

Climate tipping as a noisy bifurcation: a predictive technique

J. M. T. THOMPSON

*School of Engineering, Fraser Noble Building, Kings College, University of Aberdeen,
Aberdeen AB24 3UE, UK and Department of Applied Mathematics and Theoretical Physics,
Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road,
Cambridge CB03 0WA, UK*

AND

JAN SIEBER*

*Department of Mathematics, University of Portsmouth, Lion Gate Building, Lion Terrace,
Portsmouth PO1 3HF, UK*

*Corresponding author: jan.sieber@port.ac.uk

[Received on 3 November 2009; revised on 26 October 2010;
accepted on 16 November 2010]

It is often known, from modelling studies, that a certain mode of climate tipping (of the oceanic thermohaline circulation, for example) is governed by an underlying fold bifurcation. For such a case, we present a scheme of analysis that determines the best stochastic fit to the existing data. This provides the evolution rate of the effective control parameter, the variation of the stability coefficient and the path itself and its tipping point. By assessing the actual effective level of noise in the available time series, we are then able to make probability estimates of the time of tipping. This new technique is applied, first, to the output of a computer simulation for the end of greenhouse Earth about 34 million years ago when the climate tipped from a tropical state into an icehouse state with ice caps. Second, we use the algorithms to give probabilistic tipping estimates for the end of the most recent glaciation of the Earth using actual geological ice-core data.

Keywords: climate tipping; slow passage through saddle-node bifurcation; time-series analysis.

1. Introduction

One concern of the UN Climate Change Conference in [Copenhagen \(2009\)](#) was the prediction of future climate change, subject to a variety of carbon dioxide emission scenarios. A particularly alarming feature of any such prediction would be a sudden and (perhaps) irreversible abrupt change called a *tipping point* ([Lenton *et al.*, 2008](#); [Scheffer, 2009](#)). Such events are familiar in nonlinear dynamics, where they are called (dangerous) bifurcations at which one form of behaviour becomes unstable and the system jumps rapidly to a totally different *steady state*. Many tipping points, such as the switching on and off of ice ages, are well documented in paleoclimate studies over millions of years of the Earth's history.

There is currently much interest in examining climatic tipping points to see if it is feasible to predict them in advance using time series data derived from past behaviour. Assuming that tipping points may well be governed by a bifurcation in an underlying dynamical system, recent work looks for a slowing down of intrinsic transient responses within the data, which is predicted to occur before most bifurcational instabilities ([Held & Kleinen, 2004](#); [Livina & Lenton, 2007](#)). This is done, for example, by determining the *propagator*, which is estimated via the correlation between successive elements of

the time series, in a window sliding along the time series. This propagator is a measure for the linear stability. It should increase to unity at tipping.

Many trial studies have been made on climatic computer models where an arbitrary time, $t = T$, can be chosen to represent ‘today’. The challenge is then to predict a future critical time, $t^C > T$, at which the model will exhibit a tipping instability, using only the time history of some variable (average sea temperature, say) generated by the model before time T . The accuracy of this prediction can then be assessed by comparing it to the actual continued response of the simulation beyond time T . In some cases, these trials have been reasonably successful.

Much more challenging, and potentially convincing, is to try to predict real ancient climate tipplings, using their preceding geological data. The latter would be, for example, re-constituted time series provided by ice cores, sediments, etc. Using this data, the aim would be to see to what extent the actual tipping could have been accurately predicted in advance.

One past tipping point that has been analysed in this manner (Livina & Lenton, 2007) is the end of the Younger Dryas event, about 11,500 years ago, when the Arctic warmed by 7°C in 50 years. These authors used a time series derived from Greenland ice-core paleotemperature data. A second such study (one of eight made by Dakos *et al.*, 2008), using data from tropical Pacific sediment cores, gives an excellent prediction for the end of ‘greenhouse’ Earth about 34 million years ago when the climate tipped from a tropical state into an icehouse state.

After a review of recent research on tipping, we study in Section 5 the saddle-node normal form with a drifting normal form parameter and with additive Gaussian noise to determine the probability (or rate) of early noise-induced escape from the potential well depending on the drift speed of the normal form parameter and the noise amplitude.

We show how one can extract the relevant normal form quantities from a given time series. This allows one to adjust predictions of tipping events based on the propagator to take into account the probability of early escape. We demonstrate how it is possible to estimate the probability of noise-induced escape from the potential well using two time series: one is an output from a simple stochastic model and the other is a paleotemperature record from ice-core data. Since the propensity of the system to escape early from its potential well depends only on the order of magnitude of the ratio between drift speed and noise amplitude, we expect our estimates to be reasonably robust.

This prediction science is very young, but the above trials on paleoevents seem very encouraging, and we describe some of them more fully below. The prediction of *future* tipping points, vital to guide decisions about geo-engineering, will benefit from the experience drawn from these trials and will need the high quality data currently being recorded worldwide by climate scientists today.

2. Tipping of the climate and its sub-systems

2.1 Tipping points

Work at the beginning of this century which sets out to define and examine *climate tipping* (Rahmstorf, 2001; Lockwood, 2001; National Research Council, 2002; Alley *et al.*, 2003; Rial *et al.*, 2004) focused on abrupt climate change: namely, when the Earth system is forced to cross some threshold, triggering a sudden transition to a new state at a rate determined by the climate system itself and (usually) faster than the cause, with some degree of irreversibility. Recently, the Intergovernmental Panel on Climate Change (2007) made some brief remarks about abrupt and rapid climate change, while Lenton *et al.* (2008) have sought to define these points more rigorously. The physical mechanisms underlying these tipping points are typically internal positive feedback effects of the climate system (thus, a certain propensity for saddle-node bifurcations).

2.2 Tipping elements

In principle, a climate tipping point might involve simultaneously many features of the Earth system, but it seems that many tipping points might be strongly associated with just one fairly well-defined sub-system. These *tipping elements* are well-defined sub-systems of the climate, which work (or can be assumed to work) fairly independently, and are prone to sudden change. In modelling them, their interactions with the rest of the climate system are typically expressed as a control parameter (or forcing) that varies slowly over time.

Recently, Lenton *et al.* (2008) have listed nine tipping elements that they consider to be primary candidates for future tipping due to human activities and as such have relevance to political decision making at Copenhagen (2009) and beyond. These elements, their possible outcomes, and Lenton's assessment of whether their tipping might be associated with an underlying bifurcation are

- (1) loss of Arctic summer sea ice (possible bifurcation);
- (2) collapse of the Greenland ice sheet (bifurcation);
- (3) loss of the West Antarctic ice sheet (possible bifurcation);
- (4) shut-down of the Atlantic thermohaline circulation (THC; fold bifurcation);
- (5) increased amplitude or frequency of the El Niño Southern Oscillation (some possibility of bifurcation);
- (6) switch-off of the Indian summer monsoon (possible bifurcation);
- (7) changes to the Sahara/Sahel and West African monsoon, perhaps greening the desert (possible bifurcation);
- (8) loss of the Amazon rainforest (possible bifurcation) and
- (9) large-scale dieback of the Northern Boreal forest (probably not a bifurcation).

The analysis and prediction of tipping points of climate sub-systems is currently being pursued in several streams of research, and we should note in particular the excellent book by Marten Scheffer about tipping points in 'Nature and Society', which includes ecology and some climate studies (Scheffer, 2009).

3. Bifurcations and their precursors

3.1 Generic bifurcations of dissipative systems

The great revolution of nonlinear dynamics over recent decades has provided a wealth of information about the bifurcations that can destabilize a slowly evolving system like the Earth's climate. These bifurcations are defined as points during the slow variation of a 'control' parameter at which a qualitative topological change of behaviour is observed in the multi-dimensional phase space of the system.

The Earth's climate is what dynamicists would call a *dissipative system*, and for this, the bifurcations that can be typically encountered under the variation of a single control parameter are classified into three types: *safe*, *explosive* and *dangerous* (Thompson *et al.*, 1994; Thompson & Stewart, 2002).

The safe bifurcations, such as the supercritical Hopf bifurcation, exhibit a continuous supercritical growth of a new attractor path with no fast jump or enlargement of the attracting set. They are determinate with a single outcome even in the presence of small noise and generate no hysteresis with the path retraced on reversal of the control sweep. The explosive bifurcations are less common phenomena lying intermediate between the safe and dangerous types: we simply note here that, like the safe bifurcations, they do not generate any hysteresis. The dangerous bifurcations are typified by the simple fold

(saddle-node bifurcation) at which a stable path increasing with a control parameter becomes unstable as it curves back towards lower values of the control and by the subcritical bifurcations. They exhibit the sudden disappearance of the current attractor, with a consequential sudden jump to a new attractor (of any type). They can be indeterminate in outcome, depending on the topology of the phase space, and they always generate hysteresis with the original path not reinstated on control reversal.

Any of these three bifurcation types could in principle underlie a climate tipping point. But it is the dangerous bifurcations that will be of major concern, giving as they do a sudden jump to a different steady state with hysteresis, so that the original steady state will not be re-instated even if the controlling cause is itself reversed. So any future climatic tipping to a warmer steady state may be irreversible: a subsequent reduction in CO₂ concentration will not (immediately or perhaps ever) restore the system to its pre-tipping condition.

3.2 Time series analysis of incipient bifurcations

Most of the bifurcations in dissipative systems, including the static and cyclic folds that are the most likely to be encountered in climate studies, have the following useful precursor (see [Thompson & Stewart, 2002](#) for more details). The stability and attracting strength of the current steady state is becoming steadily weaker and weaker in one mode as the bifurcation point is approached. This implies that under inevitable noisy disturbances, transient motions returning to the attractor will become slower and slower: in the limit, the rate of decay of the transients decreases linearly to zero along the path.

[Held & Kleinen \(2004\)](#) and [Livina & Lenton \(2007\)](#) have recently presented algorithms that are able to detect incipient saddle-node bifurcations from time series of dynamical systems. Both methods estimate the linear decay rate (LDR) towards a quasistationary equilibrium that is assumed to exist and to drift towards a saddle-node bifurcation. Typical test data for the algorithms come either from geological records or from output of climate models. Both algorithms (*degenerate fingerprinting* by [Held & Kleinen, 2004](#) and *detrended fluctuation analysis* (DFA) by [Livina & Lenton, 2007](#)) have to make assumptions about the process underlying the recorded time series that are generally believed to be sensible for the tipping elements listed by [Lenton *et al.* \(2008\)](#). First, one has to assume that the process is a dynamical system close to a stable equilibrium that drifts only slowly but is perturbed by (random) disturbances. The second assumption is that the system is effectively 1D, that is, the equilibrium of the undisturbed system is strongly stable in all directions except a single critical one. Quantitatively, this means that one assumes the presence of three well-separated time scales, expressed as rates:

$$\kappa_{\text{drift}} \ll \kappa_{\text{crit}} \ll \kappa_{\text{stab}}.$$

Here, κ_{drift} is the average drift rate of those quantities that the algorithm treats as a parameter, e.g. freshwater forcing in studies of the THC (the global heat- and salinity-driven conveyor belt of oceanic water). The rate κ_{crit} is the rate with which a small disturbance in state space relaxes back to equilibrium. The rate κ_{stab} is the decay rate of all other *non-critical* modes. We note that the drift rate κ_{drift} becomes *larger* than κ_{crit} once the drifting parameter is very close to its bifurcation value. Third, one assumes that the disturbances are small in the sense that the relaxation to equilibrium is governed mostly by the LDRs (this implies, for example, that the potential well in which the dynamical system can be imagined to be sitting is approximately symmetric in the critical direction).

The basic procedure proposed by [Held & Kleinen \(2004\)](#) consists of three steps, given a time series (t_k, z_k) of measurements z_k at possibly unevenly spaced time points t_k .

1. Interpolation. Choose a stepsize Δt satisfying

$$\kappa_{\text{crit}} \ll \frac{1}{\Delta t} \ll \kappa_{\text{stab}}$$

and interpolate such that the spacing in time is uniform. Now, one has a new time series

$$z_{k,\text{new}} = z(k \Delta t),$$

evenly spaced in time.

2. Detrending. Remove the slow drift of the equilibrium by subtracting a slowly moving average. For example, choose a Gaussian kernel

$$G_k(t) = \frac{1}{\sqrt{2\pi}d} \exp\left(-\frac{1}{2} \frac{(t - k \Delta t)^2}{d^2}\right)$$

of bandwidth d satisfying

$$\kappa_{\text{drift}} \ll \frac{1}{d} \ll \kappa_{\text{crit}}$$

and subtract the average

$$Z(k \Delta t) = \frac{\sum_{i=1}^N G_k(i \Delta t) z_i}{\sum_{i=1}^N G_k(i \Delta t)}$$

of z_k over the kernel. The result of this is a time series

$$y_k = z_k - Z(k \Delta t),$$

which fluctuates around zero and can be considered as stationary on time scales shorter than $1/\kappa_{\text{drift}}$.

3. Fit LDR in sliding window. One assumes that the remaining time series, y_k , can be modelled by a stable scalar linear mapping disturbed by noise, a so-called AR(1) model

$$y_{k+1} = c_k y_k + \theta \eta_k, \quad (3.1)$$

where $\theta \eta_k$ is the instance of a random disturbance of amplitude θ at time $k \Delta t$ and c_k is the *propagator*, related to κ_{crit} at time $k \Delta t$ via

$$c_k = \exp(-\kappa_{\text{crit},k} \Delta t).$$

If one assumes that the disturbances η_k have a normal distribution and are independent from each other and that c_k is nearly constant on time scales shorter than $1/\kappa_{\text{drift}}$, one can choose a sliding window size $w = 2m + 1$ satisfying

$$\kappa_{\text{drift}} \ll \frac{1}{w \Delta t}$$

and determine the propagator c_k by an ordinary least-squares fit of

$$y_{j+1} = c_k y_j$$

over the set of indices $j = k - m, \dots, k + m$. An estimate for the noise amplitude θ can be obtained from the standard deviation of the residual of the linear least-squares fit:

$$\theta_k = \text{stdev} \left([y_j - c_k y_j]_{j=k-m}^{k+m} \right).$$

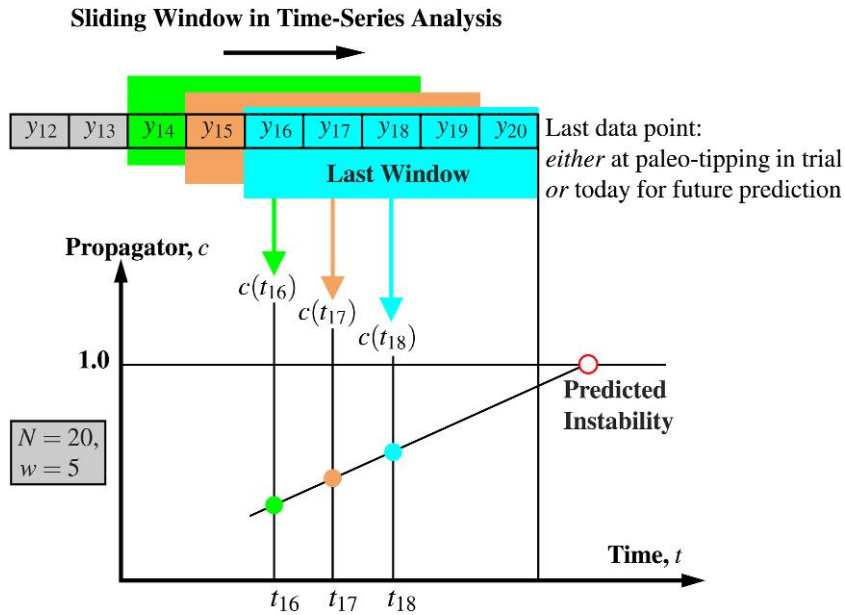


FIG. 1. Illustration of sliding windows used in degenerate fingerprinting proposed by Held & Kleinen (2004). The extracted autocorrelation coefficient c_k is an estimate for the propagator at the mid-time of the sliding windows. Extrapolation is required for prediction.

This process will stop when the front end of the last window hits the last data point. Note that when using paleodata, the end of the analysed time series should be chosen before the tipping point that is the object of the investigation. This choice is essential to prevent data, spurious to our predictions of the pre-tipping behaviour, entering the analysis: it also makes a complete analogy with any attempt to predict future tipping points from data terminating today. Figure 1 illustrates this requirement on the sliding windows (with a time series of length $N = 20$ for illustration). It also shows a simple linear extrapolation that one might make in order to predict where tipping occurs. In this manner, Held & Kleinen (2004) obtain the so-called propagator graph of the estimated c_k versus its central time as illustrated in Fig. 1. On this graph, c is expected to head towards $+1$ at any incipient bifurcation. In other words, the slowing down of the relaxation from disturbances along the times series can serve as an early warning signal for an imminent bifurcation (Dakos *et al.*, 2008). We should finally note that having used a first-order mapping in (3.1) and employed autocorrelation techniques, the propagator c is often called the first-order autoregressive coefficient and written as ARC(1). The prediction based on DFA, as proposed by Livina & Lenton (2007), also reconstructs the propagator c but does so via the scaling exponent of the variance of the (detrended) time series. For a more complete description of the time series techniques employed see the recent review by Thompson & Sieber (2010). Also, see Corsi & Taranto (2007) for a discussion of similar problems in power systems engineering (the prediction of fold bifurcations leading to voltage collapse in energy networks).

4. Review of recent work

4.1 First prediction of an ancient tipping

The first prediction of an ancient climate tipping event using preceding geological data is due to Livina & Lenton (2007) who tested their DFA on the rapid warming of the earth that occurred about 11,500 years

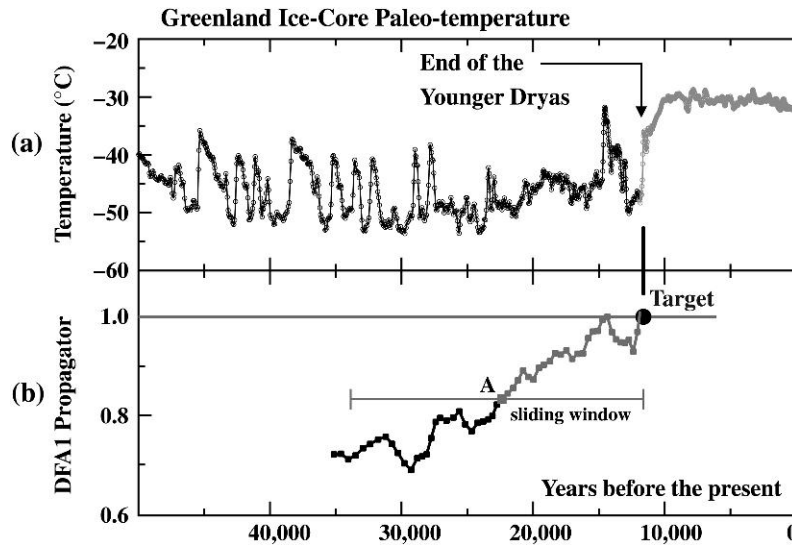


FIG. 2. Results of Livina & Lenton (2007): (a) Greenland ice-core (GISP2) paleotemperature with an unevenly spaced record, visible in the varying density of symbols on the curve. The total number of data points is $N = 1586$. (b) the DFA1 propagator is calculated in sliding windows of length $w = 500$ points and mapped into the middle points of the windows. A typical sliding window ending near the tipping is shown.

ago at the end of the so-called Younger Dryas event (analysing Greenland ice-core paleotemperature data, which is available from 50,000 years ago to the present).

This Younger Dryas event (Houghton, 2004) was a curious cooling just as the Earth was warming up after the last ice age, as is clearly visible, e.g. in records of the oxygen isotope $\delta^{18}\text{O}$ in Greenland ice. It ended in a dramatic tipping point, 11,500 years ago, when the Arctic warmed by 7°C in 50 years. Its behaviour is thought to be linked to changes in the THC. This ‘conveyor belt’ is driven by the sinking of cold salty water in the North and can be stopped if too much fresh-melt makes the water less salty and so less dense. At the end of the ice age when the ice sheet over North America began to melt, the water first drained down the Mississippi basin into the Gulf of Mexico. Then, suddenly, it cut a new channel near the St. Lawrence river to the North Atlantic. This sudden influx of fresh water cut off part of the ocean conveyor belt, the warm Atlantic water stopped flowing North, and the Younger Dryas cooling was started. It was the ‘re-start’ of the circulation that could have ended the Younger Dryas at its rapid tipping point, propelling the Earth into the warmer Pre-Boreal era.

The results of Livina & Lenton (2007) are shown in Fig. 2, where their propagator (based on DFA) is seen heading towards its critical value of +1 at about the correct time. Note, though, that from a prediction point of view, the propagator graph should end at point A when the estimation window reaches the tipping point. In this example, extracting the propagator is particularly challenging because the data set was comparatively small (1586 points) and unevenly spaced.

4.2 Systematic study of eight ancient tippings

In a more recent paper, Dakos *et al.* (2008) systematically estimated a propagator stability coefficient from reconstructed time series of real paleodata preceding eight ancient tipping events. These are

- (a) the end of the greenhouse Earth about 34 million years ago when the climate tipped from a tropical state (which had existed for hundreds of millions of years) into an icehouse state with ice caps, using data from tropical Pacific sediment cores;
- (b) the end of the last glaciation, and the ends of three earlier glaciations, drawing data from the Antarctica Vostok ice core;
- (c) the Bølling–Allerød transition that was dated about 14,000 years ago, using data from the Greenland GISP2 ice core;
- (d) the end of the Younger Dryas event about 11,500 years ago, as discussed in Section 4.1, but drawing not on the Greenland ice core, but rather on data from the sediment of the Cariaco basin in Venezuela and
- (e) the desertification of North Africa when there was a sudden shift from a savanna-like state with scattered lakes to a desert about 5,000 years ago, using the sediment core from ODP Hole 658C, off the west coast of Africa.

In all the cases studied by Dakos *et al.* (2008), the propagator c_k as extracted by degenerate fingerprinting was shown to exhibit a statistically significant increase (corresponding to a slowing down of the relaxation) prior to the tipping transition. Dakos *et al.* (2008) also demonstrated that their principal result, the statistically significant increase of c_k , is robust with respect to variations in smoothing kernel bandwidth d , sliding window length w and the interpolation procedure.

5. Noise-induced systematic bias of extrapolated prediction

The intention behind the development of the time series analysis algorithms goes beyond statistical evidence of an increasing LDR: the goal of both algorithms is to *predict the time* (or probability) of the tipping event from the observational data before the event takes place. This is more challenging and suffers from additional uncertainties. Apart from the dependence of the value of c_k on algorithm parameters (e.g. the sliding window length w), for a prediction of the time of tipping, we have to extrapolate. This implies that we have to assume that the underlying control parameter drifts with nearly constant speed during the recorded time series. This is often not the case in the study of sub-systems of the climate when the control parameter is determined by the dynamics of another, coupled, sub-system. Even if the control parameter drifts with constant speed, for prediction, we have to assume in addition *transversality*, that is, the control parameter has to vary the unfolding parameter of the normal form of the saddle-node bifurcation nearly linearly.

Figure 3 shows two time series (a and c) and the corresponding time series of extracted estimates for the propagator c_k (b and d). Time series (a) is the output of a model simulation for a transition to an icehouse Earth and is taken from Dakos *et al.* (2008). The model as presented by Dakos *et al.* (2008) is a scalar stochastic ordinary differential equation where a control parameter is varied linearly in time and the system is known to encounter a saddle-node bifurcation (originally, the model was developed by Fraedrich, 1978; see supplement of Dakos *et al.*, 2008). Time series (b) shows the propagator c_k extracted from time series (a) using degenerate fingerprinting. Time series (c) is a snapshot of temperatures before the end of the last glaciation, 20,000 years ago. The data are taken from Petit *et al.* (1999); the window of the snapshot is identical to Fig. 1(I) in Dakos *et al.* (2008). The estimated propagator c_k for time series (c) is shown in diagram (d). The graphs in diagrams (b) and (d) are qualitatively the same as in Dakos *et al.* (2008) but differ slightly quantitatively. This is likely due to minor variations between the algorithm parameters that we used and the algorithm parameters used by Dakos *et al.* (2008). To provide a visual clue about the level of uncertainty in the time series of c_k , the diagrams (b) and (d)

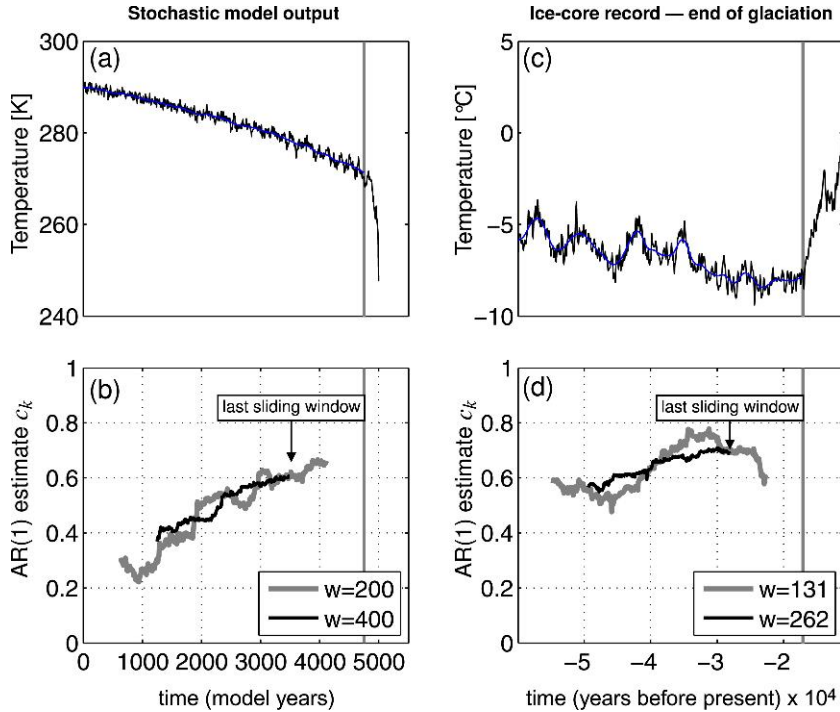


FIG. 3. LDR estimates for the output of a model of transition to icehouse earth (a and b) and for an geological temperature record of the end of the recent glaciation (c and d). See *Petit et al. (1999)* and main text for description. Data source: <ftp://ftp.ncdc.noaa.gov/pub/data/paleo/icecore/antarctica/vostok/deutnat.txt>.

show estimates extracted using two different window size parameters w . Despite the uncertainty, two features of the time series of propagators c_k are discernible: first, as studied in detail in *Dakos et al. (2008)*, c_k is increasing. Second, linear extrapolation does not match if we expect the tipping to occur at the extrapolated time for $c = 1$ (which would correspond to the critical value of the propagator). In both cases, the tipping occurs earlier, and at least for the model output (which is known to drift through a saddle-node bifurcation), the bias of the linear extrapolation is systematic. There are two competing effects that might determine the systematic bias. On the one hand, the equilibrium starts to drift rapidly (square-root like) when the drifting parameter approaches the saddle-node bifurcation; on the other hand, random disturbances may kick the dynamical system out of the shrinking basin of attraction prematurely. The numerical study in the following sections quantifies these two effects.

5.1 Rate of noise-induced escape from basin near saddle-node

We take the normal form of a saddle-node bifurcation, perturb it by adding Gaussian white noise and let the control parameter a drift with speed ϵ (the climate model of *Fraedrich (1978)* is exactly of this type):

$$dx = [a - x^2]dt + \sigma dW_t, \quad (5.1)$$

$$da = -\epsilon dt. \quad (5.2)$$

The perturbation W_t (sometimes written $W(t)$ to avoid double subscripts) is a Wiener process, which is defined by two properties:

1. $W_0 = 0$ and
2. for every sequence of time points $0 \leq t_1 \leq t_2 \leq \dots \leq t_k$, all increments $W(t_{i+1}) - W(t_i)$ are independent random variables with Gaussian distribution of zero mean and variance $t_{i+1} - t_i$ (and, thus, standard deviation $\sqrt{t_{i+1} - t_i}$).

The coefficient σ controls the amplitude of the noise variance added to the slowly drifting equilibrium. If we freeze the drifting (set $\epsilon = 0$) and set the noise amplitude to zero ($\sigma = 0$), then the dynamics of (5.1) corresponds to the dynamics of an overdamped particle in a potential well of the shape

$$U(x) = -ax + x^3/3.$$

This dynamics of (5.1) with a fixed a and no noise has a stable equilibrium X_s at \sqrt{a} and an unstable equilibrium X_u at $-\sqrt{a}$. Correspondingly, the potential well U has a (local) minimum at $X_s = \sqrt{a}$ and a hilltop (local maximum) at $X_u = -\sqrt{a}$. For x going to $-\infty$, the potential U falls off to $-\infty$, and for x going to $+\infty$, it increases to $+\infty$. The differential equation (5.2) for a governs the drifting of the control and has the solution $a(t) = a(0) - \epsilon t$.

We say that a trajectory $x(t)$ ‘escapes’ if it reaches $-\infty$ (typically in finite time). In numerical tests, one can detect if the trajectory crosses a fixed negative threshold x_{th} from which it is unlikely to return to the well (for example, set x_{th} at some fixed distance below X_u). We are interested in finding the cumulative escape probability $P_{\text{esc}}(a)$ for a trajectory to escape before the drifting control parameter has reached the value a . If the stochastic process (5.1)–(5.2) starts with a sufficiently large $a(0)$ and if ϵ is sufficiently small, then this probability P_{esc} depends only weakly on the initial distribution of x as long as this initial distribution is concentrated inside the potential well.

The perturbation W_t is a random disturbance (noise) as if x is the position of a particle that has a temperature. If the modulus of the noise is orders of magnitude smaller than the height of the barrier (the potential difference between X_s and X_u , which is $4a^{3/2}/3$), then the escape rate formulas developed for chemical reaction rates (Kramers’ escape rate, see Hänggi *et al.*, 1990) can be applied. As we are interested in the transition of a through zero, the standard reaction rate theory is not applicable. Analytical formulas for jumping times in periodic potentials from one period to the next and arbitrary noise are given in Malakhov & Pankratov (1996). See also Fischer & Imkeller (2006) for an analysis of stochastic resonance between two slowly varying potential wells.

5.1.1 Re-scaling of parameters. In the noise-perturbed drifting normal form (5.1)–(5.2), we can re-scale time t , x , a and ϵ such that the noise amplitude σ is equal to 1. Exploiting that $dW_{\rho s} = \sqrt{\rho}dW_s$, and setting

$$t_{\text{old}}\sigma^{2/3} = t_{\text{new}}, \quad \sigma^{-2/3}x_{\text{old}} = x_{\text{new}}, \quad \sigma^{-4/3}a_{\text{old}} = a_{\text{new}} \quad \text{and} \quad \sigma^{-2}\epsilon_{\text{old}} = \epsilon_{\text{new}}, \quad (5.3)$$

we obtain a stochastic process for the re-scaled quantities that is of the same form as the original noise-perturbed normal form (5.1)–(5.2) (except that the parameter σ has been absorbed as a unit):

$$dx = [a - x^2]dt + dW_t, \quad (5.4)$$

$$da = -\epsilon dt. \quad (5.5)$$

In the rescaled coordinates, the node (the local minimum of the well) and the saddle (the barrier) are at

$$X_{u,\text{new}} = \sigma^{-2/3} X_{u,\text{old}} \quad \text{and} \quad X_{s,\text{new}} = \sigma^{-2/3} X_{s,\text{old}},$$

respectively.

5.1.2 Escape rates for the saddle-node normal form with noise. For sufficiently small drift speeds ϵ of the control parameter a , we can approximate the escape probability $P_{\text{esc}}(a)$ in the drifting system (5.4)–(5.5) using quantities of the frozen system (drift speed $\epsilon = 0$). A useful quantity is the ‘escape rate’ k_0 . For fixed $a > 0$ in the noise-perturbed normal form (5.4), the escape rate $k_0(a)$ can be defined as follows:

1. set a large ensemble (size N) of initial values x inside the potential well (e.g. $x = X_s$);
2. evolve the noise-perturbed normal form (5.4) and measure the fraction $r(t)$ of instances that have reached $-\infty$ during the previous unit time interval:

$$r(t) = \lim_{h \rightarrow 0} \lim_{N \rightarrow \infty} \frac{1}{h} \frac{\text{number of instances reaching } -\infty \text{ during } [t, t+h]}{\text{number of instances still finite at time } t} \text{ and}$$

3. after a transient this fraction $r(t)$ converges to a constant k_0 .

This convergence is achieved only in the limit $N \rightarrow \infty$ due to depletion for finite ensembles. In numerical calculations, one can replace $-\infty$ with a finite threshold $x_{\text{th}} \ll X_u$. One can avoid depletion of the finite ensemble by re-initializing the escaped instance to an initial value that is randomly selected from the remaining (non-escaped) ensemble (‘putting the particle back into the well’).

Figure 4 shows this escape rate k_0 as a function of a as a grey curve with circles. Note that the x -axis corresponds to the re-scaled a after transformation (5.3). An escape rate $k_0 \approx 0.1$ at $a\sigma^{-4/3} = 0.4$ means that approximately 10% of the realizations cross the threshold $x_{\text{th}} = -5$ to escape per unit time interval during the solution of the saddle-node normal form (5.1) with $a\sigma^{-4/3} = 0.4$.

5.1.3 Quasistatic approximation of the cumulative escape probability. Using the escape rate $k_0(a)$ for the frozen problem ($\epsilon = 0$ in (5.4)–(5.5)), we can approximate the cumulative probability P_{esc} of escape for a trajectory of the saddle-node normal form with drifting control parameter. If we assume that escape is irreversible and denote by $p(t)$ the probability that the trajectory has not escaped until time t , then we have the relation

$$p(t+h) = [1 - (k_0(a(t)) + O(\epsilon))h]p(t) \tag{5.6}$$

for small time steps $h > 0$. The factor $1 - k_0(a(t))h$ is the probability that the trajectory escapes during the time interval $[t, t+h]$ if we approximate the slowly changing variable a by its left-end value $a(t)$ in this interval. Relation (5.6) expresses that a trajectory will not escape until time $t+h$ if it has not escaped until time t and it does not escape during the interval $[t, t+h]$. Letting h go to zero and ignoring the slow drift of a during the time interval $[t, t+h]$ of order $O(\epsilon)$, we obtain the differential equation

$$\frac{d}{dt} p(t) = -k_0(a(t))p(t),$$

which has the solution

$$p(t) = \exp\left(\int_0^t -k_0(a(s))ds\right) \tag{5.7}$$

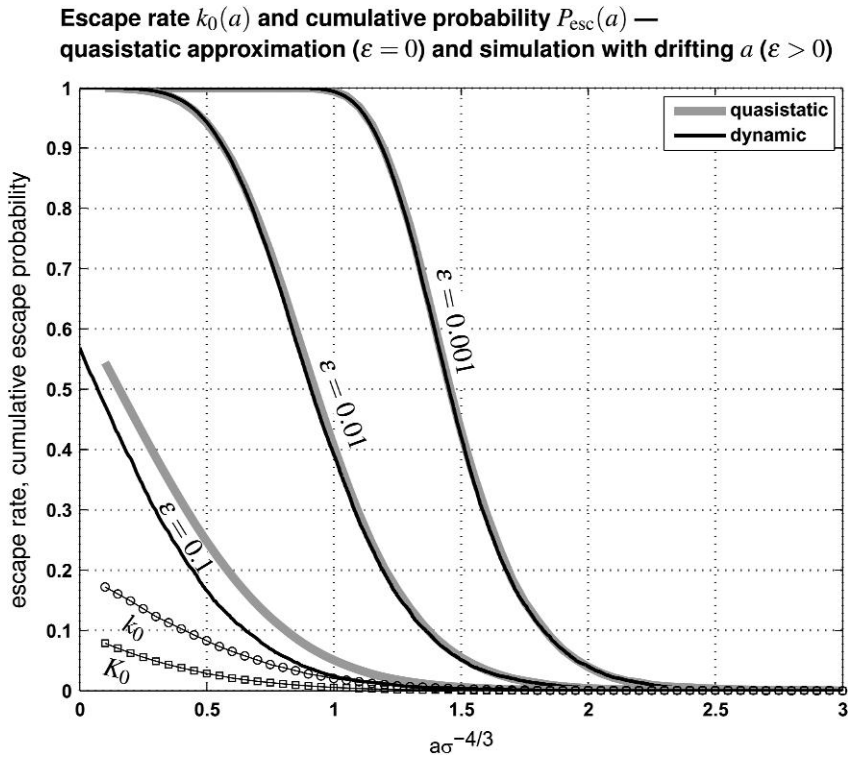


FIG. 4. Escape rate k_0 from potential well for fixed a (in the original parameters of (5.1)–(5.2)). The function $K_0(a)$ is the integral of k_0 , as defined in (5.8), integrated from a to 3. The curves $\epsilon = 0.1$, $\epsilon = 0.01$ and $\epsilon = 0.001$ are the cumulative probabilities for escape from the potential well for dynamically decreasing depth of the potential well before a is reached. The quantity ϵ indicates the rate of change of the potential in the unscaled parameters. For comparison the quasistatic estimate for the cumulative probability of escape is indicated using grey curves (see also Fig. 5).

if we start with an initial distribution concentrated in the potential well ($p(0) = 1$). We want to find the approximate cumulative probability $P_{\text{esc}}(a)$ of escape before the control has drifted to a certain value a , so we substitute $a = a(0) - \epsilon t$ into expression (5.7) for $p(t)$:

$$P_{\text{esc}}(a) = 1 - \exp(-K_0(a)/\epsilon) \quad \text{where } K_0(a) = \int_a^{a(0)} k_0(a') da'. \quad (5.8)$$

This approximation (5.8) for the escape probability assumes that ϵ is small, that the escape is irreversible (which is accurate if the threshold x_{th} is sufficiently negative), and that $a(0) \gg 1$ (which makes P_{esc} nearly independent of the initial distribution of x).

Figure 4 shows the quasistatic approximation of the cumulative escape probability P_{esc} for $\epsilon = 0.1, 0.01$ and 0.001 as grey curves. Superimposed are numerical observations of the cumulative escape probability P_{esc} for the normal form with drifting parameter a (5.4)–(5.5) as black curves. In a simulation of (5.4)–(5.5), one can approximate the cumulative escape probability P_{esc} until time step t_n with the




FIG. 5. Percentiles of the probability of escape for dynamically varying a (in black) in the original parameters of (5.1)–(5.2). For fixed ϵ the $x\%$ curve shows at which a $x\%$ of the realizations have escaped. For comparison, the grey curves show the estimate obtained using the quasistatic approximation. The depth of the potential well (and, thus, a) decreases in time.

help of a recursion for the probability $1 - P_{\text{esc}}$ of not escaping

$$1 - P_{\text{esc}}(t_{n+1}) = \frac{N - N_{\text{esc}}(t_n)}{N} (1 - P_{\text{esc}}(t_n)),$$

where $N_{\text{esc}}(t_n)$ is the number of realizations that escape at time step t_n and N is the overall ensemble size. We keep the overall ensemble size N (of non-escaped realizations) constant by re-initializing every escaped realization to a random non-escaped instance.

One can see that if the control parameter a drifts slowly ($\epsilon \ll 1$), the cumulative escape probability increases sharply from nearly 0 to 1 in a range of a of about length 1 (e.g. between $a = 1$ and $a = 2$ for $\epsilon = 0.001$).

Figure 5 shows the percentiles of the cumulative escape probabilities systematically for ϵ ranging between 0.001 and 0.5. Again, we superimpose black curves showing the percentiles of the cumulative probability observed during a simulation of the normal form with drifting control parameter (5.4)–(5.5).

We draw two conclusions from the results shown in Figs 4 and 5:

1. the quasistatic approximation (5.8) of the cumulative escape probability is quantitatively accurate to order ϵ . For large ϵ (~ 0.1), the effect of the dynamic drifting of the control parameter delays the escape slightly. Naturally, this effect is much weaker than the delay in exchange of stability observed in slow passages through Hopf or Pitchfork bifurcations (see Kuske, 1999; Baer *et al.*, 1989; Su *et al.*, 2004 for studies that quantify also the effect of noise).
2. Random disturbances make early escape probable as soon as the ratio of drift speed ϵ and variance of the disturbance σ^2 becomes small in the original parameters of saddle-node normal form

with drift and noise (5.1)–(5.2). The percentiles in Fig. 5 quantify this effect after the axes have been scaled using transformation (5.3): $\epsilon \mapsto \epsilon\sigma^{-2}$, $a \mapsto a\sigma^{-4/3}$.

5.1.4 Control of accuracy in stochastic simulations. The ensemble size for the numerical simulation was $N = 400$, and the integration was performed with the Euler–Maruyama scheme (which is of order 1 for this problem) using stepsize $h = 0.01$. Control calculations using different method parameters give results that are visually indistinguishable from Fig. 4. For control, we varied (one-by-one) the stepsize ($h = 0.05$), the threshold for a realization to count as escaped ($x_{\text{th}} = -10$) and the ensemble size ($N = 800$). We also varied the re-initialization strategy: alternatively, we chose the new position of a particle after escape according to a Gaussian distribution with mean $X_s = \sqrt{a}$ (the bottom of the well) and variance σ^2/\sqrt{a} . This would be the stationary distribution obtained for the linearization of the rescaled normal form (5.4) in its equilibrium X_s . This alternative re-initialization works by construction only for $a > 0$.

Figures 4 and 5 both show either long-time limits (such as $k_0(a)$ in Fig. 4) or cumulative quantities: $K_0(a)$ and P_{esc} . These quantities can be approximated more accurately with relatively small ensemble sizes and simple ensemble integration than probability densities. The density of the escape probability is the time derivative of P_{esc} , which would be a much ‘noisier’ function of time (or a) for finite randomly drawn ensembles than P_{esc} unless one applies more sophisticated numerical methods and uses larger ensemble sizes (Kuske, 1999, 2000).

For large a numerical simulations of the stochastic differential equation (5.4) become inefficient because escape events become extremely rare. However, this case is covered analytically by classical chemical reaction rate theory (Hänggi *et al.*, 1990):

$$k_0(a) \sim \frac{2\sqrt{a}}{\pi} \exp(-2\sqrt{a}) \quad \text{for } a \gg 1, \quad (5.9)$$

where the expression on the right-hand side contains the dominant terms of a in the pre-factor and in the exponent. Thus, replacing the upper limit of integration $a(0)$ by $+\infty$ in the approximation (5.8) of P_{esc} gives only a small change.

5.2 Extraction of noise-induced escape probability from time series

In order to estimate if noise-induced early escape plays a role, one needs to assume that the time series z_k is generated by a system with a parameter that drifts and approaches a saddle-node bifurcation. We denote the drifting parameter by a in a manner that the saddle-node bifurcation is at $a = 0$ and equilibria of the deterministic system with fixed a exist for $a > 0$. We introduce the parameter Z_0 as the z -value of the equilibrium at the saddle-node parameter $a = 0$. If we assume that the underlying process is subject to an additive (Gaussian) perturbation of amplitude σ_z , then z_k can be written as the measurements of the system

$$dz = \left[qa(t) - \frac{(z - Z_0)^2}{q} + O((z - Z_0)^3) \right] dt + \sigma_z dW_t \quad (5.10)$$

at times $k \Delta t$. The parameter q in (5.10) measures the width of the parabola $a(z)$ of equilibria of (5.10) for zero noise amplitude ($\sigma_z = 0$) and ignoring third-order powers of $z - Z_0$:

$$a(z) = \frac{(z - Z_0)^2}{q^2}.$$

If q is negative, then the stable side of the parabola is below Z_0 and if q is positive, then the stable side is above Z_0 . From the general form (5.10) of a system near a saddle-node, one can obtain the normal form

$$dx = [a(t) - x^2]dt + \sigma dW_t \quad (5.11)$$

by ignoring third-order powers of $z - Z_0$ and introducing the re-scaled variable

$$x = \frac{z - Z_0}{q}.$$

The noise amplitude scales correspondingly:

$$\sigma = \frac{\sigma_z}{q}. \quad (5.12)$$

We demonstrate how to extract a (crude) estimate of the state and the parameters of (5.10) from a time series z_k as shown for the model output in Fig. 3(a,b). As long as $a(t)$ is large, the potential well is deep such that the Gaussian perturbation $\sigma_z dW_t$ is unlikely to kick $z(t)$ out of the well of the stable equilibrium. Choosing Δt (the time spacing of the measurements z_k) as our time unit, we first revert the estimates c_k and θ_k obtained from the degenerate fingerprinting into estimates of the parameters $\kappa(t)$ and $\sigma(t)$ appearing in the Ornstein–Uhlenbeck process obtained from linearizing (5.10) in its stable equilibrium at $q\sqrt{a}$:

$$dz = -\kappa(t)zdt + \sigma_z(t)dW_t.$$

The estimates κ_k for $\kappa(k\Delta t)$ and σ_k for $\sigma(k\Delta t)$ are

$$\kappa_k = -\log c_k, \quad (5.13)$$

where the estimate c_k is calculated using the fitting procedure in Section 3.2 and

$$\sigma_{z,k} = \theta_k \sqrt{\frac{2\kappa_k}{(1 - c_k^2)}} \quad (5.14)$$

(see Gillespie, 1992), where θ_k is the standard deviation of the detrended time series from the AR(1) estimate. The estimate κ_k is an approximation of the LDR

$$\kappa(k\Delta t) = 2\sqrt{a(k\Delta t)} \quad (5.15)$$

as long the time series z_k is near the bottom of the potential well. Thus, we can estimate $a(k\Delta t)$ for the times $k\Delta t$ as a_k :

$$a_k = \frac{\kappa_k^2}{4}.$$

Consequently, estimates of the parameters a and σ_z are byproducts of the AR(1) estimate to obtain c_k . The only remaining unknown quantity that one needs to convert to the normal form (5.11) is the scaling factor q . If we drop third-order terms of $z - Z_0$, fix a at a_k , and consider only the deterministic part ($\sigma_z = 0$), then the equilibrium Z_k of (5.10) satisfies $\sqrt{a_k}q - Z_k + Z_0 = 0$, which is equivalent to

$$\frac{\kappa_k}{2}q + Z_0 = Z_k. \quad (5.16)$$

We see that q is the ratio between the slope of $\kappa_k/2$ (for which we have an estimate) and the slope of the equilibrium state that the time series z_k fluctuates around. An estimate for Z_k has also been obtained during the degenerate fingerprinting procedure as the kernel average of z_k (also named Z_k in Section 3.2). Thus, an estimate for q is the ratio between the mean slope of Z_k and the mean slope of κ_k .

In our analysis of the normal form with drift and noise, we considered the case where the parameter a drifts with a uniform (small) speed ϵ . This is unrealistic for paleoclimate records and in models whenever the parameter a is driven by the output of another sub-system, for example, if a is freshwater forcing as in Rahmstorf (2000). In order to predict how the parameter a continues to drift beyond the final sliding window, we assume that the process driving a is stationary. So,

$$a(t) = a(0) - \int_0^t \epsilon(s) ds, \quad (5.17)$$

where ϵ is not constant but a random variable with a stationary probability distribution. The distribution has been estimated for the period of time where κ is available: ϵ_k are the increments between successive a estimates:

$$\epsilon_k = \frac{a_k - a_{k+1}}{\Delta t}.$$

These ϵ_k form an empirical sample of the distribution of $\epsilon(s)$ from which we can draw to calculate $a(t)$ using (5.17) without estimating any further parameters. If the mean of ϵ exists and is bigger than zero, then $a(t)$ will reach 0 almost surely, resulting in a probability distribution

$$P_a(t) = 1 - P(a(s) > 0 \text{ for all } s < t),$$

which is the probability that the random variable a reaches its critical value $a = 0$ before time t . The probability $P_{\text{esc}}(t)$ of a trajectory escaping before time t can now be estimated using a direct numerical simulation of

$$\begin{aligned} dx &= [a - x^2] dt + \sigma dW_t, \\ da &= -\epsilon(t) dt. \end{aligned}$$

Alternatively, if the mean of ϵ is sufficiently small, one can use the quasistatic approximation

$$P_{\text{esc}}(t) = 1 - E \left[\exp \left(\int_0^t k_0(a(s)) ds \right) \right],$$

where the escape rate $k_0(a)$ (which exists only for positive a) is given in Fig. 4 and $E[x]$ is the mean of the random variable x .

Figure 6(b) and (c) show the quantities κ_k , a_k and $\sigma_{z,k}$ for the time series z_k of the stochastic model output in Fig. 6(a) (the same series z_k as in Fig. 3(a)). Figure 6(b) also shows the cumulative probability function $P_a(t)$ for a crossing the critical value 0 before t (dashed curve) and its quartiles (dashed threshold lines), and the probability $P_{\text{esc}}(t)$ for escape before time t (solid curve) and its quartiles (solid threshold lines). The threshold for a trajectory counting as escaped was set at $x_{\text{th}} = -5$. Table 1 lists the values of q as determined by relation (5.16) and $\sigma = \sigma_z/q$. Also, the table lists the empirical mean value of the random variable ϵ and its empirical standard deviation (the distribution of ϵ is only moderately skewed). The mean of ϵ/σ^2 indicates if the parameter drift is rapid or slow in rescaled normal

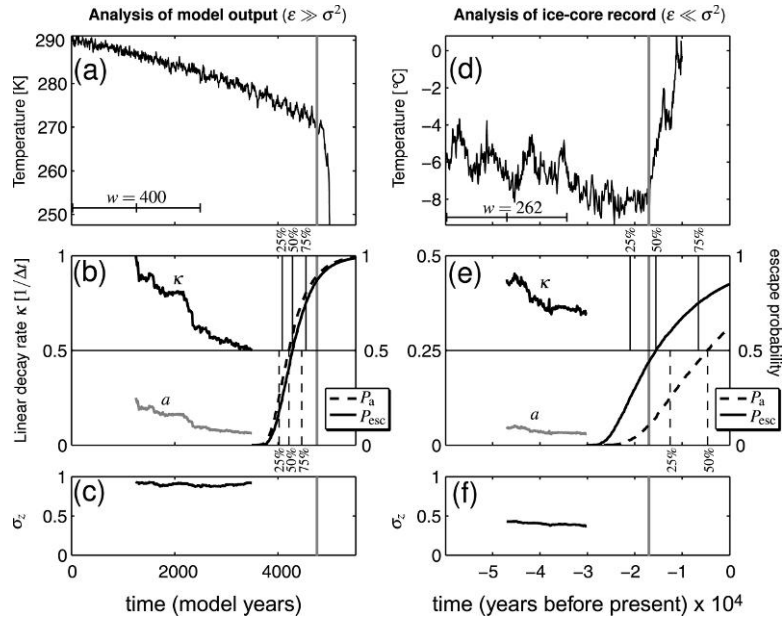


FIG. 6. Prediction of saddle-node bifurcation and early escape for two time series shown in Dakos *et al.* (2008). Diagram (a) shows the stochastic model output also shown in Fig. 3(a). All quantities have been obtained only from data prior to the cut-off time indicated by the grey vertical line in diagram (a). Diagram (b) shows the linear decay rate κ_k (as extracted by degenerate fingerprinting with window size w , compare to Fig. 3(b)) and the corresponding normal form parameter a . It also contains the cumulative probability functions for reaching the critical value $a = 0$ and for escape, P_a and P_{esc} on the right axis. The graph in diagram (c) is the estimated noise amplitude σ extracted from the time series. The diagrams (d–f) discuss the time series shown in Fig. 3(c), which is a snapshot from the paleoclimate record of Petit *et al.* (1999) at the end of the last glaciation using the identical procedure to the one applied to the time series in diagram (a).

form coordinates. Figure 6(e) and (f) show the quantities κ , a , P_{esc} , P_a and σ_z for the ice-core record in Fig. 6(d) (the same series z_k as in Fig. 3(c)).

We observe that the parameter drift in the model output (see Fig. 6(b)) is fast compared to the noise level such that early escape plays no role. Table 1 also shows that the ratio between drift and variance is large, which confirms the visual impression that drift dominates the noise. The probability distribution for escape is even shifted to the right such that trajectories are expected to escape ‘later’ than the drifting system parameter a reaches its critical value.

We note that the actual escape of the observed instance from Fig. 6(a) occurs relatively late (at the 90% percentile). This shifts down (to below 75%) for shorter sliding windows in the degenerate fingerprinting procedure. Both probability distributions are relatively symmetric and concentrated in a range of approximately 1000 model years.

In contrast, the time series of the ice-core record, shown in Fig. 6(d), has a slowly drifting parameter compared to the noise level (note that one cannot be certain that the underlying mechanism for the apparent tipping is indeed a passage through a saddle-node). Consequently, trajectories of the estimated saddle-node normal form system escape significantly earlier than a reaches its critical value: the cumulative probability P_{esc} is shifted to the left of P_a . For example, the median (50%) time for escape is 1.1×10^4 years before the median time for reaching the critical value $a = 0$. We note that both distributions, P_a and P_{esc} are highly skewed, having a long tail at the right end. This makes the expected

TABLE 1 *Numerical values of estimates of parameters that are not visible in Fig. 6*

Parameter	Model output Fig. 6(a–c)	Ice-core record Fig. 6(d–f)
Mean ϵ	2.2×10^{-4}	5.2×10^{-6}
Standard deviation ϵ	1.1×10^{-3}	5.9×10^{-5}
q	621	13.5
σ	1.4×10^{-3}	3.0×10^{-2}
Mean $\epsilon\sigma^{-2}$	105	5.9×10^{-3}

escape time (mean escape time) as a point estimate sensitive to small perturbations such as measurement uncertainty or a different choice of method parameters in the degenerate fingerprinting procedure, whereas the median times are comparatively robust. A systematic analysis, including the dependence on the cut-off time (the grey vertical line in all panels of Fig. 6), is a topic for future work.

6. Conclusion

Methods to identify incipient climate tipping using time series analysis have recently been developed by Held & Kleinen (2004) and Livina & Lenton (2007). These methods have been tested on model outputs and paleoclimate data by Livina & Lenton (2007) and Dakos *et al.* (2008). Time series analysis should be seen as a complement to the huge modelling efforts that are invested into the analysis and prediction of climate changes. While Dakos *et al.* (2008) could demonstrate that the characteristic quantity extracted from the time series, the propagator, indeed increases (as it should according to bifurcation theory) using this quantity for prediction is much more challenging. The extraction of the propagator from a time series makes assumptions that are, for geological records, difficult to check: separation of time scales between parameter drift, decay of critical mode and decay of stable modes or the nearly constant-speed approach of the underlying control parameter towards its critical value.

We studied another source of uncertainty, which exists even if the underlying deterministic dynamics is drifting with constant speed close to a saddle-node bifurcation and the estimated LDR is accurate. Namely, the escape of the dynamics from the potential well around the stable equilibrium can be either premature or delayed, depending on the ratio between parameter drift speed and the amplitude of the random disturbances. We derived an approximate (semi-analytic) formula that is valid in the quasistatic limit (that is, the parameter drifts sufficiently slowly compared to the noise amplitude). We found that the early escape effect vanishes if the drift is more rapid (drift speed $\epsilon \approx 0.3$ in normal form scaling).

We also demonstrated how one can estimate this effect from time series data, using what we might call a ‘fold for incipient tipping’ method. This was tested for two examples, the output of a stochastic model (a case of rapid drift) and an ice-core record (a case of slow drift or large noise amplitude). We plan to study the consistency of the proposed escape prediction for time series more systematically in the future. This should be done by generating time series instances from saddle-node normal forms with drifting parameter and noise, predicting escape distributions from these time series, and then comparing the predictions to the original distributions shown in Fig. 5.

Our estimates for early escape from a potential well are all stated in probabilistic terms, which is appropriate if one treats the disturbances as random. Normal form based estimates such as ours may also be useful when studying tipping in climate models because running simulations for large ensembles of realizations in sophisticated climate models is expensive. If the modelled scenario is close to tipping and

deterministic calculations reveal the presence of a saddle-node bifurcation, then the diagrams in Figs 4 and 5 help to give estimates for the escape probability depending on the speed of parameter drift and the amplitude of the disturbances.

Finally, we note that our scenario corresponds to the escape from a potential well of an overdamped particle (this appears to be the relevant case, for example, in ocean circulation models, see [Dijkstra et al., 2004](#)). The weakly damped case,

$$\frac{d^2x}{dt^2} + \gamma \frac{dx}{dt} = -U(x) + \sigma \zeta(t),$$

where γ is a small coefficient determining the amount of damping ($0 < \gamma \ll 1$), U is the potential and ζ is the disturbance (for example, a random Gaussian noise) has also been covered by reaction rate theory (the classical theory treats the case where the noise amplitude σ is much smaller than the potential barrier, see [Hänggi et al., 1990](#) for a review).

REFERENCES

- ALLEY, R. B., MAROTZKE, J., NORDHAUS, W. D., OVERPECK, J. T., PETEET, D. M., PIELKE, R. A., PIERREHUMBERT, R. T., RHINES, P. B., STOCKER, T. F., TALLEY, L. D. & WALLACE, J. M. (2003) Abrupt climate change. *Science*, **299**, 2005–2010.
- BAER, S. M., ERNEUX, T. & RINZEL, J. (1989) The slow passage through a Hopf bifurcation: delay, memory effects, and resonance. *SIAM J. Appl. Math.*, **49**, 55–71.
- COPENHAGEN (2009) *United Nations Climate Change Conference*, 7–18 December 2009. Copenhagen.
- CORSI, S. & TARANTO, G. N. (2007) Voltage instability—the different shapes of the nose. *iREP Symposium—Bulk Power System Dynamics and Control—VII, Revitalizing Operational Reliability*. Charleston, SC: IEEE.
- DAKOS, V., SCHEFFER, M., VAN NES, E. H., BROVKIN, V., PETOUKHOV, V. & HELD, H. (2008) Slowing down as an early warning signal for abrupt climate change. *Proc. Natl. Acad. Sci. U.S.A.*, **105**, 14308–14312.
- DIJKSTRA, H. A., TE RAA, L. & WEIJER, W. (2004) A systematic approach to determine thresholds of the ocean's thermohaline circulation. *Tellus Ser A: Dyn. Meteorol. Oceanogr.*, **56**, 362–370.
- FISCHER, M. & IMKELLER, P. (2006) Noise-induced resonance in bistable systems caused by delay feedback. *Stoch. Anal. Appl.*, **24**, 135–194.
- FRAEDRICH, K. (1978) Structural and stochastic analysis of a zero-dimensional climate system. *Q. J. R. Meteorol. Soc.*, **104**, 461–474.
- GILLESPIE, D. T. (1992) *Markov Processes: An Introduction for Physical Scientists*. New York: Academic Press.
- HÄNGGI, P., TALKNER, P. & BORKOVEC, M. (1990) Reaction-rate theory: fifty years after Kramers. *Rev. Mod. Phys.*, **62**, 251–341.
- HELD, H. & KLEINEN, T. (2004) Detection of climate system bifurcations by degenerate fingerprinting. *Geophys. Res. Lett.*, **31**, L23207.
- HOUGHTON, J. T. (2004) *Global Warming: The Complete Briefing*. Cambridge, UK: Cambridge University Press.
- INTERGOVERNMENTAL PANEL ON CLIMATE CHANGE (2007) *Climate Change 2007, Contribution of Working Groups I–III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, I The Physical Science Basis, II Impacts, Adaptation and Vulnerability, III Mitigation of Climate Change*. Cambridge, UK: Cambridge University Press.
- KUSKE, R. (1999) Probability densities for noisy delay bifurcations. *J. Stat. Phys.*, **96**, 797–816.
- KUSKE, R. (2000) Gradient-particle solutions of Fokker–Planck equations for noisy delay bifurcations. *SIAM J. Sci. Comput.*, **22**, 351–367.
- LENTON, T. M., HELD, H., KRIEGLER, E., HALL, J. W., LUCHT, W., RAHMSTORF, S. & SCHELLNHUBER, H. J. (2008) Tipping elements in the Earth's climate system. *Proc. Natl. Acad. Sci. U.S.A.*, **105**, 1786–1793.

- LIVINA, V. N. & LENTON, T. M. (2007) A modified method for detecting incipient bifurcations in a dynamical system. *Geophys. Res. Lett.*, **34**, L03712.
- LOCKWOOD, J. (2001) Abrupt and sudden climatic transitions and fluctuations: a review. *Int. J. Climatol.*, **21**, 1153–1179.
- MALAKHOV, A. & PANKRATOV, A. L. (1996) Influence of thermal fluctuations on time characteristics of a single Josephson element with high damping exact solution. *Phys. C: Supercond.*, **269**, 46–54.
- NATIONAL RESEARCH COUNCIL (2002) *Abrupt Climate Change: Inevitable Surprises*. Washington, DC: National Academic Press.
- PETIT, J. R., JOUZEL, J., RAYNAUD, D., BARKOV, N. I., BARNOLA, J.-M., BASILE, I., BENDER, M., CHAPPELLAZ, J., DAVIS, M., DELAYGUE, G., DELMOTTE, M., KOTLYAKOV, V. M., LEGRAND, M., LIPENKOV, V. Y., LORIUS, C., PÉPIN, L., RITZ, C., SALTZMAN, E. & STIEVENARD, M. (1999) Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature*, **399**, 429–436.
- RAHMSTORF, S. (2000) The thermohaline ocean circulation: a system with dangerous thresholds? *Clim. Change*, **46**, 247–256.
- RAHMSTORF, S. (2001) Abrupt climate change. *Encyclopaedia of Ocean Sciences* (J. Steele, S. Thorpe & K. Turekian, eds). London: Academic Press, p. 1–6.
- RIAL, J. A., PIELKE, R. A., BENISTON, M., CLAUSSEN, M., CANADELL, J., COX, P., HELD, H., DE NOBLET-DUCOUDRÉ, N., PRINN, R., REYNOLDS, J. F. & SALAS, J. D. (2004) Nonlinearities, feedbacks and critical thresholds within the Earth's climate system. *Clim. Change*, **65**, 11–38.
- SCHEFFER, M. (2009) *Critical Transitions in Nature and Society*. Princeton, NJ: Princeton University Press.
- SU, J., RUBIN, J. & TERMAN, D. (2004) Effects of noise on elliptic bursters. *Nonlinearity*, **17**, 133–157.
- THOMPSON, J. M. T. & SIEBER, J. (2010) Predicting climate tipping points. *Geo-Engineering Climate Change: Environmental necessity or Pandora's Box?* (B. Launder & J. M. T. Thompson eds), Chapter 3. Cambridge, UK: Cambridge University Press, pp. 50–83.
- THOMPSON, J. M. T. & STEWART, H. B. (2002) *Nonlinear Dynamics and Chaos*, 2nd edn. Chichester, UK: Wiley.
- THOMPSON, J. M. T., STEWART, H. B. & UEDA, Y. (1994) Safe, explosive, and dangerous bifurcations in dissipative dynamical systems. *Phys. Rev. E*, **49**, 1019–1027.