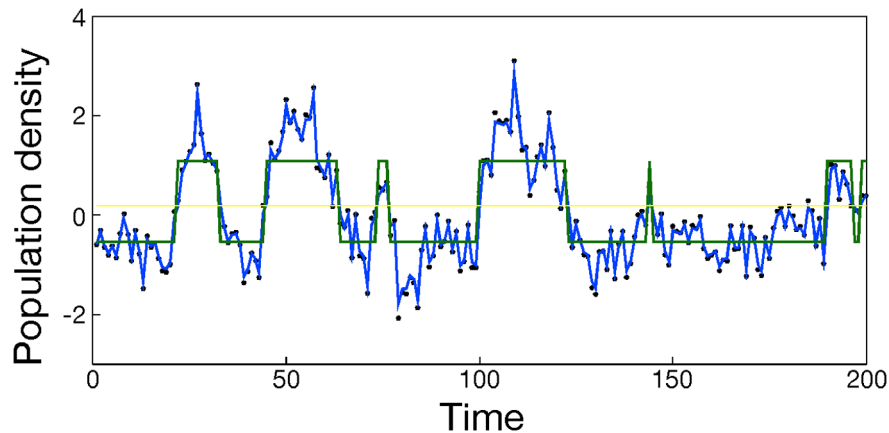


Fig. 6. Illustration of a SETARSS(1) model fit to a time series simulated from the harvesting model of May (1977) in which all parameters are fixed and stochasticity causes the population $x(t)$ to switch between domains of attraction to two stable points (Fig. 1B). (A) Differences between the fitted trajectory (blue line) and the simulated data (black dots) are attributed in the model to measurement error. The green lines give the estimates of $b_0(t)$ and $b_0'(t)$, and the yellow line gives the threshold c which separates the two AR(1) processes. Parameter estimates are: $b_0 = -0.537$, $b_0' = 1.09$, $b_1 = 0.444$, $b_1' = 0.706$, $c = 0.184$, $\sigma_\varepsilon = 0.469$, and $\sigma_x = 0.216$, and the log likelihood = -155.05 .

tions, but instead list them as speculations. Many of these potential advantages are shared with other model-based approaches (Clarke and Signorino 2010, Livina et al. 2010, Carpenter and Brock 2011). We strongly recommend that multiple methods be applied to any data set, and that simulations be performed to compare the potential advantages of the methods in the context of the data set.

A first potential advantage is that the approaches we have presented are model-based, with the analyses producing a structurally simple linear autoregressive model with statistically fitted parameters. Autoregressive models can avoid challenges and subjective decisions that are inherent in metric-based approaches, such as deciding whether and how to filter or detrend data, and choosing the best rolling window for calculating the metric (Dakos et al. 2008), as both filtering and rolling window size may have a strong influence on whether a metric predicts a regime shift (Lenton 2011, Dakos et al. 2012, Lenton et al. 2012). For time-varying AR(p) and VAR(1) models (Eqs. 2, 4), detrending is performed de facto by allowing b_0 to change through time (although simple detrending can also be used; see *Methods*). The rate at which b_0 changes, and hence the effective magnitude of detrending,

is given by the maximum likelihood estimates of the parameter σ_0^2 and therefore is determined as part of the overall fitting procedure. Thus, the approach gives statistical descriptions of the different properties of the time series (changes in the location b_0 , changes in the dynamics around b_0 given by b_i , $i \geq 1$). How these properties are used to infer regime shifts is then left up to the researcher. For example, in some models (e.g., Fig. 1A) the deterministic equilibrium of the process changes more rapidly as a regime shift is approached, and a researcher could decide to use this information to infer a regime shift. This carries the risk, however, that changes in the location are caused by some environmental trend that has no effect on the dynamics and therefore should not be used to infer a regime shift. Our approach has been to use only information about the dynamics around the (possibly temporally varying) mean of the process to infer a regime shift.

Second, our model-based approaches use maximum likelihood estimation, thereby admitting a full range of statistical inference, including hypothesis testing, confidence intervals on coefficients, and model selection. Statistical inference makes it possible to directly compare between alternative hypotheses, such as whether or not

the coefficients that govern dynamics in a an $AR(p)$ vary through time. For metric-based approaches, the only way to test whether trends in a metric are statistically significant is to perform simulations, comparing the change in the metric computed from the data with the distribution of changes in the metric computed from data simulated under the null hypothesis that the dynamics of the system do not change (i.e., non-bifurcating systems) (Dakos et al. 2012). The difficulty with this approach is that it is often not clear how to formulate an appropriate null model to be used to simulate data, and the chances of rejecting the null hypothesis will depend on decisions made about how the null model is constructed. For example, a natural null model might be a time-invariant $AR(1)$, yet this would be inappropriate for data with cyclic dynamics (e.g., Fig. 1C) for which an $AR(p)$ model would be more appropriate. Furthermore, the presence of measurement error would make an $ARMA(p,p)$ a better null model than an $AR(p)$ (Staudenmayer and Buonaccorsi 2005, Ives et al. 2010). In our model-based approach, determining the appropriate null model is more straightforward and statistically coherent, because we can compare the same model with and without time-varying coefficients, or with and without a threshold between alternative stable states.

Third, as a state-space model, measurement error can be included as either user-defined input or, as we have illustrated here, a parameter estimated in the model. Measurement error introduces short-term (non-autocorrelated) variation to time-series data, yet critical slowing down by definition involves autocorrelated variation. In effect, in our models measurement error is smoothed out of the time series during the fitting with the Kalman filter (Harvey 1989). Factoring out measurement error should better reveal the long-term patterns of variation, leading to fewer false negatives in the identification of impending regime shifts. While this in general is likely to be true, measurement error could have more complicated effects, especially if the magnitude of measurement error varied with the values in the time series or showed systematic changes through time; these complications would best be handled by, if possible, obtaining independent estimates of the measurement error and incorporating these directly into the model

(i.e., not estimating measurement error in the time-series analysis). Metric-based approaches do not account for measurement error, which is likely to lead to false negatives in the identification of early warning signals (Carpenter and Brock 2010).

Fourth, our approach should have good statistical power, because it incorporates the full information about the time series. This being said, the power of any statistical method will depend on the length and the underlying signal of a time series. As we illustrated (Fig. 3), power analyses can and should be conducted to assess the results of any analyses, and to compare the power of multiple approaches in the context of the data set in hand.

DISCUSSION

In this paper, we illustrated how locally linear state-space models can be used to measure changes in dynamics prior to regime shifts and to identify alternative stable states in time series. We showed that a rise in the dominant eigenvalue of a fitted time-varying autoregressive model, $TVARSS(p)$, can be used as an early warning signal in time series approaching a regime shift through a fold or a Neimark-Sacker bifurcation (Figs. 2, 4, 5). We also used a threshold autoregressive model, $SETARSS(p)$, to identify alternative states in a time series that flickers between alternative stable points (Fig. 6). These model-based approaches have several potential advantages over metric-based approaches that currently dominate analyses of time series for regime shifts and alternative stable states.

Although the approaches we presented are model-based, they are nonetheless simple, general models of time series. If sufficient information exists about a system, a potentially better strategy is to fit system-specific, nonlinear, biologically explicit models in order to investigate the possibility of alternative states (e.g., Ives et al. 2008, Schooler et al. 2011). This typically requires time series that are rich in dynamical patterns taken from systems that are biologically well-understood. In the absence of such requirements, or when performing a comparative analysis of the dynamics of many time series (e.g., Turchin and Taylor 1992, Murdoch et al. 2002, Ziebarth et al. 2010), simpler, non-system-

specific models are appropriate.

Although locally linear state-space models have advantages over existing metric-based approaches, they are not immune to problems that infect all approaches for detecting early warning signals (Scheffer et al. 2009, Hastings and Wysham 2010, Carpenter and Brock 2011). State-space models cannot distinguish between strong nonlinear changes in dynamics that lead to regime shifts from similarly strong nonlinear changes that do not. Furthermore, they will be difficult to apply to very short or very noisy time series. As there is no universal solution to the challenges of analyzing complex time series, multiple approaches should still be applied when trying to detect regime shifts.

Although we have presented locally linear state-space models as tools to identify regime shifts and alternative stable states, this is only a subset of the problems to which they can be applied. The time-varying models should be capable of identifying any type of change in the dynamical characteristics of time series, and the threshold models could be formulated to identify regions of state space that differ dynamically. Thus, they may prove useful additions to the general arsenal of tools currently available for time-series analyses.

ACKNOWLEDGMENTS

This work was prompted by the workshop “Practical Methods for Analysis of Early-Warnings for Regime Shifts” organized by M. Scheffer, S. R. Carpenter, and VD at the Santa Fe Institute, October 2011. Funding was provided in part by NSF-DEB-0816613 to ARI. VD was supported by a Spinoza (NWO) Award, a European Research Council Advanced Grant both awarded to M. Scheffer, and a Rubicon grant awarded to VD. We thank Carl Boettiger, John Drake, and an anonymous reviewer for very helpful suggestions.

LITERATURE CITED

- Akaike, H. 1973. Information theory and an extension of the maximum likelihood principle. Pages 267–281 in B. N. Petrov and F. Csaki, editors. Second International Symposium on Information Theory. Akademiai Kiado, Budapest, Hungary.
- Box, G. E. P., G. M. Jenkins, and G. C. Reinsel. 1994. Time series analysis: forecasting and control. Third edition. Prentice Hall, Englewood Cliffs, New Jersey, USA.
- Carpenter, S. and W. Brock. 2011. Early warnings of unknown nonlinear shifts: a nonparametric approach. *Ecology* 92:2196–2201.
- Carpenter, S. R. 2003. Regime shifts in lake ecosystems: patterns and variation. International Ecology Institute, Oldendorf/Luhe, Germany.
- Carpenter, S. R. and W. A. Brock. 2010. Early warnings of regime shifts in spatial dynamics using the discrete Fourier transform. *Ecosphere* 1:art10.
- Carpenter, S. R., J. J. Cole, M. L. Pace, R. Batt, W. A. Brock, T. Cline, J. Coloso, J. R. Hodgson, J. F. Kitchell, D. A. Seekell, L. Smith, and B. Weidel. 2011. Early warnings of regime shifts: a whole-ecosystem experiment. *Science* 332:1079–1082.
- Clarke, K. A. and C. S. Signorino. 2010. Discriminating methods: tests for non-nested discrete choice models. *Political Studies* 58:368–388.
- Dakos, V., M. Scheffer, E. H. van Nes, V. Brovkin, V. Petoukhov, and H. Held. 2008. Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences USA* 105:14308–14312.
- Dakos, V., E. H. van Nes, P. D’Odorico, and M. Scheffer. 2012. Robustness of variance and autocorrelation as indicators of critical slowing down. *Ecology* 93:264–271.
- Drake, J. M. and B. D. Griffen. 2010. Early warning signals of extinction in deteriorating environments. *Nature* 467:456–459.
- Guttal, V. and C. Jayaprakash. 2008. Changing skewness: an early warning signal of regime shifts in ecosystems. *Ecology Letters* 11:450–460.
- Harvey, A. C. 1989. Forecasting, structural time series models and the Kalman filter. Cambridge University Press, Cambridge, UK.
- Hastings, A. and D. B. Wysham. 2010. Regime shifts in ecological systems can occur with no warning. *Ecology Letters* 13:464–472.
- Held, H. and T. Kleinen. 2004. Detection of climate system bifurcations by degenerate fingerprinting. *Geophysical Research Letters* 31. doi: 10.1029/2004GL020972
- Ives, A. R., K. C. Abbott, and N. L. Ziebarth. 2010. Analysis of ecological time series with ARMA(p, q) models. *Ecology* 91:858–871.
- Ives, A. R., B. Dennis, K. L. Cottingham, and S. R. Carpenter. 2003. Estimating community stability and ecological interactions from time-series data. *Ecological Monographs* 73:301–330.
- Ives, A. R., A. Einarsson, V. A. A. Jansen, and A. Gardarsson. 2008. High-amplitude fluctuations and alternative dynamical states of midges in Lake Myvatn. *Nature* 452:84–87.
- Ives, A. R. and V. A. A. Jansen. 1998. Complex dynamics in stochastic tritrophic models. *Ecology* 79:1039–1052.

- Lade, S. J. and T. Gross. 2012. Early warning signals for critical transitions: A generalized modeling approach. *PLoS Computational Biology* 8. doi: 10.1371/journal.pcbi.1002360
- Lenton, T. M. 2011. Early warning of climate tipping points. *Nature Climate Change* 1:201–209.
- Lenton, T. M., V. N. Livina, V. Dakos, E. H. van Nes, and M. Scheffer. 2012. Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness. *Philosophical Transactions of the Royal Society a, Mathematical Physical and Engineering Sciences* 370:1185–1204.
- Livina, V. N., F. Kwasiok, and T. M. Lenton. 2010. Potential analysis reveals changing number of climate states during the last 60 kyr. *Climate of the Past* 6:77–82.
- May, R. M. 1974. *Stability and complexity in model ecosystems*. Second edition. Princeton University Press, Princeton, New Jersey, USA.
- May, R. M. 1977. Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature* 269:471–477.
- Murdoch, W. W., B. E. Kendall, R. M. Nisbet, C. J. Briggs, E. McCauley, and R. Bolser. 2002. Single-species models for many-species food webs. *Nature* 417:541–543.
- Nicholson, A. J. and V. A. Bailey. 1935. The balance of animal populations. *Proceedings of the Zoological Society of London* 3:551–598.
- Scheffer, M., J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. van Nes, M. Rietkerk, and G. Sugihara. 2009. Early-warning signals for critical transitions. *Nature* 461:53–59.
- Scheffer, M., S. Carpenter, J. A. Foley, C. Folke, and B. Walker. 2001. Catastrophic shifts in ecosystems. *Nature* 413:591–596.
- Schooler, S. S., B. Salau, M. H. Julien, and A. R. Ives. 2011. Alternative stable states explain unpredictable biological control of salvinia in Kakadu. *Nature* 470:86–89.
- Seekell, D. A., S. R. Carpenter, and M. L. Pace. 2011. Conditional heteroscedasticity as a leading indicator of ecological regime shifts. *American Naturalist* 178:442–451.
- Self, S. G. and K. Y. Liang. 1987. Asymptotic properties of maximum-likelihood estimators and likelihood ratio tests under non-standard conditions. *Journal of the American Statistical Association* 82:605–610.
- Staudenmayer, J. and J. R. Buonaccorsi. 2005. Measurement error in linear autoregressive models. *Journal of the American Statistical Association* 100:841–852.
- Stram, D. O. and J. W. Lee. 1994. Variance-components testing in the longitudinal mixed effects model. *Biometrics* 50:1171–1177.
- Strogatz, S. H. 1994. *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering*. Perseus Books, Cambridge, Massachusetts, USA.
- Tong, H. 1990. *Nonlinear time series analysis: a dynamical systems approach*. Oxford University Press, Oxford, UK.
- Turchin, P. and A. D. Taylor. 1992. Complex dynamics in ecological time series. *Ecology* 73:289–305.
- van Nes, E. H. and M. Scheffer. 2007. Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *American Naturalist* 169:738–747.
- Wissel, C. 1984. A universal law of the characteristic return time near thresholds. *Oecologia* 65:101–107.
- Ziebarth, N. L., K. C. Abbott, and A. R. Ives. 2010. Weak population regulation in ecological time series. *Ecology Letters* 13:21–31.

SUPPLEMENTAL MATERIAL

SUPPLEMENT

Matlab source code for performing TVARSS (p), TVVARSS (1), and SETARSS (p) model fitting (*Ecological Archives* C003-006-S1).