

A new approach for rapid detection of nearby thresholds in ecosystem time series

Stephen R. Carpenter, William A. Brock, Jonathan J. Cole and Michael L. Pace

S. R. Carpenter (srcarpen@wisc.edu), Center for Limnology, Univ. of Wisconsin, Madison, WI 53706, USA. – W. A. Brock, Dept of Economics, Univ. of Wisconsin, Madison, WI 53706, USA and: Dept of Economics, Univ. of Missouri, Columbia, MO 65211, USA. – J. J. Cole, Cary Inst. of Ecosystem Studies, Millbrook, NY 12545, USA. – M. L. Pace, Dept of Environmental Sciences, Univ. of Virginia, Charlottesville, VA 22904, USA.

Massive changes to ecosystems sometimes cross thresholds from which recovery can be difficult, expensive and slow. These thresholds are usually discovered in post hoc analyses long after the event occurred. Anticipating these changes prior to their occurrence could give managers a chance to intervene. Here we present a novel approach for anticipating ecosystem thresholds that combines resilience indicators with Quickest detection of change points. Unlike existing methods, the Quickest detection method is updated every time a data point arrives, and minimizes the time to detect an approaching threshold given the users' tolerance for false alarms. The procedure accurately detected an impending regime shift in an experimentally manipulated ecosystem. An ecosystem model was used to determine if the method can detect an approaching threshold soon enough to prevent a regime shift. When the monitored variable was directly involved in the interaction that caused the regime shift, detection was quick enough to avert collapse. When the monitored variable was only indirectly linked to the critical transition, detection came too late. The procedure is useful for assessing changes in resilience as ecosystems approach thresholds. However some thresholds cannot be detected in time to prevent regime shifts, and surprises will be inevitable in ecosystem management.

Recognition of an incoming nuclear missile is a case where fast detection is needed to address an impending significant event. During the Cold War, scientists developed a family of methods for quickest detection of change points, such as appearance of a missile on radar or a submarine on sonar (Shiryaev 2010). These methods, known as Quickest detection (QD) methods, employ a likelihood ratio for two models, one for the status quo ('no submarine') and one for a new and different situation ('submarine detected'). At each time series step, the model likelihoods and their ratio are updated. When the updated likelihood ratio exceeds a detection threshold, an alarm is triggered. QD methods minimize the time to detection of a change point, given the user's tolerance for false alarms. They are widely used to screen for changes in military activities, industrial processes, and financial systems (Lai 1995) and recently have been used for early detection of cyber attacks (Tartakovsky et al. 2006).

In ecology, a growing literature addresses indicators of approaching thresholds, or impending loss of resilience, in ecosystems. As an ecosystem moves toward a critical threshold, the stability basin flattens out so that perturbations have larger effects, recovery from perturbations takes longer, and variance increases (Carpenter and Brock 2006). Rising variance is directly related to slower recovery from disturbance as the ecosystem approaches the threshold

(Scheffer et al. 2009). In addition, variance can increase when shocks cause brief excursions into a different stability basin, a phenomenon known as flickering (Brock and Carpenter 2010, Wang et al. 2012). Rising variance and autocorrelation prior to critical transitions is well-established in models and has been demonstrated for laboratory microcosms and whole ecosystems (Dakos et al. 2008, Drake and Griffen 2010, Carpenter et al. 2011, Dai et al. 2012, Seekell et al. 2012, Wang et al. 2012, Pace et al. 2013). It is important to note that rising variance and autocorrelation occur for many, but not all large transitions in ecosystems (Scheffer et al. 2009). For example, a completely external disturbance such as an earthquake, volcanic eruption, tsunami or asteroid strike is not expected to have ecological early warnings. Less obviously, some types of critical transitions do not provide early warnings (Scheffer et al. 2009, Hastings and Wysham 2010). On the other hand, some massive ecological regime shifts that are not critical transitions are preceded by rising variance and autocorrelation (Kéfi et al. 2012). Nonetheless, many important kinds of ecosystem regime shifts, such as those involved in eutrophication, fisheries collapse, rangeland degradation, desertification and others involve critical transitions that can be announced by early warning signals (Scheffer et al. 2012).

Statistical time-series analyses of the distance to thresholds have used several different approaches (Dakos et al. 2012). Approaches that reconstruct the potential surface (Livina and Lenton 2007) are data-intensive but reveal the shape of the alternate states when data are sufficient. In contrast, non-parametric methods to estimate variance or eigenvalues make no assumption about the underlying process, and also require rather large data sets (Brock and Carpenter 2012). Other approaches involve comparing two or more models for the observed time series using likelihood statistics (Boettiger and Hastings 2012b, Ives and Dakos 2012). Conditional heteroscedasticity, a persistent form of variance in time series, rises to statistically significant levels as a critical threshold is approached (Seekell et al. 2012). All of these methods have important uses (Dakos et al. 2012) but none are designed for decision making ‘online’ meaning automated, instantaneous and in real time. In online decision systems, statistics are updated recursively each time a data point is collected. In ecology, automated sensor systems that record data at frequencies ranging from seconds to hours are increasingly common (Porter et al. 2012). An alarm is triggered when a statistical threshold is crossed, and a decision maker acts on the alarm to change some aspect of ecosystem management. QD methods are designed for use online. In the case of early warning indicators of critical transitions, this paper uses QD methods to detect the statistical signatures of an approaching critical transition, and not the transition itself.

For online decision making to be useful, the alarm and subsequent management action must be capable of changing the ecosystem to avoid an unwanted outcome. If an approaching threshold is detected in an ecosystem under stress, and the stress is mitigated, does the ecosystem recover or is the threshold crossed anyway? Previous simulation-model studies of lake eutrophication (Contamin and Ellison 2009) and management of fish habitat and harvest (Biggs et al. 2009) show that both recovery and collapse are possible.

Application of QD methods to problems of ecosystem thresholds thus raises two questions. How quickly can the QD methods discern the changing indicators? If a change is detected, can management interventions prevent the regime shift, or will the inertia of ecosystem dynamics carry the ecosystem past the threshold? Here we address these questions using field observations of an ecosystem going through a regime shift and an ecosystem model of critical transitions in a food web.

Methods

Quickest detection

The QD method employs two probability densities: $f(x_i)$ which is the density of the i th observation of x in the original state of the system, and $g(x_i)$ which is the density of the i th observation of x after the early warning has arrived. The time series x_i can be measurements of any ecological variable, or any statistic such as the variance or autocorrelation of an ecological variable. The likelihood ratio in a given time step is

$$\Lambda_i = \frac{g(x_i)}{f(x_i)} \quad (1)$$

The procedure (Polunchenko and Tartakovsky 2012)(section 4, Eq. 22 and 23) is based on the stopping time Ψ_A

$$\Psi_A = \inf (n \geq 1 : R_n \geq A) \quad (2)$$

In plain English, n is the first time step where R_n exceeds A , and Ψ_A is a random variable with values taken by this time step. Infimum or ‘inf’ is the smallest value of positive non-zero n that has R_n greater than or equal to A . Here A is the detection threshold. A can be chosen using the expected number of time steps to a false positive, given the user’s tolerance of false alarms (Pollak and Tartakovsky 2009). In practice, we selected A to be within the range of A values where time to first alarm is not sensitive to A , according to sensitivity analyses presented in the Results.

The detection statistic (Shiryayev–Roberts statistic) R_n is given by the recursion

$$R_n = (1 + R_{n-1})\Lambda_n \quad (3)$$

Thus R_n is a likelihood ratio that is updated using the new observation that arrives at each time step. The derivation of Eq. 3 subject to Eq. 2 is explained in Shiryayev (2010) and Polunchenko and Tartakovsky (2012). The method is optimal in the sense that it minimizes the average delay to detection of an event, where the average time to false alarm is greater than some specified bound (Pollak and Tartakovsky 2009)(section 2, theorem 1 and surrounding text).

This process, called the Quickest detection or Shiryayev–Roberts procedure, requires three user-defined components: the detection threshold A and the densities $f(x_i)$ and $g(x_i)$. We define these components on a case-by-case basis below. For sequential monitoring, an ‘alarm’ is recorded every time R_n exceeds A . R_n is reset to 1 in the time step following the alarm (Polunchenko and Tartakovsky 2012).

Field time series

To evaluate the behavior of the QD method with field data, we used the 2008–2011 daily chlorophyll time series for a whole-lake experiment described by Carpenter et al. (2011). In the field experiment, a critical transition was triggered by gradually adding individuals of a top predator species (largemouth bass *Micropterus salmoides*) to a lake food web that was dominated by prey fishes at the beginning of the experiment. The critical transition resulted from a well-known trophic cascade in which largemouth bass became dominant, prey fish became rare, zooplankton biomass shifted to large-bodied animals and chlorophyll concentration declined (Carpenter and Kitchell 1993). The manipulated ecosystem (Peter Lake) was monitored in parallel with an undisturbed reference ecosystem (Paul Lake) that was dominated by largemouth bass throughout the experiment. Previous post hoc analyses indicate that early warnings began during summer 2009 and the regime shift was complete or nearly so by late summer of 2010 (Carpenter et al. 2011, Seekell et al. 2012, Pace et al. 2013).

The QD method was applied to the time series of standard deviations of log-transformed daily chlorophyll computed using 28 day rolling windows (Carpenter et al. 2011). We set $A = 1000$, which was chosen to be within the range where the time to first alarm is not very sensitive to A (Results). The likelihood of the manipulated lake's observed standard deviations was calculated for two distributions. One is the distribution of standard deviations in the original state, and the other is the distribution in the new state. These two sets of probabilities are $f(x_i)$ and $g(x_i)$, respectively (Eq. 1). Functions for the probability densities were the normal distributions

$$\begin{aligned} f(S_{t,M}) &: N(\mu = S_{t,C} - \Delta_f, \sigma_f) \\ g(S_{t,M}) &: N(\mu = S_{t,C} - \Delta_g, \sigma_g) \end{aligned} \quad (4)$$

Normal probability plots indicated that the normal distribution was a reasonable approximation. The means (μ) of distributions for the two states are calculated by making adjustments (Δ_f and Δ_g) to the standard deviation of chlorophyll in the reference lake ($S_{t,C}$). Pre-manipulation data showed that the standard deviation of chlorophyll in the manipulated lake ($S_{t,M}$) was less than that of the reference lake ($S_{t,C}$) (Carpenter et al. 2011, Supplementary material Appendix A1 Table A2). During the critical transition, the standard deviation of chlorophyll in the manipulated lake increased relative to the reference lake as expected from theory and models (Carpenter et al. 2008, 2011, Scheffer et al. 2009). Thus Δ_f adjusts μ to approximate the standard deviation of chlorophyll in the original state of the manipulated lake, and Δ_g adjusts μ to approximate the standard deviation of chlorophyll in the new state of the manipulated lake. To estimate the parameters of the $f()$ function, Δ_f and σ_f were calculated for the first 50 data points, and the estimate of σ was rounded upward to allow for wider fluctuations in the future. Thus we set $\Delta_f = 0.3$ and $\sigma_f = \sigma_g = 0.1$. We arbitrarily set $\Delta_g = 0.12$ to represent a substantial increase in the standard deviation.

We chose this specification of the densities $f(x_i)$, $g(x_i)$ based upon experience with computer experiments and ease of implementation on field data. The references for the QD method discuss general issues of choosing $f(x_i)$ and $g(x_i)$ (Pollak and Tartakovsky 2009, Polunchenko and Tartakovsky 2012). For interested readers we present a second analysis based on autocorrelation in Supplementary Materials Appendix A1.

The field data were used to evaluate the sensitivity of the QD method to choice of A and parameters of the $f(x_i)$ and $g(x_i)$ functions. The day of first alarm was computed for ranges of A , Δ_f and Δ_g and σ_g .

Model time series

Even if QD triggers an alarm before a tipping point is reached, an immediate management intervention may not prevent the regime shift. To test whether the alarm from QD comes soon enough to prevent a regime shift, we used a food web model. In the model, a top predator fish (or piscivore) at the higher of two alternate stable equilibria is harvested at gradually increasing rates. If the harvest rate crosses

a threshold, then the piscivore biomass collapses to a low stable state. In our simulations, the harvest rate was gradually increased until an alarm occurs, and then the harvest rate is dropped to zero. To evaluate the success of the intervention, we then determined whether the piscivore biomass remained at the high stable state or dropped to the low stable state.

The model is also capable of simulating the critical transition caused by stocking adult piscivores starting from a low population of piscivores, similar to the whole-lake experiment of Carpenter et al. (2011). However, preventing this critical transition is not usually relevant to management, whereas preventing overharvest is a common problem in fisheries management. Moreover, the critical transition caused by stocking adult piscivores is not likely to be reversible by simply ending the stocking. Reproduction by adult piscivores already stocked has the potential to carry the ecosystem past the critical point. Therefore, simulation of harvest management provides a more stringent test of the early warning system.

The model (Supplementary Material Appendix A2) represented ecosystem changes in a lake food web as harvest rates of the piscivore (top predator) are increased gradually (Carpenter et al. 2008). The model represents biomass dynamics of adult and juvenile piscivores, prey fish, zooplankton and phytoplankton. It was calibrated using 14 years of data from whole-lake experiments in five lakes (Carpenter and Kitchell 1993, Carpenter et al. 2001). The model exhibits rising variance and autocorrelation before the critical transition (Carpenter et al. 2008, Brock and Carpenter 2012).

When the harvest rate of top predators is low, the ecosystem has a single stable point with high biomass of top predators (Carpenter et al. 2008, Brock and Carpenter 2012). When the harvest rate is increased past a critical level from below, the ecosystem becomes bistable with two stable points at low and high biomass of top predators. We studied the performance of the QD method as harvest rates were gradually increased toward the critical point leading to a transition to the alternate bistable condition.

We investigated the performance of the rolling-window standard deviations of two monitored variables, prey fish and phytoplankton biomass. Prey fish are directly involved in the interaction that creates the bifurcation, whereas phytoplankton are two trophic levels away from the bifurcation. In practice, it is difficult to measure high-frequency time series of prey fish biomass but relatively easy to measure high-frequency time series of phytoplankton. Therefore it is useful to know whether phytoplankton-based indicators can provide early warnings. Window length was 60 time steps.

We also evaluated the possibility of reversing a regime shift by detecting a loss of resilience and then immediately decreasing the harvest rate of top predators in the model. For reversal experiments, we used the QD procedure with

$$\begin{aligned} f(S_t) &: N(\mu_f, \sigma_f) \\ g(S_t) &: N(\mu_g, \sigma_g) \end{aligned} \quad (5)$$

where S_t is the time series of rolling window standard deviations for prey fish or phytoplankton. For reversal experiments

based on the prey fish standard deviation, $\mu_f = 0.03$, $\sigma_f = 0.03$, $\mu_g = 0.06$, and $\sigma_g = 0.03$ for the simulations reported below. For reversal experiments based on the phytoplankton standard deviation, we set $\mu_f = 0.2$, $\sigma_f = 0.2$, $\mu_g = 0.4$, and $\sigma_g = 0.1$. All computations were performed in R ver. 2.15.0 downloaded 28 May 2012 (<www.r-project.org/>).

Results

Time series of daily \log_{10} transformed chlorophyll concentrations measured in Paul (reference) and Peter (manipulated) lakes are presented in Fig. 1A. Rolling-window standard deviation of chlorophyll was elevated in year 2 (2009) and the first half of year 3 (2010) (Fig. 1B). In concert with the elevated standard deviation, the Shiryayev–Roberts statistic (Eq. 2) spikes frequently during year 2 and

the early part of year 3, and the first alarm occurs early in year 2 (Fig. 1C). The first alarm indicates that the early warning has arrived and the threshold is approaching (‘the missile has been detected’). This occurred on day 130 (Fig. 1C), after the first bass addition on day 57 (green line, Fig. 1B) but before the bass additions on days 137 and 171. The subsequent alarms indicate that the threshold is still nearby (‘the missile is still there’). The regime shift was complete by day 293 (Fig. 1C, gray line). The field experiment was designed to hold the ecosystem close to the threshold for as long as possible, while still moving slowly toward the threshold. It is difficult to state precisely when the new bass-dominated regime began, but we observed many young-of-year bass in late 2009 and their high survivorship over the winter of 2009–2010. This recruitment of bass was a key event in the completion of the trophic cascade.

In a management setting, action would be taken after the first alarm if tolerance for the pending regime shift was low.

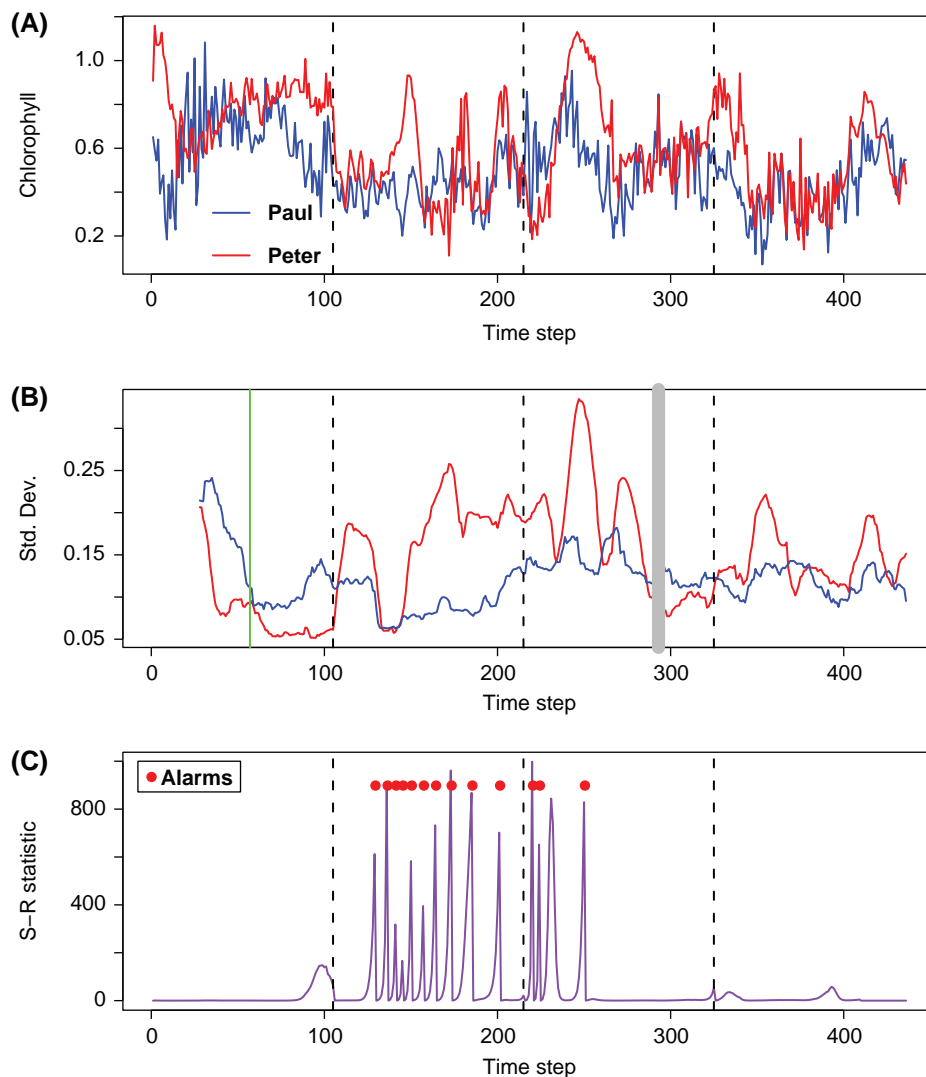


Figure 1. Time series from the whole-lake experiment from summer stratified seasons of 2008–2011 organized as contiguous days from the start of the experiment. Vertical dashed lines denote breaks between years. (A) Log₁₀ of daily chlorophyll concentration in Paul (reference, blue) and Peter (manipulated, red) lakes. (B) Rolling window (28 day) standard deviation of chlorophyll in Paul (reference, blue) and Peter (manipulated, red) lakes. The vertical green line shows the time of first bass addition and the vertical gray stripe shows the approximate time when the regime shift was completed. (C) Shiryayev–Roberts statistic and days of alarms (red dots).

Analysis of the rolling-window autocorrelation time series showed the first alarm at time step 144 (Supplementary material Appendix A1), which is similar to the time step of first alarm (130) shown by the standard deviation.

Time series of chlorophyll, pH, dissolved oxygen, zooplankton biomass and planktivorous fishes exhibit statistical signatures of critical transition, including elevated variance and autocorrelation, in years 2 and 3, consistent with the pattern shown here (Carpenter et al. 2011, Seekell et al. 2012, Pace et al. 2013). These papers show that the ecosystem regime shift was complete by late in year 3 (about day 300 on the x-axes of Fig. 1). Thus the QD procedure discerns a loss of resilience as the regime shift begins and well before the regime shift is complete.

Over broad ranges of the parameters, the day of first alarm is rather constant (Fig. 2). The time of first alarm is especially important, because the first alarm would trigger action in a management situation. For values of the alarm threshold A from $10^{2.5}$ to 10^9 , the first alarm occurs near day 130, or early in the second year (2009) (Fig. 2A). It is possible to set A high enough that no alarms occur (not shown).

Note that R will increase by one unit per time step if the original and new states of the system are identical (i.e. $\Lambda = 1$). This suggests that A should be set to be greater

than the number of time steps that will be analyzed. For this data set the number of time steps was 436, $\log_{10}(436) \approx 2.6$, near the breakpoint of the curve in Fig. 2A. Therefore the day of the first alarm was roughly similar for values of A greater than the number of time steps in this data set.

For values of Δ_f between zero and about 0.04 the first alarm occurred around day 130, i.e. early in year 2 (Fig. 2B). In applications of QD, Δ_f and σ_f can be estimated from data obtained early in the time series.

For values of Δ_g above about 0.1, the first alarm occurred around day 130, i.e. early in year 2, and increased slowly as Δ_g increased (Fig. 2C). The day of first alarm was similar for values of σ_g between 0.01 and 0.3 (Fig. 2D). In applications of QD, simulations could be used to estimate the ranges of Δ_g and σ_g for which the alarm time was relatively insensitive to parameter choice.

The QD procedure based on daily observations of chlorophyll provided an early warning. However the field data cannot be used to determine whether the regime shift could have been reversed following detection. We used a simulation model to evaluate the possibility of reversing a regime shift.

In the food web model, standard deviation of prey fishes rises steeply as the ecosystem passes through the critical

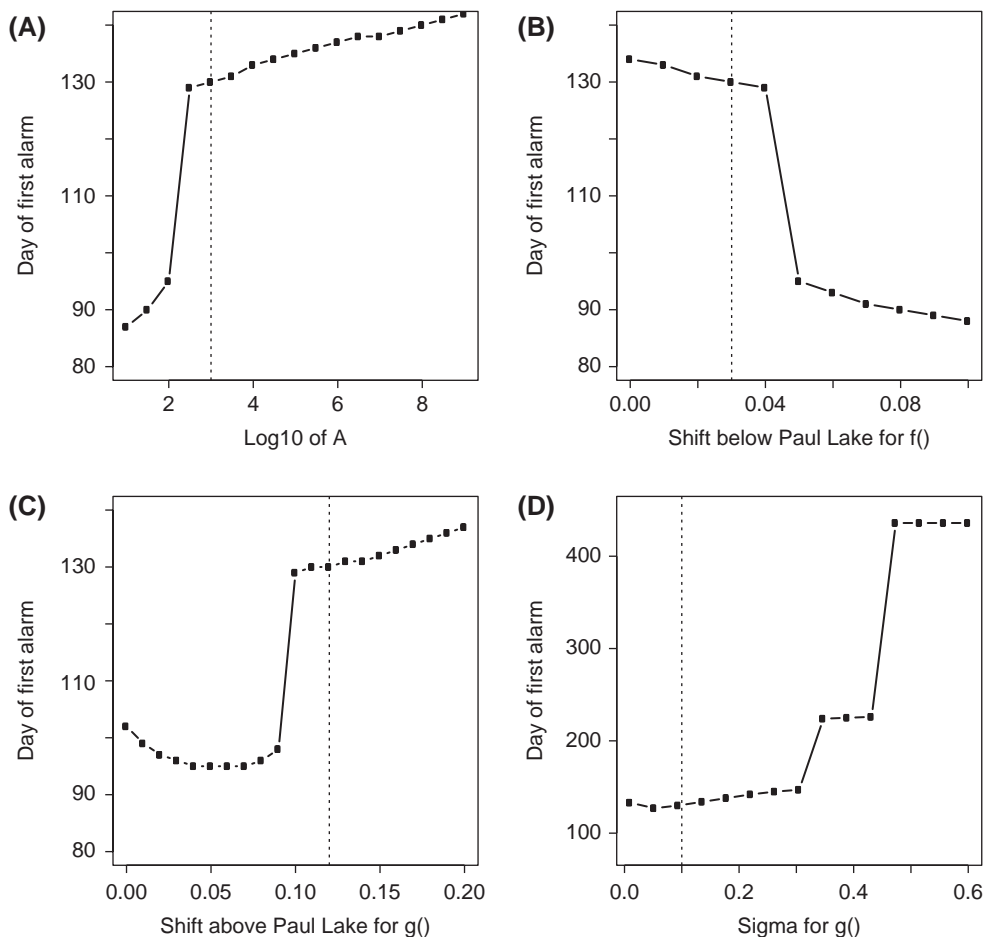


Figure 2. Day of first alarm versus values of the parameters of the quickest detection algorithm. Vertical dashed lines indicate nominal value of each parameter. (A) Day of first alarm versus A , the detection threshold. (B) Day of first alarm versus Δ_f , the shift below the Paul Lake value in $f()$. (C) Day of first alarm versus Δ_g , the shift above the Paul Lake value in $g()$. (D) Day of first alarm versus value of σ used in $g()$.

transition due to the gradual increase in harvest rate of piscivores (Fig. 3A). This behavior is consistent with previous analyses of this model (Carpenter et al. 2008, Brock and Carpenter 2012). Management interventions were simulated by decreasing harvest rate to zero after the first alarm. When the resilience indicator was computed from rolling-window standard deviation of the prey fish, the decrease of harvest rate to zero prevented the collapse of the piscivore biomass and rise of the prey fish biomass (Fig. 3B). In contrast, when the resilience indicator was computed from rolling-window standard deviation of the phytoplankton the decrease of harvest rate to zero failed to prevent the collapse of the piscivore biomass and rise of the prey fish biomass (Fig. 3C).

In the model, the prey fish are directly involved in the nonlinear interaction that creates the critical transition, whereas the phytoplankton are two trophic levels removed from this key interaction (Carpenter et al. 2008). Additional noise terms (added to zooplankton and phytoplankton dynamics) affect the phytoplankton but are unrelated to the variance created by the critical transition. Thus the standard deviation of the prey fish sounds the alarm relatively early in the time series (Table 1). The alarm based on phytoplankton statistics is delayed by about 150 time steps and comes too late to prevent the collapse of the piscivore population.

Discussion

The rolling-window standard deviation combined with the Shiryayev–Roberts statistic provided signals of resilience loss in field data and in models. Sensitivity analysis shows that the timing of alarms depends only weakly on the parameters over broad ranges. Furthermore the method is optimal in the following sense: it minimizes the average detection delay subject to a lower bound on the average time to false positive alarms (Polunchenko and Tartakovsky 2012, pp. 662–663).

The QD method is designed for online management in real time, in which statistics are continually updated as new data are observed, an alarm is issued as soon as a test criterion is exceeded, and management action follows immediately. In this respect, the QD approach differs from the methods employed so far to measure early warnings or resilience of ecosystems (Dakos et al. 2012). The QD procedure is likelihood-based and therefore depends on computing the likelihood for two alternative models. It could easily be adapted for use with more sophisticated and specific time series models for particular ecosystem regime shifts, such as the models used by Carpenter and Pace (1997), Carpenter (2002), Boettiger and Hastings (2012a) and Ives and Dakos (2012).

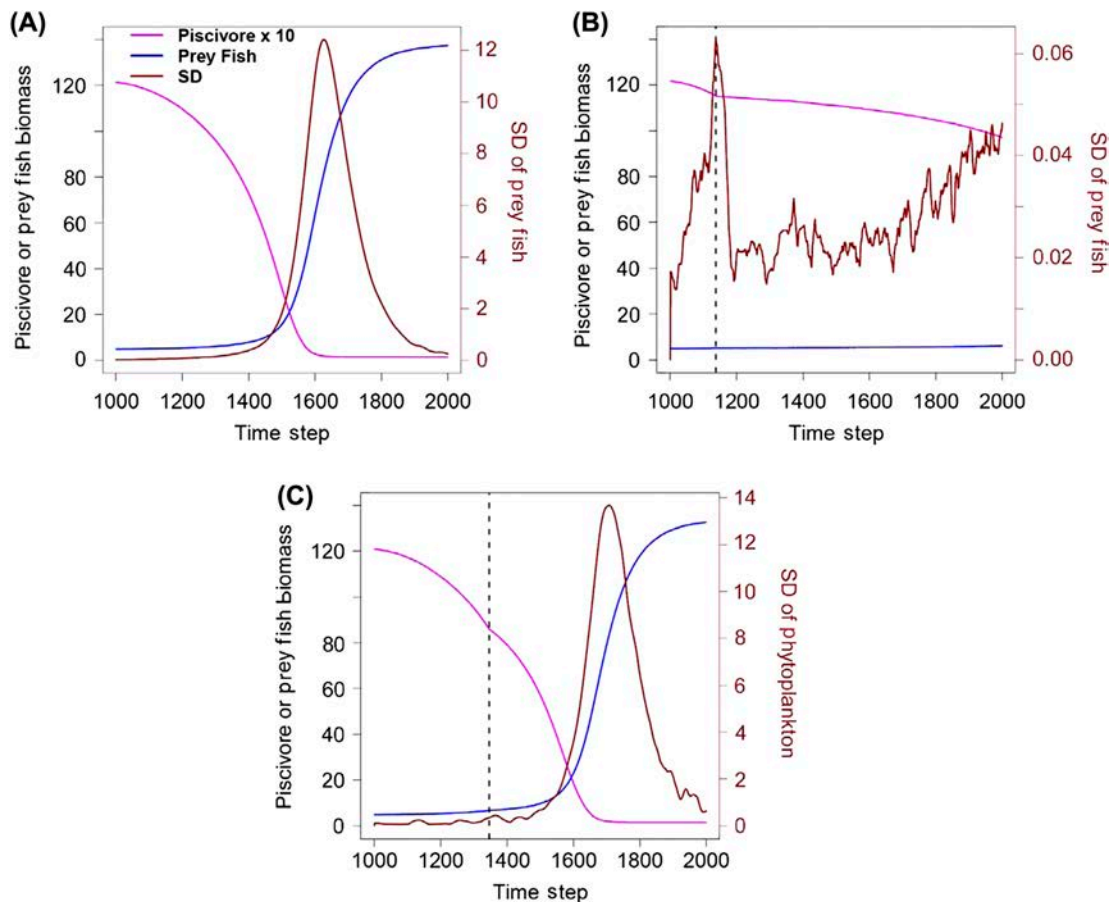


Figure 3. Simulated time series using the food web model. Each panel presents piscivore biomass ($\times 10$), prey fish biomass, and rolling window standard deviation of a resilience indicator. Vertical dashed lines in (B) and (C) show the time step when piscivore harvest drops to zero based on day of first alarm. (A) Baseline situation with no intervention. (B) Case where intervention is triggered by the rising standard deviation of the prey fish. (C) Case where the intervention is triggered by the rising standard deviation of the phytoplankton.

Table 1. Statistics for reversal experiments using two indicators, rolling window standard deviation of prey fish or rolling window standard deviation of phytoplankton. 100 independent simulations of 2000 time steps were computed for each indicator. Mean (with standard deviation in parentheses) of the 100 simulations are reported for day of first alarm, piscivore biomass at step 2000, and planktivore biomass at step 2000. Day of first alarm is measured from the start of the experiment in contiguous summer stratified seasons.

Indicator	Day of first alarm	Piscivore biomass at step 2000	Planktivore biomass at step 2000
Rolling window SD of prey fish	1149 (10)	11.1 (0.2)	5.3 (0.1)
Rolling window SD of phytoplankton	1299 (86)	1.6 (3.7)	109 (46)

Our application of the QD method is different from the original use of the method, which was inspired by problems where the change from $f(x_i)$ to $g(x_i)$ occurs abruptly (Shiryaev 2010). In contrast, the literature on early warning of ecosystem thresholds developed in the context of emerging bifurcations, i.e. where the signal develops gradually through time (Carpenter and Brock 2006). We use the QD method to identify a change in a statistical indicator, not the underlying ecological threshold. Recent work by Ludkovski (2012) and Han (2010) on public health surveillance problems employ similar methods where the signal gradually propagates through space. Since Ludkovski and Han exploit spatial information as well as time series information their work suggests a potentially valuable extension to spatial dynamics in ecology.

One challenge in applying the QD method lies in establishing the distribution of $g(x_i)$ in the absence of information about potential alternate states. Our approach was to modify $f(x_i)$ (which is fitted to initial data) by increasing the mean to account for the expected shift and increasing the variance to account for our uncertainty. While arbitrary, this approach is simple and based on data. Sensitivity analysis shows that outcomes are affected only weakly by the assumptions over a wide range. In other cases, ecological experience with systems in a variety of conditions should allow development of a $g(x_i)$ distribution that is consistent with observed transitions. For example, eutrophication and trophic cascades have been observed in many lakes and provide a basis for estimating $g(x_i)$. In still other cases, models may be used to simulate a variety of alternative distributions and evaluate their efficacy relative to a tolerance for false positives or false negatives in detecting loss of resilience.

In our examples, the first alarm was not a false alarm or a delayed alarm; it occurred before the ecological threshold was reached, and the threshold was in fact crossed. Nonetheless the alarm may come too late, in a situation where managers attempt to avoid crossing the threshold. The model results show that the threshold could be avoided only when the monitored variable was directly involved in the critical transition. When prey fish standard deviation was the monitored variable, alarms occurred early and intervention prevented the shift in the dominant fish. When phytoplankton standard deviation was monitored, alarms occurred later and intervention failed to prevent the regime shift. In this model, the signal of the critical transition is generated in the fish compartments. The phytoplankton biomass is subject to additional sources of variance unrelated to the critical transition, and therefore it takes longer to discern the change in standard deviation and send an alarm. The momentum of ecosystem change leads to collapse of the piscivores even though the harvest rate is reduced sharply. Note, however, that in the analyzed field example phytoplankton

warnings were evident long before the regime shift. In this field experiment the goal was to push the system to the alternate, piscivore-dominated state so we cannot address if the phytoplankton warnings came too late to prevent the transition.

Our findings corroborate the conclusions of previous studies that have considered the prevention of regime shifts in ecosystems. Using a similar model of fish interactions, Biggs et al. (2009) found that piscivore stocks could be preserved by harvest interventions under some conditions but not others. In a model of lake eutrophication, Contamin and Ellison (2009) found a similar result – eutrophication could be prevented by prompt action under certain conditions. Using a model for the North Atlantic thermohaline circulation, Kleinen et al. (2003) found that climate warming due to greenhouse gas emission probably could not be reversed in time to prevent a critical transition in ocean physics. So far all tests of the use of resilience indicators to prevent collapse of ecosystems have employed models. Experimental field tests of the indicators at the scale of whole ecosystems are needed to confirm or refute the possibility of using resilience indicators to prevent regime shifts. Quickest detection offers a flexible method that can easily be adapted for such tests in complex field situations.

The complex relationships of changing environmental drivers, ecosystem responses, thresholds, and the possibility of regime shifts remains an open topic of scientific research. QD is only one tool among many for investigating resilience and the possibility that statistical changes can indicate approaching thresholds of ecosystem change. We are far from having reliable tools for detecting regime shifts for ecosystem management. Not all regime shifts are caused by critical transitions (Carpenter 2003), not all critical transitions have early warnings (Scheffer et al. 2012), and, as shown in our study, not all early warnings come soon enough for management interventions to prevent the regime shift. Surprises will always be with us. Nonetheless, emerging statistical tools have the potential to reveal new insights about some important kinds of ecosystem regime shifts.

Acknowledgements – We thank R. Batt and the referees for helpful comments on the manuscript. R. Batt, T. Cline, J. Coloso, J. Kurzweil, R. Johnson, D. Seekell, L. Smith and G. Wilkinson contributed to the field results reported here. This work was supported by the NSF.

References

- Biggs, R. et al. 2009. Turning back from the brink: Detecting an impending regime shift in time to avert it. – *Proc. Natl Acad. Sci. USA* 106: 826–831.

- Boettiger, C. and Hastings, A. 2012a. Early warning signals and the prosecutor's fallacy. – *Proc. R. Soc. B* 279: 4734–4739.
- Boettiger, C. and Hastings, A. 2012b. Quantifying limits to detection of early warning for critical transitions. – *J. R. Soc. Interface* 9: 2527–2539.
- Brock, W. A. and Carpenter, S. R. 2010. Interacting regime shifts in ecosystems: implication for early warnings. – *Ecol. Monogr.* 80: 353–367.
- Brock, W. A. and Carpenter, S. R. 2012. Early warnings of regime shift when the ecosystem structure is unknown. – *PLoS ONE* 7: e45586.
- Carpenter, S. R. 2002. Ecological futures: building an ecology of the long now. – *Ecology* 83: 2069–2083.
- Carpenter, S. R. 2003. Regime shifts in lake ecosystems: pattern and variation. – Ecology Institute.
- Carpenter, S. R. and Kitchell, J. F. (eds) 1993. Trophic cascades in lakes. – Cambridge Univ. Press.
- Carpenter, S. R. and Pace, M. L. 1997. Dystrophy and eutrophy in lake ecosystems: implications of fluctuating inputs. – *Oikos* 78: 3–14.
- Carpenter, S. R. and Brock, W. A. 2006. Rising variance: a leading indicator of ecological transition. – *Ecol. Lett.* 9: 311–318.
- Carpenter, S. R. et al. 2001. Trophic cascades, nutrients and lake productivity: whole-lake experiments. – *Ecol. Monogr.* 71: 163–186.
- Carpenter, S. R. et al. 2008. Leading indicators of trophic cascades. – *Ecol. Lett.* 11: 128–138.
- Carpenter, S. R. et al. 2011. Early warnings of regime shifts: a whole-ecosystem experiment. – *Science* 332: 1079–1082.
- Contamin, R. and Ellison, A. M. 2009. Indicators of regime shifts in ecological systems: what do we need to know and when do we need to know it. – *Ecol. Appl.* 19: 799–816.
- Dai, L. et al. 2012. Generic indicators for loss of resilience before a tipping point leading to population collapse. – *Science* 336: 1175–1177.
- Dakos, V. et al. 2008. Slowing down as an early warning signal for abrupt climate change. – *Proc. Natl Acad. Sci. USA* 105: 14308–14312.
- 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.
- Drake, J. M. and Griffen, B. D. 2010. Early warning signals of extinction in deteriorating environments. – *Nature* 467: 456–459.
- Han, S. W. 2010. Efficient change detection methods for bio and health care surveillance. – School of Industrial and Systems Engineering, Georgia Inst. of Technology.
- Hastings, A. and Wysham, D. B. 2010. Regime shifts in ecological systems can occur with no warning. – *Ecol. Lett.* 13: 464–472.
- Ives, A. R. and Dakos, V. 2012. Detecting dynamical changes in nonlinear time series using locally linear state-space models. – *Ecosphere* 3: art58.
- Kéfi, S. et al. 2012. Early warning signals also precede non-catastrophic transitions. – *Oikos* 122: 641–648.
- Kleinen, T. et al. 2003. The potential role of spectral properties in detecting thresholds in the Earth system: application to the thermohaline circulation. – *Ocean Dynamics* 53: 53–63.
- Lai, T. L. 1995. Sequential changepoint detection in quality control and dynamical systems. – *J. R. Stat. Soc. B* 57: 613–658.
- Livina, V. N. and Lenton, T. M. 2007. A modified method for detecting incipient bifurcations in a dynamical system. – *Geophys. Res. Lett.* 34: L03712.
- Ludkovski, M. 2012. Bayesian quickest detection in sensor arrays. – *Sequential Anal.* 31: 481–504.
- Pace, M. L. et al. 2013. Zooplankton provide early warnings of a regime shift in a whole-lake manipulation. – *Limnol. Oceanogr.* 58: 525–532.
- Pollak, M. and Tartakovsky, A. G. 2009. Optimality of the Shiryaev–Roberts procedure. – *Stat. Sin.* 19: 1729–1739.
- Polunchenko, A. and Tartakovsky, A. G. 2012. State of the art in sequential change-point detection. – *Meth. Comput. Appl. Probabil.* 14: 649–683.
- Porter, J. H. et al. 2012. Staying afloat in the sensor data deluge. – *Trends Ecol. Evol.* 27: 121–129.
- Scheffer, M. et al. 2009. Early-warning signals for critical transitions. – *Nature* 461: 53–59.
- Scheffer, M. et al. 2012. Anticipating critical transitions. – *Science* 338: 344–348.
- Seekell, D. et al. 2012. Conditional heteroskedasticity forecasts regime shift in a whole-ecosystem experiment. – *Ecosystems* 15: 741–747.
- Seekell, D. A. et al. 2012. Conditional heteroskedasticity as a leading indicator of ecological regime shifts. – *Am. Nat.* 178: 442–451.
- Shiryaev, A. N. 2010. Quickest detection problems: fifty years later. – *Sequential Anal.* 29: 345–385.
- Tartakovsky, A. G. et al. 2006. A novel approach to detection of intrusions in computer networks via adaptive sequential and batch-sequential change-point detection methods. – *IEEE Trans. Signal Processing* 54: 3372–3382.
- Wang, R. et al. 2012. Flickering gives early warning signals of a critical transition to a eutrophic lake state. – *Nature* 492: 419–422.

Supplementary material (available online as Appendix oik-00539 at <www.oikosoffice.lu.se/appendix>). Appendix A1-A2.

Carpenter, S. R., Brock W. A., Cole J. J. and Pace, M. L. 2013. A new approach for rapid detection of approaching thresholds in ecosystem time series. – *Oikos* 000: 000–000.

Appendix A1

Example using autocorrelation

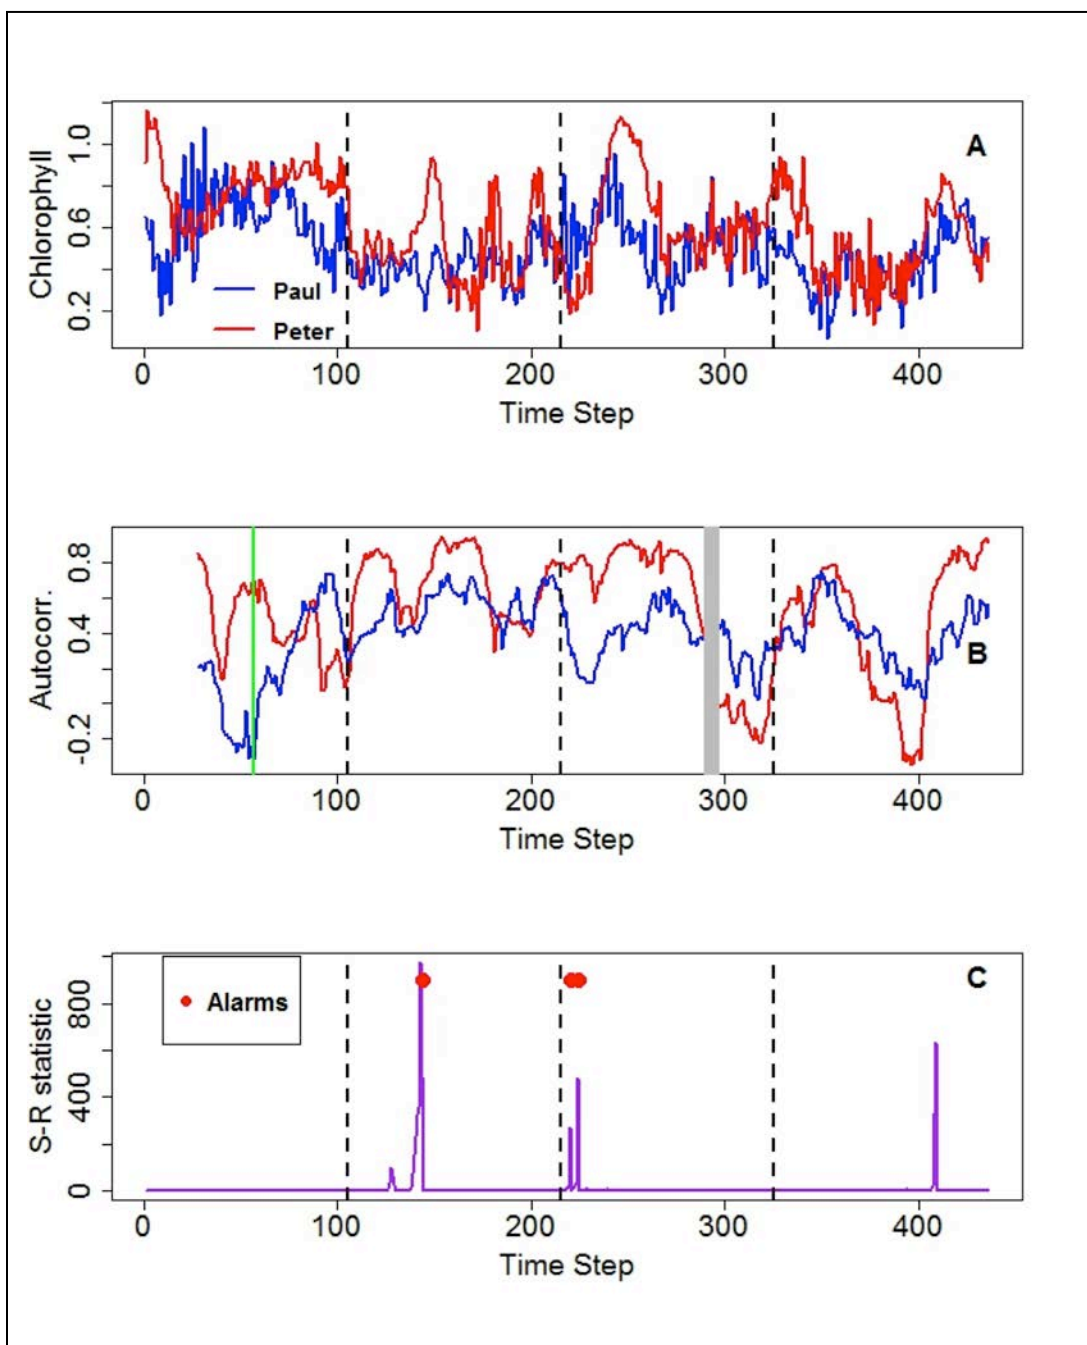
The QD method was applied to time series of lag-1 autocorrelations for the whole-lake experiment computed using the same 28 day rolling windows used for variance in the main text. The detection threshold $A = 1000$ as in the main text. Functions for the original state $f(\alpha_t)$ and new state $g(\alpha_t)$, where α_t is the rolling-window autocorrelation at time step t , were

$$\begin{aligned} f(\alpha_{t,M}) &: N(\mu = \alpha_{t,C}, \sigma = 0.2) \\ g(\alpha_{t,M}) &: N(\mu = 1, \sigma = 0.04) \end{aligned} \tag{A1.1}$$

As in the main text, the subscript M stands for the manipulated lake and C stands for the reference lake. In the original state $f(\alpha_t)$, the mean autocorrelation in the manipulated lake is assumed equal to that of the reference lake, with a standard deviation equal to the observed standard deviation of the autocorrelation of the reference lake. In the altered state $g(\alpha_t)$ the mean autocorrelation of the reference lake is assumed equal to one, the theoretical value for the early warning, and the standard deviation was assumed equal to zero.

Results (Fig. A1) show that autocorrelation remained close to one for much of the second and third years (Fig. A1 panel B). The time step of first alarm was 144 (Fig. A1 panel C), somewhat later than the time step of first alarm for variance which was 130.

Figure A1. Time series from the whole-lake experiment from summer stratified seasons of 2008-2011. Vertical dashed lines denote breaks between years. (A) Log₁₀ of daily chlorophyll concentration in Paul (reference, blue) and Peter (manipulated, red) lakes. (B) Rolling window (28 day) lag-1 autocorrelation of chlorophyll in Paul (reference, blue) and Peter (manipulated, red) lakes. The vertical green line shows the time of first bass addition and the vertical gray stripe shows the approximate time when the regime shift was completed. (C) Shiryayev-Roberts statistic and days of alarms (red dots).



Appendix A2

Food-web model

This appendix presents the equations of the food web model based closely on the presentation by Brock and Carpenter (2012). Parameter values are presented in Carpenter et al. (2008) except where stated otherwise. Dynamics of adult piscivores A follow

$$\frac{dA}{dt} = s_J J - qEA - m_A A + \sigma_A \frac{dW_A}{dt} \quad (\text{A2.1})$$

where s_J is the rate coefficient for maturation of juveniles (J) adults ($s_J = 0.5 \text{ y}^{-1}$); q is catchability; E is angler effort; m_A is the mortality rate coefficient for A ($m_A = 0.5 \text{ y}^{-1}$); and σ_A is the noise magnitude.

Dynamics of juveniles J follow

$$\frac{dJ}{dt} = \left(f_A + \sigma_J \frac{dW_J}{dt} \right) A - c_{JA} JA - \frac{c_{JF} v F J}{h + v + c_{JF} F} - s_J J \quad (\text{A2.2})$$

f_A is fecundity of adults; c_{JA} is a cannibalism coefficient; c_{JF} is predation on juveniles by planktivores; v is the vulnerability coefficient; h is the hiding coefficient.

Planktivore F dynamics follow

$$\frac{dF}{dt} = D_F (F_R - F) - c_{FA} FA + \sigma_F \frac{dW_F}{dt} \quad (\text{A2.3})$$

D_F is an exchange coefficient between the refuge and open water; F_R is the refuge density of planktivores; c_{FA} is consumption coefficient of planktivores by piscivores; σ_F is the noise magnitude.

Herbivore (zooplankton) H dynamics follow

$$\frac{dH}{dt} = D_H (H_R - H) + \alpha c_{HP} HP - c_{HF} HF + \sigma_H \frac{dW_H}{dt} \quad (\text{A2.4})$$

D_H is an exchange coefficient between refuge and open water; H_R is the refuge density of herbivores; α is the conversion coefficient between phytoplankton and herbivores; c_{HP} is a grazing coefficient; c_{HF} is the predation coefficient for planktivory; σ_H is the noise magnitude.

Phytoplankton (P) dynamics follow

$$\frac{dP}{dt} = r_P L \gamma(I_0, P) P - mP - c_{PH} HP + \sigma_P \frac{dW_P}{dt} \quad (\text{A2.5})$$

r_P is a growth coefficient; L is phosphorus loading; γ is the growth function described in Carpenter et al. (2008); m is non-grazing mortality; c_{PH} is the grazing coefficient; σ_P is the noise magnitude.

The growth function γ imposes self-shading on the phytoplankton which imposes density-dependency on their growth (Carpenter et al. 2008). Because surface irradiance I_0 is constant in these simulations, the dynamic impact of γ is stabilizing, through density-dependence of P dynamics. Details of the growth function γ are presented in Carpenter et al. (2008) and are not essential for understanding results presented here.

Results

Brock, W. A. and Carpenter, S. R. 2012. Early warnings of regime shift when the ecosystem structure is unknown. - PLoS ONE 7: e45586.

Carpenter, S. R. et al. 2008. Leading indicators of trophic cascades. - Ecol. Lett. 11: 128-138.