Slowing down as an early warning signal for abrupt climate change

Vasilis Dakos*, Marten Scheffer*†, Egbert H. van Nes*, Victor Brovkin‡§, Vladimir Petoukhov‡, and Hermann Held‡

*Department of Aquatic Ecology and Water Quality Management, Wageningen University, P.O. Box 47, 6700 AA, Wageningen, The Netherlands; and ‡Potsdam Institute for Climate Impact Research, P.O. Box 601203, D-14412 Potsdam, Germany

Edited by Stephen R. Carpenter, University of Wisconsin, Madison, WI, and approved July 16, 2008 (received for review March 11, 2008)

In the Earth's history, periods of relatively stable climate have often been interrupted by sharp transitions to a contrasting state. One explanation for such events of abrupt change is that they happened when the earth system reached a critical tipping point. However, this remains hard to prove for events in the remote past, and it is even more difficult to predict if and when we might reach a tipping point for abrupt climate change in the future. Here, we analyze eight ancient abrupt climate shifts and show that they were all preceded by a characteristic slowing down of the fluctuations starting well before the actual shift. Such slowing down, measured as increased autocorrelation, can be mathematically shown to be a hallmark of tipping points. Therefore, our results imply independent empirical evidence for the idea that past abrupt shifts were associated with the passing of critical thresholds. Because the mechanism causing slowing down is fundamentally inherent to tipping points, it follows that our way to detect slowing down might be used as a universal early warning signal for upcoming catastrophic change. Because tipping points in ecosystems and other complex systems are notoriously hard to predict in other ways, this is a promising perspective.

catastrophic shifts | critical slowing down | autocorrelation | alternative stable states | tipping point

The relative constancy of the climate over the past 10,000 years is exceptional in view of the large variability found in reconstructions of almost all periods before. Particularly noteworthy in the records of past climate dynamics are occasional sharp transitions from one state to another. Such transitions happened at various time scales (1). For instance, ≈34 million years ago the earth changed suddenly from the tropical state in which it had been for hundreds of millions of years to a state with ice caps, a shift known as the greenhouse–icehouse transition (2, 3) (Fig. 1*A*). A prominent feature of the climate cycles that followed is the abrupt termination of most glacial periods (4) (Fig. 1 *C*, *E*, *G*, and *I*). Zooming in on a finer time scale shows that there are sharp shifts too. A well known example is the Younger Dryas period, when, just after the recovery from the last glacial maximum, the climate at Greenland relapsed to very cold conditions for many centuries and then suddenly jumped back to a >10° warmer state (5) (Fig. 1*M*). An even more recent abrupt climate shift is the sudden shift of North Africa from a savanna-like state with scattered lakes to a desert state ≈5,000 years ago (6) (Fig. 1*O*).

Proposed explanations for these and other examples of abrupt climate change usually invoke the existence of thresholds in external conditions where the climate system is particularly sensitive, or even has a tipping point (7), similar to that of a canoe where one leans over too much to one side. In models such tipping points correspond to bifurcations (8) where, at a critical value of a control parameter, an attractor becomes unstable, leading to a shift to an alternative attractor. The underlying mechanism causing such extreme sensitivity at particular thresholds is typically a positive feedback. The earth system is notoriously riddled with such positive feedbacks (9–11). Unfortunately, the explanations for abrupt climatic change in the past remain rather hypothetical because they are difficult to test. Even if the proposed mechanisms seem plausible, our capacity to model these systems accurately is too limited to conclude with reasonable certainty that tipping points are involved. This is particularly worrisome in view of the possibility of hitting on a tipping point as current climate change proceeds. Although most climate scientists would acknowledge that possibility, we are simply unable to predict if and when future climate change might bring us to a critical threshold (1). Even though climate models are rapidly improving, the chances that we will soon be able to predict potential tipping points with sufficient accuracy seem negligible. A similar situation exists in ecology where the existence of thresholds for catastrophic shifts has been shown for a range of systems (12), but prediction of such shifts has remained elusive.

In the face of our limited mechanistic insight it would be invaluable to have another way to find out whether past abrupt climate change was related to the crossing of critical thresholds, and to know whether parts of our current climate system may be approaching such a threshold. A possible clue that we explore here is to use the theoretical finding that, as a rule, dynamical systems become "slow" when a critical point is approached as conditions are gradually changing. In technical terms, the mechanism is that the maximum real part of the eigenvalues of the Jacobian matrix tends to zero as a bifurcation point is approached (13). As a result the dynamical system becomes increasingly slow in recovering from small perturbations (13–15).

Although an ideal way to test whether a system is slowing down (15) would be to study its response to small experimental perturbations, this is obviously of little use for analyzing past climate change. An alternative is to interpret fluctuations in the state of a system as its responds to natural perturbations. Slowing down should then simply be reflected as a decrease in the rates of change in the system, and therefore, as an increase in the short-term autocorrelation in the time series (16). Various authors have elaborated methods to detect slowing down associated with a shift in model-generated time series of the thermohaline circulation (17–19). Kleinen *et al.* (17) analyzed spectral properties, and Held and Kleinen (18) focused on autocorrelation as a statistic to detect slowing down before the transition. Livina and Lenton (19) suggested an approach inspired by a technique for detecting long-term memory in a time series. Despite the interest in this field, so far, no significant signs of slowing down before a shift have been shown on real data.

Author contributions: V.D., M.S., and E.H.v.N. designed research; V.B., V.P., and H.H. contributed new reagents/analytic tools; V.D. analyzed data; and V.D. and M.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

†To whom correspondence should be addressed. E-mail: marten.scheffer@wur.nl.

§Present address: Max Planck Institute for Meteorology, Bundesstrasse 55, 20146 Hamburg, Germany.

This article contains supporting information online at www.pnas.org/cgi/content/full/0802430105/DCSupplemental.

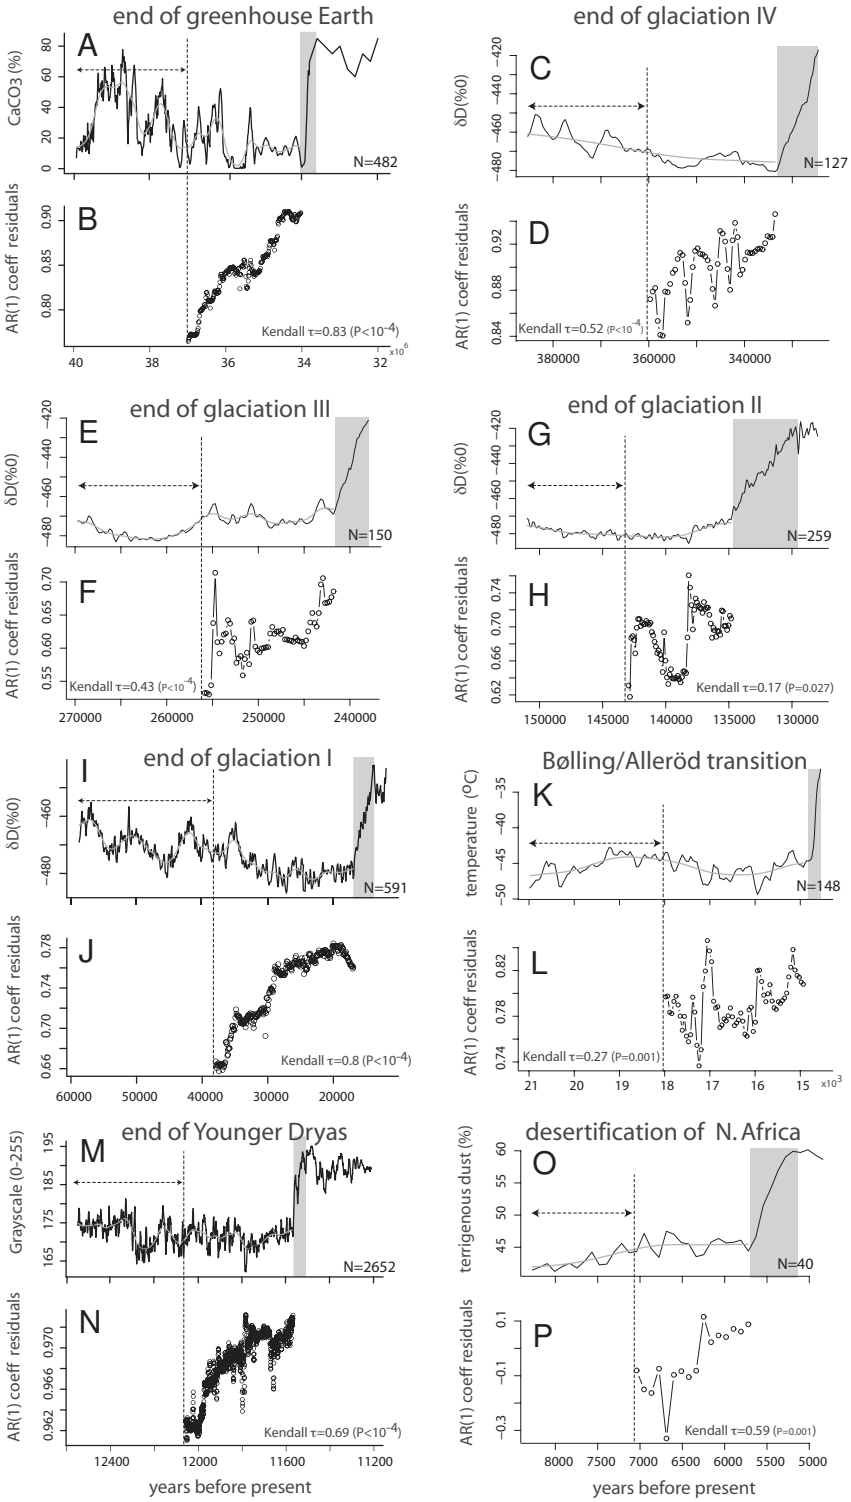

Fig. 1. Eight reconstructed time series of abrupt climate shifts in the past. (*A*) The end of the greenhouse Earth, (*M*) the end of the Younger Dryas, (*K*) the Bølling-Alleröd transition, (*O*) the desertification of North Africa, (*I*) the end of the last glaciation, and (*G*, *E*, and *F*) the ends of earlier glaciations. In all cases the dynamics of the system slow down before the transition, as revealed by an increasing trend in autocorrelation (*B*, *D*, *F*, *H*, *J*, *L*, *N*, and *P*). The gray bands identify transition phases. The arrows mark the width of the moving window used to compute slowness. The smooth gray line through the time series is the Gaussian kernel function used to filter out slow trends. Data in *A* come from tropical Pacific sediment core records, data in *M* are from the Cariaco basin sediment, data in *K* come from the Greenland GISP2 ice core, data in *O* from the sediment core ODP Hole 658C off the west coast of Africa, and data presented in *C*, *E*, *G*, and *I* are from the Antarctica Vostok ice core (additional details are in supporting information (SI) Table S1 and Fig. S1).

Here, we analyze the change in autocorrelation in time series of eight ancient events of abrupt climate change reconstructed from geological records (Fig. 1; see *Methods*) to examine whether the climate system slows down when a critical threshold is approached. Because we are interested in the possibility of using such information as an early warning signal, we used only

data from before the actual transition (Fig. 1, shaded bands) to scan for slowing down. Details of the time series and the identification of the period before the shift can be found in Table S1. Techniques of data processing are described in *Methods*.

Results and Discussion

Evidence for Critical Slowing Down. In all examples of abrupt climate change we analyzed, autocorrelation showed an increase in the period before the shift (Fig. 1 *B*, *D*, *F*, *H*, *J*, *L*, *N*, and *P*), suggesting that these climate systems did indeed slow down before the abrupt change, as expected theoretically for systems approaching a tipping point. All of the trends were significant as measured by the Kendall rank correlation coefficient τ, but the strength of the correlation varied among cases. There was a marked increase in slowing down before the end of the greenhouse Earth (Fig. 1*B*), the end of the Younger Dryas (Fig. 1*N*), and the end of glaciation I (Fig. 1*J*). Autocorrelation moderately increased before the end of glaciation IV, glaciation III, and the desertification of North Africa (Fig. 1 *D*, *F*, and *P*), whereas the end of the Bølling-Alleröd (Fig. 1*L*) and the end of glaciation II (Fig. 1*H*) showed weak signs of slowing down. We explored the likelihood that our method would find such results by chance, that is, without an underlying critical slowing down causing the pattern, by studying the occurrence of trends in computer-generated surrogate time series (see *Methods*). The approach was to generate large numbers of randomized time series with characteristics similar to the analyzed stretches of climate series before episodes of abrupt change, and see in how many cases our analysis would find an increase in autocorrelation by chance. These analyses (see *SI Text*, Table S2, and Fig. S3) indicated that the probability of finding the increase in autocorrelation detected in the data by chance is very low for the three transitions that showed the strongest slowing down (end of greenhouse Earth, end of Younger Dryas, and end of glaciation I). These records have the most detailed data (all >450 data points). The other time series are much less detailed (all <150 data points), and our surrogate data analyses suggest higher probabilities of finding the observed trends by chance in those cases. The lower number of points in some of the series obviously makes the results less reliable, not only because of the small number of points *per se*, but also because the resolution can be insufficient to capture the short-term autocorrelation. This is especially so in the case of the desertification of North Africa, where the points are spaced almost a century apart, which may well be too short to capture the interactive dynamics of vegetation and monsoon supposed to drive the dynamics. The scarcity of points in the record ($N_p = 30$ before the transition) results in residuals of alternating positive and negative values and in estimates of the autoregressive coefficient α_1 that show a negative autocorrelation (Fig. 1*P*).

To weigh the combined uncertainties, and look at the overall picture, we computed the probability of finding the complete set of *P* values by chance, by using Fisher's combined probability approach. This combined probability appears to be very small ($P < 0.003$) irrespective of the approach taken to generate surrogate data (Table S2).

Robustness of Results to the Choice of Methods. The results obviously depend on choices made in the data processing (see *Methods*). Two important parameters are the bandwidth used in the function for filtering and the size of the sliding window used to compute the autocorrelation. We performed an extensive analysis of the sensitivity of outcomes to the choice of these parameters for our three longest time series. The results indicate that the observed increase in autocorrelation before the climate shifts is a rather robust outcome. Actually, this analysis shows that we could have obtained more significant trends by tailoring these parameters for the specific series (Fig. S4). We also explored whether interpolation used to generate equidistant data for the time series analyses might have caused spurious trends in autocorrelation (*SI Text*, Fig. S1, Fig. S2). Estimates of autocorrelation on the noninterpolated data gave approximately similar results (Table S3) (see also *SI Text*).

Comparison to Model Results. Approaching the problem from a different angle, to check whether the theoretically predicted critical slowing down may indeed be expected to be visible from autoregressive coefficients in climate data, we also used our methods to analyze simulation results from climate models that were slowly driven across a known threshold (Fig. 2). The models deal with three quite different systems: the North African paleo-monsoon system, the thermo-haline circulation, and the earth temperature as affected by the ice-albedo feedback. Model details and references are given in *Methods* and in the *SI Text*. In all cases our indicator picked up an increase in slowness, comparable to that found in the geological records. Also, the results of bootstrap analyses and sensitivity analyses applied to model results are comparable to those from our climate datasets (Fig. S3, Fig. S4, Table S2). This lends further support to the idea that the patterns detected in the data do indeed correspond to critical slowing down as predicted by the theory.

Perspectives. It may seem rather surprising that all cases of sharp climate shifts we analyzed were announced well before they happened by changes in the pattern of fluctuations. Indeed, our bootstrap analysis shows that approximately half of the positive trends in autocorrelation may well have arisen by chance (the desertification of North Africa, the Bølling-Alleröd transition and the end of glaciations II and III). Nonetheless, our analyses also show that the combined probability of finding these trends is extremely low. Furthermore, the close similarity to what can be shown in climate models suggests that the patterns in the data may indeed represent the slowing down of a system approaching a tipping point.

Our results have profound implications for climate science. So far, support for the idea that tipping points can be the explanation for dramatic climatic shifts in the past has been based on models of specific mechanisms. Although compelling cases have been built, there is always considerable uncertainty because it is simply very difficult to prove what had been the mechanism behind such events in the far past. The slowing down that our analysis suggests does not point to any specific mechanism. Rather, it is a universal property of systems approaching a tipping point. Therefore, it represents an independent line of evidence, complementing model-based approaches, suggesting that tipping points exist in the climate system. Clearly, this is an important insight because it implies that, in principle, internal feedback can propel the climate system through an episode of rapid change once a critical threshold is reached.

Obviously, detection of critical slowing down has two faces. In hindsight it may help to tease out whether past dynamics may be explained by the existence of critical thresholds. With respect to predicting future climate change, it may give us an indication of whether we are entering a situation in which the parts of the earth system may amplify rather than buffer human-induced climate change. Clearly, there are challenges and limitations. Long time series of sufficient quality are needed, and resolution needs to be sufficient to capture the characteristic time scale of the internal dynamics of the system. Similarly, good detrending is challenging but critically important, because unfiltered trends may lead to patterns in autocorrelation that are not related to the system's dynamical response to perturbations we wish to probe. An important fundamental limitation we should keep in mind is that slowing down will only occur if the system is moving gradually toward a threshold. Therefore, transitions caused by a sudden large disturbance without a preceding gradual loss of

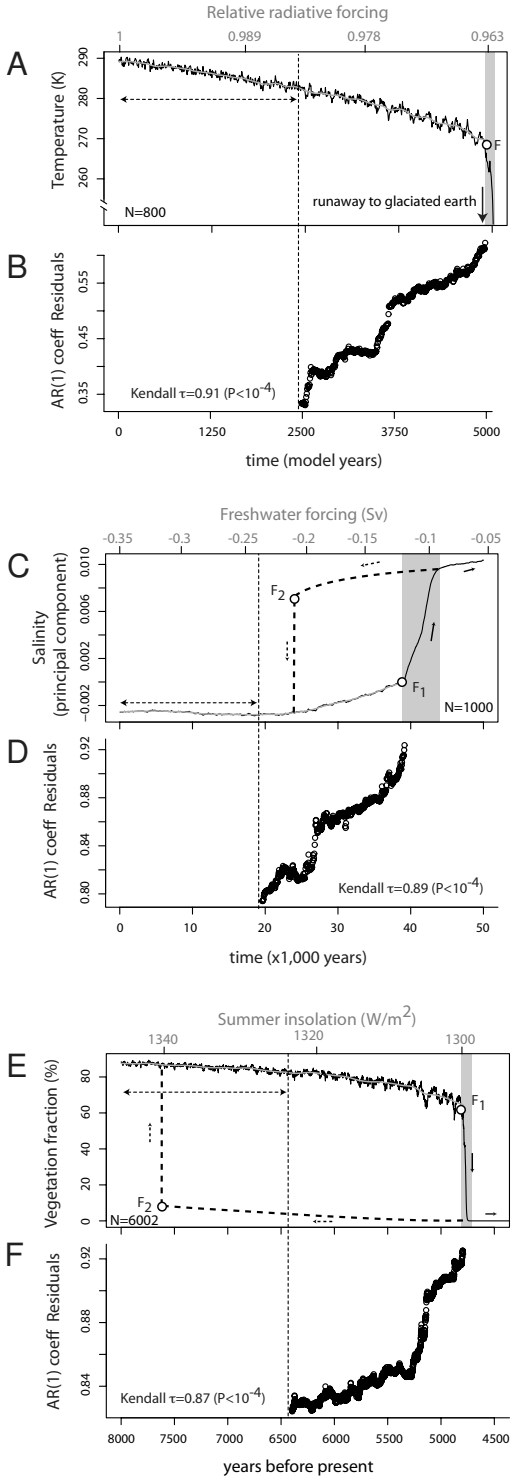

Fig. 2. Three simulated abrupt climate transitions. Transition to an icehouse Earth (*A*), collapse of the thermo-haline circulation (*C*), and desertification of North Africa (*E*) (see *SI Text* for details on simulations). As in the reconstructed real dynamics, the transition is preceded by slowing down as revealed by increased autocorrelation (*B*, *D*, and *F*). The gray bands identify the transition phases. The arrows mark the width of the sliding window used to compute slowness. The smooth gray line through the time series is the Gaussian kernel function used to filter out slow trends. All models pass a fold bifurcation *F* as a control parameter is slowly changing (relative radiation, freshwater forcing, and insolation, respectively). In the case of the ocean circulation and desertification model (*C* and *E*), there are also alternative attractors present implying hysteresis (dashed line), if the change in the control variable would be reversed on the shift. Points F_1 and F_2 are saddle-node bifurcation points.

resilience will not be announced by slowing down. Certainly, current trends in atmospheric carbon are rather fast compared with the dynamics of ice caps and ocean heat contents, and fluctuations of such variables may therefore not show detectable slowing down on century scales. By contrast, slowing down could possibly be detected in faster subsystems that might have tipping points such as regional atmospheric circulation patterns. In view of our current inability to predict potential abrupt climate shifts (1), having slowing down as a clue for detecting whether such parts of the climate system may be approaching a threshold is a marked step forward in projecting future climatic changes.

Putting our results in an even wider perspective, it is important that slowing down is a universal property of systems approaching a tipping point. This implies that our techniques might in principle be used to construct operational early warning systems for critical transitions in a wider range of complex systems where tipping points are suspected to exist, ranging from disease dynamics and physiology to social and ecological systems.

Methods

Data Sources. We used examples of climatic transitions that have been widely interpreted as significant shifts in the climate record and for which underlying positive feedback has been suggested as mechanistic explanation. We have not preselected the examples on the basis of preliminary results from our own analyses. The time series used represent climate data proxies derived from different sources. All were downloaded from the World Data Center for Paleoclimatology, National Geophysical Data Center, Boulder, Colorado (http://www.ncdc.noaa.gov/paleo/data.html). The terrigenous dust record was accessed from the Lamont Doherty Earth Observatory Institute of Columbia University, New York (http://www.ldeo.columbia.edu/~peter/Resources/data.html). Full details on the data records used are given in Table S1.

Data Selection. We used only points in the record that correspond to the period before the transition (Table S1). The exact transition points were determined by eye and all were approximately equal to those cited in the original articles where the records appeared. We have chosen the transition points conservatively, in the sense that we avoided including points that were part of the transition itself. This is important, because, due to increased serial correlation, inclusion of such points would bias the estimate of our slowing-down indicator. In a few cases, double values for the same date occurred in the original files. Those were averaged to provide a single value for each chronology.

Interpolation. We used linear interpolation to transform the climate records to time series with equidistant data. This allows us to use the time series analysis approaches suggested earlier for detecting slowing down (17, 18, 20) on real reconstructed climate records.

Detrending. To filter out long trends and to achieve stationarity we subtracted a Gaussian kernel smoothing function from the data and used the remaining residuals for the estimation of the autoregressive coefficient at lag 1. We chose a bandwidth in such a way that we do not overfit while still removing the long-term trends visible in the records. The same treatment was applied also to the simulated time series and the original records without interpolated points (see Table S3). Fig. S2 shows the interpolated, filtered time series and the resulting residual time series of Fig. 1 *A*, *M*, and *I* for visual inspection.

Autocorrelation. The autocorrelation at lag 1 was computed by fitting an autoregressive model of order 1 (AR1 model of the form $x_{t+1} = \alpha_1 x_t + \varepsilon_t$, by an ordinary least-squares (OLS) fitting method) applied on the data points within a sliding window of fixed size up to the transition point. In each case we took a sliding window of half the size of the interpolated time series. We tested for evidence of slowing down by estimating the nonparametric Kendal rank-correlation τ statistic on the estimates of the autoregressive coefficients α_1 (details in *SI Text*).

Model-Generated Time Series. We used a stochastic 1D energy balance climate model forced by relative incoming radiation to simulate data of ocean temperature that reflect a transition to an icehouse earth (21) (Fig. 2*A*). The thermo-haline circulation dynamics is generated by the CLIMBER-2 climate model of intermediate complexity (Fig. 2*C*). The data series on desertification in Western North Africa (Fig. 2*E*) was produced by using a stochastic version

of the climate box model (22) forced by reconstructed solar irradiance and atmospheric CO_2 concentration. See *SI Text* for the model details.

Surrogate Data. For each time series we tested the likelihood of obtaining our computed trend statistics (Kendall's τ rank correlation) by chance by using 1,000 surrogate time series of the same length as the filtered simulated and real data in three different ways. First, we bootstrapped our datasets by shuffling the original residual time series and picking data with replacement to generate surrogate records of similar distribution (mean and variance). Second, we produced a surrogate time series that had the same Fourier spectrum and amplitudes as the original sets (23). Last, we created surrogate datasets produced by an autoregressive model of order 1 with the same variance, mean and autocorrelation at lag 1 with the residuals time series starting from the same initial value as in the original series (23). For each surrogate set, we computed the trend detection statistic. We then calculated the probability that our estimates of the trend statistic would be observed by chance as the fraction of the 1,000 surrogate series scoring the same value or a higher one. The probability distributions for the model and data trend statistic as well as details on how we produced the surrogate sets are summarized in Table S2 and Fig. S3.

ACKNOWLEDGMENTS. We thank the IPL Institute for providing the opportunity for interdisciplinary cooperation, Steve Carpenter and William Brock for insightful discussions on early warning signals in time series, Cajo ter Braak for statistical advice, and the two anonymous reviewers whose comments greatly improved the manuscript. This work was funded by the Netherlands Organization for Scientific Research (V.D.).

1. Alley RB, *et al.* (2003) Abrupt climate change. *Science* 299:2005–2010.
2. Tripati A, Backman J, Elderfield H, Ferretti P (2005) Eocene bipolar glaciation associated with global carbon cycle changes. *Nature* 436:341–346.
3. Kump LR (2005) Palaeoclimate—Foreshadowing the glacial era. *Nature* 436:333–334.
4. Petit JR, *et al.* (1999) Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature* 399:429–436.
5. Clark PU, Pisias NG, Stocker TF, Weaver AJ (2002) The role of the thermohaline circulation in abrupt climate change. *Nature* 415:863–869.
6. deMenocal P, Ortiz J, Guilderson T, Sarnthein M (2000) Coherent high- and low-latitude climate variability during the holocene warm period. *Science* 288:2198–2202.
7. Lenton TM, *et al.* (2008) Tipping elements in the Earth's climate system. *Proc Natl Acad Sci USA* 105:1786–1793.
8. Kuznetsov YA (1995) *Elements of Applied Bifurcation Theory* (Springer, New York).
9. Lawton J (2001) Earth system science. *Science* 292:1965.
10. Woodwell GM, *et al.* (1998) A simulation model to illustrate feedbacks among resource consumption, production, and factors of production in ecological-economic systems. *Climatic Change* 40:495–518.
11. Rial JA, *et al.* (2004) Nonlinearities, feedbacks and critical thresholds within the earth's climate system. *Climatic Change* 65:11–38.
12. Scheffer M, Carpenter SR, Foley JA, Folke C, Walker B (2001) Catastrophic shifts in ecosystems. *Nature* 413:591–596.
13. Strogatz SH (1994) *Nonlinear Dynamics and Chaos with Applications to Physics, Biology, Chemistry and Engineering* (Westview, Reading, MA).
14. Wissel C (1984) A universal law of the characteristic return time near thresholds. *Oecologia* 65:101–107.
15. van Nes EH, Scheffer M (2007) Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *Am Nat* 169:738–747.
16. Ives AR (1995) Measuring resilience in stochastic-systems. *Ecol Monogr* 65:217–233.
17. Held H, Kleinen T (2004) Detection of climate system bifurcations by degenerate fingerprinting. *Geophys Res Lett* 31:L23207.
18. 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 Dynam* 53:53–63.
19. Livina VN, Lenton TM (2007) A modified method for detecting incipient bifurcations in a dynamical system. *Geophys Res Lett* 34:L23207.
20. Ives AR, Dennis B, Cottingham KL, Carpenter SR (2003) Estimating community stability and ecological interactions from time-series data. *Ecol Monogr* 73:303–330.
21. Fraedrich K (1978) Structural and stochastic analysis of a zero-dimensional climate system. *Q J R Meteor Soc* 104:461–474.
22. Brovkin V, Claussen M, Petoukhov V, Ganopolski A (1998) On the stability of the atmosphere-vegetation system in the Sahara/Sahel region. *J Geophys Res* 103:31613–31624.
23. Theiler J, Eubank S, Longtin A, Galdrikian B, Farmer JD (1992) Testing for nonlinearity in time-series—The method of surrogate data. *Physica D* 58:77–94.

Supporting Information

Dakos *et al.* 10.1073/pnas.0802430105

SI Text

Derivation of Model-Simulated Data. Our simulated data presented in Fig. 2 come from three climate models of different complexity.

(*i*) We used a simple one-dimensional climate model (1, 2) to simulate a transition from a greenhouse to an icehouse Earth (Fig. 2*A*). The model has temperature, T, as the only state variable that represents the average temperature of an ocean on a spherical planet subjected to radiative heating (2) according to the equation:

$$\frac{dT}{dt} = \frac{1}{c}\left\{ -\varepsilon\sigma T^4 + \frac{1}{4}\mu I_0 bT + \frac{1}{4}\mu I_0(1-a)\right\}$$
$$\text{with } a_p = a - bT \qquad \textbf{[S1]}$$

where ε is effective emissivity, μ is relative intensity of solar radiation, I_0 is solar irradiance, c is a constant thermal inertia, and a_p is the planetary albedo. Parameters a and b define a linear feedback between ice and albedo variability and temperature. In this simple climate system, there is one internal equilibrium of nonglacial conditions, which, when I_0 drops below a certain threshold, there is a runaway effect to ice climate through a fold bifurcation.

We extended the deterministic skeleton of the model by including a stochastic term following the general form of a stochastic differential equation:

$$dx = f(x,\theta)dt + \sigma(x)dW, \qquad \textbf{[S2]}$$

where x is the state variable, f is the deterministic part of the model that depends on the control parameter θ, and σ scales the amount of noise that is introduced in the model with dW, a Wiener process. In this climate model, we used as control parameter the relative radiation μ. We produced a time series by decreasing the control parameter μ linearly with time from 1 to 0.9524, allowing a transition from a warm to a cold climate. We used σ equal to 0.003 (applied multiplicatively on the state variable) and all of the rest of the parameter values as they appear in ref. 2. We changed the original time scale of the model (= 1 sec) by rescaling time with a factor of $\delta = 20 \times 10^6$ (new time scale = 0.6342 years). Simulations were performed in MATLAB v.7.1.0246 by using an Euler–Murayama method to solve the stochastic equation with Ito calculus.

(*ii*) The thermo-haline circulation model simulation presented here is produced from the CLIMBER-2 model (3, 4), which is a coupled climate model of intermediate complexity. The ocean component originates from the module by Wright and Stocker (5). A freshwater forcing at 44° northern latitude is applied; the average forcing is superimposed with a Gaussian white noise time series. The 50,000 years transient run sees a linear increase in atmospheric CO_2 from 280 ppm to 800 ppm, implying an increased average freshwater forcing.

(*iii*) The deterministic climate model of the desertification of North Africa (Fig. 2*C*) (6) was extended by accounting for the synoptic component w_{syn} of vertical velocity w at the top of the planetary boundary layer:

$$w = w_m + w_h + w_{\text{syn}}$$
$$w_{\text{syn}} = \frac{K_T}{H_0}\left(k_{ts}^w \frac{\max(0, T_L - T_{cr})}{T_L^0 - T_{cr}} + k_{sl}^w \frac{T_L - T_B}{T_L^0 - T_B^0}\right)(1 + \xi(0,\sigma_w)) \qquad \textbf{[S3]}$$

and the synoptic component U_{syn} of the Hadley circulation potential U:

$$U = U_0 + U_{\text{syn}}$$
$$U_{\text{syn}} = k_{\text{syn}}^U U_0\left(1 + \frac{T_L - T_B}{T_L^0 - T_B^0}\xi(0,\sigma_U)\right) \qquad \textbf{[S4]}$$

which allows for the contribution from the synoptic-scale baroclinic and barotropic atmospheric eddies with characteristic time scales from 2 to 10 days. w_m is the vertical velocity in the mean monsoon circulation and w_h is the vertical velocity associated with the mean Hadley circulation, U_0 is the mean Hadley circulation potential, T_B and T_L are surface air temperature at the southern box boundary and over land, respectively, T_B^0 and T_L^0 are their reference values, K_T is a vertical macroeddy diffusion coefficient in the free troposphere, H_0 is a scale height for the atmospheric density, $\xi(0, \sigma_w)$ and $\xi(0, \sigma_U)$ are normally distributed stochastic variables with zero mean and variances σ_w and σ_U, respectively, and k_{ts}^w, k_{sl}^w and k_{syn}^U are model parameters that reflect partial contributions from the corresponding physical processes.

The terms $\frac{K_T}{H_0}k_{ts}^w\frac{\max(0,T_L - T_{cr})}{T_L^0 - T_{cr}}(1 + \xi(0,\sigma_w))$ and $\frac{K_T}{H_0}k_{sl}^w\frac{T_L - T_B}{T_L^0 - T_B^0}(1 + \xi(0,\sigma_w))$ in Eq. **S3** describe the components of the synoptic-scale vertical velocity perturbation attributed to tropical storms and squall lines, respectively. These parameterizations assume that tropical storms form when the temperature exceeds a critical threshold T_{cr} [assumed to be 26°C (7)], whereas the squall lines are mainly generated because of the lower troposphere wind shear in the African Easterly Jet associated with a temperature gradient $T_L - T_B$ between Sahara and the Gulf of Guinea (8). Parameters k_{ts}^w and k_{sl}^w were assigned 0.2 and 0.8, respectively, which reflects partial contributions to the synoptic-scale variability from the tropical storms and squall lines based on the empirical data (9, 10). The value of the variance σ_w was assigned 0.1 (11).

The synoptic term of the Hadley circulation potential (Eq. **S4**) includes a contribution from the synoptic variability, $k_{\text{syn}}^U U_0$, due to the synoptic-scale perturbations of the zonally averaged wind, and from the term associated with the local fluctuations of the Hadley circulation, which is assumed to be proportional to the local horizontal temperature gradient, $k_{\text{syn}}^U U_0\frac{T_L - T_B}{T_L^0 - T_B^0}\xi(0,\sigma_U)$. Parameters k_{syn}^U and σ_U were set equal to 0.05 and 0.1, respectively, based on empirical data (11).

Derivation of Paleoclimate Proxy Data. Because we were interested in measuring slowing down before the transition, we restricted our analysis to the period just before the transition in both simulated and proxy records. The exact parts of the original time series that we selected for our analysis, together with the size of the original record and data sources, are presented in supporting information (SI) Table S1. Because the exact selection of the part of the record is critical for the outcome of our analysis, we were careful to avoid points that were part of the transition. Because of increased serial correlation as the transition trend begins, including such points would bias the estimate of the AR(1) coefficient.

Data Analyses: Interpolation, Detrending, and Estimation of Autocorrelation at Lag 1. We applied the same analyses both to the simulated data and the real paleoclimate proxy records. Because the available paleoclimate data were of unequal density, we used linear interpolation to make our records equidistant (Table S3). However, as indicated in the main text, interpolation can create spurious trends in autocorrelation. A positive trend in autocorrelation could occur as an artifact of interpolation if the density of points would decrease toward the shift (and, hence, the role of interpolation would increase). Therefore, we checked the evolution of the time intervals in the original records and compared them with the equidistant time intervals of the interpolated records (Fig. S1). In general, the time intervals of the interpolated datasets were rather similar to the time intervals in the original time series close to the transition. Only in Fig. S1*a* (the end of greenhouse Earth), did the time between subsequent data points decrease toward the shift. However, this happened at the very end, and can therefore not be responsible for the long-term increasing autocorrelation trend detected. In any case, as shown in the next section, we analyzed the sensitivity of our results to interpolation systematically for all time series.

We removed slow trends in the original records by applying a Gaussian kernel smoothing function [based on the Nadaraya–Watson kernel regression estimate (12)] over the interpolated record before the transition and subtracted it from the interpolated record to obtain the residual time series (Fig. S2). The choice of the size of the bandwidth is important in this process. We picked bandwidths such that we do not overfit our data but yet filter out the slower trends in the records.

In ref. 13 changes in power spectra were used as an indicator for the proximity to thresholds, tested in a 1D model of the thermohaline circulation of the North Atlantic. The spectrum is equivalent to the full autocorrelation function. Here, according to ref. 14, the spatial dynamics become degenerate at the transition, leading to the observation of the critical mode in arbitrary generic 1D time series. By assuming time-scale separation at the bifurcation, we can use only the first entry of the autocorrelation function, that is, the lag-1 information; in that sense, the current method is more parsimonious. To calculate the autocorrelation at lag 1, which is an estimator of the slowing down of the system (15), we fitted an autoregressive model of order 1 (AR1) on data that are included within a sliding window of half the size of the record before the transition. The AR1 *ansatz* (14) is of the form $x_{t+1} = \alpha_1 x_t + \varepsilon_t$, fitted by an ordinary least-squares method (OLS) with Gaussian random error ε_t. Note that we calculated no intercept, because we are fitting the AR1 model on the detrended residuals with mean zero. Although there have been modifications to the AR1 *ansatz* (16) where the authors used detrended fluctuation analysis (DFA), we used the degenerate fingerprinting approach (14) because of its most direct relation to generic bifurcations and straightforward applicability.

Finally, to determine the evolution of the AR1 estimates before the transition we used the nonparametric Kendall τ rank correlation coefficient to check against the null hypothesis of randomness for a sequence of measurements against time (17).

All analyses were implemented in MATLAB v7.1.0246 (Mathworks) and in Rv2.4.1 (R project for Statistical Computing). Specifically, we used (*i*) for the linear interpolation, the function *interp1* (MATLAB); (*ii*) for the detrending of the records, the function *ksmooth* (R); (*iii*) for the estimates of the autoregressive coefficients, the function *ar.ols* (R); and (*iv*) for the calculation of the trend statistic, the function *cor.test* (R) for the Kendall τ correlation statistics together with the *P* values (two-tailed with $\alpha = 0.05$).

Effect of Interpolating on the Results. We also explored the estimates of our trend statistic on the original records without interpolating missing values. Obviously, working with nonequidistant data violates the basic assumptions behind time series analysis. However, we pursued it only as an extra check on the robustness of our results. Thus, we treated the original time series as equidistant, we removed the slow trends (using the same bandwidth for the Gaussian filter as we did in their interpolated counterparts), and we estimated the autoregressive coefficients at lag 1 within a sliding window of half the size of the time series. In all eight cases, our positive trend from the interpolated records are similar to those from the original time series (Table S3). In all cases the trend statistic was of the same order of magnitude, and although there were three cases where interpolated records yielded a stronger increase in the AR(1) coefficient than the noninterpolated ones (end of greenhouse Earth, end of Younger Dryas, and end of glaciation II), there were three other cases where the opposite was observed (end of glaciation I, III, and Bølling–Alleröd) and one in which there was no real difference (glaciation IV).

Analysis of Surrogate Time Series. To test for the likelihood of obtaining estimates of trend statistics by randomness, we created surrogate time series by three different ways.

(*i*) We bootstrapped our datasets by reshuffling the order of the detrended original time series and by picking data with replacement to generate surrogate records of similar probability distribution (mean and variance) (18) (H_0 1).

(*ii*) We produced surrogate time series with the same autocorrelations and the same probability distribution as the data, to test against the H_0 hypothesis that our datasets are a realization of a Gaussian linear stochastic process (19, 20) (H_0 2). We did this by replicating data of the same Fourier spectrum and amplitudes as of the original set using the MATLAB function *generate_iAAFT* (21).

(*iii*) To test against the H_0 hypothesis that the data are produced by a colored-noise process with similar variance, mean, and autocorrelation at lag 1 with the original detrended time series (22) (H_0 3), we generated surrogate sets by an AR1 model $x_{t+1} = \alpha_1 x_t + \alpha_0 + \sigma\varepsilon_t$, where $\alpha_1 = A(1)$, $\sigma^2 = v(1 - \alpha_1^2)$, $\alpha_0 = \mu(1 - \alpha_1)$, with v the variance, μ the mean, $A(1)$ the autocorrelation at lag 1 from the residual time series (estimated by using function *acf* as implemented in *R*), and σ a scaling factor for the Gaussian random error ε_t.

We estimated the probability that our estimates of the trend statistic would be observed by chance as the fraction of the 1,000 surrogate series scoring the same value or a higher one. Specifically for the Kendall τ, we estimated this probability as the number of cases in which the statistic was equal to or higher than the estimate of the original record, $P(\tau \geq \tau^*)$. We also estimated the combined probability for observing the trend statistic estimate in each the H_0 hypotheses test by chance. For this, we used the Fisher's combined probability test (23) to estimate the X^2 statistic, given by:

$$X^2_{2k} = -2\sum_{i=1}^{k} \ln(p_i) \qquad \text{[S5]}$$

where k is the amount of tests (here, $k = 8$) and P is the probability estimated for each H_0 hypotheses test (Table S2). The combined probability for the X^2 statistic was given by a χ^2 distribution with $2k$ degrees of freedom.

The probability estimates for the model and data trend statistic under the three different H_0 hypotheses are shown in Table S2. The probability of, by chance, acquiring a similar trend estimate as in the original record differs from case to case. In the case of the models, the probabilities were consistently very low ($P < 0.05$). Similarly low probabilities were estimated in the records of the transitions of the greenhouse Earth, the Younger

Dryas and the glaciation I (Fig. S3). In the shorter time series the probabilities of finding the observed trends by chance is much higher. Nonetheless the combined probability of finding positive trends in all eight data series is obviously very low (lower row Table S2).

Robustness Against Choice of Window Size and Filtering Resolution. The results of our analyses are obviously influenced by the standard deviation (defined by bandwidth size) used in the kernel function for filtering and the size of the sliding window used to compute autocorrelation. In the latter there is a trade-off between time resolution and reliability of the estimate. Smaller windows allow one to track short-term changes in autocorrelation. However, the small number of data points in the window makes the estimate of autocorrelation less reliable. The filtering poses another trade-off. A too-wide filter does not remove slow trends that may lead to spurious autocorrelation. Especially, at the ends of the time series the deviation becomes obvious if a too-wide kernel size is used. A too-narrow filter removes the short-term fluctuations that we intend to study for signs of slowing down. A systematic sensitivity analysis for our three longest time series and the model results indicate that the results are quite robust, and that actually we could have obtained more significant trends by tuning the parameters for the specific series (Fig. S4).

1. Fraedrich K (1979) Catastrophes and resilience of a zero-dimensional climate system with ice-albedo and greenhouse feedback. *Q J R Meteorol Soc* 105:147–167.
2. Fraedrich K (1978) Structural and stochastic analysis of a zero-dimensional climate system. *Q J R Meteorol Soc* 104:461–474.
3. Ganopolski A, *et al.* (2001) CLIMBER-2: A climate system model of intermediate complexity. Part II: model sensitivity. *Clim Dyn* 17:735–751.
4. Petoukhov V, *et al.* (2000) CLIMBER-2: A climate system model of intermediate complexity. Part I: model description and performance for present climate. *Clim Dyn* 16:1.
5. Stocker TF, Wright DG, Mysak LA (1992) A zonally averaged, coupled ocean atmosphere model for paleoclimate studies. *J Clim* 5:773–797.
6. Brovkin V, Claussen M, Petoukhov V, Ganopolski A (1998) On the stability of the atmosphere-vegetation system in the Sahara/Sahel region. *J Geophys Res* 103:31613–31624.
7. Gray WM (1968) Global view of origin of tropical disturbances and storms. *Mon Weather Rev* 96:669–679.
8. Cook KH (1999) Generation of the African easterly jet and its role in determining West African precipitation. *J Clim* 12:1165–1184.
9. Grist JP, Nicholson SE (2001) A study of the dynamic factors influencing the rainfall variability in the West African Sahel. *J Clim* 14:1337–1359.
10. Joseph E, *et al.* (2003) Paper presented at the First International GPM GV Requirements Workshop, Cosenors House, Abington, UK, 4–7 November 2003.
11. Petoukhov V, Eliseev AV, Klein R, Oesterle H (2007) On statistics of the free-troposphere synoptic component: an evaluation of skewnesses and mixed third-order moments contribution to the synoptic-scale dynamics and fluxes of heat and humidity. *Tellus A,* 10.1111/j.1600-0870.2007.00276.
12. Hastie TJA, Tibshirani RJ (1990) Generalized Additive Models, Monographs on statistics and applied probability (Chapman & Hall, London).
13. 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.
14. Held H, Kleinen T (2004) Detection of climate system bifurcations by degenerate fingerprinting. *Geophys Res Lett* 31:L23207.
15. van Nes EH, Scheffer M (2007) Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *Am Nat* 169:738–747.
16. Livina VN, Lenton TM (2007) A modified method for detecting incipient bifurcations in a dynamical system. *Geophys Res Lett* 34:L23207.
17. Mann HB (1945) Nonparametric tests against trend. *Econometrica* 13:245–259.
18. Efron B, Tibshirani R (1986) Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Stat Sci* 1:54–77.
19. Schreiber T, Schmitz (1996) A improved surrogate data for nonlinearity tests. *Phys Rev Lett* 77:635.
20. Schreiber T, Schmitz A (2000) Surrogate time series. *Physica D Nonlinear Phenomena* 142:346–382.
21. Gautama T, Mandic DP, Van Hulle MM (2004) A novel method for determining the nature of time series. *IEEE Trans Biomed Eng* 51:728–736.
22. Theiler J, Eubank S, Longtin A, Galdrikian B, Farmer JD (1992) Testing for nonlinearity in time-series—The method of surrogate data. *Physica D* 58:77–94.
23. Sokal RR, Rohlf FJ (1995) *Biometry: The principles and Practice of Statistics in Biological Research* (Freeman, New York), 3rd Ed.
24. Tripati A, Backman J, Elderfield H, Ferretti P (2005) Eocene bipolar glaciation associated with global carbon cycle changes. *Nature* 436:341–346.
25. Hughen KA, Southon JR, Lehman SJ, Overpeck JT (2000) Synchronous radiocarbon and climate shifts during the last deglaciation. *Science* 290:1951–1954.
26. Alley RB (2000) The Younger Dryas cold interval as viewed from central Greenland. *Q Sci Rev* 19:213–226.
27. deMenocal P, Ortiz J, Guilderson T, Sarnthein M (2000) Coherent high- and low-latitude climate variability during the holocene warm period. *Science* 288:2198–2202.
28. Petit JR, *et al.* (1999) Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature* 399:429–436.

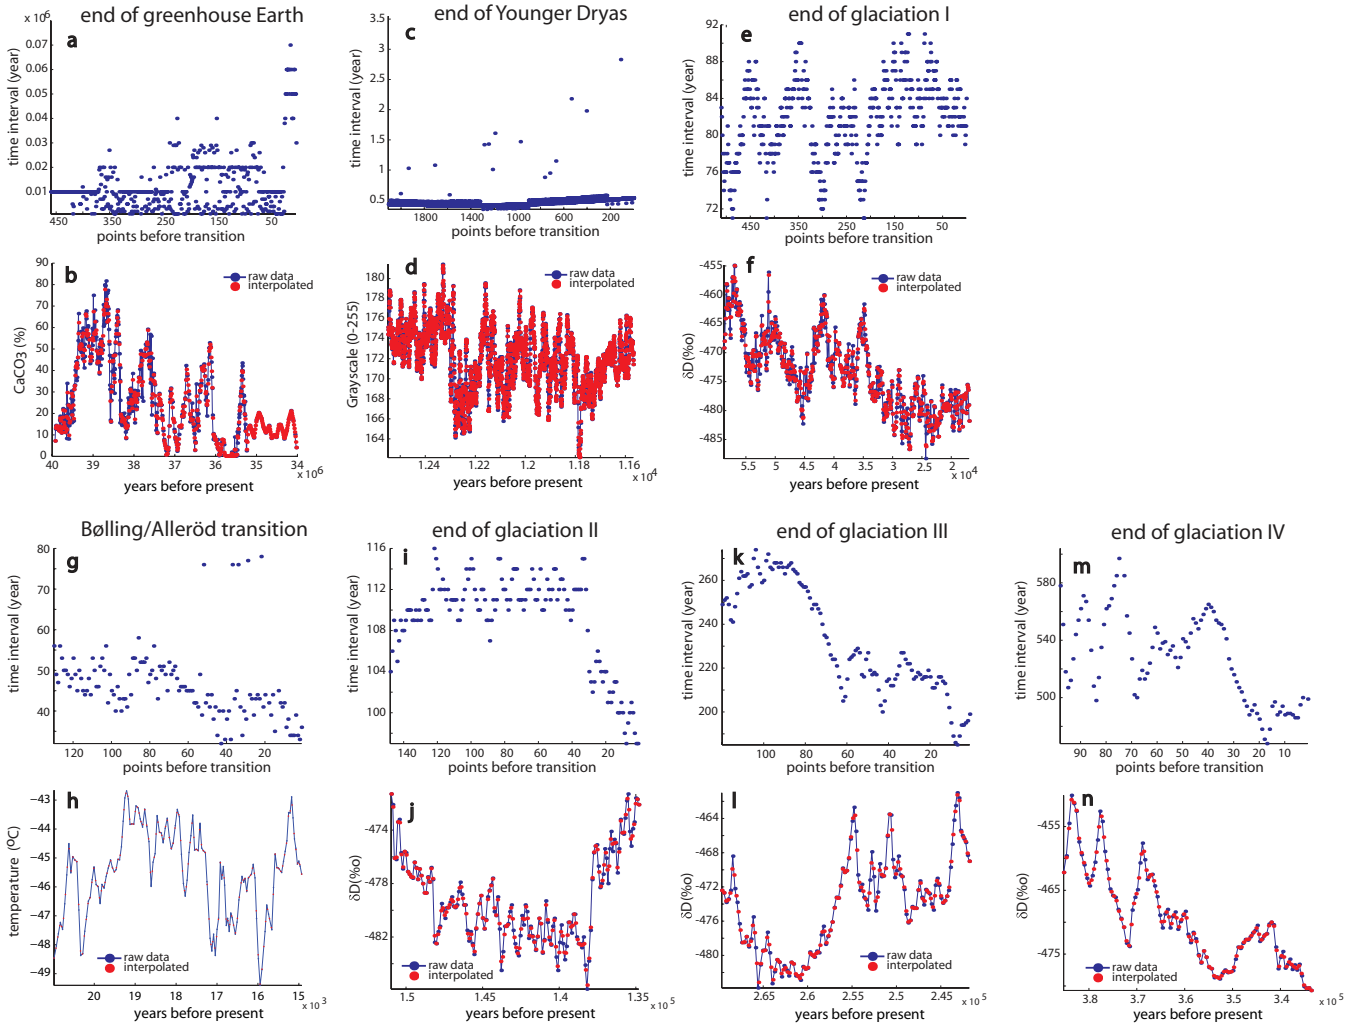

Fig. S1. Plots of time intervals and interpolated data. (*a, c, e, g, i, k,* and *m*) Plots of time intervals over proximity to the transition in the original records shown in the main text (Fig. 1). Time intervals are estimated in time units of years by differencing the times at which the data were dated. Smaller time intervals indicate increased density of points. In all cases, except from *a*, the time intervals close to the transitions in the original records are of similar magnitude to the time interval of the interpolated datasets (compare with Table S3). (*b, d, f, h, j, l,* and *n*) Plots of the interpolated data (red points) superimposed on the original records (blue points).

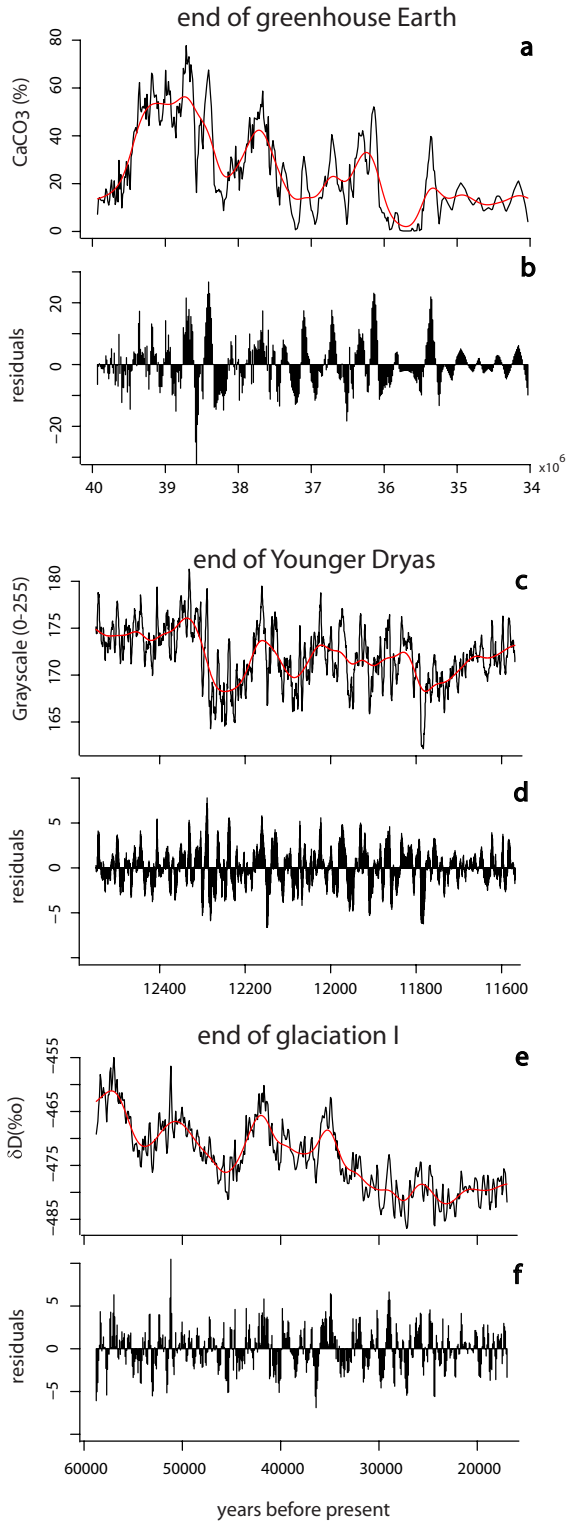

Fig. S2. Effects of filtering on the records shown in Fig. 1 in the main text. (*a*, *c*, and *e*) Data points before the transition and the Gaussian kernel filter used for detrending. (*b*, *d*, and *f*) The residual time series after subtracting the trend (red line). Records shown correspond to the records depicted in Fig. 1 in the main text.

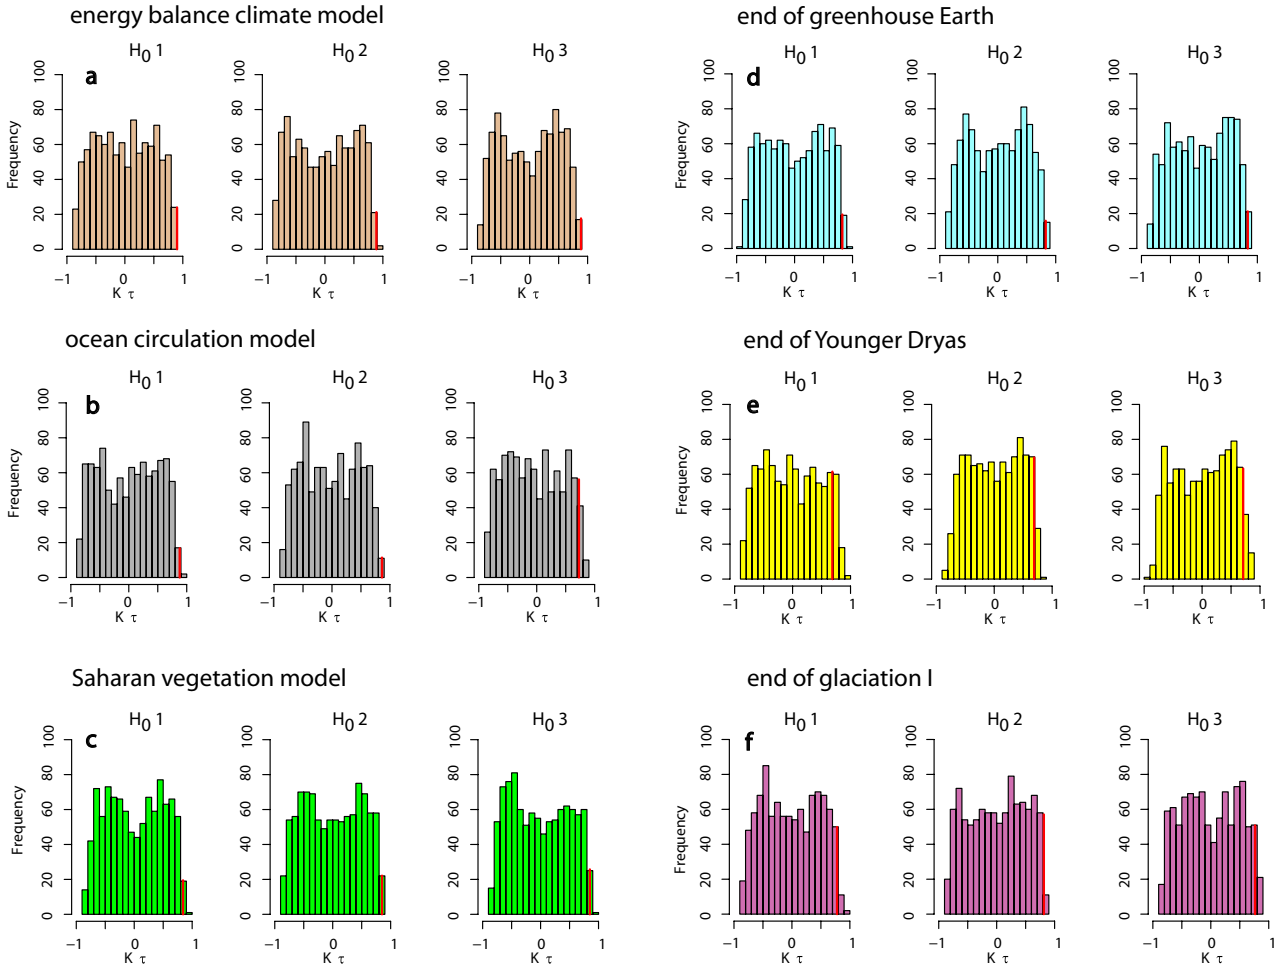

Fig. S3. Probability distributions of the estimated trend statistic (Kendall's τ) of the ranked order test, under three alternative H_0 hypotheses for a set of 1,000 surrogate time series. Under H_0 1, datasets are generated after bootstrapping the residual time series records, under H_0 2 new datasets are produced with similar distribution and Fourier spectra as the residual time series, and under H_0 3 the surrogate time series have been produced from an autoregressive model with similar autocorrelation at lag 1, mean, and variance as in the residual records. Red lines indicate the limit over which the surrogate trend statistic is higher than the trend statistic of the original residual records. From this subset, only values of significance *P* equal to or higher than the original record are used to estimate the likelihood of acquiring trend statistic estimates of similar magnitude.

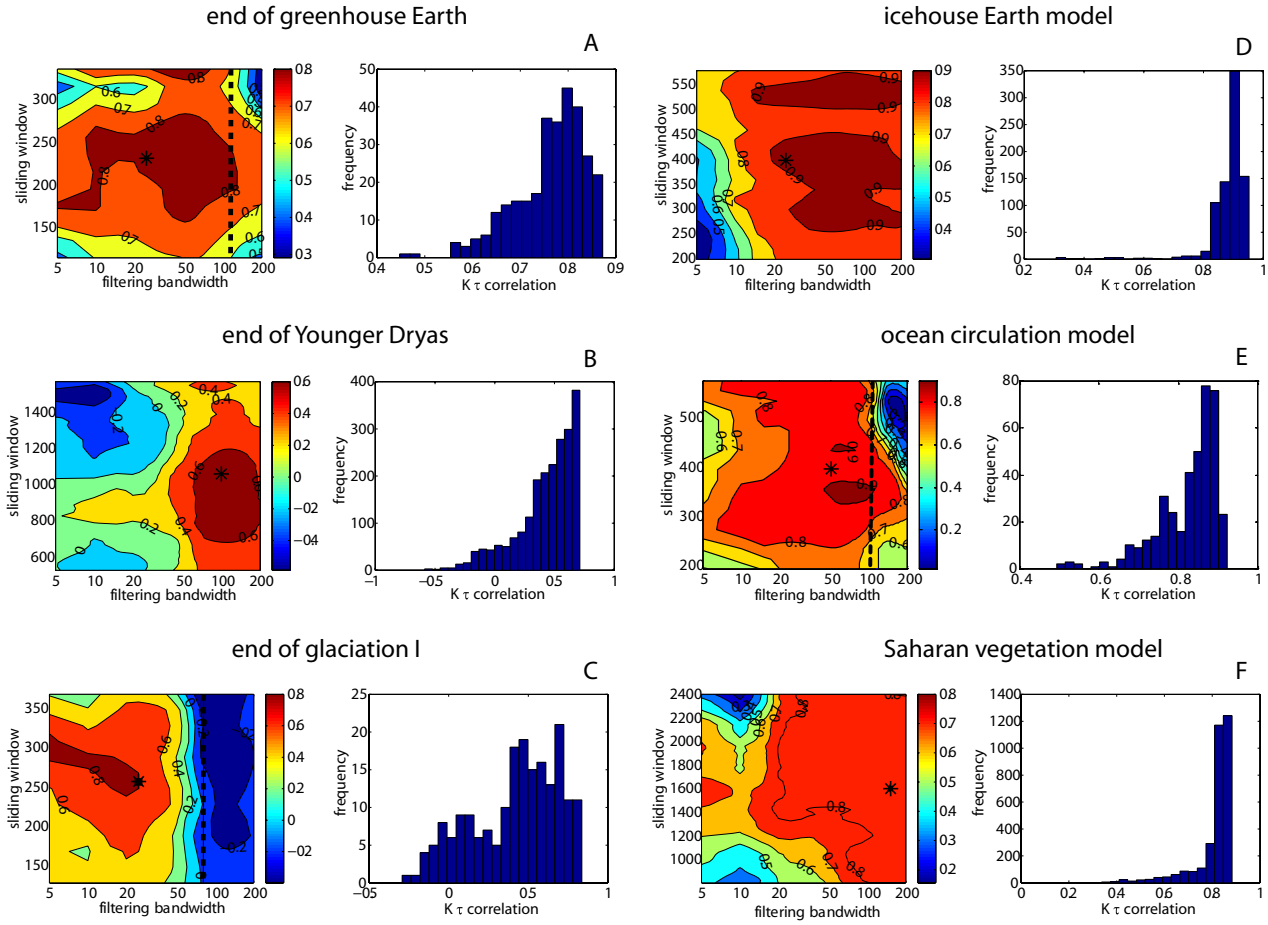

Fig. S4. Contour plots show the effect of the width of the sliding window and the width of the kernal filter on the observed trend in autocorrelation for the climate original data and simulated results shown in the main text as measured by the Kendall's τ. Stars indicate the parameter choice used in the analyses presented in the text. We visually inspected the fit of the kernel filter line to the data and demarcated where the kernel width becomes too large to follow the trend in the data (dashed vertical lines in some contour plots). The histograms give the frequency distribution of the trend statistic for the scanned parameter area to the left hand side of the dashed line.

Table S1. Paleoclimate records together with their origin, the proxy measured, the approximate range, and transition threshold used in the analyses, as well as the dataset citation and the original reference of the record

Paleo record	Origin	Climate proxy (units)	Time range, yrs BP	Time of transition	*N*	Dataset	Original reference
End of greenhouse Earth	ODP tropical Pacific core 1218	$CaCO_3$ (%)	$(39–32) \times 10^6$	34×10^6	482	*	24
Bølling-Alleröd transition	GISP2 ice core	Temperature (°C)	21,000–14,600	15,000	147	†	25
End of Younger Dryas	Cariaco basin core PL07–58PC	Grayscale (0–255)	12,500–11,200	11,600	2652	‡	26
Desertification of North Africa	ODP Hole 658C	Terrigenous dust (%)	8,300–4,800	7,500	40	§	27
End of glaciation I	Vostok ice core	d2H (%)	58,800–12,000	17,000	591	¶	28
End of glaciation II	Vostok ice core	d2H (%)	151,000–128,000	135,000	258	¶	28
End of glaciation III	Vostok ice core	d2H (%)	270,000–238,000	242,000	149	¶	28
End of glaciation IV	Vostok ice core	d2H (%)	385,300–324,600	334,100	126	¶	28

*Tripati A, *et al.* (2005). Eocene Greenhouse-Icehouse Transition Carbon Cycle Data. IGBP PAGES/World Data Center for Paleoclimatology Data Contribution Series no. 2005-056. NOAA/NGDC Paleoclimatology Program, Boulder CO.

†Alley R (2004) GISP2 Ice Core Temperature and Accumulation Data. IGBP PAGES/World .Data Center for Paleoclimatology Data Contribution Series no. 2004-013. NOAA/NGDC Paleoclimatology Program, Boulder CO.

‡Hughen K, *et al.* (2000) Cariaco Basin 2000 Deglacial 14C and Grey Scale Data, IGBP PAGES/World Data Center A for Paleoclimatology Data Contribution Series no. 2000-069.NOAA/NGDC Paleoclimatology Program, Boulder CO.

§deMenocal PB, *et al.* (2001) Holocene Variations in Subtropical Atlantic SST. IGBP PAGES/World Data Center A for Paleoclimatology Data Contribution Series no 2001-054. NOAA/NGDC Paleoclimatology Program, Boulder CO.

¶Petit JR, *et al.* (2001) Vostok Ice Core Data for 420,000 Years, IGBP PAGES/World Data Center for Paleoclimatology Data Contribution Series no. 2001-076. NOAA/NGDC Paleoclimatology Program, Boulder CO.

Table S2. Probability of acquiring the estimated values for the trend statistic (Kendall's τ) of the original and simulated residual time series under three alternative H_0 hypotheses for a set of 1,000 surrogate time series

(N = 1000 surrogate sets)	H_0 1	H_0 2	H_0 3
Original record (residuals)	Kendall τ	Kendall τ	Kendall τ
End of greenhouse Earth	0.014**	0.004**	0.011**
End of Younger Dryas	0.086*	0.03**	0.055*
End of glaciation I	0.013**	0.011**	0.021**
Bølling–Alleröd transition	0.367	0.340	0.332
End of glaciation II	0.402	0.397	0.386
End of glaciation III	0.247	0.235	0.234
End of glaciation IV	0.186	0.043**	0.125
Desertification of North Africa	0.140	0.165	0.091*
Fisher's combined probability	*0.002847*	*0.000206*	*0.001278*
Simulated record (residuals)			
Energy balance climate model	$<10^{-4}$**	0.002**	$<10^{-4}$**
Saharan vegetation model	0.002**	0.001**	0.006**
Ocean circulation model	0.003**	$<10^{-4}$**	$<10^{-4}$**

Under H_0 1, datasets are generated after bootstrapping, under H_0 2 new data sets are produced with similar distribution and Fourier spectra as the residual time series, and under H_0 3 the surrogate time series have been produced from a autoregressive model with similar autocorrelation at lag 1, mean, and variance as in the residual records. *, $P \leq 0.1$. **, $P \leq 0.05$. In italics, the combined probability for obtaining the estimated probabilities for each hypothesis is provided.

Table S3. Summary of trend statistic for the original (noninterpolated) and interpolated paleo records, and their probabilities (*P*)

Record	*N* points original/ interpolated	Bandwidth size	Original K τ (*P*)	Interpolated K τ (*P*)
End of greenhouse Earth	461/462	25	0.5 ($<10^{-4}$)	0.83 ($<10^{-4}$)
End of Younger Dryas	2,110/2,111	100	0.34 ($<10^{-4}$)	0.69 ($<10^{-4}$)
End of glaciation I	512/513	25	0.85 ($<10^{-4}$)	0.8 ($<10^{-4}$)
Bølling–Alleröd transition	131/132	25	0.37 ($<10^{-4}$)	0.27 (0.001)
End of glaciation II	149/150	25	0.08 (0.31)	0.17 (0.27)
End of glaciation III	121/122	10	0.67 ($<10^{-4}$)	0.43 ($<10^{-4}$)
End of glaciation IV	99/100	50	0.51 ($<10^{-4}$)	0.52 ($<10^{-4}$)
Desertification of North Africa	88/88*	10	0.58 (0.001)	0.58 (0.001)

*In the case of the desertification of North Africa the original data were already interpolated.