PREDICTING CLIMATE TIPPING AS A NOISY BIFURCATION: A REVIEW

J. MICHAEL T. THOMPSON

Department of Applied Mathematics & Theoretical Physics,
Cambridge University, Centre for Mathematical Sciences,
Wilberforce Road, Cambridge, CB3 0WA, UK

School of Engineering (Sixth Century Professor),
Aberdeen University, UK
jmtt@ucl.ac.uk

JAN SIEBER

Department of Mathematics, University of Portsmouth,
Portsmouth, PO1 3HF, UK
jan.sieber@port.ac.uk

There is currently much interest in examining climatic tipping points, to see if it is feasible to predict them in advance. Using techniques from bifurcation theory, recent work looks for a slowing down of the intrinsic transient responses, which is predicted to occur before an instability is encountered. This is done, for example, by determining the short-term autocorrelation coefficient ARC(1) in a sliding window of the time-series: this stability coefficient should increase to unity at tipping. Such studies have been made both on climatic computer models and on real paleoclimate data preceding ancient tipping events. The latter employ reconstituted time-series provided by ice cores, sediments, etc., and seek to establish whether the actual tipping could have been accurately predicted in advance. One such example is the end of the Younger Dryas event, about 11 500 years ago, when the Arctic warmed by 7°C in 50 yrs. A second 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, using data from tropical Pacific sediment cores. This prediction science is very young, but some encouraging results are already being obtained. Future analyses will clearly need to embrace both real data from improved monitoring instruments, and simulation data generated from increasingly sophisticated predictive models.

Keywords: Climate tipping; bifurcation prediction; time-series analysis.

1. Introduction

Predicting the future climate is now a major challenge to the world, as witnessed by the recent Copenhagen Conference and its sequels. In studying changes to the Earth's climate, perhaps the most important feature to watch out for, and try to anticipate, is a so-called tipping point at which the climate makes a sudden, and often irreversible, change. Major events of this type are well documented in geological records, striking examples being the on-and-off switching of prehistoric ice ages, as illustrated in Fig. 1. The current reason for concern is the apparently coordinated increase of the average global temperature and the percentage of carbon dioxide in the atmosphere, as illustrated in Fig. 2. Many scientists believe, firstly, that the
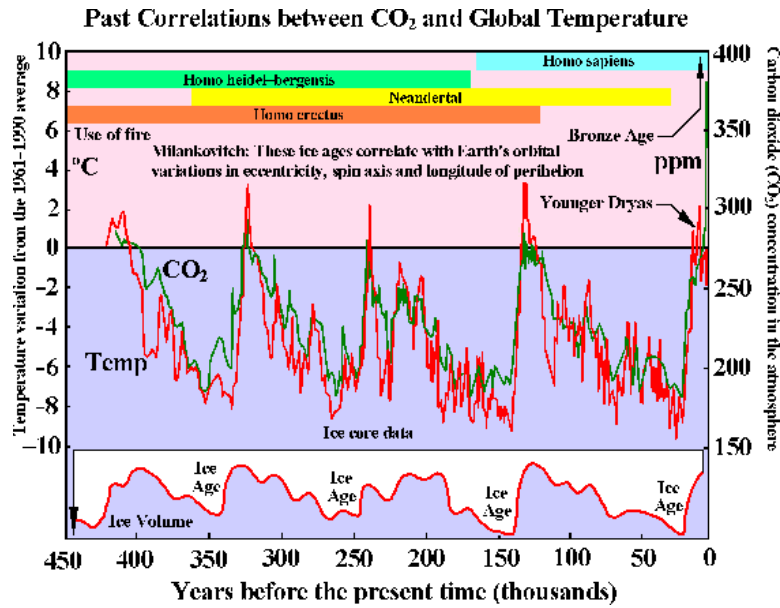
Fig. 1.   Correlation between the CO_2 concentration and temperature during the prehistoric ice ages, showing the variation of the ice volume and human development.

rise in CO_2 concentration is due to human activity (especially the burning of fossil fuels), and secondly that it is this rise in CO_2 levels that is causing the rise in temperature. The concern, then, is that this temperature rise (which itself would have serious consequences for agriculture and flooding) might cause a climate tipping of major proportions.

Fig. 2.   Correlation between the CO_2 concentration and temperature during recent times, illustrating the concern about global warming. Based on a composite picture from the Woods Hole Research Centre.

The analysis and prediction of tipping points, often focused on climate subsystems, is currently being pursued in several streams of research, and we should note in particular the excellent book by Scheffer [2009], *Critical Transitions in Nature and Society*, which includes ecological and climatic studies. Some brief remarks about abrupt and rapid climate change were made by the Intergovernmental Panel on Climate Change [IPCC, 2007], and more recently Lenton *et al.* [2008] have sought to define these points rigorously.

Table 1 shows a list of subsystem candidates proposed by Lenton *et al.* [2008], and the possible effects of their tipping on the global climate. All of these subsystems have strong internal positive feedback mechanisms. Thus, they have a certain propensity for tipping and are susceptible to input (human or otherwise). We should note that Lenton's first element is the shrinking of the Arctic summer sea ice, which is progressing rapidly and is causing considerable concern. The changes over a 23 year span are shown in two NASA photographs in Fig. 3. As column 2 of Table 1 shows, the primary deterministic mechanisms behind several of the listed tipping events are so-called *bifurcations*, special points in the control parameter space (see columns 4 and 5) at which the deterministic part of the dynamical system governing the climate changes qualitatively (for example, the currently attained steady state disappears).
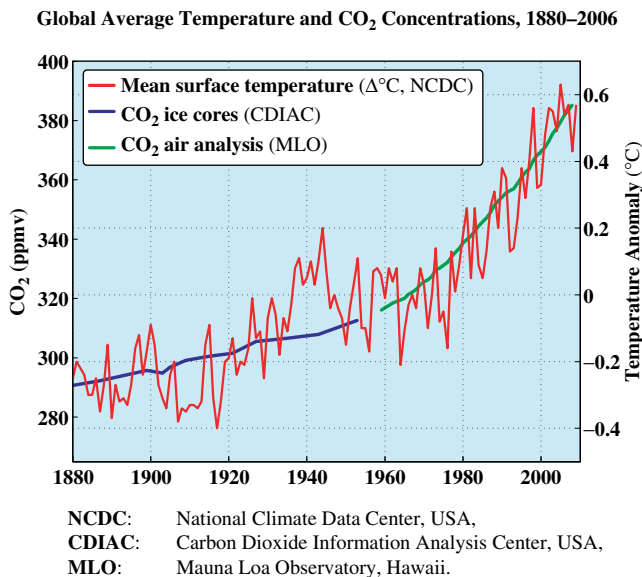
Table 1. Summary of Lenton's Tipping Elements, namely climate subsystems that are likely to be candidates for future tipping with relevance to political decision making. In column 2, the possibility of there being an underlying bifurcation is indicated as follows: black = high, grey = medium, white = low. Notice that in column 4 EEP denotes the Eastern Equatorial Pacific and in the last column ITCZ denotes the Intertropical Convergence Zone. This list will be discussed in greater detail in Sec. 5.

| Tipping Element		Feature, F (Change)	Control Parameter, μ	μ_{crit}	Global Warming	Transition Time, T	Key Impacts
Arctic summer sea-ice	(grey)	Areal extent $(-)$	Local ΔT_{air}, ocean heat transport	??	$+0.5$ to $+2°C$	~ 10 yrs (rapid)	Amplified warming, ecosystem change
Greenland ice sheet (GIS)	(black)	Ice volume $(-)$	Local ΔT_{air}	$\sim +3°C$	$+1$ to $+2°C$	>300 yrs (slow)	Sea level $+2$ to $+7$ m
West antarctic ice sheet (WAIS)	(grey)	Ice volume $(-)$	Local ΔT_{air} or, less ΔT_{ocean}	$+5$ to $+8°C$	$+3$ to $+5°C$	>300 yrs (slow)	Sea level $+5$ m
Atlantic thermohaline circulation	(black)	Overturning $(-)$	Freshwater input to North Atlantic	$+0.1$ to $+0.5$ Sv	$+3$ to $+5°C$	~ 100 yrs (gradual)	Regional cooling, sea level, ITCZ shift
El Niño Southern oscillation	(grey)	Amplitude $(+)$	Thermocline depth, sharpness in EEP	??	$+3$ to $+6°C$	~ 100 yrs (gradual)	Drought in SE Asia and elsewhere
Indian summer monsoon (ISM)	(grey)	Rainfall $(-)$	Planetary albedo over India	0.5	N/A	~ 1 yr (rapid)	Drought, decreased carrying capacity
Sahara/Sahel and W. African monsoon	(grey)	Vegetation fraction $(+)$	Precipitation	100 mm/yr	$+3$ to $+5°C$	~ 10 yrs (rapid)	Increased carrying capacity
Amazon rain-forest	(grey)	Tree fraction $(-)$	Precipitation, dry season length	1100 mm/yr	$+3$ to $+4°C$	~ 50 yrs (gradual)	Biodiversity loss, decreased rainfall
Boreal forest	(white)	Tree fraction $(-)$	Local ΔT_{air}	$\sim +7°C$	$+3$ to $+5°C$	~ 50 yrs (gradual)	Change in type of the ecosystem
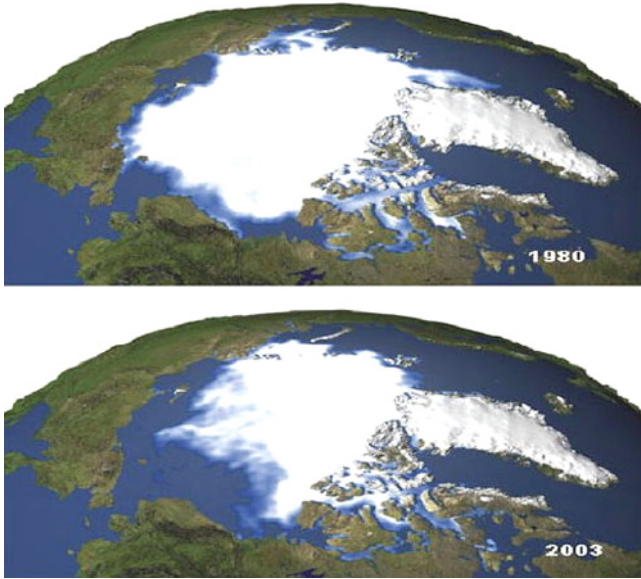
Fig. 3.   Two NASA satellite photographs, showing the reduction of Arctic snow and ice cover over a 23 yrs interval. *Source*: http://nasascience.nasa.gov/images/about-us/accomplishments/YIR2004_arctic.jpg.

In Sec. 3 we review possible bifurcations and classify them into three types, *safe*, *explosive* and *dangerous*. Almost universally these bifurcations have a precursor: in at least one mode, all feedback effects cancel at the linear level, which means that the system is slowing down, and the local (or linear) decay rate (LDR) to the steady state decreases to zero.

Most of the relevant research is devoted to creating climate models from first principles, tuning and initializing these models by assimilating geological data, and then running simulations of these models to predict the future. Climate models come in varying degrees of sophistication and realism, more complex ones employing up to 3×10^8 variables [Dijkstra, 2008]. Predictions do not rely solely on a single "best model" starting from the "real initial conditions". Typically, all qualified models are run from ensembles of initial conditions and then a statistical analysis over all generated outcomes is performed [IPCC, 2007].

An alternative to the *model and simulate* approach (and in some sense a short-cut) is to realize that mathematically some of the climate-tipping events correspond to *bifurcations* (see Sec. 3 for a discussion), and then to use time-series analysis techniques to extract precursors of these bifurcations directly from observational data. This method

still benefits from the modeling efforts because simulations generated by predictive models allow analysts to hone their prediction techniques on masses of high quality data, with the possibility of seeing whether they can predict what the computer eventually displays as the outcome of its run. Transferring these techniques to *real data* from the Earth itself is undoubtedly challenging. Still, bifurcation predictions directly from real time-series will be a useful complement to modeling from first principles because they do not suffer from all the many difficulties of building and initializing reliable computer models. Our review discusses the current state of bifurcational predictions in climate time-series, focusing on methods introduced by Held and Kleinen [2004] and Livina and Lenton [2007]. Held and Kleinen analyze the collapse of the global conveyor belt of oceanic water, the thermohaline circulation (THC). This conveyor is important, not only for the water transport, *per se*, but because of the heat and salt that it redistributes.

The paper by Livina and Lenton [2007] is particularly noteworthy in that it includes what seems to be the first bifurcational predictions using real data, namely the Greenland ice-core paleo-temperature data spanning the time from 50 000 years ago to the present. The unevenly spaced data comprised 1586 points and their DFA-propagator (this quantity reaches +1 when the local decay rate vanishes; see Sec. 4.1) was calculated in sliding windows of length 500 data points. The results are shown in Fig. 4, and the rapid warming at the end of the Younger Dryas event, around 11 500 yrs before the present is anticipated by an upward trend in the propagator, which is heading towards its critical value of +1 at about the correct time. With the data set running over tens of thousands of years, this study should be seen primarily as an estimate of the end of the last glaciation, rather than the Younger Dryas event itself. The sliding window that ends near the tipping is highlighted, and we note that (as we emphasize at the end of Sec. 4.1), from a prediction point of view, the propagator estimates would end at point A. The grey propagator curve beyond A uses time-series points beyond the tipping point, which would not normally be available: in any event, they should not be used, because they contaminate the grey results with data from a totally different climatic state.

In a second notable paper, Dakos *et al.* [2008] systematically estimated the LDR for real data in
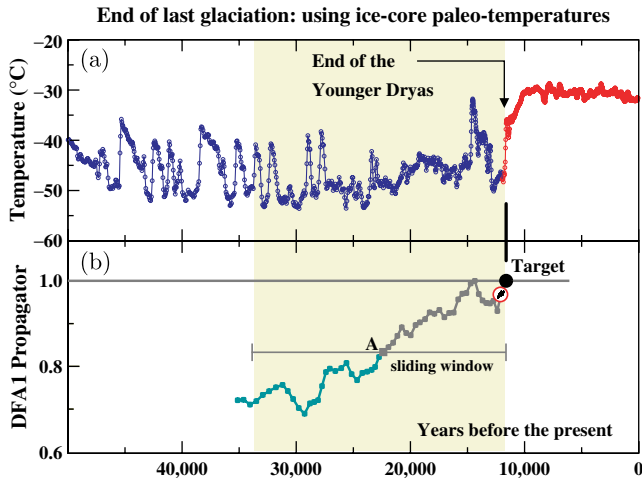
End of last glaciation: using ice-core paleo-temperatures

Fig. 4. Results of Livina and Lenton [2007] for the end of the last glaciation (a) Greenland ice-core (GISP2) paleo-temperature with an unevenly spaced record, visible in the varying density of symbols on the curve. The total number of data points is $N = 1586$. In (b) the DFA1-propagator is calculated in sliding windows of length 500 points and mapped into the middle points of the windows. The results of a second and much more local study by Dakos *et al.* [2008] (that we shall be discussing in Fig. 14) are highlighted by the red circle.

their analysis of eight ancient tipping events via reconstructed time-series. 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–Alleröd transition which 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 when the Arctic warmed by 7°C in 50 yrs, drawing on data from the sediment of the Cariaco basin in Venezuela. This examines at a much shorter time-scale, and with different data, the transition of Fig. 4,

(e) the desertification of North Africa when there was a sudden shift from a savanna-like state with scattered lakes to a desert about 5000 years ago, using the sediment core from ODP Hole 658C, off the west coast of Africa.

In all of these cases, the dynamics of the system are shown to slow down before the transition. This

slow-down was revealed by a short-term autocorrelation coefficient, ARC(1), of the time-series which examines to what extent a current point is correlated to its preceding point. It gives an estimate of the LDR, and is expected to increase towards unity at an instability, as described in Sec. 4.

2. Climate Models as Dynamical Systems

Thinking about modeling is a good introduction to the ideas involved in predicting climate change, so we will start from this angle. Now, to an applied mathematician, the Earth's climate is just a very large dynamical system that evolves in time. Vital elements of this system are the Earth itself, its oceans and atmosphere, and the plants and animals that inhabit it (including, of course, ourselves). In summary, the five key components are often listed succinctly as atmosphere, ocean, land, ice and biosphere. Arriving as external stimuli to this system are sunlight and cosmic rays, etc.: these are usually viewed as driving forces, often just called *forcing*. In modeling the climate we need not invoke the concepts of quantum mechanics (for the very small) or relativity theory (for the very big or fast).

So one generally considers a system operating under the deterministic rules of classical physics, employing, for example, Newton's Laws for the forces, and their effects, between adjacent large blocks of sea water or atmosphere. A block in the atmosphere might extend 100 km by 100 km horizontally and 1 km vertically, there being perhaps 20 blocks stacked vertically over the square base: for example, in a relatively low resolution model, Selten *et al.* [2004] used blocks of size 3.75° in latitude and longitude with 18 blocks stacked vertically in their simulation. (For current high resolution models see [IPCC, 2007].) So henceforth in this section, we will assume that the climate has been modeled primarily as a large deterministic dynamical system evolving in time according to fixed rules. For physical, rather than biological entities, these rules will usually relate to adjacent (nearest-neighbor) objects at a particular instant of time (with no significant delays or memory effects). It follows that our climate model will have characteristics in common with the familiar mechanical systems governed by Newton's laws of motion. From a given set of starting conditions (positions and velocities of all the components, for example), and external deterministic *forcing* varying in a prescribed fashion with

time, there will be a unique outcome as the model evolves in time. Plotting the time-evolution of these positions and velocities in a conceptual multidimensional *phase space* is a central technique of dynamical systems theory.

Despite the unique outcome, the results of chaos theory remind us that the response may be essentially unknown over time-scales of interest because it can depend with infinite sensitivity on the starting conditions (and on the numerical approximations used in a computer simulation). To ameliorate this difficulty, weather and climate forecasters now often make a series of parallel simulations from an ensemble of initial conditions which are generated by adding different small perturbations to the original set: and they then repeat all of this on different models. This ensemble approach, pioneered by Tim Palmer and others, is described by Buizza *et al.* [1998] and Sperber *et al.* [2001].

Mechanical systems are of two main types. First is the idealized closed *conservative* (sometimes called *Hamiltonian*) system in which there is no input or output of energy, which is therefore conserved. These can be useful in situations where there is very little "friction" or energy dissipation, such as when studying the orbits of the planets. A conservative system, like a pendulum with no friction at the pivot and no air resistance, tends to move for ever: it does not exhibit transients, and does not have any *attractors*. Second, is the more realistic dissipative system where energy is continuously lost (or dissipated). An example is a real pendulum which eventually comes to rest in the hanging-down position, which we call a *point attractor*. A more complex example is a damped pendulum driven into resonance by steady harmonic forcing from an AC electromagnet: here, after some irregular transient motion, the pendulum settles into a stable "steady" oscillation, such as a *periodic attractor* or a *chaotic attractor*. In general, a dissipative dynamical system will settle from a complex transient motion to a simpler attractor as the time increases towards infinity. These attractors, the stable steady states of the system, come in four main types: the point attractors, the periodic attractors, the quasi-periodic (toroidal) attractors and the chaotic attractors [Thompson & Stewart, 2002].

Climate models will certainly not be conservative, and will dissipate energy internally, though they also have some energy input: they can be reasonably expected to have the characteristics of the well-studied dissipative systems of (for example) engineering mechanics, and are, in particular, well known to be highly nonlinear.

3. Concepts from Bifurcation Theory

A major component of nonlinear dynamics is the theory of bifurcations, these being points in the slow evolution of a system at which qualitative changes or even sudden jumps of behavior can occur.

In the field of dissipative dynamics *codimension-1* bifurcations are those events that can be "typically" encountered under the slow sweep of a single control parameter. A climate model will often have (or be assumed to have) such a parameter under the quasi-static variation of which the climate is observed to gradually evolve on a "slow" time-scale. Slowly varying parameters are external influences that vary on geological time-scales, for example, the obliquity of the Earth's orbit. Another common type of slowly varying parameter occurs if one models only a subsystem of the climate, for example, oceanic water circulation. Then the influence of an interacting subsystem (for example, freshwater forcing from melting ice sheets) acts as a parameter that changes slowly over time.

An encounter with a bifurcation during this evolution will be of great interest and significance, and may give rise to a dynamic jump on a much faster time-scale. A complete list of the (typical) codimension-1 bifurcations, to the knowledge of the authors at the time of writing, is given by Thompson and Stewart [2002]. It is this list of local and global bifurcations that is used to populate Tables 2–5. The technical details and terminology of these tables need not concern the general reader, but they do serve to show the vast range of bifurcational phenomena that can be expected even in the simplest nonlinear dynamical systems, and certainly in climate models.

A broad classification of the codimension-1 *attractor bifurcations* of dissipative systems into *safe*, *explosive* and *dangerous* forms [Thompson *et al.*, 1994] is illustrated in Tables 2–4 and Fig. 5, while all are summarized in Table 5 together with notes on their precursors. It must be emphasized that these words are used in a technical sense. Even though in general the safe bifurcations are often literally safer than the dangerous bifurcations, in certain contexts this may not be the case. In particular, the safe bifurcations can still be in a literal

Table 2.   Safe bifurcations. These include the supercritical forms of the local bifurcations and the less well-known global "band merging". The latter is governed by a saddle-node event on a chaotic attractor. Alternative names are given in brackets.

Safe Bifurcations	
(a) Local Supercritical Bifurcations	
1. Supercritical Hopf	Point to cycle
2. Supercritical Neimark-Sacker (secondary Hopf)	Cycle to torus
3. Supercritical Flip (period-doubling)	Cycle to cycle
(b) Global Bifurcations	
4. Band Merging	Chaos to chaos

These bifurcations are characterized by the following features:

SUBTLE: continuous supercritical growth of new attractor path
SAFE: no fast jump or enlargement of the attracting set
DETERMINATE: single outcome even with small noise
NO HYSTERESIS: path retraced on reversal of control sweep
NO BASIN CHANGE: basin boundary remote from attractors
NO INTERMITTENCY: in the responses of the attractors

sense very dangerous: as when a structural column breaks at a "safe" buckling bifurcation!

Note carefully here that when talking about bifurcations we use the word "local" to describe events that are essentially localized in phase space. Conversely we use the word "global" to describe events that involve distant connections in phase space. With this warning, there should be no chance of confusion with our use, elsewhere, of the word "global" in its common parlance as related to the Earth.

In Tables 2–4 we give the names of the bifurcations in the three categories, with alternative names given in parentheses. We then indicate the change in the type of attractor that is produced by the bifurcation, such as a point to a cycle, etc. Some of

the attributes of each class (safe, explosive or dangerous) are then listed at the foot of each table. Among these attributes, the concept of a *basin* requires some comment here. In the multidimensional phase space of a dissipative dynamical system (described in Sec. 2) each attractor, or stable state, is surrounded by a region of starting points from which a displaced system would return to the attractor. The set of all these points constitutes the *basin of attraction*. If the system were displaced to, and then released from any point outside the basin, it would move to a different attractor (or perhaps to infinity). Basins also undergo changes and bifurcations, but for simplicity of exposition in this brief review we focus on the more common attractor bifurcations. Notice, though, that the "basin

Table 3.   Explosive bifurcations. These are less common global events, which occupy an intermediate position between the safe and dangerous forms. Alternative names are given in brackets.

Explosive Bifurcations	
5. Flow Explosion (omega explosion, SNIPER)	Point to cycle
6. Map Explosion (omega explosion, mode-locking)	Cycle to torus
7. Intermittency Explosion: Flow	Point to chaos
8. Intermittency Explosion: Map (temporal intermittency)	Cycle to chaos
9. Regular-Saddle Explosion (interior crisis)	Chaos to chaos
10. Chaotic-Saddle Explosion (interior crisis)	Chaos to chaos

These bifurcations are characterized by the following features:

CATASTROPHIC: global events, abrupt enlargement of attracting set
EXPLOSIVE: enlargement, but no jump to remote attractor
DETERMINATE: with single outcome even with small noise
NO HYSTERESIS: paths retraced on reversal of control sweep
NO BASIN CHANGE: basin boundary remote from attractors
INTERMITTENCY: lingering in old domain, flashes through the new

Table 4.   Dangerous bifurcations. These include the ubiquitous folds where a path reaches a smooth maximum or minimum value of the control parameter, the subcritical local bifurcations, and some global events. They each trigger a sudden jump to a remote "unknown" attractor. In climate studies these would be called tipping points, as indeed might other nonlinear phenomena. Alternative names are given in brackets.

Dangerous Bifurcations	
(a) Local Saddle-Node Bifurcations	
11. Static Fold (saddle-node of fixed point)	from Point
12. Cyclic Fold (saddle-node of cycle)	from Cycle
(b) Local Subcritical Bifurcations	
13. Subcritical Hopf	from Point
14. Subcritical Neimark-Sacker (secondary Hopf)	from Cycle
15. Subcritical Flip (period-doubling)	from Cycle
(c) Global Bifurcations	
16. Saddle Connection (homoclinic connection)	from Cycle
17. Regular-Saddle Catastrophe (boundary crisis)	from Chaos
18. Chaotic-Saddle Catastrophe (boundary crisis)	from Chaos

These bifurcations are characterized by the following features:

CATASTROPHIC: sudden disappearance of attractor
DANGEROUS: sudden jump to new attractor (of any type)
INDETERMINACY: outcome can depend on global topology
HYSTERESIS: path not reinstated on control reversal
BASIN: tends to zero (b), attractor hits edge of residual basin (a, c)
NO INTERMITTENCY: but critical slowing in global events

boundary collision" discussed by Scheffer [2009] in connection with the population dynamics of fish eating zooplankton eating phytoplankton is simply our *saddle connection* of Table 4.

In Fig. 5 we have *schematically* illustrated three bifurcations that are *codimension-1*, meaning that they can be typically encountered under the variation of a single control parameter, μ, which is here plotted horizontally in the left column. The response, q, is plotted vertically. To many applied mathematicians, the most common (safe) bifurcation is what is called the supercritical pitchfork or stable-symmetric point of bifurcation [Thompson & Hunt, 1973]. This was first described by Euler [1744] in his classic analysis of the buckling of a slender elastic column, and is taught to engineering students as "Euler buckling" in which the load carried by the column is the control parameter. Poincaré [1885] explored a number of applications in astro-physics. In this event, the trivial *primary* equilibrium path on which the column has no lateral deflection ($q = 0$), becomes unstable at a critical point, C, where $\mu = \mu_{\mathrm{crit}}$. Passing vertically though C, and then curving towards increasing μ, is a stable *secondary* equilibrium path of deflected states, the so-called post-buckling path. The existence of (stable) equilibrium states at values of

$\mu > \mu_{\mathrm{crit}}$ is why we call the bifurcation a *supercritical* pitchfork. In contrast, many shell-like elastic structures exhibit a *dangerous* bifurcation with an (unstable) post-buckling path that curves towards decreasing values of the load, μ, and is accordingly called a *subcritical* pitchfork. These two pitchforks are excellent examples of safe and dangerous bifurcations, but they do not appear in our lists because they are not codimension-1 events in generic systems. That the bifurcation of a column is not codimension-1 manifests itself by the fact that a perfectly straight column is not a typical object; any real column will have small imperfections, lack of straightness being the most obvious one. These imperfections round off the corners of the intersection of the primary and secondary paths (in the manner of the contours of a mountain-pass), and destroy the bifurcation in the manner described by catastrophe theory [Poston & Stewart, 1978; Thompson, 1982]. We shall see a subcritical pitchfork bifurcation in a schematic diagram of the THC response due to Rahmstorf [2000] in Fig. 10. This is only observed in very simple (nongeneric) models and is replaced by a fold in more elaborate ones.

It is because of this lack of typicality of the pitchforks that we have chosen to illustrate the safe and dangerous bifurcations in Fig. 5 by other
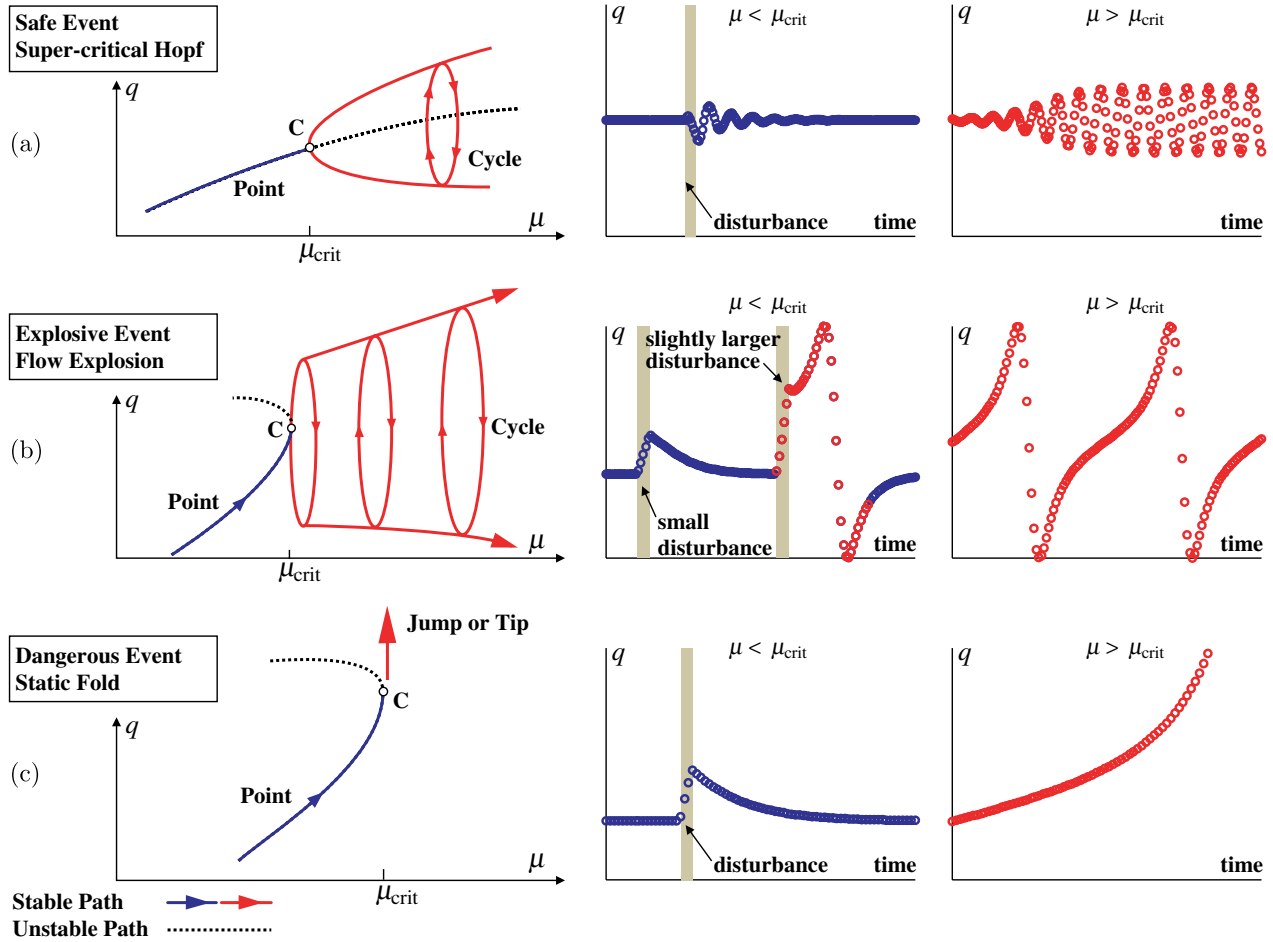
Fig. 5.   Schematic illustration of the three bifurcation types. On the left the control parameter, μ, is plotted horizontally and the response, q, vertically. The middle column shows the time-series of a response to small disturbances if $\mu < \mu_{crit}$. On the right we show how the system drifts away from its previously stable steady state if $\mu > \mu_{crit}$. The different types of events are (from top to bottom) safe (a), explosive (b) and dangerous (c).

(codimension-1) bifurcations. As a safe event, we show in Fig. 5(a) the supercritical Hopf bifurcation. This has an equilibrium path increasing monotonically with μ whose point attractor loses its stability at C in an oscillating fashion, throwing off a path of stable limit cycles which grow towards increasing μ. This occurs, for example, at the onset of vibrations in machining, and triggers the aerodynamic flutter of fins and ailerons in aircraft. Unlike the pitchfork, this picture is not qualitatively changed by small perturbations of the system.

As our explosive event, we show in Fig. 5(b) the *flow explosion* involving a saddle-node (fold) on a limit cycle. Here the primary path of point attractors reaches a vertical tangent, and a large oscillation immediately ensues. As with the supercritical Hopf, all paths are refollowed on reversing the sweep of the control parameter μ: there is no hysteresis.

Finally, as our dangerous event in Fig. 5(c), we have chosen the simple *static fold* (otherwise known as a saddle-node bifurcation), which is actually the most common bifurcation encountered in scientific applications: and we shall be discussing one for the THC in Sec. 6.1. Such a fold is in fact generated when a perturbation rounds off the (untypical) subcritical pitchfork, revealing a sharp *imperfection sensitivity* notorious in the buckling of thin aerospace shell structures [Thompson & Hunt, 1984]. In the fold, an equilibrium path of stable point attractors being followed under increasing μ folds smoothly backwards as an unstable path towards decreasing μ as shown. Approaching the turning point at μ_{crit} there is a gradual loss of attracting strength, with the *local decay rate* (LDR) of transient motions (see Sec. 4) passing directly through zero with progress along the arc-length of the path.

This makes its variation with μ parabolic, but this fine distinction seems to have little significance in the climate tipping studies of Secs. 6 and 7. Luckily, in these studies, the early decrease of LDR is usually identified long before any path curvature is apparent. As μ is increased through μ_{crit} the system finds itself with no equilibrium state nearby, so there is inevitably a fast dynamic jump to a remote attractor of any type. On reversing the control sweep, the system will stay on this remote attractor, laying one end-foundation for a possible hysteresis cycle.

We see immediately from these bifurcations that it is primarily the dangerous forms that will correspond to, and underlie, the climate tipping points that concern us here. (Though if, for example, we adopt Lenton's relatively relaxed definition of a tipping point based on time-horizons (see Sec. 5), even a safe bifurcation might be the underlying trigger.) Understanding the bifurcational aspects will be particularly helpful in a situation where some quasi-stationary dynamics can be viewed as an equilibrium path of a mainly-deterministic system, which may nevertheless be stochastically perturbed by noise. We should note that the dangerous bifurcations are often indeterminate in the sense that the remote attractor to which the system jumps often depends with infinite sensitivity on the precise manner in which the

bifurcation is realized. This arises (quite commonly and typically) when the bifurcation point is located exactly on a fractal basin boundary [McDonald *et al.*, 1985; Thompson, 1992, 1996]. In a model, repeated runs from slightly varied starting conditions would be needed to explore all the possible outcomes.

Table 5 lists the precursors of the bifurcations from Tables 2–4 that one would typically use to determine if a bifurcation is nearby in a (mostly) deterministic system. One imagines the currently observed steady state to be perturbed by a small "kick" or sudden noise. Since the steady state is still stable, the system relaxes back to it. This relaxation decays exponentially proportional to $\exp(\lambda t)$ where t is the time and λ (a negative quantity in this context) is the critical eigenvalue of the destabilizing mode [Thompson & Stewart, 2002]. The local decay rate, LDR (called κ in Sec. 4), is the negative of λ.

Defined in this way, a positive LDR tending to zero quantifies the "slowing of transients" as we head towards an instability. We see that the vast majority (though not all) of the typical events display the useful *precursor* that the local decay rate, LDR, vanishes at the bifurcation (although the decay is in some cases oscillatory). Under light stochastic noise, the variance of the critical mode will correspondingly exhibit a divergence

Table 5. List of all codimension-1 bifurcations of continuous dissipative dynamics, with notes on their precursors. Here S, E and D are used to signify the safe, explosive and dangerous events respectively. LDR is the local decay rate, measuring how rapidly the system returns to its steady state after a small perturbation. Being a linear feature, the LDR of a particular type of bifurcation is not influenced by the sub- or super-critical nature of the bifurcation.

Precursors of Codimension-1 Bifurcations		
Supercritical Hopf	S: point to cycle	LDR $\to 0$ linearly with control
Supercritical Neimark	S: cycle to torus	LDR $\to 0$ linearly with control
Supercritical flip	S: cycle to cycle	LDR $\to 0$ linearly with control
Band merging	S: chaos to chaos	separation decreases linearly
Flow explosion	E: point to cycle	Path folds. LDR $\to 0$ linearly along path
Map explosion	E: cycle to torus	Path folds. LDR $\to 0$ linearly along path
Intermittency expl: flow	E: point to chaos	LDR $\to 0$ linearly with control
Intermittency expl: map	E: cycle to chaos	LDR $\to 0$ as trigger (fold, flip, Neimark)
Regular interior crisis	E: chaos to chaos	lingering near impinging saddle cycle
Chaotic interior crisis	E: chaos to chaos	lingering near impinging chaotic saddle
Static fold	D: from point	Path folds. LDR $\to 0$ linearly along path
Cyclic fold	D: from cycle	Path folds. LDR $\to 0$ linearly along path
Subcritical Hopf	D: from point	LDR $\to 0$ linearly with control
Subcritical Neimark	D: from cycle	LDR $\to 0$ linearly with control
Subcritical flip	D: from cycle	LDR $\to 0$ linearly with control
Saddle connection	D: from cycle	period of cycle tends to infinity
Regular exterior crisis	D: from chaos	lingering near impinging saddle cycle
Chaotic exterior crisis	D: from chaos	lingering near impinging accessible saddle

proportional to the reciprocal of the LDR. The LDR precursor certainly holds, with monotonic decay, for the static fold which is what we shall be looking at in Sec. 6.1 in the collapse of the North Atlantic thermohaline circulation. The fact, noted in Table 5, that close to the bifurcation some LDRs vary linearly with the control, while some vary linearly along the (folding) path is a fine distinction that may not be useful or observable in climate studies.

The outline of the codimension-1 bifurcations that we have just presented applies to dynamical flows which are generated by *continuous* systems where time changes smoothly as in the real world, and as in those computer models that are governed by differential equations. There are closely analogous theories and classifications for the bifurcations in the dynamics of *maps* that govern (for example) *iterated* systems, where time changes in finite steps. It is these analogous theories that will be needed when dealing with experimental data sets from ice cores, etc., as we shall show in the following section. Meanwhile the theory for discrete time data, has direct relevance to the possibility of tipping points in parts of the *biosphere* where time is often best thought of in generations or seasons; in some populations, such as insects, one generation disappears before the next is born.

The equivalent concept that we shall need for analyzing discrete-time data is as follows. The method used in our examples from the recent literature (in Secs. 6 and 7) is to search for an underlying linearized deterministic map of the form

$$y_{n+1} = cy_n$$

which governs the critical slowing mode of the transients. This equation represents exponential decay when the eigenvalue of the mapping, c, is less than one, but exponential growth when c is greater than one. So corresponding to LDR dropping to zero, we shall be expecting c to increase towards unity.

4. Analysis of Time Series Near Incipient Bifurcations

Time series of observational data can help to predict incipient bifurcations in two ways. First, climate models, even if derived from first principles, require initial conditions on a fine mesh and depend on parameters (for example, the effective reradiation coefficient from the Earth's land surface). Both, initial conditions and parameters, are often not measurable directly but must be extracted indirectly by fitting the output of models to training data. This process is called data *assimilation*. The alternative is to skip the modeling step and search for precursors of incipient dangerous bifurcations directly in a monitored time-series. A typical example of an observational time-series is shown (later) in the upper part of Fig. 13. The time-series clearly shows an abrupt transition at about 34 million years before the present (BP). One of the aims of time-series analysis would be to predict this transition (and, ideally, its time) from features of the time-series *prior to the transition*. In this example one assumes that the system is in an equilibrium-like state which then disappears in a static fold, 34 million years BP. According to Table 5 the LDR tends to zero as we approach such a bifurcation.

A decreasing LDR corresponds to a slowing down of small-scale features in the time-series which one can expect to be visible in many different ways. If it is possible to apply small pulse-like disturbances (or one knows that this type of disturbance has been present during the recording) the LDR is observable directly as the recovery rate from this disturbance (this was suggested for ecological systems by van Nes and Scheffer [2007]). However, natural disturbances that are typically present are noise-induced fluctuations around the equilibrium. The noise that we have in mind might come from outside the system, or might be high frequency oscillations within the system which effectively act as noise on the larger and slower fundamental motions. From either source, these noisy fluctuations on short time-scales can be used to extract information about a decrease of the LDR. For example, the power spectrum of the noisy time-series shifts toward lower frequencies. This reddening of the spectrum was analyzed and tested by Kleinen *et al.* [2003] as an indicator of a decrease of the LDR using the box models by Stommel [1961], and by Biggs *et al.* [2009] in a fisheries model. Carpenter and Brock [2006] found that a decreasing LDR causes an increasing variance of the stationary temporal distributions in their study of stochastic ecological models. Also in studies of ecological models, Guttal and Jayaprakash [2008a, 2008b] found that increasing higher-order moments (such as skewness) of the temporal distribution can be a reliable early warning signal for a regime shift, as well as increasing higher-order moments of spatial distributions. Making the step from temporal to spatial distributions is of interest because advancing technology may be able to increase the accuracy of measured

spatial distributions more than measurements of temporal distributions (which require data from the past).

4.1. *Autoregressive modeling and detrended fluctuation analysis*

Held and Kleinen [2004] used the noise-induced fluctuations on the short-time-scale to extract information about the LDR using auto-regressive (AR) modeling. See [Box & Jenkins, 1994] for a textbook on statistical forecasting. In order to apply AR modeling to unevenly spaced, drifting data from geological records, Dakos *et al.* [2008] interpolated and detrended the time-series. We outline the procedure of Dakos *et al.* [2008] in more detail for the example of a single-valued time-series that is assumed to follow a slowly drifting equilibrium of a deterministic, dissipative dynamical system disturbed by noise-induced fluctuations.

(1) **Interpolation:** If the time spacing between measurements is not equidistant (which is typical for geological time-series) then one interpolates (for example, linearly) to obtain a time-series on an equidistant mesh of time steps Δt. The following steps assume that the time step Δt satisfies $1/\kappa \gg \Delta t \gg 1/\kappa_i$ where κ is the LDR of the time-series and κ_i are the decay rates of other, noncritical, modes. For example, Held and Kleinen [2004] found that $\Delta t = 50$ years fits roughly into this interval for their tests on simulations (see Fig. 11). The result of the interpolation is a time-series x_n of values approximating measurements on a mesh t_n with time steps Δt.

(2) **Detrending:** To remove the slow drift of the equilibrium one finds and subtracts the slowly moving average of the time-series x_n. One possible choice is the average $X(t_n)$ of the time-series x_n taken for a Gaussian kernel of a certain bandwidth d. The result of this step is a time-series $y_n = x_n - X(t_n)$ which fluctuates around zero as a *stationary time-series*. Notice that $X(t_n)$ is the smoothed curve in the upper part of Fig. 13.

(3) **Fit LDR in moving window:** One assumes that the remaining time-series, y_n, can be modeled approximately by a stable scalar linear mapping, the so-called AR(1) model, disturbed by noise

$$y_{n+1} = cy_n + \sigma\eta_n$$

where $\sigma\eta_n$ is the instance of a random error at time t_n and c (the mapping eigenvalue, sometimes called the propagator) is the correlation between successive elements of the time-series y_n. In places we follow other authors by calling c the *first-order autoregressive coefficient*, written as ARC(1). We note that under our assumptions c is related to the LDR, κ, via $c = \exp(\kappa\Delta t)$. If one assumes that the propagator, c, drifts slowly and that the random error, $\sigma\eta_n$, is *independent and identically distributed* (i.i.d.) sampled from a normal distribution then one can obtain the optimal approximation of the propagator c by an ordinary least-squares-fit of $y_{n+1} = cy_n$ over a moving time-window $[t_{m-k} \cdots t_{m+k}]$. Here the window length is $2k$, and the estimation of c will be repeated as the center of the window, given by m, moves through the field of data, as illustrated in Fig. 6. The solution c_m of this least-squares fit is an approximation of $c(t_m) = \exp(\kappa(t_m)\Delta t)$ and, thus, also gives an approximation of the LDR, $\kappa(t_m)$, at the middle of the window. The evolution of the propagator c is shown in the bottom of Figs. 11–14. Finally, if one wants to make a prediction about the time t_f at which the static fold occurs, one has to extrapolate a fit of the propagator time-series $c(t_m)$ to find the time t_f such that $c(t_f) = 1$.

The AR(1) model is only suitable to find out whether the equilibrium approaches a bifurcation or not. It is not able to distinguish between possible types of bifurcation as listed in Table 5. Higher
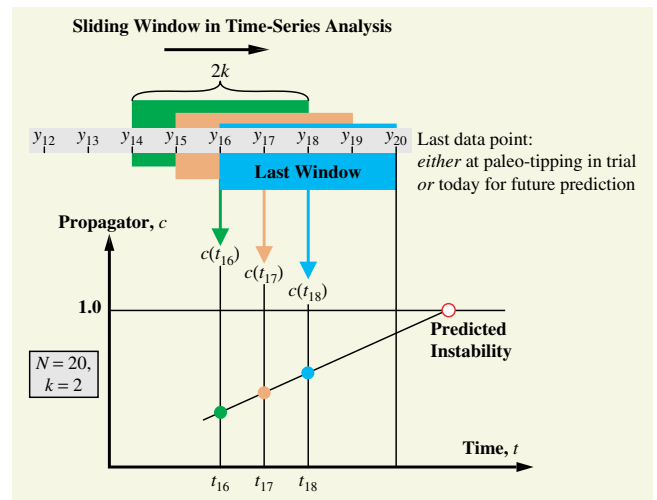

Fig. 6.   Illustration of the sliding window of length $2k$ moving along the time-series and reaching the last data point.

order AR models can be reconstructed. For the data presented by Dakos *et al.* [2008] these higher-order AR models confirm that, first, the first-order coefficient really is dominant, and, second, that this coefficient is increasing before the transition.

Livina and Lenton [2007] modified step 3 of the AR(1) approach of Held and Kleinen [2004], aiming to find estimates also for shorter time-series with a long range memory using *detrended fluctuation analysis* (DFA; originally developed by Peng *et al.* [1994] to detect long-range correlation in DNA sequences). For DFA one determines the variance $V(k)$ of the cumulated sum of the detrended time-series y_n over windows of size k and fits the relation between $V(k)$ and k to a power law: $V(k) \sim k^\alpha$. The exponent α approaches $3/2$ when the LDR of the underlying deterministic system decreases to zero. The method of Livina and Lenton [2007] was tested for simulations of the GENIE-1 model and on real data for the Greenland ice-core paleo-temperature (GISP2) data spanning the time from $50\,000$ years ago to the present. Extracting bifurcational precursors such as the ARC(1) propagator from the GISP2 data is particularly challenging because the data set is comparatively small (1586 points) and unevenly spaced. Nevertheless, the propagator estimate extracted via Livina and Lenton's detrended fluctuation analysis shows not only an increase but its intersection with unity would have predicted the rapid transition at the end of the Younger Dryas accurately. See [Lenton *et al.*, 2009] for further discussion of the GENIE simulations.

Both methods, AR analysis and DFA analysis, can in principle be used for predictions of tipping induced by a static fold that are nearly independent of the methods and the (arbitrary) parameters used. When testing the accuracy of predictions on model-generated or real data one should note the following two points.

First, assign the ARC(1) estimate to the time in the middle of the moving time window for which it has been fitted. Dakos *et al.* [2008] shifted the time argument of their ARC(1) estimate to the end point of the fitting interval because they were not concerned with accurate prediction (see Sec. 4.2).

Second, use only those parts of the time-series $c(t)$ that were derived from data prior to the onset of the transition. We can illustrate this using Fig. 4. The time interval between adjacent data points used by Livina and Lenton [2007] and shown in Fig. 4(a) is not a constant. The length of the sliding window in which the DFA1 propagator is repeatedly

estimated is likewise variable. However, we show in Fig. 4(b) a typical length of the window, drawn as if the right-hand leading edge of the window had just reached the tipping point. For this notional window, the DFA1 result would be plotted in the center of the window at point A. Since in a real prediction scenario we cannot have the right-hand leading edge of the window passing the tipping point, the DFA1 graph must be imagined to terminate at A. Although when working with historical or simulation data it is possible to allow the leading edge to pass the tipping point (as Livina and Lenton had done) the results after A become increasingly erroneous from a *prediction point of view* because the desired results for the *pretipping* DFA1 are increasingly contaminated by the spurious and irrelevant behavior of the temperature graph after the tip.

Finally, we note that the disturbances $\sigma \eta_n$ do not have to be i.i.d. random variables. The underlying local decay rate causes a correlation between subsequent measurements for any disturbance without autocorrelation. In this sense the random noise assumed to be present in the AR(1) model is merely a representative for disturbances that are present in the climate system. In fact, the precise assumption underlying the AR(1) analysis is the presence of three well separated time-scales. One time-scale, on which small fluctuations of the complex climate system occur (these fluctuations are represented by the noise), is *fast*. The second time-scale, on which disturbances decay, is *intermediate* (this characteristic time corresponds to the inverse of the LDR away from the bifurcation point). Finally, the time-scale on which the bifurcation parameter drifts is comparatively *slow*.

4.2. *Comments on predictive power*

Ultimately, methods based on AR modeling have been designed to achieve quantitative predictions, giving an estimate of when tipping occurs with a certain confidence interval (similar to Fig. 11). We note, however, that Dakos *et al.* [2008], which is the most systematic study applying this analysis to geological data, make a much more modest claim: the propagator $c(t)$ (and, hence, the estimated LDR) shows a statistically significant increase prior to each of the eight tipping events they investigated (listed in the introduction). Dakos *et al.* [2008] applied statistical rank tests to the propagator $c(t_n)$ to establish statistical significance. In the procedures of Sec. 4.1 one has to choose a number of

method parameters that are restricted by *a priori* unknown quantities, for example, the step size Δt for interpolation, the kernel bandwidth d, and the window length, $2k$. A substantial part of the analysis in Dakos *et al.* [2008] consisted of checking that the observed increase of c is largely independent of the choice of these parameters, thus, demonstrating that the increase of c is not an artefact of their method.

The predictions one would make from the ARC(1) time-series, $c(t)$, are, however, not as robust on the quantitative level (this will be discussed for two examples of Dakos *et al.* [2008] in Sec. 7). For example, changing the window length $2k$ or the kernel bandwidth d shifts the time-series of the estimated propagator horizontally and vertically: even a shift by 10% corresponds to a shift for the estimated tipping by possibly thousands of years. Also the interpolation step size Δt (interpolation is necessary due to the unevenly spaced records and the inherently non-discrete nature of the time-series) may cause spurious autocorrelation.

Another difficulty arises from an additional assumption one has to make for accurate prediction: the underlying control parameter is drifting (nearly) linearly in time during the recorded time-series. Even this assumption is not sufficient. A dynamical system can nearly reach the tipping point under gradual variation (say, increase) of a control parameter but turn back on its own if the parameter is increased further. The only definite conclusion one can draw from a decrease of the LDR to a small value is that generically there should exist a perturbation that leads to tipping. For a recorded time-series, this perturbation may simply not have happened. The term "generic" means that certain second-order terms in the underlying nonlinear deterministic system should have a substantially larger modulus than the vanishing LDR [Thompson & Stewart, 2002]. This effect may lead to false positives when testing predictions using past data even if the AR models are perfectly accurate and the assumptions behind them are satisfied.

Another problem affecting the quantitative accuracy of predictions is the possibility of noise-induced escape from the basin of attraction before the tipping point is reached. This leads to a systematic bias of a prediction that extrapolates the AR(1) propagator to estimate the time at which it reaches unity. The probability of early escape can be expressed in terms of the relation between noise level and drift speed of the bifurcation parameter.
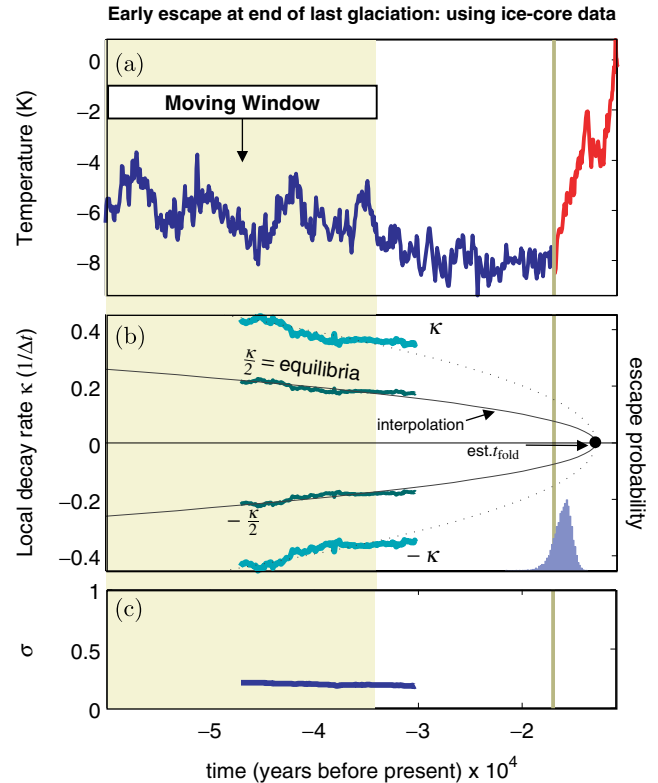

Fig. 7. Estimated probability for early escape from the stable node based on the AR(1) analysis of the ice-core record from the end of the last glaciation [Petit *et al.*, 1999]. (a) Shows the original time-series, (b) the estimated local decay rate per time step, κ, the corresponding equilibrium positions for the saddle-node normal form, and the interpolation estimate for the critical time t fold. The inset in (b) shows the probability distribution for escape. (c) Shows the estimate for the nondimensionalized noise-level.

Figure 7 shows a quantitative estimate of this effect as studied by Thompson and Sieber [2011]. The original time-series in Fig. 7(a) is an ice-core record of the end of the last glaciation from Petit *et al.* [1999], which is part of the study by Dakos *et al.* [2008]. Figure 7(b) shows the local decay rate κ, as extracted by AR(1) analysis. If one assumes that the underlying deterministic system has a control parameter that approaches its critical value for a saddle-node with linear speed one can extract the saddle-node normal form parameters using the estimate for κ. For example, in the normal form the position of the node would be at $\kappa/2$, and the position of the saddle equilibrium would be at $-\kappa/2$, as shown in Fig. 7(b). Interpolation between saddles and nodes gives an estimate for the critical time t_{fold}. An order-of-magnitude estimate of the (nondimensionalized) noise level σ, shown in Fig. 7(c), then allows an estimate of the probability

distribution for escape over time (see [Thompson & Sieber, 2011] for details). This distribution is shown as a small inset in Fig. 7(b), and it clearly shows that early escape plays a role whenever the noise level is large compared to the drift speed of the control parameter.

The effects listed above all conspire to restrict the level of certainty that can be gained from predictions based on time-series. Note, though, that from a geo-engineering point of view [Launder & Thompson, 2010], these difficulties may be of minor relevance because establishing a decrease of the LDR is of the greatest interest in its own right. After all, the LDR is the primary direct indicator of sensitivity of the climate to perturbations (such as geo-engineering measures).

5. Lenton's Tipping Elements

Work at the beginning of this century which set 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 transition to a new state at a rate determined by the climate system itself and faster than the cause, with some degree of irreversibility. As we noted in Sec. 3, this makes the *tipping points* essentially identical to the dangerous bifurcations of nonlinear dynamics.

As well as tipping points, the concept has arisen of *tipping elements*, these being well-defined subsystems of the climate which work (or can be assumed to work) fairly independently, and are prone to sudden change. In modeling them, their interactions with the rest of the climate system are typically expressed as a forcing that varies slowly over time.

Recently, Lenton *et al.* [2008] made a critical evaluation of policy-relevant tipping elements in the climate system that are particularly vulnerable to human activities. To do this, they built on the discussions and conclusions of a recent international workshop entitled "Tipping Points in the Earth System" held at the British Embassy, Berlin, which brought together 36 experts in the field. Additionally, they conducted an expert elicitation from 52 members of the international scientific community to rank the sensitivity of these elements to global warming.

In their work, they used the term tipping element to describe a subsystem of the Earth system that is at least subcontinental in scale, and can be switched into a qualitatively different state by small perturbations. Their definition is in some ways broader than that of some other workers because they wished to embrace the following: nonclimatic variables; cases where the transition is actually slower than the anthropogenic forcing causing it; cases where a slight change in control may have a qualitative impact in the future without any abrupt change however. To produce their short list of key climatic tipping elements, summarized in Table 1 (in the introduction) and below, Lenton *et al.* [2008] considered carefully to what extent they satisfied the following four conditions guaranteeing their relevance to international decision-making meetings such as Copenhagen [2009], the daughter of Kyoto.

Condition 1

There is an adequate theoretical basis (or past evidence of threshold behavior) to show that there are parameters controlling the system that can be combined into a single control μ for which there exists a critical control value μ_{crit}. Exceeding this critical value leads to a qualitative change in a crucial system feature after prescribed times.

Condition 2

Human activities interfere with the system such that decisions taken within an appropriate *political time horizon* can determine whether the critical value for the control, μ_{crit}, is reached.

Condition 3

The time to observe a qualitative change plus the time to trigger it lie within an *ethical time horizon* which recognizes that events too far away in the future may not have the power to influence today's decisions.

Condition 4

A significant number of people care about the expected outcome. This may be because (i) it affects significantly the overall mode of operation of the Earth system, such that the tipping would modify the qualitative state of the whole system, or (ii) it would deeply affect human welfare, such that the tipping would have impacts on many people, or (iii) it would seriously affect a unique feature of the biosphere. In a personal communication, Tim Lenton kindly summarized his latest views as to which of these are likely

to be governed by an underlying bifurcation. They are listed in the headings as follows.

1. *Arctic summer sea-ice*: *Possible bifurcation*

If the area covered by ice decreases, less solar energy (insolation) is reflected, resulting in increasing temperature and, thus, a further decrease in ice coverage. So area coverage has a strong positive feedback, and may exhibit bistability with perhaps multiple states for ice thickness. The instability is not expected to be relevant to Southern Ocean sea-ice because the Antarctic continent covers the region over which it would be expected to arise [Morales Maqueda *et al.*, 1998]. Some researchers think a summer ice-loss threshold, if not already passed, may be very close and a transition could occur well within this century. However Lindsay and Zhang [2005] are not so confident about a threshold, and Eisenman and Wettlaufer [2009] argue that there is probably no bifurcation for the loss of seasonal (summer) sea-ice cover: but there may be one for the year-round loss of ice cover. See also [Winton, 2006]. The decline of the summer sea ice is illustrated in Fig. 8.

2. *Greenland ice-sheet: Bifurcation*

Ice-sheet models generally exhibit multiple stable states with nonlinear transitions between them [Saltzman, 2002], and this is reinforced by paleodata. If a threshold is passed, the IPCC [2007] predicts a time-scale that is greater than 1000 years for a collapse of the sheet. However, given the uncertainties in modeling, a lower limit of 300 yrs is conceivable [Hansen, 2005].

3. *West Antarctic ice-sheet: Possible bifurcation*

Most of the West Antarctic ice-sheet (WAIS) is grounded below sea level and could collapse if a retreat of the grounding-line (between the ice sheet and the ice shelf) triggers a strong positive feedback. The ice sheet has been prone to collapse, and models show internal instability. There are occasional major losses of ice in the so-called Heinrich events. Although the IPCC [2007] has not quoted a threshold, Lenton estimates a range that is accessible this century. Note that a rapid sea-level rise (of greater than one meter per century) is more likely
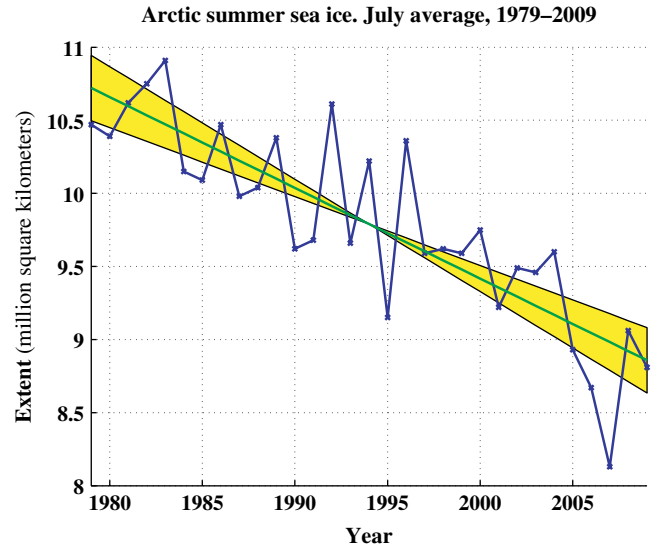

Fig. 8. Decline in the Arctic summer sea ice since 1979. The monthly July average is plotted for each year. The sea ice extent (blue) is derived from satellite images measuring the extent of ocean covered by sea ice at any concentration greater than 15%. As the satellite images do not capture the region around the North pole, this region is assumed to be covered in this data. The trend, a decrease of 3.2% per decade, is shown by a dark green line, together with its 95% confidence interval (yellow region, centered at the middle of the time period). *Source* is [Fetterer *et al.*, 2002].

to come from the WAIS than from the Greenland ice-sheet.

4. *Atlantic thermohaline circulation: Fold bifurcation*

A shut-off in Atlantic thermohaline circulation can occur if sufficient freshwater enters in the North to halt the density-driven North Atlantic Deep Water formation. Such THC changes played an important part in rapid climate changes recorded in Greenland during the last glacial cycle [Rahmstorf, 2002]: see Sec. 7 for predictive studies of the Younger Dryas tipping event. As described in Sec. 6.1, a multitude of mathematical models, backed up by past data, show the THC to exhibit bistability and hysteresis with a fold bifurcation (see Fig. 10 and discussion in Sec. 6.1). Since the THC helps to drive the Gulf Stream, a shut-down would significantly affect the climate of the British Isles.

5. *El Niño Southern Oscillation: Some possibility of bifurcation*

The El Niño Southern Oscillation (ENSO) is the most significant ocean-atmosphere mode of climate

variability, and it is susceptible to three main factors: the zonal mean thermocline depth, the thermocline sharpness in the eastern equatorial Pacific (EEP), and the strength of the annual cycle and hence the meridional temperature gradient across the equator [Guilyardi, 2006]. So increased ocean heat uptake could cause a shift from present day ENSO variability to greater amplitude and/or more frequent El Niños [Timmermann *et al.*, 1999]. Recorded data suggests switching between different (self-sustaining) oscillatory regimes: however, it could be just noise-driven behavior, with an underlying damped oscillation.

6. Indian summer monsoon: Possible bifurcation

The Indian Summer Monsoon (ISM) is driven by a land-to-ocean pressure gradient, which is itself reinforced by the moisture that the monsoon carries from the adjacent Indian Ocean. This moisture-advection feedback is described by Zickfeld *et al.* [2005]. Simple models of the monsoon give bistability and fold bifurcations, with the monsoon switching from "on" and "off" states. Some data also suggest more complexity, with switches between different chaotic oscillations.

7. Sahara/Sahel and West African monsoon: Possible bifurcation

The monsoon shows jumps of rainfall location even from season to season. Such jumps alter the local atmospheric circulation, suggesting multiple stable states. Indeed past greening of the Sahara occurred in the mid-Holocene and may have occurred rapidly in the earlier Bølling–Alleröd warming. Work by de Menocal *et al.* [2000] suggests that the collapse of vegetation in the Sahara about 5000 years ago occurred more rapidly than could be attributed to changes in the Earth's orbital features. A sudden increase in green desert vegetation would of course be a welcome feature for the local population, but might have unforeseen knock-on effects elsewhere.

8. Amazon rainforest: Possible bifurcation

In the Amazon basin, a large fraction of the rainfall evaporates causing further rainfall, and for this reason simulations of Amazon deforestation typically generate about 20–30% reductions in precipitation [Zeng *et al.*, 1996], a lengthening of the dry season, and increases in summer temperatures [Kleidon & Heimann, 2000]. The result is that it would be difficult for the forest to re-establish itself, suggesting that the system may exhibit bistability.

9. Boreal forest: Probably not a bifurcation

The Northern or Boreal forest system exhibits a complex interplay between tree physiology, permafrost, and fire. Climate change could lead to large-scale dieback of these forests, with transitions to open woodlands or grasslands [Lucht *et al.*, 2006; Joos *et al.*, 2001]. Based on limited evidence, the reduction of the tree fraction may have characteristics more like a quasi-static transition than a real bifurcation.

6. Predictions of Tipping Points in Models

6.1. *Shutdown of the Thermohaline Circulation (THC)*

We choose to look, first, at the thermohaline circulation because it has been thoroughly examined over many years in computer simulations, and its bifurcational structure is quite well understood.

The remarkable global extent of the THC is illustrated in Fig. 9. In the Atlantic it is closely related to, and helps to drive, the North Atlantic Current (including the Drift), and the Gulf Stream: so its variation could significantly affect the climate of the British Isles and Europe. It exhibits multistability and can switch abruptly in response to gradual changes in forcing which might arise from global warming. Its underlying dynamics are summarized *schematically* in Fig. 10 adapted from the paper by Rahmstorf *et al.* [2005], which itself drew on the classic paper of Stommel [1961]. This shows the response, represented by the overturning strength of the circulation (q), versus the forcing control, represented by the fresh water flux (from rivers, glaciers, etc.) into the North Atlantic, (μ). The suggestion is that anthropogenic (man-induced) global warming may shift this control parameter, μ, past the fold bifurcation at a critical value of $\mu = \mu_{\text{crit}}$ ($= 0.2$ in this highly schematic diagram). The hope is that by tuning a climate model to available climatological data we could determine μ_{crit} from that model, thereby throwing some light on the possible tipping of the real climate element.

The question of where the tipping appears in models has been addressed in a series of papers
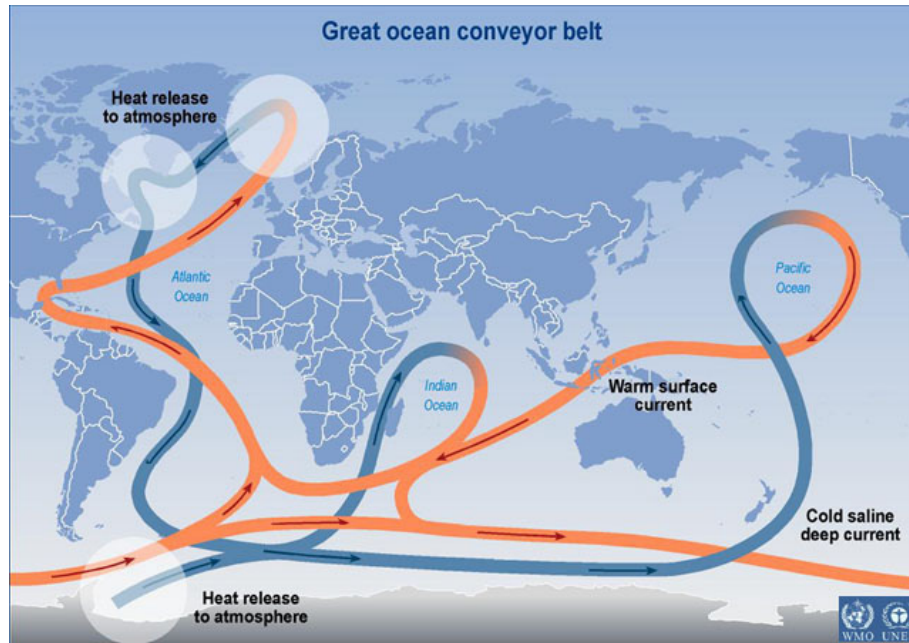
Fig. 9.   The thermohaline circulation (THC), often called the global conveyor, is the major oceanic current of the Earth. It includes warm surface currents, which sink in the polar regions to become cold and saline deep currents as shown. Figure reproduced courtesy of the World Meteorological Office (WMO).

by Dijkstra and Weijer [2003, 2005] Dijkstra *et al.* [2004], and Huisman *et al.* [2009] using a hierarchy of models of increasing complexity. The simplest model is a box model consisting of two connected boxes of different temperatures and salinity representing the North Atlantic at low and high latitudes. For this box model it is known that two stable equilibria coexist for a large range of freshwater-forcing. The upper end of the model hierarchy is a full global ocean circulation model.

Using this high-end model, Dijkstra and Weijer [2005] applied techniques of numerical bifurcation analysis to delineate two branches of stable steady-state solutions. One of these had a strong northern overturning in the Atlantic while the other had hardly any northern overturning, confirming qualitatively the sketch shown in Fig. 10. Finally, Huisman *et al.* [2009] discovered four different flow regimes of their computer model. These they call the Conveyor (C), the Southern Sinking (SS), the Northern Sinking (NS) and the Inverse Conveyor (IC), which appear as two disconnected branches of solutions, where the C is connected with the SS and the NS with the IC. The authors argue that these findings show, significantly, that the parameter volume for which multiple steady states exist is greatly increased.

An intuitive physical mechanism for bistability is the presence of two potential wells (at the

bottom of each is a stable equilibrium) separated by a saddle, which corresponds to the unstable equilibrium. Applying a perturbation then corresponds to a temporary alteration of this potential energy landscape. Dijkstra *et al.* [2004] observed that this
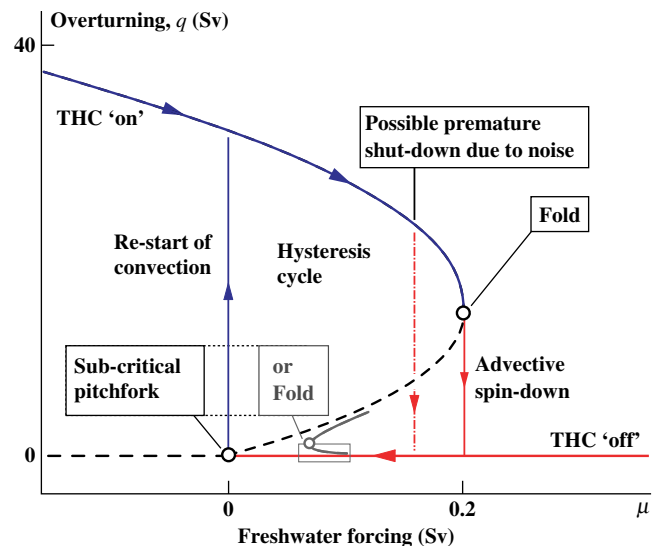

Fig. 10.   A schematic diagram of the thermohaline response showing the two bifurcations and the associated hysteresis cycle [Rahmstorf, 2000]. The subcritical pitchfork bifurcation will be observed in very simple models, but will be replaced by a fold in more elaborate ones: see, for example, Fig. 12(b). Note that 1Sv is 10^6 cubic metres per second, which is roughly the combined flow rate of all rivers on Earth.

picture is approximately true for ocean circulation if one takes the average deviation of water density (as determined by salinity and temperature) from the original equilibrium as the potential energy. They showed, first for a box model and then for a global ocean circulation model, that the potential energy landscape of the unperturbed system defines the basins of attraction fairly accurately. This helps engineers and forecasters to determine whether a perturbation (for example, increased freshwater influx) enables the bistable system to cross from one basin of attraction to the other.

Concerning the simple box models of the THC, we might note their similarity to the atmospheric convection model in which Lorenz [1963] discovered the chaotic attractor: this points to the fact that we must expect chaotic features in the THC and other climate models. See [Dijkstra, 2008] for a summary of the current state of ocean modeling from a dynamical systems point of view, and, for example, [Tziperman *et al.*, 1994; Tziperman, 1997] for how predictions of ocean models connect to full global circulation models. Building on these modeling efforts, ongoing research is actively trying to predict an imminent collapse at the main fold seen in the models (for example, Fig. 10) from bifurcational precursors in time-series. Held and Kleinen [2004] used the local decay rate (LDR; described earlier in Sec. 4 and in Table 5) as the diagnostic variable that they think is most directly linked to the distance from a bifurcation threshold. They demonstrated its use to predict the shutdown of the North Atlantic thermohaline circulation using the oceanic output of CLIMBER2, a predictive coupled model of intermediate complexity [Petoukhov *et al.*, 2000]. They made a 50 000 yrs transient run with a linear increase in atmospheric CO_2 from 280 to 800 parts per million (ppm), which generated within the model an increase in the fresh water forcing which was perturbed stochastically. This run resulted in the eventual collapse of the THC as shown in Fig. 11.

In Fig. 11(a) the graph (corresponding approximately to the schematic diagram of Fig. 10) is fairly linear over much of the time-scale: there is no adequate early prediction of the fold bifurcation in terms of path curvature. The graph of Fig. 11(b) shows the variation of the *first-order autoregressive coefficient* or *propagator*, ARC(1) which is described in Sec. 4. Unlike the response diagram of $q(t)$, the time-series of ARC(1), although noisy, allows a fairly good prediction of the imminent collapse
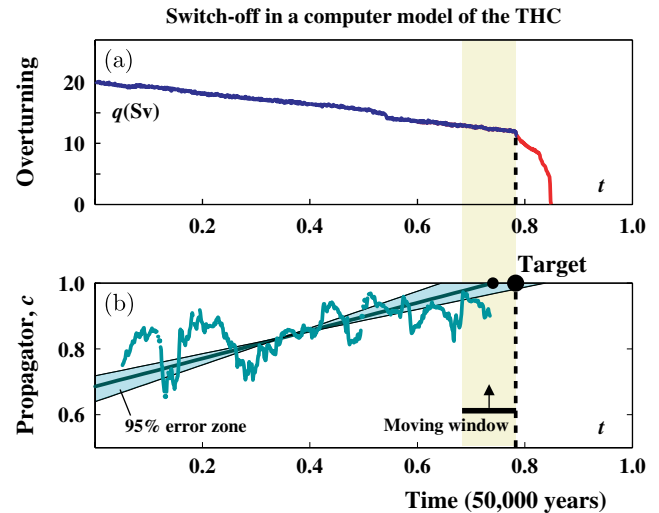

Fig. 11. Results of [Held & Kleinen, 2004] which give a good prediction of the collapse of the thermohaline circulation induced by a four-fold linear increase of CO_2 over 50 000 yrs in a model simulation. Collapse present at $t \approx 0.8$ in (a) is predicted to occur when the propagator, $c =$ ARC(1), shown in (b), or its linear fit, reaches $+1$.

using the linear fit drawn: the fairly steady rise of ARC(1) towards its critical value of $+1$ is indeed seen over a very considerable time-scale. Notice that the linear fit is surrounded by a 95% zone, giving probability bounds to the collapse time. These bounds emphasize that much more precise predictions will be needed before they can be used to guide policy on, for example, whether to implement geo-engineering proposals.

6.2. *Global glaciation and desertification of Africa*

Along side their extensive studies of past climatic events using real paleo-data, Dakos *et al.* [2008] also made some model studies as illustrated in Fig. 12. For these, and subsequent figures, the number of data points, N, is quoted in the captions.

In pictures of this type it is worth observing that there seems to be no agreed way of plotting the estimated auto-correlation coefficient. Held and Kleinen [2004] and Livina and Lenton [2007] plotted ARC(1) at the center of the moving window in which it was determined. Meanwhile Dakos *et al.* [2008] plotted ARC(1) at the final point of this window. Here, we have redrawn the results from the latter article by shifting the ARC(1) back by half the length of the sliding window, bringing the graphs into the format of Held and Kleinen [2004] and Livina and Lenton [2007]. This is important *whenever*

Run-away to Glaciated Earth   **Collapse of Thermohaline Circulation**   **Desertification of North Africa**
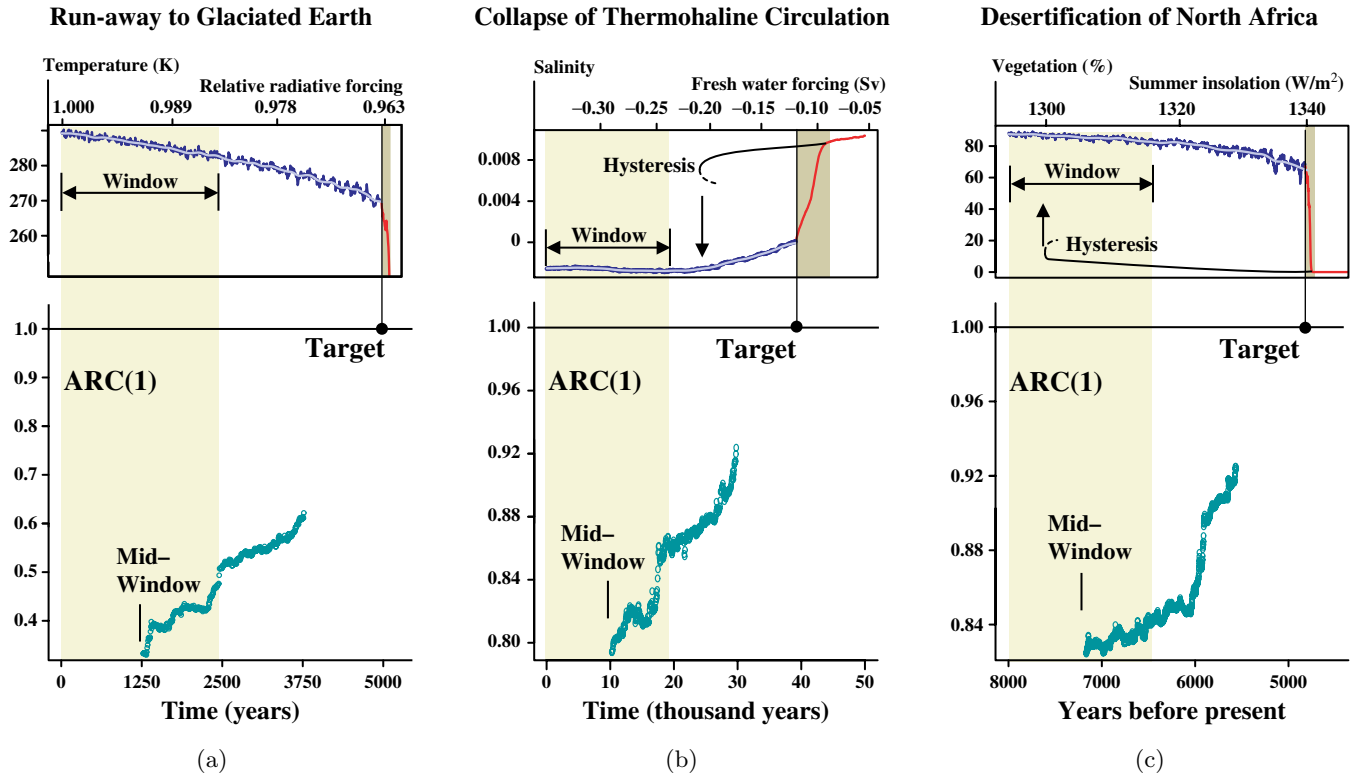

Fig. 12.   Results of Dakos *et al.* [2008] for three examples based on predictive models. (a) Run-away to glaciated Earth ($N = 800$), (b) collapse of thermohaline circulation ($N = 1000$), (c) desertification of North Africa ($N = 6002$). Notice the notional hysteresis loops sketched on (b) and (c). These pictures have been redrawn as mid-window plots.

the intention is to make a forward extrapolation to a target, as we are doing here (see Sec. 4.1). This forward extrapolation can be made by any appropriate method. In fact, approaching (close to) an underlying fold bifurcation, ARC(1) will vary linearly along the solution path, but parabolically with the control parameter: this parabolic effect will only be relevant if the upper solution path is already curving appreciably, which is not the case in most of the present examples displayed here.

7.   Predictions of Ancient Tippings

We have already presented the results of Livina and Lenton [2007] on the ending of the last glaciation (related to the Younger Dryas event) using Greenland ice-core data in Fig. 4 of Sec. 1. Here we turn to Dakos *et al.* [2008] who present a systematic analysis of eight ancient climate transitions. They show that prior to all eight of these transitions the ARC(1) propagator c extracted from the time-series of observations (as described in Sec. 4) shows a statistically significant increase, thus, providing evidence that these ancient transitions indeed correspond to fold-like tipping events. We show in the

following subsections the results of Dakos *et al.* [2008] for two of these events (leaving out the statistical tests).

7.1.   *The greenhouse to ice-house tipping*

We show first in Fig. 13 their study of the greenhouse-icehouse tipping event that happened about 34 million years ago. The time-series in Fig. 13(a) is the data, namely the calcium carbonate ($CaCO_3$) content from tropical Pacific sediment cores. The smooth central line is the Gaussian kernel function used to filter out slow trends. The graph in Fig. 13(b) shows the two plots of ARC(1) that are described in Sec. 6.2, and we notice that the mid-window projection is very close to the target, namely the known tipping point from the paleo-data.

7.2.   *End of the Younger Dryas event*

To put things in perspective, Fig. 14 shows a less-well correlated example from the Dakos paper, this one for the end of the Younger Dryas event using
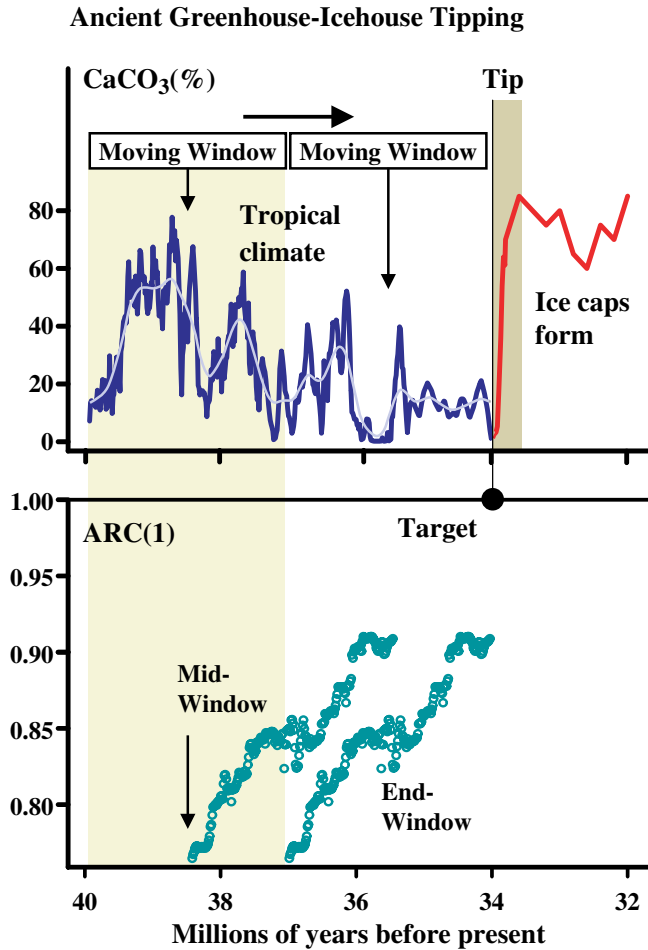
Ancient Greenhouse-Icehouse Tipping

Fig. 13. The ancient greenhouse to icehouse tipping with $N = 482$ data points. This is one of the best correlations obtained by Dakos *et al.* [2008] in their work on eight recorded tipping points. Here the sediments containing $CaCO_3$ were laid down 30–40 million years ago. Redrawn from Dakos *et al.* [2008], as described in the text.

the grayscale from the Cariaco basin sediments in Venezuela. 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, for example, in records of the oxygen isotope $\delta^{18}O$ in Greenland ice. It ended in a dramatic tipping point, about 11 500 years ago, when the Arctic warmed by $7°C$ in 50 yrs. Its behavior is thought to be linked to changes in the thermohaline circulation. As we have seen, 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 restart of the circulation that could have ended the Younger Dryas at its rapid tipping point, propelling the Earth into the warmer Pre-Boreal era.

We might note, here, that the sudden cutting of a new water channel, switching off the THC at the main fold of Fig. 10, may lie outside the scope of any prediction based on a preceding time-series. This would be especially true if the rush of water was, relatively speaking, very sudden and fast so that the control parameter (fresh water forcing) was quickly ramped past the main fold. So predicting the onset of the Younger Dryas might not be possible. Meanwhile, let us assume that the ending of the Younger Dryas was intimately associated with the switch-on of the THC close to the subcritical bifurcation of Fig. 10. Now, we might suppose that the control parameter were moving fairly slowly backwards towards the underlying subcritical bifurcation, making prediction more feasible. However, thinking physically about the processes involved, one could easily imagine that for a good precursor one would need a time-series, not of a temperature, but of an oceanic flow rate.

In Fig. 14(b), we see that the (mid-window) plot of the propagator ARC(1) gives a fairly inadequate prediction of the tipping despite its statistically significant increase. A possible cause for this discrepancy might be the violation of the central assumption underlying the extraction of ARC(1): before tipping the system is supposed to follow a slowly drifting equilibrium disturbed by noise-induced fluctuations. Note that the ARC(1) being close to its critical value $+1$ does not necessarily mean that the underlying deterministic system is close to a bifurcation, and that due to the detrending procedure the fitted ARC(1) will always be slightly less than $+1$.

We might note finally that a very recent paper on the Younger Dryas event by Bakke *et al.* [2009] presents high-resolution records from two sediment cores obtained from Lake Kråkenes in western Norway and the Nordic seas. Multiple proxies from the former show signs of rapid alternations between glacial growth and melting. Simultaneously, sea temperature and salinity show an alternation related to the ice cover and the inflow of warm, salty North Atlantic waters. The suggestion is that there was a rapid flipping between two

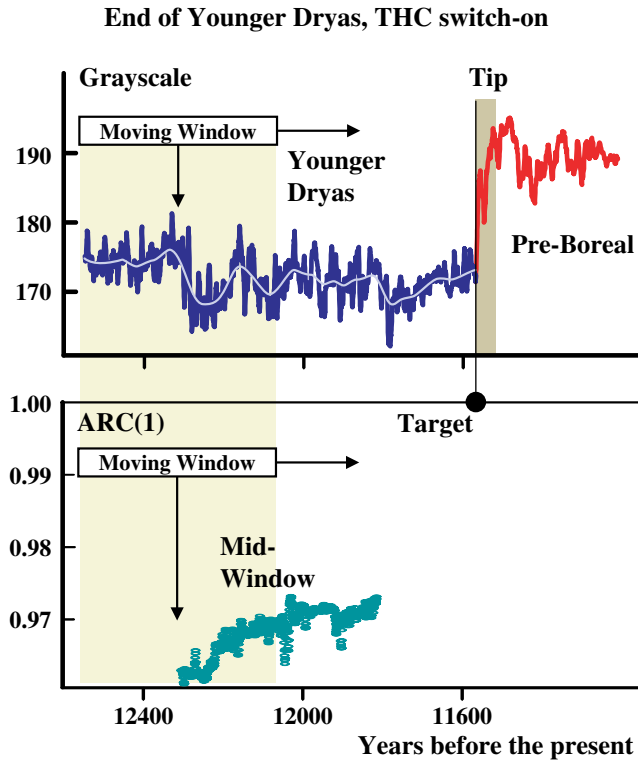### End of Younger Dryas, THC switch-on

Fig. 14.   A second illustration taken from Dakos *et al.* [2008] for the end of the Younger Dryas event using the grayscale from basin sediment in Cariaco, Venezuela. The whole of the above ARC(1) graph fits in the red circle of Fig. 4.

states before the fast tip at the end of Younger Dryas which created the permanent transition to an interglacial state. This strengthens the suspicion that the deterministic component of the dynamics behind the time-series in Fig. 14(a) is not near a slowly drifting equilibrium. It will be interesting to see if any useful time-series analyses can be made of this rapid fluttering action.

8.   Concluding Remarks

Our illustrations give a snapshot of very recent research showing the current status of predictive studies. They show that tipping events, corresponding mathematically to dangerous bifurcations, pose a likely threat to the current state of the climate because they cause rapid and irreversible transitions. Also, there is evidence that tipping events have been the mechanism behind climate transitions of the past. Model studies give hope that these tipping events are predictable using time-series analysis: when applied to real geological data from past events prediction is often remarkably good but is not always reliable. With today's and tomorrow's

vastly improved monitoring, giving times-series that are both longer (higher N) and much more accurate, reliable estimates can be confidently expected. However, if a system has already passed a bifurcation point it may be too late to do anything useful, because an irreversible transition might be already underway.

Techniques from nonlinear dynamical systems enter the modeling side of climate prediction at two points. First, in data assimilation, which plays a role in the tuning and updating of models, the assimilated data is often Lagrangian (for example, it might come from drifting floats in the ocean). It turns out that optimal starting positions for these drifters are determined by stable and unstable manifolds of the vector field of the phase-space flow [Kuznetsov *et al.*, 2003]. Second, numerical bifurcation-tracking techniques for large-scale systems have become applicable to realistic large-scale climate models [Huisman *et al.*, 2009]. More generally, numerical continuation methods have been developed (for example, LOCA by Salinger *et al.* [2002]) that are specifically designed for the continuation of equilibria of large physical systems. These general methods appear to be very promising for the analysis of tipping points in different types of deterministic climate models. These developments will permit efficient parameter studies where one can determine directly how the tipping event in the model varies when many system parameters are changed simultaneously. This may become particularly useful for extensive scenario studies in geo-engineering. For example, Dijkstra *et al.* [2004] demonstrated how bifurcation diagrams can help to determine which perturbations enable threshold-crossing in the bistable THC system, and Biggs *et al.* [2009] studied how quickly perturbations have to be reversed to avoid jumping to coexisting attractors in a fisheries model.

Furthermore, subtle microscopic nonlinearities, currently beyond the reach of climate models, may have a strong influence on the large spatial scale. For example, Golden [2009] observed that the permeability of sea ice to brine drainage changes drastically (from impermeable to permeable) when the brine volume fraction increases across the five percent mark. This microscopic tipping point may have a large-scale follow-on effect on the salinity of sea water near the arctic, and thus, the THC. Incorporating microscopic nonlinearities into the macroscopic picture is a challenge for future modeling efforts.

Concerning the techniques of time-series analysis, two developments in related fields are of interest. First, theoretical physicists are actively developing methods of time-series analysis that take into account unknown nonlinearities, allowing for short term predictions even if the underlying deterministic system is chaotic [Kantz & Schreiber, 2003]. These methods permit, to a certain extent, the separation of the deterministic, chaotic, component of the time-series from the noise (see also [Takens, 1981]). As several of the tipping events listed in Table 1 involve chaos, nonlinear time-series analysis is a promising complement to the classical linear analysis.

Second, much can perhaps be learned from current predictive studies in the related field of theoretical ecology, discussing how higher-order moments of the noise-induced distributions help to detect tipping points. See Sec. 4 for a brief description and [Biggs *et al.*, 2009] for a recent comparison between indicators in a fisheries model.

Acknowledgments

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.

Bakke, J., Lie, O., Heegaard, E., Dokken, T., Haug, G. H., Birks, H., Dulski, P. & Nilsen, T. [2009] "Rapid oceanic and atmospheric changes during the Younger Dryas cold period," *Nature Geosci.* **2**, 2005–2010.

Biggs, R., Carpenter, S. R. & Brock, W. A. [2009] "Turning back from the brink: Detecting an impending regime shift in time to avert it," *Proc. Nat. Acad. Sci. USA* **106**, 826–831.

Box, G. E. P. & Jenkins, G. M. [1994] *Time Series Analysis: Forecasting and Control*, 3rd edition (Prentice Hall PTR, Upper Saddle River, NJ, USA).

Buizza, R., Petroliagis, T., Palmer, T., Barkmeijer, J., Hamrud, M., Hollingsworth, A., Simmons, A. & Wedi, N. [1998] "Impact of model resolution and ensemble size on the performance of an ensemble prediction system," *Quart. J. Roy. Meteorol. Soc.* **124**, 1935–1960.

Carpenter, S. R. & Brock, W. A. [2006] "Rising variance: A leading indicator of ecological transition," *Ecol. Lett.* **9**, 311–318.

Copenhagen [2009] *United Nations Climate Change Conference*, 7–18 Dec. 2009.

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. Nat. Acad. Sci. USA* **105**, 14308–14312.

de Menocal, P., Ortiz, J., Guilderson, T., Adkins, J., Sarnthein, M., Baker, L. & Yarusinsky, M. [2000] "Abrupt onset and termination of the African Humid Period: Rapid climate responses to gradual insolation forcing," *Quat. Sci. Rev.* **19**, 347–361.

Dijkstra, H. A. & Weijer, W. [2003] "Stability of the global ocean circulation: The connection of equilibria within a hierarchy of models," *J. Mar. Res.* **61**, 725–743.

Dijkstra, H. A., te Raa, L. & Weijer, W. [2004] "A systematic approach to determine thresholds of the ocean's thermohaline circulation," *Tellus. Series A, Dyn. Meteorol. Oceanogr.* **56**, 362–370.

Dijkstra, H. A. & Weijer, W. [2005] "Stability of the global ocean circulation: Basic bifurcation diagrams," *J. Phys. Oceanogr.* **35**, 933–948.

Dijkstra, H. A. [2008] *Dynamical Oceanography* (Springer, NY).

Eisenman, I. & Wettlaufer, J. S. [2009] "Nonlinear threshold behavior during the loss of Arctic sea ice," *Proc. Nat. Acad. Sci. USA* **106**, 28–32.

Euler, L. [1744] *Methodus Inveniendi Lineas Curvas Maximi Minimive Proprietate Gaudentes* (*Appendix, De Curvis Elasticis*) (Marcum Michaelem Bousquet, Lausanne and Geneva).

Fetterer, F., Knowles, K., Meier, W. & Savoie, M. [2002] "Sea ice index," (updated 2010) National Snow and Ice Data Center (Boulder, Colorado) http://nsidc.org/data/seaice_index.

Golden, K. [2009] "Climate change and the mathematics of transport in sea ice," *Not. AMS* **56**, 562–584.

Guilyardi, E. [2006] "El Niño–mean state seasonal cycle interactions in a multi-model ensemble," *Clim. Dyn.* **26**, 329–348.

Guttal, V. & Jayaprakash, C. [2008a] "Changing skewness: An early warning signal of regime shifts in ecosystems," *Ecol. Lett.* **11**, 450–460.

Guttal, V. & Jayaprakash, C. [2008b] "Spatial variance and spatial skewness: Leading indicators of regime shifts in spatial ecological systems," *Theoret. Ecol.* **2**, 3–12.

Hansen, J. E. [2005] "A slippery slope: How much global warming constitutes 'dangerous anthropogenic interference'?" *Clim. Change* **68**, 269–279.

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* (Cambrige University Press, Cambridge, UK).

Huisman, S. E., Dijkstra, H. A., von Der Heydt, A. & de Ruijter, W. P. M. [2009] "Robustness of multiple equilibria in the global ocean circulation," *Geophys. Res. Lett.* **36**, L01610.

IPCC [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*, eds. Solomon, S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, K. B., Tignor, M. & Miller, H. L. (Cambridge University Press, Cambridge, UK).

Joos, F., Prentice, I. C., Sitch, S., Meyer, R., Hooss, G., Plattner, G.-K., Gerber, S. & Hasselmann, K. [2001] "Global warming feedbacks on terrestrial carbon uptake under the intergovernmental panel on climate change (IPCC) emission scenarios," *Glob. Biogeochem. Cycles* **15**, 891–907.

Kantz, H. & Schreiber, T. [2003] *Nonlinear Time Series Analysis*, 2nd edition (Cambridge University Press, Cambridge, UK).

Kleidon, A. & Heimann, M. [2000] "Assessing the role of deep rooted vegetation in the climate system with model simulations: Mechanism, comparison to observations and implications for Amazonian deforestation," *Clim. Dyn.* **16**, 183–199.

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.

Kuznetsov, L., Ide, K. & Jones, C. K. R. T. [2003] "A method for assimilation of Lagrangian data," *Monthly Weather Rev.* **131**, 2247–2260.

Launder, B. & Thompson, J. M. T. (eds.) [2010] *Geo-Engineering Climate Change: Environmental Necessity or Pandora's Box?* (Cambridge University Press, Cambridge, UK).

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. Nat. Acad. Sci. USA* **105**, 1786–1793.

Lenton, T. M., Myerscough, R. J., Marsh, R., Livina, V. N., Price, A. R., Cox, S. J. & the GENIE Team [2009] "Using GENIE to study a tipping point in the climate system," *Philos. Trans. Series A, Math. Phys. Engin. Sci.* **367**, 871–884.

Lindsay, R. W. & Zhang, J. [2005] "The thinning of Arctic sea ice, 1988–2003: Have we passed a tipping point?" *J. Climate* **18**, 4879–4894.

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. Climat.* **21**, 1153–1179.

Lorenz, E. N. [1963] "Deterministic nonperiodic flow," *J. Atmosph. Sci.* **20**, 130–141.

Lucht, W., Schaphoff, S., Erbrecht, T., Heyder, U. & Cramer, W. [2006] "Terrestrial vegetation redistribution and carbon balance under climate change," *Carbon Balance and Management* **1**, 6.

McDonald, S. W., Grebogi, C., Ott, E. & Yorke, J. A. [1985] "Fractal basin boundaries," *Physica D: Nonlin. Phen.* **17**, 125–153.

Morales Maqueda, M. A., Willmott, A. J., Bamber, J. L. & Darby, M. S. [1998] "An investigation of the small ice cap instability in the Southern Hemisphere with a coupled atmosphere-sea ice-ocean-terrestrial ice model," *Clim. Dyn.* **14**, 329–352.

National Research Council [2002] *Abrupt Climate Change: Inevitable Surprises* (Natl Acad Press, Washington, DC).

Peng, C.-K., Buldyrev, S. V., Havlin, S., Simons, M., Stanley, H. E. & Goldberger, A. L. [1994] "Mosaic organization of DNA nucleotides," *Phys. Rev. E* **49**, 1685–1689.

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.

Petoukhov, V., Ganopolski, A., Brovkin, V., Claussen, M., Eliseev, A., Kubatzki, C. & Rahmstorf, S. [2000] "CLIMBER-2: A climate system model of intermediate complexity. Part I: Model description and performance for present climate," *Clim. Dyn.* **16**, 1–17.

Poincaré, H. [1885] "Sur l'équilibre d'une masse fluide animée d'un mouvement de rotation," *Acta Mathematica* **7**, 259–380.

Poston, T. & Stewart, I. [1978] *Catastrophe Theory and its Applications* (Pitman, London, UK).

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*, eds. Steele, J., Thorpe, S. & Turekian, K. (Academic Press, London), pp. 1–6.

Rahmstorf, S. [2002] "Ocean circulation and climate during the past 120 000 years," *Nature* **419**, 207–214.

Rahmstorf, S., Crucifix, M., Ganopolski, A., Goosse, H., Kamenkovich, I., Knutti, R., Lohmann, G., Marsh, R., Mysak, L. A., Wang, Z. & Weaver, A. J. [2005] "Thermohaline circulation hysteresis: A model intercomparison," *Geophys. Res. Lett.* **32**, L23605.

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.

Salinger, A. G., Bou-Rabee, N. M., Pawlowski, R. P., Wilkes, E. D., Burroughs, E. A., Lehoucq, R. B. & Romero, L. A. [2002] "LOCA 1.0 Library of continuation algorithms: Theory and implementation manual," Sandia National Laboratories, SAND2002-0396.

Saltzman, B. [2002] *Dynamical Paleoclimatology* (Academic Press, London, UK).

Scheffer, M. [2009] *Critical Transitions in Nature and Society* (Princeton University Press, Princeton, USA).

Selten, F. M., Branstator, G. W., Dijkstra, H. A. & Kliphuis, M. [2004] "Tropical origins for recent and future Northern Hemisphere climate change," *Geophys. Res. Lett.* **31**, L21205.

Sperber, K. R., Brankovic, C., Déqué, M., Frederiksen, C. S., Graham, R., Kitoh, A., Kobayashi, C., Palmer, T., Puri, K., Tennant, W. & Volodin, E. [2001] "Dynamical seasonal predictability of the Asian summer monsoon," *Monthly Weather Rev.* **129**, 2226–2248.

Stommel, H. [1961] "Thermohaline convection with two stable regimes of flow," *Tellus* **13**, 224–230.

Takens, F. [1981] "Detecting strange attractors in turbulence," *Dynamical Systems and Turbulence*, eds. Rand, D. A. & Young, L. S. (Springer, Berlin & NY), Chapter 21, pp. 366–381.

Thompson, J. M. T. & Hunt, G. W. [1973] *A General Theory of Elastic Stability* (Wiley, London).

Thompson, J. M. T. [1982] *Instabilities and Catastrophes in Science and Engineering* (Wiley, Chichester, UK).

Thompson, J. M. T. & Hunt, G. W. [1984] *Elastic Instability Phenomena* (Wiley, Chichester, UK).

Thompson, J. M. T. [1992] "Global unpredictability in nonlinear dynamics: Capture, dispersal and the indeterminate bifurcations," *Physica D: Nonlin. Phen.* **58**, 260–272.

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.

Thompson, J. M. T. [1996] "Danger of unpredictable failure due to indeterminate bifurcation," *Zeitschrift für angewandte Mathematik und Mechanik* **76**, 205–208.

Thompson, J. M. T. & Stewart, H. B. [2002] *Nonlinear Dynamics and Chaos*, 2nd edition (Wiley, Chichester, UK).

Thompson, J. M. T. & Sieber, J. [2010] "Predicting climate tipping points," *Geo-Engineering Climate Change*: *Environmental Necessity or Pandora's Box?*, eds. Launder, B. & Thompson, J. M. T. (Cambridge University Press, Cambridge, UK), Chapter 3, pp. 50–83.

Thompson, J. M. T. & Sieber, J. [2011] "Climate tipping as a noisy bifurcation: A predictive technique," *IMA J. Appl. Math.* **76**, 27–46.

Timmermann, A., Oberhuber, J., Bacher, A., Esch, M., Latif, M. & Roeckner, E. [1999] "Increased El Niño frequency in a climate model forced by future greenhouse warming," *Nature* **398**, 694–697.

Tziperman, E., Toggweiler, J. R., Bryan, K. & Feliks, Y. [1994] "Instability of the thermohaline circulation with respect to mixed boundary conditions: Is it really a problem for realistic models?" *J. Phys. Oceanogr.* **24**, 217–232.

Tziperman, E. [1997] "Inherently unstable climate behavior due to weak thermohaline ocean circulation," *Nature* **386**, 592–595.

van Nes, E. H. & Scheffer, M. [2007] "Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift," *The American Naturalist* **169**, 738–747.

Winton, M. [2006] "Does the Arctic sea ice have a tipping point?" *Geophys. Res. Lett.* **33**, L23504.

Zeng, N., Dickinson, R. E. & Zeng, X. [1996] "Climatic impact of Amazon deforestation — a mechanistic model study," *J. Clim.* **9**, 859–883.

Zickfeld, K., Knopf, B., Petoukhov, V. & Schellnhuber, H. J. [2005] "Is the Indian summer monsoon stable against global change?" *Geophys. Res. Lett.* **32**, L15707.