

Research



Cite this article: Bury TM, Bauch CT, Anand M. 2020 Detecting and distinguishing tipping points using spectral early warning signals. *J. R. Soc. Interface* **17**: 20200482. <http://dx.doi.org/10.1098/rsif.2020.0482>

Received: 17 June 2020
Accepted: 7 September 2020

Subject Category:
Life Sciences—Mathematics interface

Subject Areas:
computational biology, environmental science

Keywords:
critical transition, early warning signal, population dynamics, power spectrum, bifurcation

Author for correspondence:
T. M. Bury
e-mail: thomas.bury@mcgill.ca

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.5125994>.

Detecting and distinguishing tipping points using spectral early warning signals

T. M. Bury^{1,2}, C. T. Bauch¹ and M. Anand²

¹Department of Applied Mathematics, University of Waterloo, Waterloo, Ontario, Canada ON N2L 3G1

²School of Environmental Sciences, University of Guelph, Guelph, Ontario, Canada ON N1G 2W1

TMB, 0000-0003-1595-9444; CTB, 0000-0001-6214-6601

Theory and observation tell us that many complex systems exhibit tipping points—thresholds involving an abrupt and irreversible transition to a contrasting dynamical regime. Such events are commonly referred to as critical transitions. Current research seeks to develop early warning signals (EWS) of critical transitions that could help prevent undesirable events such as ecosystem collapse. However, conventional EWS do not indicate the type of transition, since they are based on the generic phenomena of critical slowing down. For instance, they may fail to distinguish the onset of oscillations (e.g. Hopf bifurcation) from a transition to a distant attractor (e.g. Fold bifurcation). Moreover, conventional EWS are less reliable in systems with density-dependent noise. Other EWS based on the power spectrum (spectral EWS) have been proposed, but they rely upon spectral reddening, which does not occur prior to critical transitions with an oscillatory component. Here, we use Ornstein–Uhlenbeck theory to derive analytic approximations for EWS prior to each type of local bifurcation, thereby creating new spectral EWS that provide greater sensitivity to transition proximity; higher robustness to density-dependent noise and bifurcation type; and clues to the type of approaching transition. We demonstrate the advantage of applying these spectral EWS in concert with conventional EWS using a population model, and show that they provide a characteristic signal prior to two different Hopf bifurcations in data from a predator–prey chemostat experiment. The ability to better infer and differentiate the nature of upcoming transitions in complex systems will help humanity manage critical transitions in the Anthropocene Era.

1. Introduction

The understanding that complex systems can possess thresholds marking a sudden shift to an alternative dynamical regime has been around for a long time (e.g. in ecology [1,2]). Such a threshold may be referred to directly as a tipping point/catastrophic bifurcation, or by its inferred dynamics, a critical transition/regime shift. Predicting tipping points and their ensuing dynamics remains a significant challenge, since the observable state of a system may show little change right up until it is too late. Even where data are abundant, parametrized models based on biological principles are rarely able to pinpoint tipping points due to uncertainty in system parameters and mechanisms. However, a new wave of research is targeting stochasticity as a possible treasure trove of information on the otherwise hidden, and often surprising, dynamics of complex systems [3,4].

A significant development in this area is that of early warning signals (EWS), which are a suite of statistical metrics that are expected to undergo observable change prior to a tipping point [5,6]. Most EWS are grounded in the phenomenon of ‘critical slowing down’, which is a generic feature of local bifurcations [7]. It involves the degradation of restoring forces along some dimension of the system’s state space, resulting in a longer return time to equilibrium following a perturbation. In stochastic systems, this manifests as an increase in variance [8], higher

correlations in time [9] and space [10], lower frequencies in the power spectrum [11], and notable changes in several other statistical metrics [4]. We use the term ‘conventional EWS’ to refer to these EWS that serve as proxies for critical slowing down.

The generality of critical slowing down is both a blessing and a curse. On one hand, it allows EWS to be applied to a wide range of systems including socio-ecological [12,13], neurological [14], financial [15] and climate [16] systems. On the other hand, bifurcations come in assorted forms which each possess their own unique dynamics [17], such as a smooth transition to an intersecting state, the onset of oscillations, or an abrupt departure to a distant attractor (e.g. Transcritical, Hopf, Fold bifurcations, respectively). These local bifurcations are all accompanied by critical slowing down [7] and so cannot be distinguished by the conventional critical slowing down methodology [18]. As such, EWS that are specific to each type of bifurcation are required to predict ensuing dynamics. Moreover, external noise that is correlated or density dependent can distort EWS [19,20], so there is a need to develop metrics that are more robust to these forms of noise.

The power spectrum has shown potential as a tool for detecting bifurcations from time-series data [11,21–24] and has a rich history in time-series analysis more generally [25,26]. A seminal study in the physics literature derived analytical approximations for its behaviour prior to bifurcations of periodic orbits [21], showing how different patterns emerge depending on the type of bifurcation. Later studies investigated how changes in the power spectrum could warn of an upcoming Fold bifurcation [11], and constructed EWS to pick up changes in spectral properties such as a spectral ratio [22] and the spectral exponent [23]. We refer to EWS based on the power spectrum collectively as ‘spectral EWS’. Current spectral EWS are based on the phenomenon of spectral reddening—the movement towards lower frequencies in the power spectrum as a bifurcation approaches—and have shown promise in providing warning of upcoming Fold bifurcations. However, if the bifurcation has an underlying oscillatory component (e.g. Hopf/Flip/Neimark–Sacker), which we argue is more likely in high-dimensional systems than simplified models would suggest (electronic supplementary material, Note), then spectral reddening does not occur, and these spectral EWS cannot be expected to provide a signal. Given that the power spectrum is a complete representation of a (stationary) time series, it possesses a lot of information. The question then becomes: how does one harness this information to provide maximum information about the underlying system?

For assessing and developing EWS, the theory of stochastic processes has much to offer [3,27–29], though most studies do not go beyond the rule of thumb provided by critical slowing down. This is well illustrated by the general framework of a system of variables s that evolves in time according to

$$ds = f(s) dt + \sigma dW(t), \quad (1.1)$$

where f captures the within-system dynamics and $dW(t)$ is a vector of Wiener processes representing environmental noise with amplitudes and correlations given in the matrix σ . For relatively small noise, the dynamics about an equilibrium state are well approximated by

$$dx = Ax dt + \sigma dW(t), \quad (1.2)$$

where x is the deviation of the state from equilibrium and A is the Jacobian matrix of f at equilibrium, which contains the

local interaction terms between the variables. This process, originally studied in physics to model Brownian motion [30], is an Ornstein–Uhlenbeck process, for which general statistical properties can be derived [19,27]. Therefore, given a system that fits into this framework, analytical approximations for EWS can be derived in terms of system parameters (within A) and relative noise strengths and correlations (within σ). This way, one can move beyond generic indicators of an upcoming transition (which do not always behave as expected), towards more reliable indicators that are specific to the system being modelled.

This analytical approach was recently adopted in previous studies to investigate the behaviour of EWS specific to particular models [31,32], and different regimes of noise [20,28]. It has also been used to derive EWS approximations for a subset of local bifurcations in continuous-time [28] and analytical approximations of the power spectrum prior to continuous-time bifurcations of period orbits [21]. However, a more complete description of EWS in continuous-time is required to understand their behaviour prior to oscillatory bifurcations. This description should also include discrete-time systems, since a corresponding discrete-time model can exhibit fundamentally different dynamics (e.g. the logistic model for population growth exhibits steady-state dynamics in continuous-time, but regimes of oscillatory and chaotic dynamics in discrete-time [33]).

Here, we build on previous analytical work to derive EWS approximations for every local codimension-1 bifurcation in discrete and continuous-time systems (table 1). This provides a more complete framework for which EWS to expect preceding each type of bifurcation. We then develop two spectral EWS that are motivated by the insights from the analytical approximations. The first metric scales with the distance to the bifurcation in a favourable manner compared to conventional EWS, therefore providing a signal that is more likely to be detected. It is also more robust to density-dependent noise, and works for both oscillatory and non-oscillatory bifurcations, to which current spectral EWS are not suited. The second metric determines the class of bifurcations to which the upcoming instability belongs, allowing one to distinguish between certain types of transition. Finally, we apply the spectral EWS to model and empirical data, to demonstrate their ability to provide characteristic signals of different bifurcations.

2. Results

2.1. Insights from analytical approximations

The mathematical forms provided in table 1 reveal characteristic features of EWS that can be used to distinguish certain types of bifurcation. For example, the Fold and Hopf bifurcations are preceded by very different autocorrelation and power spectra (figure 1). The behaviour of these EWS preceding the other local codimension-1 bifurcations is shown graphically in electronic supplementary material, figure S1.

Consider lag- τ autocorrelation, which computes the correlation between data points spaced τ time units apart. The mathematical forms imply that bifurcations without an oscillatory component (Fold/Transcritical/Pitchfork) yield increasing autocorrelation for all lag times, conforming to the expected behaviour of critical slowing down (figure 1*b*). Choosing the most suitable lag time is not trivial however: smaller lag times can diminish the signal since nearby points are highly correlated even far from the bifurcation, and larger lag times

Table 1. Analytical approximations for EWS preceding each local, codimension-1 bifurcation. Approximations are for the normal form of each bifurcation [17] with additive white noise of amplitude σ . The asymptotic behaviour of the peak in the power spectrum is most easily seen for the Fold bifurcation in continuous-time, where setting $\omega = 0$ shows S_{\max} scales like $1/\lambda^2$, whereas variance scales at the slower pace of $1/\lambda$. Derivations use standard techniques from stochastic process theory [27], making the assumption of small noise (such that nonlinear terms are negligible) and quasi-stationarity (the bifurcation parameter varies sufficiently slowly). The dominant eigenvalue(s) is that of the Jacobian matrix about equilibrium. Shorthand notation includes TC, Transcritical; PF, Pitchfork, NS, Neimark–Sacker. Note that PF, Hopf, Flip and NS can be both super- and sub-critical. Expressions for the Fold, TC and PF in continuous-time are reported in [28]. Expressions for the power spectrum in continuous-time are reported in [21]. Derivations of expressions and asymptotic properties are provided in electronic supplementary material, Methods.

	bifurcation	dominant eigenvalue(s)	variance	lag- τ AC, $\rho(\tau)$	power spectrum, $S(\omega)$
continuous-time $t, \tau \in \mathbb{R}$	Fold, TC, PF	$\lambda \in \mathbb{R}, \lambda \rightarrow 0^-$	$-\frac{\sigma^2}{2\lambda}$	$e^{\lambda \tau }$	$\frac{\sigma^2}{2\pi} \left(\frac{1}{\omega^2 + \lambda^2} \right)$
	Hopf	$\lambda_{1,2} = \mu \pm i\omega_0, \mu \rightarrow 0^-$	$-\frac{\sigma^2}{2\mu}$	$e^{\mu \tau } \cos \omega_0 \tau$	$\frac{\sigma^2}{4\pi} \left(\frac{1}{(\omega - \omega_0)^2 + \mu^2} + \frac{1}{(\omega + \omega_0)^2 + \mu^2} \right)$
discrete-time $t, \tau \in \mathbb{Z}$	Fold, TC, PF	$\lambda \in \mathbb{R}, \lambda \rightarrow 1^-$	$\frac{\sigma^2}{1-\lambda^2}$	$\lambda^{ \tau }$	$\frac{\sigma^2}{2\pi} \left(\frac{1}{1+\lambda^2-2\lambda \cos(\omega)} \right)$
	Flip	$\lambda \in \mathbb{R}, \lambda \rightarrow -1^+$	$\frac{\sigma^2}{1-\lambda^2}$	$\lambda^{ \tau }$	$\frac{\sigma^2}{2\pi} \left(\frac{1}{1+\lambda^2-2\lambda \cos(\omega)} \right)$
	Neimark–Sacker	$\lambda_{1,2} = r e^{\pm i\theta}, r \rightarrow 1^-$	$\frac{\sigma^2}{1-r^2}$	$r^{ \tau } \cos(\theta\tau)$	$\frac{\sigma^2}{4\pi} \left(\frac{1}{1+r^2-2r \cos(\omega-\theta)} + \frac{1}{1+r^2-2r \cos(\omega+\theta)} \right)$

can provide a delayed signal. By contrast, bifurcations with an oscillatory component (Hopf/Flip/Neimark–Sacker) yield an increasing or decreasing trend, that depends on the relationship between the lag time τ , and the underlying period of oscillations T (figure 1e). At sufficiently low lag times ($\tau < T/4$), autocorrelation increases as the bifurcation is approached. However, for lag times near to half the period of oscillations, the autocorrelation decreases. This can be understood intuitively by noting that at a lag time of $T/2$, one is computing the correlation between peaks and troughs of the underlying frequency, which become more pronounced as the bifurcation is approached. Since these points occur on opposite sides of the trajectory mean, they possess negative correlation. This finding corroborates previous studies that have found decreasing autocorrelation in both empirical [34] and model [35] studies preceding oscillatory bifurcations. Studies that have found increasing autocorrelation prior to Hopf bifurcations [18,36] have used lag times much smaller than the period of oscillations. The analytical approximation for autocorrelation explains these previously contradictory outcomes. It also suggests that autocorrelation should be computed at multiple lag times, not just at lag-1, which is common practice in EWS studies.

The power spectrum provides perhaps the most intuitive picture of the approaching bifurcation. Prior to non-oscillatory bifurcations it moves to lower frequencies, which is a manifestation of critical slowing down, and coined ‘spectral reddening’ [11]. Prior to oscillatory bifurcations it moves to the frequency of oscillations that occur at the bifurcation. The fact that this frequency is observed in the power spectrum prior to the bifurcation is due to resonant amplification of stochasticity [37–39], which occurs in systems possessing damped oscillations subject to environmental or demographic noise.

2.2. Spectral early warning signals

We use the analytical expressions for the power spectrum to construct metrics (spectral EWS) that capture the important features that relate to the bifurcation. First, we use the peak

in the power spectrum (S_{\max}) as an indicator of bifurcation proximity. This metric increases asymptotically like σ^2/μ^2 , where σ is the external noise amplitude, and μ is the distance to the bifurcation. By contrast, variance increases like σ^2/μ , which is a slower rate of increase as the bifurcation is approached (halving the distance to the bifurcation, doubles the variance, whereas S_{\max} increases fourfold). Moreover, this scaling allows the metric to be more robust to changes in σ that may occur in systems with multiplicative or time-varying external noise. Second, we use AIC weights [40]—a metric that determines the relative parsimony of a set of models fitted to a dataset—to determine which bifurcation the measured power spectrum corresponds to. The ‘models’ here are the analytical forms for the power spectra preceding each bifurcation, and a flat power spectrum to serve as a null model (white noise). These spectral EWS should be used together as they provide complementary information about an upcoming transition. S_{\max} warns of an upcoming bifurcation, and the AIC weights provide information on the type of transition to expect by indicating which analytical form most parsimoniously fits the power spectra.

2.3. Application to a model system

We test the spectral EWS on the classical Ricker model of a logistically growing population subject to harvesting (Methods). This model exhibits different bifurcations depending on the parameter that varies [35]. An increase in harvesting rate yields a Fold bifurcation to a diminished population state, whereas increasing the intrinsic growth rate of the population (which can occur through size-selective harvesting [41]) yields a Flip (period-doubling) bifurcation to an oscillatory regime. Whereas the Fold bifurcation results in a critical transition to an alternative state, the Flip bifurcation results in a smooth, and therefore reversible transition. Given these contrasting outcomes, it is important to be able to distinguish the upcoming bifurcation in advance. The EWS are computed on bootstrapped samples from segments of the time series within a

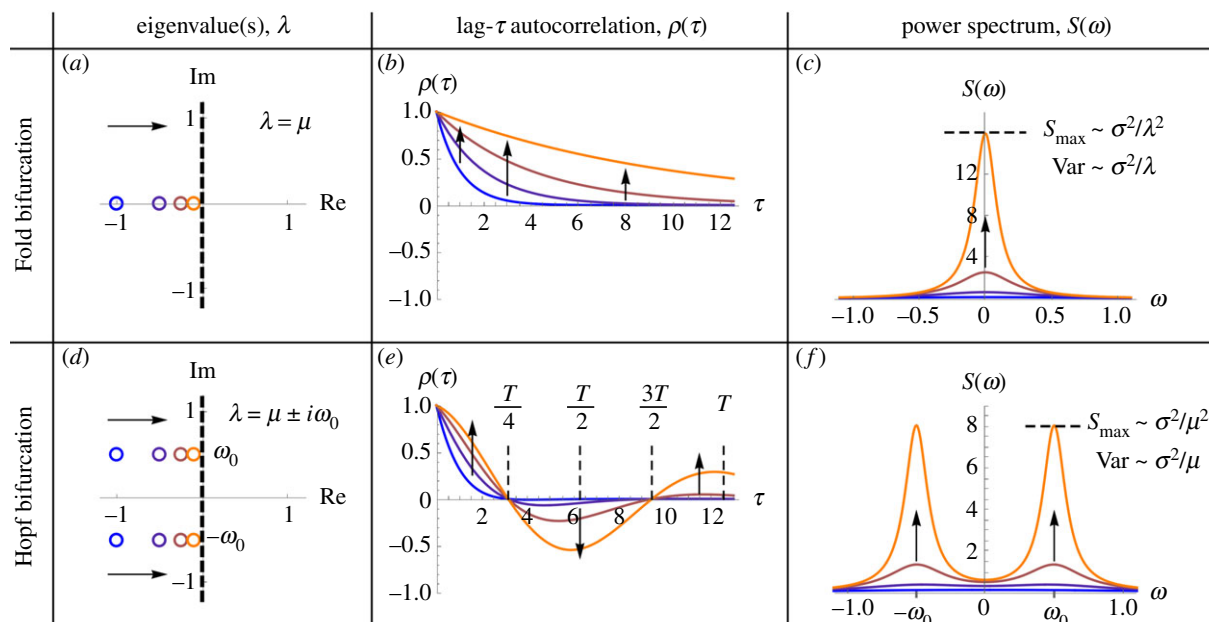


Figure 1. Contrasting analytical EWS preceding the Fold and the Hopf bifurcation. Analytical approximations for the autocorrelation function and power spectrum (table 1) are plotted at various distances to the bifurcation given by $\mu = \{-1, -0.5, -0.25, -0.1\}$, where μ is the real part of the system's dominant eigenvalue. (a) The Fold bifurcation involves a single real eigenvalue becoming positive. (b) Autocorrelation increases at all lag times, though very small lag times give a less noticeable increase, and large lag times yield a delayed increase. (c) The power spectrum becomes dominated by lower frequencies, with a peak amplitude that increases asymptotically faster than the variance (as $\mu \rightarrow 0$). (d) A Hopf bifurcation involves a complex-conjugate pair of eigenvalues obtaining positive real part. The imaginary part of the eigenvalues (ω_0) corresponds to the frequency of oscillations that occur at the bifurcation. (e) The trend of autocorrelation depends on how the lag time compares with the underlying period of oscillations ($T = 2\pi/\omega_0$). For lag times near to half the period of oscillations, one in fact observes decreasing autocorrelation. (f) The power spectrum becomes dominated by ω_0 , with the peak amplitude increasing favourably compared to variance. Parameter values are $\omega_0 = 0.5$, $\sigma = 1$. Similar figures for each type of local codimension-1 bifurcation are provided in electronic supplementary material, figure S1.

rolling window (Methods), which provides the error bars in the EWS metrics. We emphasize that these error bars are not from multiple simulations of the system, but from the single trajectory as shown. We do this since, in reality, one usually only has a single realization to work with. To verify their trends over multiple simulations, we use Kendall tau values, which serve as a measure of increasing or decreasing trend. EWS from multiple realizations of this model are provided in electronic supplementary material, figures S2–S3.

In the scenario undergoing the Fold bifurcation (figure 2), the trajectory shows a strong increasing trend in lag-1 autocorrelation, but no discernible trend in variance. The protocol of requiring increasing variance and lag-1 autocorrelation to signal a transition therefore fails in this case. This is an example of how an increase in variance does not necessarily precede transitions in systems with multiplicative noise [19]. However, since S_{\max} is a more sensitive metric to changes in bifurcation proximity than variance, it still provides a signal prior to the bifurcation. The AIC weights correctly favour the power spectrum of the Fold bifurcation, and do so by the time $t = 300$ for 98 of the 100 realizations.

The scenario undergoing the Flip bifurcation (figure 3) looks somewhat similar prior to the transition with regards to the population trend. The spectral EWS, however, tell a very different story, in agreement with the analytical approximations. Note that a Flip bifurcation from equilibrium yields oscillations of period $T = 2$ and so one expects an underlying frequency of $\omega = 2\pi/T = \pi$ prior to the transition, as observed in the power spectrum. Correspondingly, autocorrelation decreases at lag-1 and increases at lag-2 (seen analytically in figure 1e). S_{\max} shows a marked increase (stronger than that of variance), and the AIC weights favour the power spectrum

of a Hopf bifurcation, indicative of upcoming oscillations. The predicted frequency of these oscillations can be read off from the power spectrum as $\omega_0 = \pi$, in line with that of a Flip bifurcation. The spectral EWS together therefore provide a characteristic EWS of the Flip bifurcation that can be distinguished from the approaching Fold bifurcation.

To assess the predictive power of the EWS, we compute receiver operator characteristics for forced and null simulations for each bifurcation (figure 4). We find that prior to both bifurcations, S_{\max} outperforms variance as a warning signal (AUC = 0.83 versus 0.53 for the Fold bifurcation and AUC = 0.98 versus 0.96 for the Flip bifurcation). Lag-1 autocorrelation provides the strongest prediction; however, it requires prior knowledge as to whether an increase or decrease is expected, which depends on the bifurcation. In this assessment, we take decreasing lag-1 autocorrelation as indication of the transition in the Flip scenario. To assess the efficacy of the EWS at alternative parameter values, we compute EWS over all combinations of the bifurcation parameters and find that the spectral EWS distinguish the transition in each case (electronic supplementary material, figures S4–S5).

2.4. Application to empirical data

To test the spectral EWS in an empirical setting, we use chemostat data from a predator–prey experiment conducted in a previous study [42] (electronic supplementary material, figure S6). The authors showed using both experiments and a parametrized model that the system exhibits two Hopf bifurcations (H1, H2) as the dilution rate (controlling nutrient uptake) is varied (figure 5a). The observed Hopf bifurcations occur at slightly higher dilution rates than the model predicts

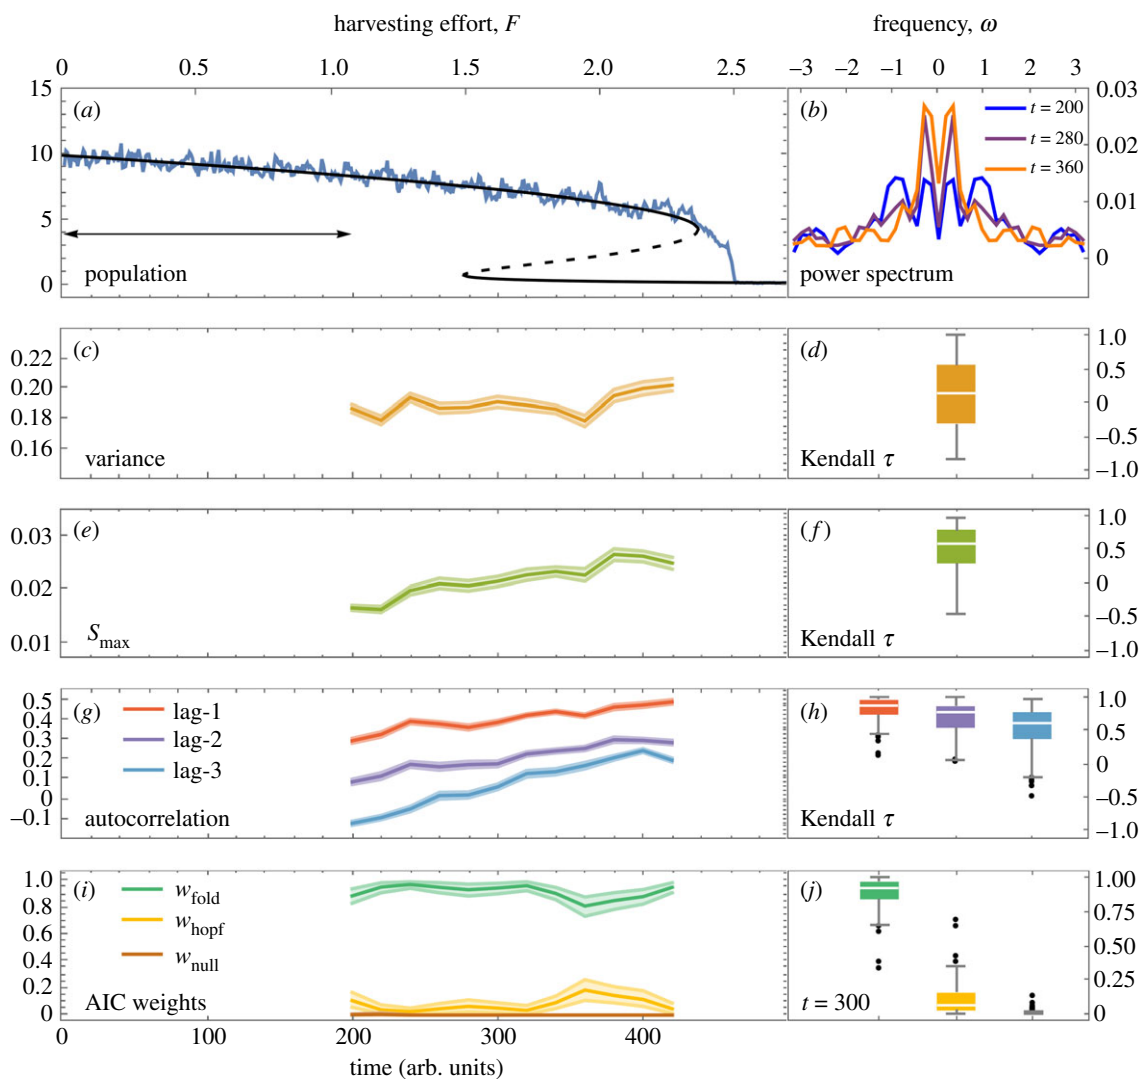


Figure 2. Characteristic EWS preceding the Fold bifurcation in the Ricker model. (a) Simulated response of population abundance (blue) to increasing harvesting effort superimposed on bifurcation diagram (black) of the deterministic model (solid, stable state; dashed, unstable state). Arrow illustrates the rolling window (40% of the time series) used for computing EWS. (b) Power spectrum becomes dominated by lower frequencies over time as expected for an approaching Fold bifurcation [11]. (c) Variance provides no useful signal due to density-dependent noise [19]. (d) Kendall tau values over 100 realizations show no consistent trend in variance. (e) S_{\max} provides an increasing trend (f) for the majority of realizations. (g) Autocorrelation provides an increasing trend, (h) with higher consistency at lower lag times. (i) AIC weights favour the Fold bifurcation. (j) At time $t = 300$, w_{fold} is dominant for 98 of 100 realizations. EWS are displayed as means with 95% confidence intervals over 100 bootstrapped samples from the time series in (a), after removing the trend with a Lowess filter. Kendall tau values are shown using box-whisker diagrams, where boxes mark the median and span the interquartile range and whiskers show the full range excluding outliers (black dots).

[42], highlighting the difficulty of locating bifurcations in real systems, and the importance of having additional predictive tools such as EWS. We compute EWS for each experimental time series and find good agreement with theoretical expectations (figure 5*b–d*). Preceding both Hopf bifurcations, the spectral EWS provide a characteristic indication of a Hopf bifurcation— S_{\max} shows an upward trend, and w_{Hopf} is the significant AIC weight. Unlike for H2, variance and S_{\max} do fluctuate prior to H1, which is likely due to the data points being more tightly spaced in the region further from the bifurcation. In addition, the power spectrum in the pre-bifurcation regime of H1 ($\delta = 0.04$) shows a dominant frequency, $\omega_0 \approx 1/3$ (electronic supplementary material, figure S8) corresponding to an underlying period of oscillations, $T = 2\pi(1/3)^{-1} \approx 19$ days, which approximates the period observed in the oscillatory regime. This suggests that in addition to determining the type of bifurcation, the power spectrum can provide an early estimate of the period of oscillations at the bifurcation.

3. Discussion

The ability to not only detect, but characterize bifurcations is essential for obtaining knowledge of upcoming qualitative changes in a system's dynamics. In this paper, we have derived analytical expressions for EWS prior to each type of local bifurcation, enabling us to construct spectral EWS that distinguish oscillatory from non-oscillatory bifurcations and provide a more sensitive warning of changes in bifurcation proximity. We demonstrated these tools with a well-known population model, showing how the onset of collapse can be distinguished from the onset of oscillations, and applied them to empirical data of a population exhibiting a Hopf bifurcation. This paper shows that spectral EWS offer complementary information to conventional EWS and should be added to the repertoire of tools for predicting tipping points in real systems.

There exist other recent developments in tipping point detection that, under appropriate circumstances, show potential to predict specific bifurcations. Eigenvalue reconstruction

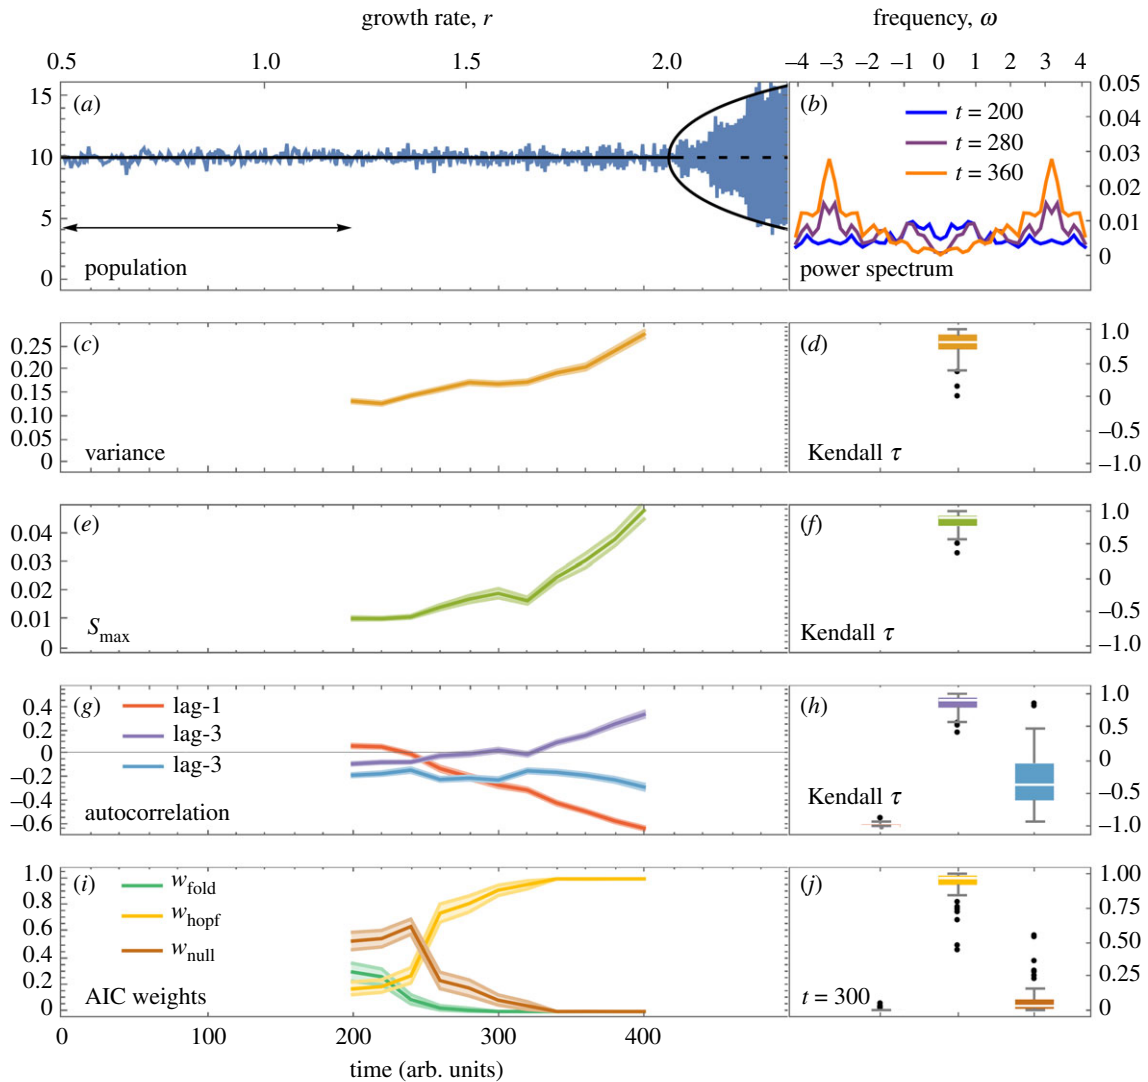


Figure 3. Characteristic EWS preceding the Flip bifurcation in the Ricker model. (a) Simulated response of population abundance (blue) to increasing growth rate superimposed on bifurcation diagram (black) of the deterministic model (solid, stable state/limit cycle; dashed, unstable state). Arrow illustrates the rolling window (40% of the time series) used for computing EWS. (b) Power spectrum becomes dominated by the frequency $\omega = \pi$, in agreement with analytical approximations prior to a Flip bifurcation (table 1). (c) Variance provides an increasing trend, and (d) is confirmed to have an increasing trend over an ensemble of 100 realizations. (e) S_{\max} yields an early warning signal with (f) a stronger increasing trend than variance. Note that while variance increases twofold, S_{\max} increases fourfold, consistent with the asymptotic relationships of these metrics with the bifurcation parameter (figure 1). (g) Autocorrelation trend depends on the relationship between lag time and underlying period of oscillations ($T = 2$). Lag-1 autocorrelation decreases, since in this case the lag is half of the period (see figure 1). (h) Higher lags yield less pronounced signals on average. (i) AIC weights favour the power spectrum of a Hopf bifurcation, indicating upcoming oscillations. (j) At time $t = 300$, w_{hopf} dominates for 98 of 100 realizations. Computation details are as in figure 2.

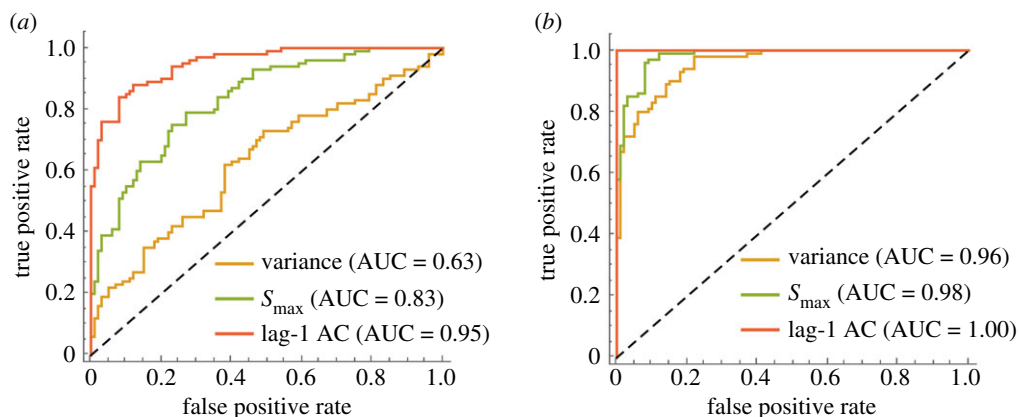


Figure 4. Receiver operator characteristics for EWS prior to the Fold (a) and Flip (b) bifurcation in the Ricker model. Area under the curve (AUC) is a measure of performance with values closer to one being better predictors of the transition. Decreasing lag-1 autocorrelation is used to predict the Flip bifurcation.

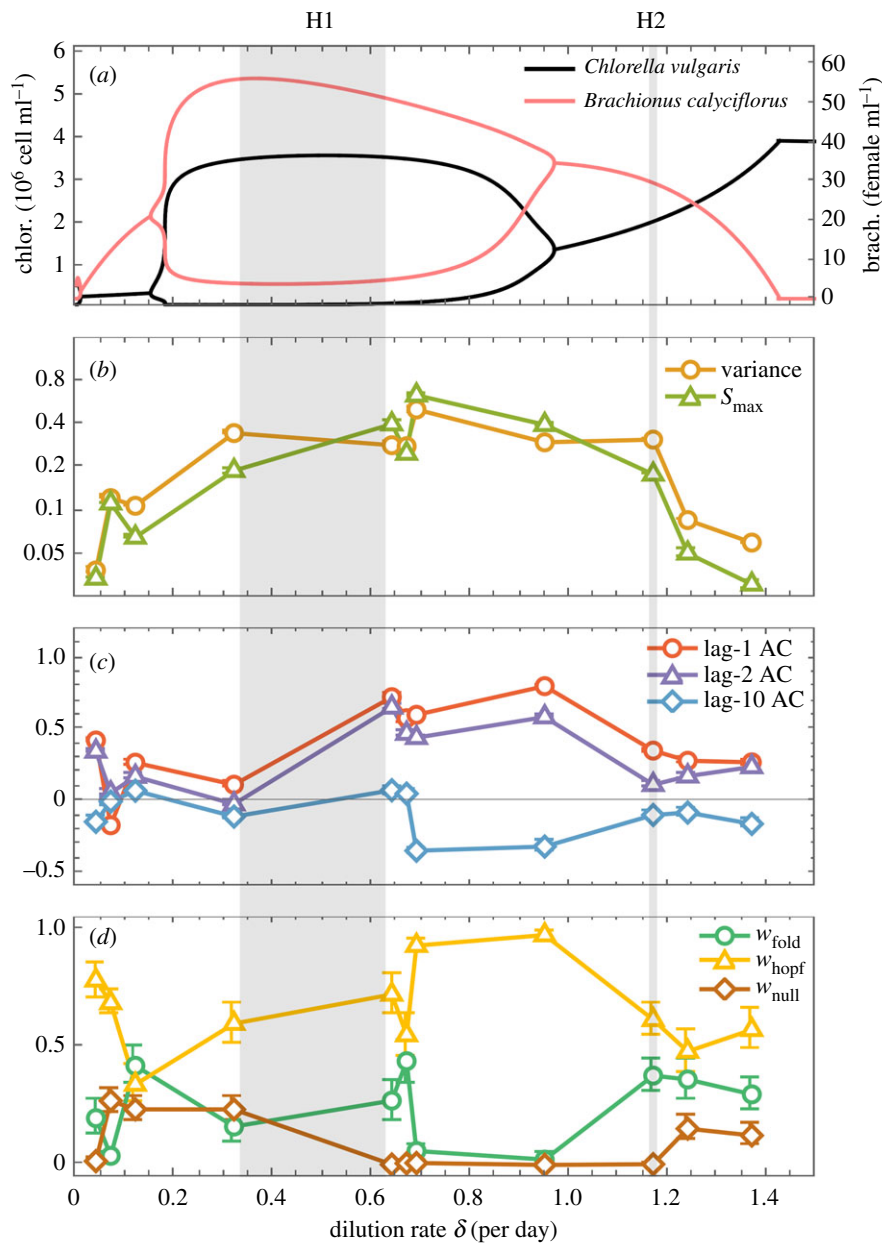


Figure 5. Characteristic early warning signals in an empirical predator–prey system. (a) Bifurcation diagram of a parametrized model [42] for the predator–prey dynamics between *Brachionus calyciflorus* and *Chlorella vulgaris*. Lines show stable states/limit cycles, indicating two Hopf bifurcations. Hopf bifurcations in the experiments (H1, H2) actually occur somewhere within the grey regions, which separate the oscillatory from equilibrium behaviour observed in the experiments [42]. (b) Variance and S_{\max} as a function of dilution rate in experiments. Increasing trends are observed in S_{\max} prior to the Hopf bifurcations. (c) Autocorrelation at low (1,2) and high (10) lag times provides no clear early warning signal. (d) The Hopf AIC weight dominates in the vicinity of the bifurcation, thus paired with S_{\max} , provides a characteristic early warning signal of an oscillatory bifurcation. Early warning signals are computed for 100 bootstrapped samples of each chemostat time series (fixed dilution rate), and displayed as means with 95% confidence intervals, across the samples and across the two species. Spectral EWS are also present in the time series of each species considered separately (electronic supplementary material, figure S7).

[43] allows one to obtain an approximation of the Jacobian matrix from time-series data. This method permits the monitoring of stability loss via a diminishing eigenvalue, and the underlying frequency via the eigenvalue's imaginary part (figure 1a,d). This approach, however, requires time-series data from multiple variables in order to make assertions on the type of bifurcation, whereas the power spectrum can be computed from an individual time series. Another approach considers the scaling of critical slowing down as a bifurcation is approached [14]; however, this requires a controlled setting whereby changes in the bifurcation parameter can be measured. Finally, promising studies in large perturbation theory [44,45] demonstrate how information on the bifurcation type and distance can be obtained by carefully monitoring a

system's recovery trajectory following a perturbation. This again requires a controlled environment, and a system whereby large perturbations are feasible.

This study uses ecological data from controlled chemostat experiments, where environmental conditions are kept fixed. Similar to other studies on EWS in experimental populations [46,47], this allows us to obtain estimates for the EWS at different values of the bifurcation parameter without requiring a rolling window. Our empirical results show that spectral EWS behave according to theoretical predictions in an empirical system. For real ecosystems however, environmental conditions vary over time, more similar to our model scenarios. Future studies should therefore test spectral EWS on high-resolution data of natural systems subject to time-varying environmental

conditions. Unfortunately, many ecological datasets are too sparse for EWS to be effective [48]. However, technological advances in measurement devices are facilitating high-frequency data collection across many scientific disciplines. In ecology, sondes that are deployed on lakes to measure chlorophyll concentration are able to transmit data every minute, helping to anticipate algal blooms [49]. In medicine, wearable devices can now provide continuous physiological data to help detect disease transitions [50]. We anticipate that natural systems for which high-frequency data are available, will be those to most benefit from spectral EWS and EWS more generally.

Aside from detecting transitions, spectral EWS could contribute to a better understanding of systems where competing models exist. For example, the transition from quiescence to the spiking of neurons in the mammalian cortex has been modelled as both a Fold and Hopf bifurcation, depending on the underlying model assumptions [14]. Computing spectral EWS of empirical data, one could discern which bifurcation the transition corresponds to, providing stronger validation for one model over the other and therefore learning something about the underlying mechanisms of the system.

The findings of this study have to be seen in light of some limitations, which should be addressed in future studies. First, the spectral EWS proposed here only distinguish between oscillatory and non-oscillatory bifurcations, and not between bifurcations within these groups. For example, a supercritical and subcritical Hopf bifurcation, while giving rise to the same spectral EWS, result in structurally different transitions. Whereas the supercritical Hopf yields a smooth, reversible transition to small oscillations, the subcritical Hopf yields a critical, irreversible transition to a distant limit cycle. Distinguishing between such bifurcations requires consideration of nonlinear terms, which this study has assumed to be negligible—a valid assumption for systems subject to small noise relative to the width of the basin of attraction. Future studies should investigate how nonlinear terms alter the power spectrum prior to a bifurcation.

Second, this study only considers transitions arising due to local codimension-1 bifurcations. However, complex systems may lose stability via global bifurcations, which are not accompanied by critical slowing down, resulting in the failure of corresponding EWS [51,52]. In such situations, EWS based on other dynamical features are required, such as critical attractor growth, which has been shown to precede interior crises of excitable systems [53]. Given that the spectral EWS here are derived from changes in local dynamics, they cannot be expected to provide warning of global bifurcations. In addition, bifurcations of a higher codimension (number of parameters that must be varied for the bifurcation to occur), such as the Bautin bifurcation, remain little explored in the EWS literature yet arise in many models (e.g. [54]).

Third, spectral EWS, as with conventional EWS, may be distorted in cases of large noise, multiple scales of correlation and sparsely sampled data [48]. Further research should investigate the robustness of spectral EWS to these factors, and how the power spectrum varies in each case. For example, one would expect correlated environmental noise to create an additional peak in the power spectrum at the characteristic frequency of the noise. Finally, the work here does not consider spatial systems, where a power spectrum can be computed in both space and time. Investigating the

behaviour of the power spectrum prior to different bifurcations in these systems may inspire further spectral EWS to predict and characterize transitions in spatial systems.

The spectral EWS developed here have several implications for science and policy. Competing hypotheses can lead to models with different types of bifurcations [14,55]. Monitoring time-series data with spectral EWS can be used to infer particular models based on their bifurcation structure, and therefore deepen our understanding of the system itself. To avoid a regime shift, action must be taken well in advance of the bifurcation [22], requiring indicators that respond quickly to changes in bifurcation proximity. Incorporating S_{\max} may improve the likelihood of a sufficiently early warning due to its high sensitivity to bifurcation proximity. Moreover, the ability to distinguish oscillatory bifurcations (Hopf/Flip/Neimark–Sacker) from non-oscillatory ones is crucial, since their inferred dynamics produce contrasting outcomes. Our work here furthers methodology to learn useful information from stochasticity [3] and offers ready-to-go tools for further application.

4. Material and methods

4.1. Analytical derivations of early warning signals

Derivations use the normal form of each local bifurcation [17] with additive white noise. For example, derivations for the Fold bifurcation in continuous-time come from the system

$$\frac{dx}{dt} = \alpha - x^2 + \sigma\xi(t), \quad (4.1)$$

where α is the bifurcation parameter, $\xi(t)$ is a white noise process and σ is the noise amplitude. The corresponding normal form system in discrete-time is

$$x_{t+1} = x_t + \alpha - x_t^2 + \sigma\epsilon_t, \quad (4.2)$$

where now ϵ_t is a normal random variable of mean zero and unit variance. For each bifurcation, we linearize the system, yielding an Ornstein–Uhlenbeck process (in continuous-time) or a vector autoregression (in discrete-time). EWS for the corresponding stationary process are then derived from first principles using standard techniques from stochastic processes [25,27] (electronic supplementary material, Methods).

4.2. Population model

We use a Ricker-type model that describes the logistic growth of a population subject to harvesting [35]. The model reads

$$N_{t+1} = N_t e^{(r(1-N_t/K)+\sigma\epsilon_t)} - F \frac{N_t^2}{N_t^2 + h^2}, \quad (4.3)$$

where N_t is the population size at time t , r is the intrinsic growth rate, K is the carrying capacity, F is the maximum rate of harvesting, h is a half-saturation constant, σ is the noise amplitude and ϵ_t is a normal random variable with zero mean and unit variance. Baseline parameters are $r = 0.75$, $K = 10$, $F = 0$, $h = 0.75$, $\sigma = 0.04$. The model exhibits a Fold bifurcation at $F = 2.36$, and a Flip (period-doubling) bifurcation at $r = 2.00$ followed by a sequence of further Flip bifurcations to chaos.

We simulate two model scenarios, similar to a previous study [35]. In one, the harvesting rate F increases linearly over $[0, 2.7]$, resulting in a Fold bifurcation. In the other, the growth rate r increases linearly over the interval $[0.5, 2.3]$ resulting in a Flip bifurcation. All other parameters remain fixed. Both scenarios are simulated for 500 time-steps. Negative population values arising from noise are reset to zero. To test statistical significance

of EWS, null trajectories are also simulated in which the bifurcation parameter is fixed.

4.3. Chemostat data analysis

Chemostat data were available for 14 different dilution rates [42] (electronic supplementary material, figure S8). The dilution rate was held constant in each of the experimental treatments and hence, as in the original study, the empirical Hopf bifurcations are considered to occur between two neighbouring dilution rates such that one dilution rate gives rise to observable oscillations and the other does not. For computing EWS, we considered only time series with greater than 25 data points, of which there were 11. The data were initially detrended with a Lowess filter with an 80 day span to account for any unintentional drift in the dilution rate.

4.4. Bootstrapping time-series data

Prior to computing EWS, we use block-bootstrapping to generate an ensemble of samples. This involves detrending the data using a Lowess filter, followed by using a rolling window to obtain overlapping segments that can be considered approximately stationary. For each segment, we generate 100 bootstrapped samples. Each sample is built by selecting blocks of the time series randomly with displacement. The block length is drawn from a geometric distribution [56], with an average large enough such that significant temporal correlations in the time series are retained.

4.5. Computing early warning signals

Conventional EWS were computed according to common practices [6]. The power spectrum was approximated using Welch's method [57], which computes multiple periodograms from overlapping segments of the time-series data and averages them. The periodogram is given by $P(k) = |\tilde{x}(k)|^2$ where

$$\tilde{x}(k) = \frac{1}{\sqrt{n}} \sum_{j=1}^n x_j e^{2\pi i(j-1)(k-1)/n} \quad (4.4)$$

is the discrete Fourier transform, and n is the number of data points in each segment. From the resulting power spectrum $S(k)$, we compute $S_{\max} = \max_k S(k)$. The AIC weights w_{fold} , w_{hopf} and w_{null} are found by fitting the canonical power spectrum forms

$$S_{\text{fold}}(\omega; \sigma, \lambda) = \frac{\sigma^2}{2\pi} \frac{1}{\omega^2 + \lambda^2}, \quad (4.5)$$

$$S_{\text{hopf}}(\omega; \sigma, \mu, \omega_0) = \frac{\sigma^2}{2\pi} \left(\frac{1}{(\omega - \omega_0)^2 + \mu^2} + \frac{1}{(\omega + \omega_0)^2 + \mu^2} \right) \quad (4.6)$$

and
$$S_{\text{null}}(\omega; \sigma) = \frac{\sigma^2}{2\pi} \quad (4.7)$$

to the measured power spectrum using a nonlinear optimization procedure, and taking appropriate ratios of the AIC scores [40]. Finer details are provided in electronic supplementary material, Methods. In addition, we have developed *ewstools*, an open-source Python package located at <https://github.com/ThomasMBury/ewstools> and available on the Python Package Index. To date, *ewstools* provides general functionality for computing conventional and spectral EWS, including auxiliary methods such as bootstrapping. We intend to continue expanding the package and welcome contributions from other Python users.

Data accessibility. The chemostat data used in this study can be obtained from Gregor Fussmann upon reasonable request.

Authors' contributions. All authors conceived of and designed the study. T.M.B. developed the EWS methodology and software, performed the analysis and wrote the first draft. All authors revised and commented on the manuscript.

Competing interests. We declare we have no competing interest.

Funding. This research was funded by NSERC Discovery grants to M.A. and C.T.B. The funders had no role in the research.

Acknowledgements. We thank Gregor Fussmann (McGill University) for providing the chemostat data and two anonymous reviewers for helpful comments on the manuscript.

References

- May RM. 1977 Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature* **269**, 471–477. (doi:10.1038/269471a0)
- Holling CS. 1973 Resilience and stability of ecological systems. *Annu. Rev. Ecol. Syst.* **4**, 1–23. (doi:10.1146/annurev.es.04.110173.000245)
- Boettiger C. 2018 From noise to knowledge: how randomness generates novel phenomena and reveals information. *Ecol. Lett.* **21**, 1255–1267. (doi:10.1111/ele.13085)
- Clements CF, Ozgul A. 2018 Indicators of transitions in biological systems. *Ecol. Lett.* **21**, 905–919. (doi:10.1111/ele.12948)
- Scheffer M *et al.* 2009 Early-warning signals for critical transitions. *Nature* **461**, 53–59. (doi:10.1038/nature08227)
- Dakos V *et al.* 2012 Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. *PLoS ONE* **7**, e41010. (doi:10.1371/journal.pone.0041010)
- Strogatz SH. 2014 *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*, pp. 44–92, 241–300. Boulder, CO: Westview press.
- Carpenter S, Brock W. 2006 Rising variance: a leading indicator of ecological transition. *Ecol. Lett.* **9**, 311–318. (doi:10.1111/j.1461-0248.2005.00877.x)
- Held H, Kleinen T. 2004 Detection of climate system bifurcations by degenerate fingerprinting. *Geophys. Res. Lett.* **31**, L23207. (doi:10.1029/2004GL020972)
- Dakos V, van Nes EH, Donangelo R, Fort H, Scheffer M. 2010 Spatial correlation as leading indicator of catastrophic shifts. *Theor. Ecol.* **3**, 163–174. (doi:10.1007/s12080-009-0060-6)
- Kleinen T, Held H, Petschel-Held G. 2003 The potential role of spectral properties in detecting thresholds in the Earth system: application to the thermohaline circulation. *Ocean Dyn.* **53**, 53–63. (doi:10.1007/s10236-002-0023-6)
- Bauch CT, Sigdel R, Pharaon J, Anand M. 2016 Early warning signals of regime shifts in coupled human–environment systems. *Proc. Natl Acad. Sci. USA* **113**, 14 560–14 567. (doi:10.1073/pnas.1604978113)
- Pananos AD, Bury TM, Wang C, Schonfeld J, Mohanty SP, Nyhan B, Salathé M, Bauch CT. 2017 Critical dynamics in population vaccinating behavior. *Proc. Natl Acad. Sci. USA* **114**, 13 762–13 767. (doi:10.1073/pnas.1704093114)
- Meisel C, Klaus A, Kuehn C, Plenz D. 2015 Critical slowing down governs the transition to neuron spiking. *PLoS Comput. Biol.* **11**, e1004097. (doi:10.1371/journal.pcbi.1004097)
- Diks C, Hommes C, Wang J. 2015 Critical slowing down as an early warning signal for financial crises? *Empir. Econ.* **57**, 1–28. (doi:10.1007/s00181-018-1527-3)
- Boers N. 2018 Early-warning signals for Dansgaard-Oeschger events in a high-resolution ice core record. *Nat. Commun.* **9**, 2556. (doi:10.1038/s41467-018-04881-7)
- Kuznetsov YA. 2013 *Elements of applied bifurcation theory*, vol. 112, pp. 79–150. Berlin, Germany: Springer Science & Business Media.
- Kéfi S, Dakos V, Scheffer M, Van Nes EH, Rietkerk M. 2013 Early warning signals also precede non-catastrophic transitions. *Oikos* **122**, 641–648. (doi:10.1111/j.1600-0706.2012.20838.x)

19. Dakos V, Van Nes EH, D'Odorico P, Scheffer M. 2012 Robustness of variance and autocorrelation as indicators of critical slowing down. *Ecology* **93**, 264–271. (doi:10.1890/11-0889.1)
20. Qin S, Tang C. 2018 Early-warning signals of critical transition: effect of extrinsic noise. *Phys. Rev. E* **97**, 0032406. (doi:10.1103/PhysRevE.97.032406)
21. Wiesenfeld K. 1985 Noisy precursors of nonlinear instabilities. *J. Stat. Phys.* **38**, 1071–1097. (doi:10.1007/BF01010430)
22. Biggs R, Carpenter SR, Brock WA. 2009 Turning back from the brink: detecting an impending regime shift in time to avert it. *Proc. Natl Acad. Sci. USA* **106**, 826–831. (doi:10.1073/pnas.0811729106)
23. Prettyman J, Kuna T, Livina V. 2018 A novel scaling indicator of early warning signals helps anticipate tropical cyclones. *EPL (Europhys. Lett.)* **121**, 10002. (doi:10.1209/0295-5075/121/10002)
24. Tzuk O, Ujjwal SR, Fernandez-Oto C, Seifan M, Meron E. 2019 Period doubling as an indicator for ecosystem sensitivity to climate extremes. *Sci. Rep.* **9**, 1. (doi:10.1038/s41598-018-37186-2)
25. Box GE, Jenkins GM, Reinsel GC. 1994 *Time series analysis: forecasting and control*, pp. 197–199. Englewood Cliffs, NJ: John Wiley & Sons.
26. Priestley MB. 1981 *Spectral analysis and time series*, vol. 1. London, UK: Academic press.
27. Gardiner CW *et al.* 1985 *Handbook of stochastic methods*, vol. 3. Berlin, Germany: Springer.
28. O'Regan SM, Burton DL. 2018 How stochasticity influences leading indicators of critical transitions. *Bull. Math. Biol.* **80**, 1–25. (doi:10.1007/s11538-018-0429-z)
29. Kuehn C. 2013 A mathematical framework for critical transitions: normal forms, variance and applications. *J. Nonlinear Sci.* **23**, 457–510. (doi:10.1007/s00332-012-9158-x)
30. Uhlenbeck GE, Ornstein LS. 1930 On the theory of the Brownian motion. *Phys. Rev.* **36**, 823. (doi:10.1103/PhysRev.36.823)
31. Negahbani E, Steyn-Ross DA, Steyn-Ross ML, Wilson MT, Sleight JW. 2015 Noise-induced precursors of state transitions in the stochastic Wilson–Cowan model. *J. Math. Neurosci. (JMN)* **5**, 9. (doi:10.1186/s13408-015-0021-x)
32. Ridolfi L, D'Odorico P, Laio F. 2015 Indicators of collapse in systems undergoing unsustainable growth. *Bull. Math. Biol.* **77**, 339–347. (doi:10.1007/s11538-013-9922-6)
33. May RM *et al.* 1974 Biological populations with nonoverlapping generations: stable points, stable cycles, and chaos. *Science* **186**, 645–647. (doi:10.1126/science.186.4164.645)
34. Gopalakrishnan E, Sharma Y, John T, Dutta PS, Sujith R. 2016 Early warning signals for critical transitions in a thermoacoustic system. *Sci. Rep.* **6**, 35310. (doi:10.1038/srep35310)
35. Dakos V, Glaser SM, Hsieh Ch, Sugihara G. 2017 Elevated nonlinearity as an indicator of shifts in the dynamics of populations under stress. *J. R. Soc. Interface* **14**, 20160845. (doi:10.1098/rsif.2016.0845)
36. Batt RD, Brock WA, Carpenter SR, Cole JJ, Pace ML, Seekell DA. 2013 Asymmetric response of early warning indicators of phytoplankton transition to and from cycles. *Theor. Ecol.* **6**, 285–293. (doi:10.1007/s12080-013-0190-8)
37. McKane AJ, Newman TJ. 2005 Predator-prey cycles from resonant amplification of demographic stochasticity. *Phys. Rev. Lett.* **94**, 218102. (doi:10.1103/PhysRevLett.94.218102)
38. Gang H, Ditzinger T, Ning C, Haken H. 1993 Stochastic resonance without external periodic force. *Phys. Rev. Lett.* **71**, 807–810. (doi:10.1103/PhysRevLett.71.807)
39. Boettiger C, Batt R. 2019 Bifurcation or state tipping: assessing transition type in a model trophic cascade. *J. Math. Biol.* **80**, 1–13. (doi:10.1007/s00285-019-01358-z)
40. Wagenmakers EJ, Farrell S. 2004 AIC model selection using Akaike weights. *Psychon. Bull. Rev.* **11**, 192–196. (doi:10.3758/BF03206482)
41. Olsen EM, Heino M, Lilly GR, Morgan MJ, Brattey J, Ernande B, Dieckmann U. 2004 Maturation trends indicative of rapid evolution preceded the collapse of northern cod. *Nature* **428**, 932–935. (doi:10.1038/nature02430)
42. Fussmann GF, Ellner SP, Shertzer KW, Hairston Jr NG. 2000 Crossing the Hopf bifurcation in a live predator-prey system. *Science* **290**, 1358–1360. (doi:10.1126/science.290.5495.1358)
43. Williamson MS, Lenton TM. 2015 Detection of bifurcations in noisy coupled systems from multiple time series. *Chaos* **25**, 036407. (doi:10.1063/1.4908603)
44. Lim J, Epureanu BI. 2011 Forecasting a class of bifurcations: theory and experiment. *Phys. Rev. E* **83**, 016203. (doi:10.1103/PhysRevE.83.016203)
45. Ghadami A, Gourgou E, Epureanu BI. 2018 Rate of recovery from perturbations as a means to forecast future stability of living systems. *Sci. Rep.* **8**, 9271. (doi:10.1038/s41598-018-27573-0)
46. Dai L, Vorselen D, Korolev KS, Gore J. 2012 Generic indicators for loss of resilience before a tipping point leading to population collapse. *Science* **336**, 1175–1177. (doi:10.1126/science.1219805)
47. Chen A, Sanchez A, Dai L, Gore J. 2014 Dynamics of a producer-free-loader ecosystem on the brink of collapse. *Nat. Commun.* **5**, 3713. (doi:10.1038/ncomms4713)
48. Perretti CT, Munch SB. 2012 Regime shift indicators fail under noise levels commonly observed in ecological systems. *Ecol. Appl.* **22**, 1772–1779. (doi:10.1890/11-0161.1)
49. Pace ML, Batt RD, Buelo CD, Carpenter SR, Cole JJ, Kurtzweil JT, Wilkinson GM. 2017 Reversal of a cyanobacterial bloom in response to early warnings. *Proc. Natl Acad. Sci. USA* **114**, 352–357. (doi:10.1073/pnas.1612424114)
50. Tyler J, Choi SW, Tewari M. 2020 Real-time, personalized medicine through wearable sensors and dynamic predictive modeling: a new paradigm for clinical medicine. *Curr. Opin. Syst. Biol.* **20**, 17–25. (doi:10.1016/j.coisb.2020.07.001)
51. Hastings A, Wysham DB. 2010 Regime shifts in ecological systems can occur with no warning. *Ecol. Lett.* **13**, 464–472. (doi:10.1111/j.1461-0248.2010.01439.x)
52. Boettiger C, Ross N, Hastings A. 2013 Early warning signals: the charted and uncharted territories. *Theor. Ecol.* **6**, 255–264. (doi:10.1007/s12080-013-0192-6)
53. Karnatak R, Kantz H, Bialonski S. 2017 Early warning signal for interior crises in excitable systems. *Phys. Rev. E* **96**, 042211. (doi:10.1103/PhysRevE.96.042211)
54. Russo L, Spiliotis K, Giannino F, Mazzoleni S, Siettos C. 2019 Bautin bifurcations in a forest-grassland ecosystem with human-environment interactions. *Sci. Rep.* **9**, 2665. (doi:10.1038/s41598-019-39296-x)
55. Webb C. 2003 A complete classification of Darwinian extinction in ecological interactions. *Am. Nat.* **161**, 181–205. (doi:10.1086/345858)
56. Politis DN, Romano JP. 1994 The stationary bootstrap. *J. Am. Stat. Assoc.* **89**, 1303–1313. (doi:10.1080/01621459.1994.10476870)
57. Welch P. 1967 The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.* **15**, 70–73. (doi:10.1109/TAU.1967.1161901)