

Early warning signals and the prosecutor's fallacy

Carl Boettiger and Alan Hastings

Proc. R. Soc. B published online 10 October 2012
doi: 10.1098/rspb.2012.2085

References

[This article cites 23 articles, 8 of which can be accessed free](#)

<http://rsjb.royalsocietypublishing.org/content/early/2012/10/09/rspb.2012.2085.full.html#ref-list-1>

P<P

Published online 10 October 2012 in advance of the print journal.

Subject collections

Articles on similar topics can be found in the following collections

[ecology](#) (1182 articles)
[environmental science](#) (175 articles)
[theoretical biology](#) (40 articles)

Email alerting service

Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click [here](#)

Advance online articles have been peer reviewed and accepted for publication but have not yet appeared in the paper journal (edited, typeset versions may be posted when available prior to final publication). Advance online articles are citable and establish publication priority; they are indexed by PubMed from initial publication. Citations to Advance online articles must include the digital object identifier (DOIs) and date of initial publication.

Early warning signals and the prosecutor's fallacy

Carl Boettiger^{1,*} and Alan Hastings²

¹Center for Population Biology, and ²Department of Environmental Science and Policy, 1 Shields Avenue, University of California, Davis, CA 95616 USA

Early warning signals have been proposed to forecast the possibility of a critical transition, such as the eutrophication of a lake, the collapse of a coral reef or the end of a glacial period. Because such transitions often unfold on temporal and spatial scales that can be difficult to approach by experimental manipulation, research has often relied on historical observations as a source of natural experiments. Here, we examine a critical difference between selecting systems for study based on the fact that we have observed a critical transition and those systems for which we wish to forecast the approach of a transition. This difference arises by conditionally selecting systems known to experience a transition of some sort and failing to account for the bias this introduces—a statistical error often known as the prosecutor's fallacy. By analysing simulated systems that have experienced transitions purely by chance, we reveal an elevated rate of false-positives in common warning signal statistics. We further demonstrate a model-based approach that is less subject to this bias than those more commonly used summary statistics. We note that experimental studies with replicates avoid this pitfall entirely.

Keywords: early warning signals; tipping point; alternative stable states; likelihood methods

1. INTRODUCTION

Mathematics . . . while assisting the trier of fact in the search of truth, must not cast a spell over him. ([1], p. 320)

In the case of *People v. Collins* 1968, the California Supreme Court considered the evidence of an expert witness described by the court as 'an instructor of mathematics at a state college', which concluded that the probability that a randomly selected individual would match the description given by the victim would be less than 1 in 12 million [1]. The prosecution had produced an individual matching the prosecutor's detailed description, and convinced by the mathematics, the lower courts had found him guilty.

The prosecution has only observed that the probability of seeing the evidence (E) they produced given a random innocent individual (I), $P(E|I)$ is very small. From this, one cannot conclude that the individual is indeed guilty, that is, that the probability the individual is innocent given the evidence $P(I|E)$ is also very small. In a city with millions of people, there might be several individuals who match the description of the evidence. Mathematically, $P(E|I)$ need not equal $P(I|E)$; instead these expressions are related by Bayes theorem,

$$P(E|I) = P(I|E) \frac{P(E)}{P(I)}, \quad (1.1)$$

$P(E) \ll 1$ and $P(I) \approx 1$, so $P(E|I) \approx P(I|E)P(E)$, and consequently we cannot conclude that $P(I|E) \ll 1$ from $P(E|I) \ll 1$. Realizing this mistake, the California Supreme Court reversed the decision, and the case became a widely recognized example of the prosecutor's fallacy [2]. Here, we explore how a similar misconception can arise from

the use of historical data to evaluate methods for detecting early warning signals of critical transitions.

Catastrophic transitions or tipping points, where a complex system shifts suddenly from one state to another, have been implicated in a wide array of ecological and global climate systems such as lake ecosystems [3], coral reefs [4], savannah [5], fisheries [6] and tropical forests [7]. Recent research has begun to identify statistical patterns commonly associated with these sudden catastrophic transitions, which could be used as an *early warning sign* to identify an approaching tipping point, which might provide managers time to react to and avert an undesirable state shift [8,9]. An array of statistical patterns associated with tipping point phenomena has been suggested for the detection of early warning signals associated with such sudden transitions. Two of the most commonly used are a pattern of increasing variance [10] and a pattern of increasing autocorrelation [11], which have been tested in both experimental manipulation [3,12–14] and historical observations [15–20].

(a) Testing patterns on historical data

Historical examples of sudden transitions taken from the paleo-climate record provide an important way to test and evaluate potential leading indicator methods, and have been widely used for this purpose [15–20]. Similarly, it has been suggested that data gathered from ecological systems such as lakes that were monitored before they experienced sudden eutrophication, or grasslands subjected to overgrazing could contain data that could help reveal when similar systems are approaching a tipping point [3].

However, testing methods for early warning signals against historical examples of transitions is susceptible to statistical mistakes that arise from selecting data conditional on that data having already exhibited a sudden

* Author for correspondence (cboettig@ucdavis.edu).

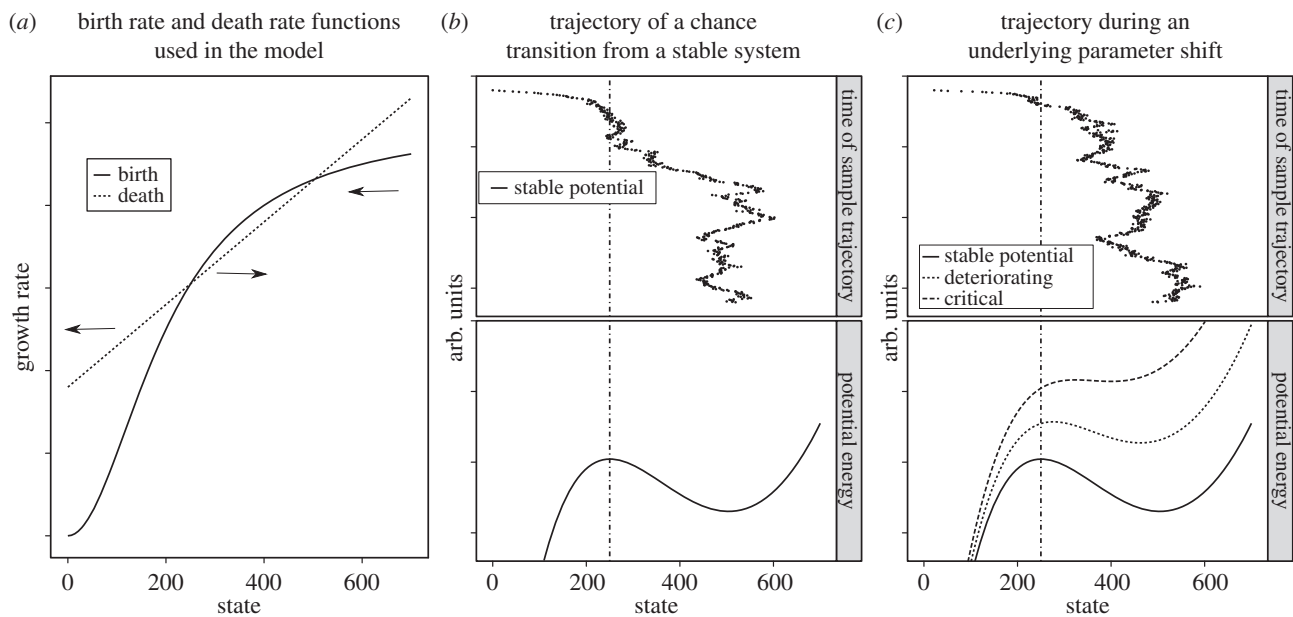


Figure 1. The prosecutor's fallacy. (a) Plot of the model functions shown in equation (2.1) with parameters $a = 180$, $K = 500$, $e = 0.5$ and $h = 200$. When the death rate is higher than the birth rate, the system dynamics drive the state (population size) to smaller values. When the birth rate is higher, the system moves right, as indicated by the arrows. (b) The potential energy is given by the negative integral of $b(n) - d(n)$, shown in the lower plot. The potential function gives an intuitive picture of the stability of a system by imagining the curve as a surface on which a ball is free to bounce across wells correspond to stable points and peaks to unstable points. While most trajectories remain near the stable well, some transition out merely by chance. An example of such a trajectory is shown in the top panel, in which time increases along the vertical axis. Though initially oscillating around the stable state, a chance excursion carries it beyond the Allee threshold (vertical dotted line). Such chance trajectories can produce the statistical patterns as observed in true critical transitions seen in (c): early warning signals are aimed at detecting systems that are slowly moving towards a tipping point or bifurcation, illustrated in the successive curves (deteriorating and critical). Top panel: an example trajectory from a simulation under this process shows the state of the system as the potential moves towards the bifurcation point. The original position of the Allee threshold is shown by the vertical dotted line (though it moves slightly as the parameter changes).

transition. A central tenant of early warning theory is that the system in question is slowly approaching a tipping point that lies some unknown distance away. If nothing is done to remedy the situation, this slow change will inevitably carry the system beyond the tipping point, which introduces a sudden, rapid transition into an undesirable state [8]. This process can be described mathematically as a *bifurcation*, in which a slowly changing parameter reaches a critical value that causes the system stability to change.

Not all sudden transitions are caused by some 'guilty' process slowly driving the system over a tipping point—the kind of process that early warning signals are designed to detect. Some systems may experience such transitions purely by chance, leaving a stable state on an extremely unlikely excursion that happens to stray to far from the stable attractor [9,18], consider this possibility in transitions that arise from analysing historical climate record. Like the evidence presented before the California Supreme Court in 1968, the chance of observing such an 'innocent' transition a priori may be very small, but when selected from a historical record of many possible transitions, this possibility can no longer be ignored.

Figure 1 shows a schematic illustrating critical transitions under each of these scenarios. In figure 1a, the system experiences a bifurcation and should contain an early warning signal. In figure 1b, a similar-looking trajectory emerges from a simulation of a stable system that should not contain a warning signal. While the simulation of the bifurcation scenario shown on the left produces a

similar transition every time, the transition shown on the right is somewhat less probable, occurring in only one per cent of simulations.

2. METHODS AND RESULTS

To investigate whether early warning signals are vulnerable to this fallacy, we simulate a system that is not driven towards a bifurcation such as in figure 1b. This simulation approach allows us to determine whether examining historical events is a valid way to test the utility of these indicators. We simulated 20 000 replicates of a stochastic individual-based birth–death process with an Allee threshold [21], which arises from positive fitness effects at low densities. Above the Allee threshold, the population returns to a positive equilibrium size, whereas below the threshold the population decreases to zero. The model can be represented as a continuous time birth–death process where births and deaths are Poisson events that depend on the current density with rates given by

$$b(n) = \frac{Kn^2}{n^2 + h^2} \quad (2.1)$$

and

$$d(n) = en + a, \quad (2.2)$$

a model with a linear death rate and density-dependent birth rate that drives the Allee effect at low densities and limits growth at high densities. In this model, n

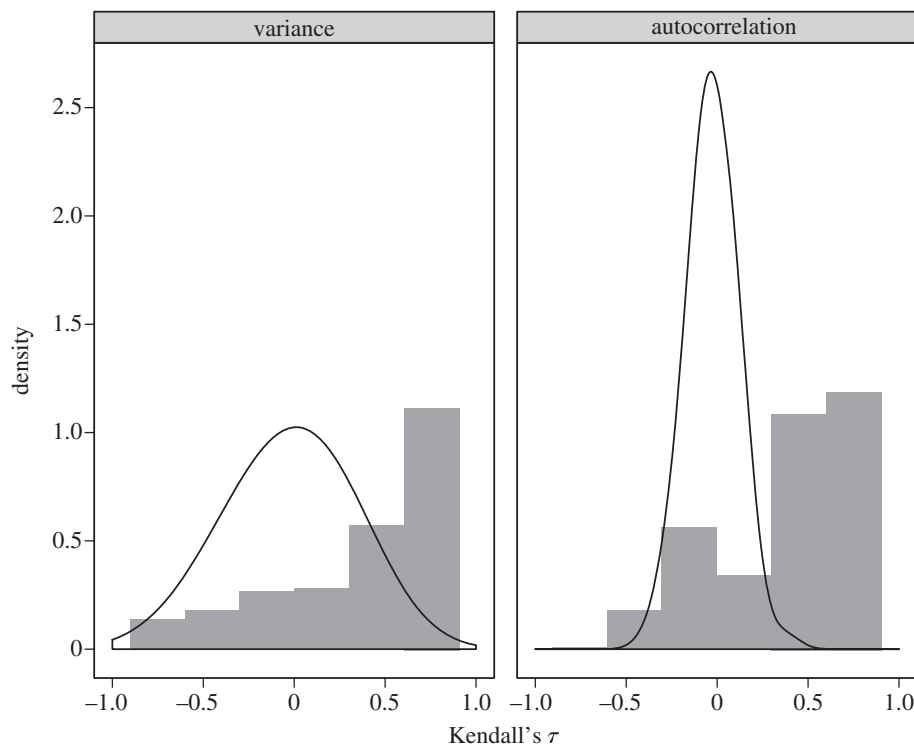


Figure 2. The distribution of the correlation statistic τ for two early warning indicators (variance, autocorrelation) on replicates conditionally selected for having collapsed by chance in simulations is shown in grey bars. Solid lines indicate the estimated density of the statistic from a random sample of the simulations (not conditional on observing a transition). Positive values of τ correspond to a pattern of an indicator increasing with time; typically taken as evidence that a system is approaching a critical transition. In these simulations, the pattern arises instead from the prosecutor's fallacy of conditional selection.

indicates the discrete number of individuals in the population, K indicates a carrying capacity as set by a limiting resource, e a per-capita death rate (the e scaling term in the birth equation allows the carrying capacity K to correspond to a positive equilibrium point), a an additional mortality imposed on the population such as harvest, h is a parameter controlling at what population size the addition of more individuals switches from conferring a positive benefit on growth from Allee interactions $n < h$ to a negative impact on growth owing to increased competition, $n > h$. The key feature of this model is the alternate stable states introduced by this effect; other functional forms for equation (2.1) could serve equally well for these simulations [22]. Although this system can be forced through a bifurcation by increasing the death rate, in these simulations, all parameters are held constant and no bifurcation occurs. Consequently, we do not anticipate an early warning signal of an approaching bifurcation.

The simulation starts from the positive equilibrium population size. Although the chance of a transition across the Allee threshold in any given time step is small, given enough time, this system will eventually experience such a rare event, driving the population extinct. We ran each replicate over 50 000 time units, sampling the system every 50 time units. In this time, window 266 of the 1000 replicates experience population collapse. To keep the examples of comparable sample size, we focus on a section of the data 500 time points prior to the system approaching the transition.

To test whether selecting systems that have experienced spontaneous transitions could bias the analysis towards false-positive detection of early warning signals (the prosecutor's fallacy), we selected replicates conditional on

having collapsed in the simulations. We then selected a window around each system that ended just before the collapse, while the population values were still above the Allee threshold. For each replicate, we calculated the most common early warning indicators, variance and autocorrelation [8,10,16], around a moving window equal to half the length of that time series.

To test for the presence of a warning signal in these indicators, we computed values of Kendall's τ for both indicators for each of the 266 replicates. Kendall's τ is a non-parametric measure of rank correlation frequently used to identify an increasing trend ($\tau > 0$ in early warning signals [16,23], defined as $\tau \times \frac{1}{2}n(n-1)$ in n observations.¹ τ takes values in $(-1,1)$. The distribution of τ values observed across these replicates is shown in figure 2. We compare the distribution of τ from all the simulations to the distribution conditioned on experiencing a chance transition to the alternative stable state. To avoid an effect of sample size, the time series are all chosen to be of the same length.

To demonstrate that the effect we observe is not unique to models with Allee effects, we provide an example of the effect arising in a discrete-time model with two non-zero stable states adapted from [4],

$$X_{t+1} = X_t \exp \left(r \left(1 - \frac{X_t}{K} \right) - \frac{a \times X_t^{Q-1}}{X_t^Q + H^Q} \right). \quad (2.3)$$

which combines a logistic growth model with a saturating predator response (see [24] for detailed discussion), shown in figure 3. Code to replicate the analysis can be found at <https://github.com/cboettig/earlywarning/tree/prosecutor/>.




Figure 3. The identical analysis from figure 2 is shown for the model in equation (2.3), using parameters $r = 0.75$, $K = 10$, $a = 1.7$, $Q = 3$ and $H = 1$. A similar statistical bias, particularly towards positive values of τ , occurs in this model as well.

For each of these replicates, we also take a model-based approach, estimating parameters for an approximate linear model of the system approaching a saddle node bifurcation, as described by Boettiger & Hastings [25],

$$dX = \sqrt{r_t}(\varphi(r_t) - X_t) dt + \sigma\sqrt{\varphi(r_t)} dB_t. \quad (2.4)$$

In this model, the parameter m describes the approach towards the saddle-node bifurcation. Estimates $m < 0$ are expected in systems approaching a bifurcation, while for stable systems, m should be approximately zero. None of the estimates across the 266 simulations differed from zero in our study; hence the model-based estimation shows no evidence of bias on data that has been selected conditional on collapse.

3. DISCUSSION

The attempts to detect early warning signs for critical transitions are based on the concept of a deteriorating environment as embodied in a changing parameter [8], which is a different kind of transition than one that is driven instead by stochasticity in an environment that is otherwise constant and exhibiting no directional change. When trying to use historical data to understand critical transitions, we often do not know which category, changing environment or simply chance, an observed large change falls into.

We have shown here that systems that undergo rare sudden transitions owing to chance look statistically different from their counterparts that do not, even though they are driven by the same stochastic process. In particular, such conditionally selected examples are more likely to show signs associated with an early warning of an approaching tipping point, such as increasing

variance or increasing autocorrelation, as measured by Kendall's τ . This increases the risk of false positives—cases in which a warning signal being tested appears to have successfully detected an underlying change in the system leading to a tipping point, when in fact the example comes instead from a stable system with no underlying change in parameters. Figure 2 shows that many of the chance crashes show values of τ that are significantly larger than those observed in the otherwise identical replicates that did not experience a chance transition, thus ‘detecting’ an underlying change in the system dynamics that is not in fact present.

(a) *Chance transitions are false positives for early warning signals*

It seems tempting to argue that this bias towards positive detection in historical examples is not problematic—each of these systems did indeed collapse; so the increased probability of exhibiting warning signals could be taken as a successful detection. Unfortunately, this is not the case. At the moment the forecast is made, these systems are not likely to transition, because they experience a strong pull towards the original stable state. A closer look at the patterns involved shows why common indicators such as autocorrelation and variance can be misleading.

As the system gets farther from its stable point, it is more likely to draw a random step that returns it towards the stable point. Despite this, there is always some probability that it will move further still; so systems that do cross the tipping point must do so rather quickly by a string of events. This pattern, clearly visible before the crashes in each of the examples in figure 1, produces a string of observations that appear more highly autocorrelated (if we are sampling the system frequently enough to

catch the excursion at all) than we observe in the rest of the fluctuations around the equilibrium. Yet, this autocorrelation comes from a chance trajectory moving quickly *away* from the stable state, not from the critical slowing down pattern in the return times to the stable state that precede a saddle-node bifurcation and motivate the early warning signal.

This longer than expected excursion results in a higher than expected variance in that window as well. Both variance and autocorrelation are calculated using a moving window over the time-series, which allows the method to pick out a pattern of change as the window moves along the sequence. If this chance excursion that precedes the crash happens to fill a significant part of the moving window, the resulting pattern will tend to show an increase in autocorrelation or variance. If the chance excursion is relatively rapid compared with the frequency at which the system is observed (spacing of the data) or the width of the moving window, the excursion may not significantly alter the general pattern. In this way, some of the events in which a crash is observed will appear to present these statistical patterns of increased variance or autocorrelation without being harbingers of approaching critical transitions.

(b) The truncation of observations

If we had a complete knowledge of the system dynamics, then we could eliminate the bias we observe here because the bias arises from the transient branch of the trajectory that crosses the threshold, and if the system were truncated at the minimum of the potential, then the effect we emphasize here would not appear. But, it is not possible to truncate the system in any practical application. The precise location of the minimum of the potential (the location of the deterministic equilibrium) is unknown. Moreover, under the hypothesis that the system is approaching a critical transition, the location of the minimum potential moves; so it cannot easily be estimated by previous observations, (see figure 1c where the equilibrium point moves in the direction of the transition). Thus, it is neither practical nor desirable to suggest that historical time series can be used by following a simple truncation rule that avoids the branch of a trajectory crossing the threshold to another basin of attraction. Exactly where a particular study will choose to truncate such a trajectory will necessarily be arbitrary without an underlying model of the process. Frequently this is done by removing the very steep, monotonic branch of the trajectory expected, once the system crosses the unstable threshold. Such an approach corresponds with our choice of termination and produces the bias we discuss here.

The examples of figure 1, though only single replicates, may be useful in illustrating these issues. Figure 1c, top panel shows a sample trajectory of a system with a parameter shift, while 1b shows a trajectory without a shift. Both trajectories become more highly autocorrelated and higher variance near the end of the time series (time increases on the y -axis in figure 1). The part of the time series following the critical transition shows a fast and monotonic trajectory to the unstable trajectory, and would usually be excluded by an analysis for warning signals in advance of the transition. No such clear pattern exists prior to the transition in figure 1b. An alternative

proposal to terminate the trajectory in (b) earlier would also risk decreasing the signal seen in (c), and would be inconsistent with the application of warning signals in the forecasting context, where there would be no such truncation.

(c) Comparing to the model-based method

In our numerical experiment, the model-based estimate of early warning signals appears more robust than the summary statistics, producing the same estimates on both the conditionally selected replicates as on a random sample of the replicates. This is a consequence of the more rigid specifications that come with a model-based approach—the pattern expected is less general than any increase in variance or autocorrelation, but instead must be one that matches its approximation of the saddle-node bifurcation. This observation highlights the difference between the pattern driving the false positive trends in increasing variance and increasing autocorrelation and the pattern anticipated in the saddle-node model. This should not however be taken as evidence that the model-based approach is immune to the bias of the prosecutor's fallacy.

(d) Importance of experimental approaches

The problem we highlight ultimately stems from the difficulty of having only a single realization with which to examine a complex problem. The only way to deal with this problem embodied is through replication, as can be done in an experimental system in laboratory manipulations such as Drake & Griffen [12], Veraart *et al.* [13] and Dai *et al.* [14] and at the scale of whole lake ecosystems in Carpenter [3]. Experimental procedures avoid the hazard of the prosecutor's fallacy by generating a complete sample of replicates, rather than by selecting a subset of cases from some larger historical sample.

This research was supported by funding from NSF grant EF 0742674 to A.H. and a Computational Sciences Graduate Fellowship from the Department of Energy grant DE-FG02-97ER25308 and NERSC Supercomputing grant DE-AC02-05CH11231 to C.B. The authors thank M. Baskett, T. A. Perkins and N. Ross for helpful comments on earlier drafts of the manuscript, and also P. Ditlevsen and an anonymous reviewer for their comments.

ENDNOTE

¹A pair of observations (x_i, y_i) and (x_j, y_j) are concordant if $x_i > x_j$ and $y_i > y_j$ or $x_i < x_j$ and $y_i < y_j$ and discordant otherwise; equalities excepted.

REFERENCES

- 1 Sullivan, p320, in *People v. Collins*, 68 Cal. 2d 319 (1968).
- 2 Thompson, W. C. & Schumann, E. L. 1987 Interpretation of statistical evidence in criminal trials: the prosecutor's fallacy and the defense attorney's fallacy. *Law Hum. Behav.* **11**, 167–187. (doi:10.1007/BF01044641)
- 3 Carpenter, J. 2011 May the best analyst win. *Science* **331**, 698–699. (doi:10.1126/science.331.6018.698)
- 4 Mumby, P. J., Hastings, A. & Edwards, H. J. 2007 Thresholds and the resilience of Caribbean coral reefs. *Nature* **450**, 98–101. (doi:10.1038/nature06252)
- 5 Kéfi, S., Rietkerk, M., Alados, C. L., Pueyo, Y., Papanastasis, V. P., Elaiach, A. & de Ruiter, P. C. 2007 Spatial vegetation patterns and imminent desertification in Mediterranean

- arid ecosystems. *Nature* **449**, 213–217. (doi:10.1038/nature06111)
- 6 Berkes, F. *et al.* 2006 Ecology. Globalization, roving bandits, and marine resources. *Science (New York, NY)* **311**, 1557–1558. (doi:10.1126/science.1122804)
 - 7 Hirota, M., Holmgren, M., Van Nes, E. H. & Scheffer, M. 2011 Global resilience of tropical forest and savanna to critical transitions. *Science* **334**, 232–235. (doi:10.1126/science.1210657)
 - 8 Scheffer, M. *et al.* 2009 Early-warning signals for critical transitions. *Nature* **461**, 53–59. (doi:10.1038/nature08227)
 - 9 Lenton, T. M. 2011 Early warning of climate tipping points. *Nat. Clim. Change* **1**, 201–209. (doi:10.1038/nclimate1143)
 - 10 Carpenter, S. R. & Brock, W. A. 2006 Rising variance: a leading indicator of ecological transition. *Ecol. Lett.* **9**, 311–318. (doi:10.1111/j.1461-0248.2005.00877.x)
 - 11 van Nes, E. H. & Scheffer, M. 2007 Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *Am. Nat.* **169**, 738–747. (doi:10.1086/516845)
 - 12 Drake, J. M. & Griffen, B. D. 2010 Early warning signals of extinction in deteriorating environments. *Nature* **467**, 456–459. (doi:10.1038/nature09389)
 - 13 Veraart, A. J., Faassen, E. J., Dakos, V., van Nes, E. H., Lürling, M. & Scheffer, M. 2011 Recovery rates reflect distance to a tipping point in a living system. *Nature* **481**, 357–359. (doi:10.1038/nature10723)
 - 14 Dai, L., Vorselen, D., Korolev, K. S. & 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)
 - 15 Livina, V. N. & Lenton, T. M. 2007 A modified method for detecting incipient bifurcations in a dynamical system. *Geophys. Res. Lett.* **34**, 1–5. (doi:10.1029/2006GL028672)
 - 16 Dakos, V., Scheffer, M., van Nes, E. H., Brovkin, V., Petoukhov, V. & Held, H. 2008 Slowing down as an early warning signal for abrupt climate change. *Proc. Natl Acad. Sci. USA* **105**, 14 308–14 312. (doi:10.1073/pnas.0802430105)
 - 17 Lenton, T. M., Livina, V. & Dakos, V. 2012 Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness. *Phil. Trans. R. Soc. A* **370**, 1185–1204. (doi:10.1098/rsta.2011.0304)
 - 18 Ditlevsen, P. D. & Johnsen, S. J. 2010 Tipping points: early warning and wishful thinking. *Geophys. Res. Lett.* **37**, 2–5. (doi:10.1029/2010GL044486)
 - 19 Guttal, V. & Jayaprakash, C. 2008 Spatial variance and spatial skewness: leading indicators of regime shifts in spatial ecological systems. *Theoret. Ecol.* **2**, 3–12. (doi:10.1007/s12080-008-0033-1)
 - 20 Thompson, J. M. T. & Sieber, J. 2010 Climate tipping as a noisy bifurcation: a predictive technique. *IMA J. Appl. Math.* **76**, 27–46. (doi:10.1093/imamat/hxq060)
 - 21 Courchamp, F., Berec, L. & Gascoigne, J. 2008 *Allee effects in ecology and conservation*. Oxford, UK: Oxford University Press.
 - 22 Scheffer, M., Carpenter, S. R., Foley, J. A., Folke, C. & Walker, B. 2001 Catastrophic shifts in ecosystems. *Nature* **413**, 591–596. (doi:10.1038/35098000)
 - 23 Dakos, V., Kéfi, S., Rietkerk, M., Nes, E. H. V. & Scheffer, M. 2011 Slowing down in spatially patterned ecosystems at the brink of collapse. *Am. Nat.* **177**, E153–E166. (doi:10.1086/659945)
 - 24 May, R. M. 1977 Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature* **269**, 471–477. (doi:10.1038/269471a0)
 - 25 Boettiger, C. & Hastings, A. 2012 Quantifying limits to detection of early warning for critical transitions. *J. R. Soc. Interface* **9**, 2527–2539. (doi:10.1098/rsif.2012.0125)