

3

Predicting climate tipping points

J. MICHAEL T. THOMPSON AND JAN SIEBER

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 years. 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, relevant to geo-engineering, will clearly need to embrace both real data from improved monitoring instruments, and simulation data generated from increasingly sophisticated predictive models.

3.1 Introduction

The geo-engineering proposals assessed in this book aim to combat global warming by proactively manipulating the climate. All authors are agreed that these are indeed risky procedures. They would only be actively pursued if all else had failed, and there was a well-researched consensus that to do nothing would lead rapidly to an environmental catastrophe of major proportions. We should note as well that if the climate is thought to have already passed a bifurcation point, one has to consider carefully whether it is in fact too late to usefully apply geo-engineering techniques, because an irreversible transition may already be under way. Biggs *et al.* (2009) study an analogous problem in their fisheries model: after tipping how fast does one have to implement measures to jump back to the previous state?

The emergence of a consensus would inevitably rely on scientific projections of future climatic events, central to which would be sudden, and usually irreversible features that are now called *tipping points*. The Intergovernmental Panel on Climate Change (IPCC 2007) made some brief remarks about abrupt and rapid climate change, but more recently Lenton *et al.* (2008) have sought to define these points rigorously.

The physical mechanisms underlying these tipping points are typically internal positive feedback effects of the climate system. Since any geo-engineering measure will have to rely strongly on natural positive feedback mechanisms to amplify its effect, the proximity to a tipping point is of real significance to the engineers planning the intervention.

Table 3.1 shows a list of candidates proposed by Lenton *et al.* (2008), and the possible effects of their tipping on the global climate. All of these subsystems of the climate have strong internal positive feedback mechanisms. Thus, they have a certain propensity for tipping and are susceptible to input (human or otherwise).

As column 2 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). In Section 3.3 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. This implies that, near tipping points, geo-engineering can be expected to be most effective (and most dangerous) because the climate system is most susceptible to disturbances.

Table 3.1 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 risk of there being an underlying bifurcation is indicated as follows: *H* = high, *M* = medium, *L* = low. This list will be discussed in greater detail in Section 3.5.

Tipping element	Risk	Feature, F (change)	Control parameter, μ	μ_{crit}	Global warming	Transition time, <i>T</i>	Key impacts
Arctic summer sea-ice	M	Areal extent (-)	Local ΔT_{air} , ocean heat transport	??	+0.5 to +2 °C	~ 10 yr (rapid)	Amplified warming, ecosystem change
Greenland ice sheet (GIS)	H	Ice volume (-)	Local ΔT_{air}	~+3 °C	+1 to +2 °C	>300 yr (slow)	Sea level +2 to +7 m
West antarctic ice sheet (WAIS)	M	Ice volume (-)	Local ΔT_{air} or, less ΔT_{ocean}	+5 to +8 °C	+3 to +5 °C	>300 yr (slow)	Sea level +5 m
Atlantic thermohaline circulation	H	Overturning (-)	Freshwater input to North Atlantic	+0.1 to +0.5 Sv	+3 to +5 °C	~100 yr (gradual)	Regional cooling, sea level, ITCZ ^a shift
El Niño Southern Oscillation	M	Amplitude (+)	Thermocline depth, sharpness in EEP ^b	??	+3 to +6 °C	~100 yr (gradual)	Drought in SE Asia and elsewhere
Indian summer monsoon (ISM)	M	Rainfall (-)	Planetary albedo over India	0.5	NA	~1 yr (rapid)	Drought, decreased carrying capacity
SaharaSahel and W - African monsoon	M	Vegetation fraction (+)	Precipitation	100 mm yr ⁻¹	+3 to +5 °C	~10 yr (rapid)	Increased carrying capacity
Amazon rainforest	M	Tree fraction (-)	Precipitation, dry season length	1100 mm yr ⁻¹	+3 to +4 °C	~50 yr (gradual)	Biodiversity loss, decreased rainfall
Boreal forest	L	Tree fraction (-)	Local ΔT_{air}	~+7 °C	+3 to +5 °C	~50 yr (gradual)	Change in type of the ecosystem

^a ITCZ, Intertropical Convergence Zone.

^b EEP, Eastern Equatorial Pacific.

The analysis and prediction of tipping points of climate subsystems is currently being pursued in several streams of research, and we should note in particular the excellent book by Marten Scheffer about tipping points in nature and society, which includes ecology and some climate studies, due to appear this year (Scheffer 2009).

Most of the research is devoted to creating climate models from first principles, tuning and initialising 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 only rely on a single 'best model' starting from the 'real initial conditions'. Typically, all qualified models are run from ensembles of initial conditions and then 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 realise that mathematically some of the climate-tipping events correspond to *bifurcations* (see Section 3.3 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 modelling 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 modelling from first principles because they do not suffer from all the many difficulties of building and initialising reliable computer models. Our review discusses the current state of bifurcation predictions in climate time-series, focussing on methods, introduced by Held and Kleinen (2004) and Livina and Lenton (2007), for the analysis of the collapse of the global conveyor belt of oceanic water, the thermohaline circulation (THC). This conveyor is important, not so much 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 Section 3.4.1) was calculated in sliding windows of length 500 data points. The results are shown in Figure 3.1, and the rapid warming at the end of the Younger Dryas event, around 11 500 years before the present, is spectacularly anticipated by an upward trend in the propagator, which is heading towards its critical value of +1 at about the correct time.

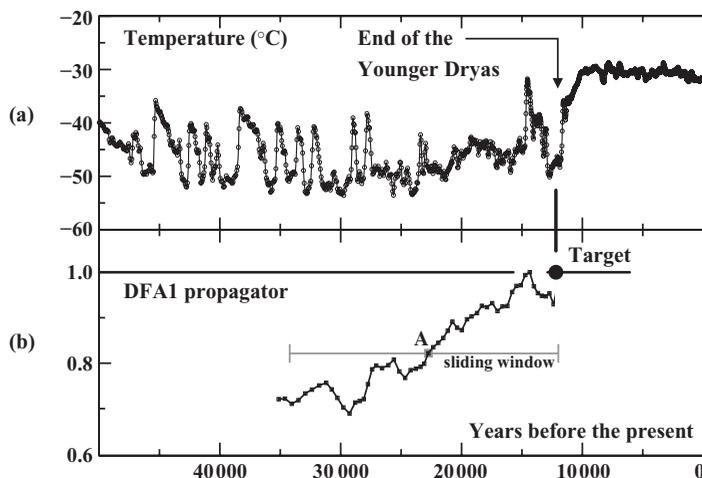


Figure 3.1 Results of Livina and Lenton (2007). (a) Greenland ice core (GISP2) paleo-temperature with an unevenly spaced record, visible in the varying density of symbols on the curve. The total number of data points is $N = 1586$. In (b) the DFA1 propagator is calculated in sliding windows of length 500 points and mapped into the middle points of the windows. A typical sliding window ending near the tipping is shown. Thus, from a prediction point of view the propagator estimates would end at point A (see remarks at the end of Section 3.4.1).

In a second notable paper, Dakos *et al.* (2008) systematically estimated the LDR for real data in 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 on data from the Antarctica Vostok ice core,
- (c) the Bølling–Allerød 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 years, drawing on data from the sediment of the Cariaco Basin in Venezuela,
- (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 slows down before the transition. This slow-down was revealed by a short-term autocorrelation coefficient, $\text{ARC}(1)$, of the time series which examines to what extent the current points are correlated to preceding points and gives an estimate of the LDR. It is expected to increase towards unity at an instability, as described in Section 3.4.

3.2 Climate models as dynamical systems

Thinking about modelling 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 modelling 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) use 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 modelled 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-neighbour) 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. See Kantz and Schreiber (2003) for the relevance of phase space to time-series analysis.

Despite the unique outcome, the results of chaos theory remind us that the response may be essentially unknowable over timescales 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 idealised 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 non-linear.

3.3 Concepts from bifurcation theory

A major component of non-linear 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 behaviour can occur.

In the field of dissipative dynamics *co-dimension-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’ timescale. Slowly varying parameters are external influences that vary on geological timescales, 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.

Table 3.2 *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 parentheses.*

(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 characterised 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

Table 3.3 *Explosive bifurcations. These are less common global events, which occupy an intermediate position between the safe and dangerous forms. Alternative names are given in parentheses.*

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 characterised 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

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 timescale. A complete list of the (typical) co-dimension-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 3.2–3.5. The technical details and terminology of these tables need not concern the general

Table 3.4 *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 parentheses.*

(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 characterised 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

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 as we see in Section 3.6.

A broad classification of the co-dimension-1 *attractor bifurcations* of dissipative systems into *safe*, *explosive* and *dangerous* forms (Thompson *et al.* 1994) is illustrated in Tables 3.2–3.4 and Figure 3.2, while all are summarised in Table 3.5 together with notes on their precursors. It must be emphasised 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 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 localised 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.

Table 3.5 *List of all co-dimension-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 supercritical nature of the bifurcation.*

Precursors of Co-dimension-1 Bifurcations		
Supercritical Hopf	S: point to cycle	LDR \rightarrow 0 linearly with control
Supercritical Neimark	S: cycle to torus	LDR \rightarrow 0 linearly with control
Supercritical flip	S: cycle to cycle	LDR \rightarrow 0 linearly with control
Band merging	S: chaos to chaos	Separation decreases linearly
Flow explosion	E: point to cycle	Path folds. LDR \rightarrow 0 linearly along path
Map explosion	E: cycle to torus	Path folds. LDR \rightarrow 0 linearly along path
Intermittency explosion: flow	E: point to chaos	LDR \rightarrow 0 linearly with control
Intermittency explosion: map	E: cycle to chaos	LDR \rightarrow 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 \rightarrow 0 linearly along path
Cyclic fold	D: from cycle	Path folds. LDR \rightarrow 0 linearly along path
Subcritical Hopf	D: from point	LDR \rightarrow 0 linearly with control
Subcritical Neimark	D: from cycle	LDR \rightarrow 0 linearly with control
Subcritical flip	D: from cycle	LDR \rightarrow 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

In Tables 3.2–3.4 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 Section 3.2) 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.

In Figure 3.2 we have *schematically* illustrated three bifurcations that are *co-dimension-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.

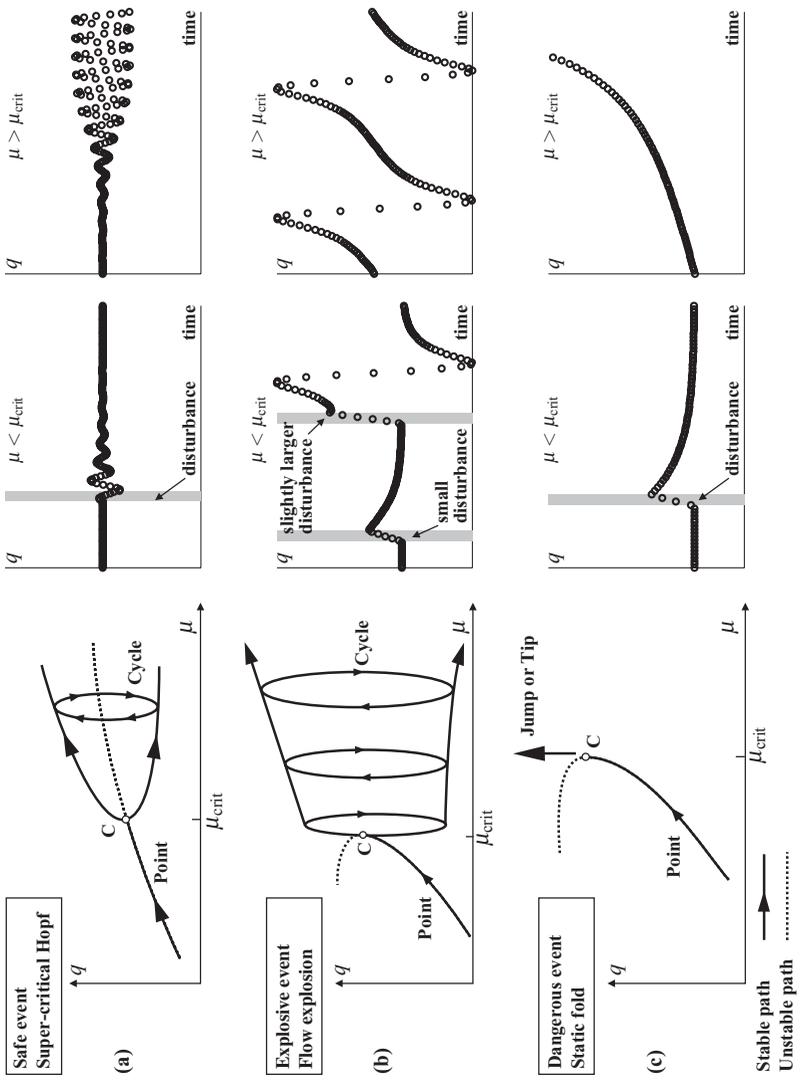


Figure 3.2 Schematic illustration of the three bifurcation types. On the left the control parameter, μ , is plotted horizontally and the response, q , vertically. The middle column shows the time series of a response to small disturbances if $\mu < \mu_{crit}$. On the right we show how the system drifts away from its previously stable steady state if $\mu > \mu_{crit}$. The different types of events are (from top to bottom) (a) safe, (b) explosive and (c) dangerous.

The response, q , is plotted vertically. To many people, 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 astrophysics. 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 $\mu = \mu_{\text{crit}}$. Passing vertically through 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 $\mu > \mu_{\text{crit}}$ 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 co-dimension-1 events in generic systems. That the bifurcation of a column is not co-dimension-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 Figure 3.3 below. This is only observed in very simple (non-generic) 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 Figure 3.2 by other (co-dimension-1) bifurcations. As a safe event, we show in Figure 3.2(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 Figure 3.2(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 re-followed on reversing the sweep of the control parameter μ : there is no hysteresis.

Finally, as our dangerous event in Figure 3.2(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 Section 3.6.1. 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 LDR of transient motions (see Section 3.4) 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 Sections 3.6–3.7. 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 the 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 Section 3.5), 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 realised. 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 3.5 lists the precursors of the bifurcations from Tables 3.2–3.4 that one would typically use to determine if a bifurcation is nearby in a (mostly) deterministic system. One perturbs the observed steady state by a small ‘kick’. As the steady state is still stable the system relaxes back to the steady state. 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 Section 3.4), 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 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 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 Section 3.6.1 in the collapse of the North Atlantic thermohaline circulation. The fact, noted in Table 3.5, 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 co-dimension-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 of dynamical *maps* governed (for example) by *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 analysing discrete-time data is as follows. The method used in our examples from the recent literature (in Sections 3.6 and 3.7) is to search for an *underlying* linearised 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 1, but exponential growth when c is greater than 1. So corresponding to LDR dropping to zero, we shall be expecting c to increase towards unity.

3.4 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 re-radiation 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 modelling 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 Figure 3.6. The time-series clearly shows an abrupt transition at about 34 million years before the present. 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 ago. According to Table 3.5 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. These fluctuations on short timescales 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 analysed 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) find 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 (2008, 2009) find 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).

3.4.1 Autoregressive modelling and detrended fluctuation analysis

Held and Kleinen (2004) use the noise-induced fluctuations on the short timescale to extract information about the LDR using autoregressive (AR) modelling (see Box *et al.* (1994) for a text book on statistical forecasting). In order to apply AR modelling 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, non-critical, modes. For example, Held and Kleinen (2004) found that $\Delta t = 50$ years fits roughly into this interval for their tests on simulations (see Figure 3.4 below). 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 Figure 3.6 below.
- (3) *Fit LDR in moving window* One assumes that the remaining time-series, y_n , can be modelled 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 $\text{ARC}(1)$. 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 $[t_{m-k} \dots t_{m+k}]$. Here the window length is $2k$, and the estimation of c will be repeated as the centre of the window, given by m , moves through the field of data. The solution c_m of this least-squares fit is an approximation of $c(t_m) = \exp(\kappa(t_m)\Delta t)$ and, thus, gives also 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 Figure 3.6. 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 is close to a bifurcation or not. It is not able to distinguish between possible types of bifurcation as listed in Table 3.5. Higher-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 modelling and DFA, can in principle be used for (nearly) model-free prediction of tipping induced by a static fold. 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) have shifted the time argument of their ARC(1) estimate to the endpoint of the fitting interval because they were not concerned with accurate prediction (see Section 3.4.2).

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 Figure 3.1. The time interval between adjacent data points used by Livina and Lenton (2007) and shown in Figure 3.1(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 Figure 3.1(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 centre 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 have done) the results after A become increasingly erroneous from a *prediction point of view*

because the desired results for the *pre-tipping* DFA1 are increasingly contaminated by the spurious and irrelevant behaviour of the temperature graph after the tip.

3.4.2 Comments on predictive power

Ultimately, methods based on AR modelling have been designed to achieve quantitative predictions, giving an estimate of when tipping occurs with a certain confidence interval (similar to Figure 3.4). 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 $c(t)$ shows a statistically significant increase prior to each of the eight tipping events they investigated (listed in the introduction); correspondingly the estimated LDR will show a statistically significant *decrease*. Dakos *et al.* (2008) applied statistical rank tests to the propagator $c(t_n)$ to establish statistical significance. In the procedures of Section 3.4.1 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 Section 3.7). 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 per cent 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 non-linear 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.

These difficulties all conspire to restrict the level of certainty that can be gained from predictions based on time series. Fortunately, from a geo-engineering point of view, 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).

3.5 Lenton's tipping elements

Work at the beginning of this century which set out to define and examine *climate tipping* (Lockwood 2001; Rahmstorf 2001; National Research Council 2002; Alley *et al.* 2003; Rial *et al.* 2004) focussed 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 Section 3.3, this makes the *tipping points* essentially identical to the dangerous bifurcations of non-linear 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 modelling 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) have 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 use the term tipping element to describe a subsystem of the Earth system that is at least sub-continental 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 wish to embrace the following: non-climatic 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 however any abrupt change. To produce their short list of key climatic tipping elements, summarised in Table 3.1 (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 behaviour) 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 are interfering 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 summarised 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* Area coverage has strong positive feedback, and may exhibit bi-stability with perhaps multiple states for ice thickness (if the area covered by ice decreases less energy from insolation is reflected, resulting in increasing temperature and, thus, decreased ice coverage). 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 (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).
- (2) *Greenland ice sheet: bifurcation* Ice-sheet models generally exhibit multiple stable states with non-linear transitions between them (Saltzman 2002), and this is reinforced by paleo-data. If a threshold is passed, the IPCC (2007) predicts a timescale of greater than 1000 years for a collapse of the sheet. However, given the uncertainties in modelling a lower limit of 300 years 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 1 metre per century) is more likely 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 Section 3.7.2 for predictive studies of the Younger Dryas tipping event. As described in Section 3.6.1, a multitude of mathematical models, backed up by past data, show the THC to exhibit bi-stability and hysteresis with a fold bifurcation (see Figure 3.3 and discussion in Section 3.6.1). 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 behaviour, 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 bi-stability 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–Allerød 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 per cent 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 bi-stability.
- (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 (Joos *et al.* 2001; Lucht *et al.* 2006). Based on limited evidence, the reduction of the tree fraction may have characteristics more like a quasi-static transition than a real bifurcation.

3.6 Predictions of tipping points in models

3.6.1 Shutdown of the thermohaline circulation

We choose to look, first, at the thermohaline circulation (THC) 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 well known. 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 multi-stability and can switch abruptly in response to gradual changes in forcing which might arise from global warming. Its underlying dynamics are summarised *schematically* in Figure 3.3 (below) 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 freshwater 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 ,

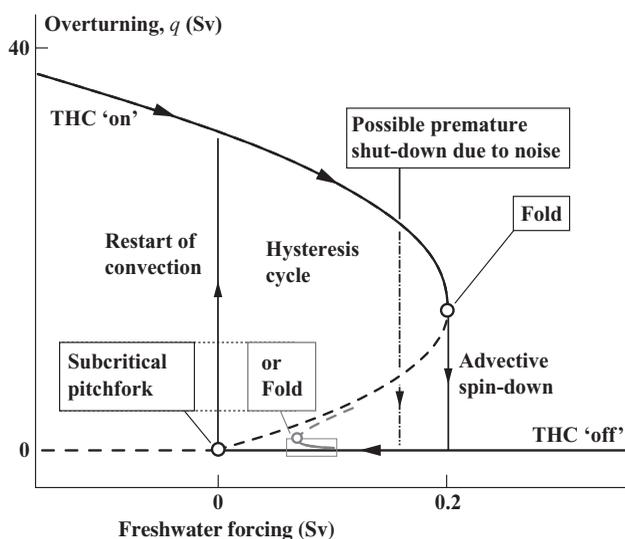


Figure 3.3 A schematic diagram of the thermohaline response showing the two bifurcations and the associated hysteresis cycle (Rahmstorf 2000). The subcritical pitchfork bifurcation will be observed in very simple models, but will be replaced by a fold in more elaborate ones: see, for example, Figure 3.5(b). Note that 1 Sv is 10^6 cubic metres per second, which is roughly the combined flow rate of all rivers on Earth.

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 shows in models has been addressed in a series of papers 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 theory 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 Figure 3.3. Finally, Huisman *et al.* (2009) have 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 bi-stability 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 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 bi-stable 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 modelling from a dynamical systems point of view, and, for example, Tziperman *et al.* (1994) and Tziperman (1997) for how predictions of ocean models connect to full global circulation models.

Building on these modelling efforts, ongoing research is actively trying to predict an imminent collapse at the fold seen in the models (for example, Figure 3.3) from bifurcational precursors in time series. Held and Kleinen (2004) use the LDR (described earlier in Section 3.4 and in Table 3.5) as the diagnostic variable that they think is most directly linked to the distance from a bifurcation threshold. They demonstrate 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 make a 50 000-years transient run with a linear increase in atmospheric CO₂ from 280 to 800 parts per million (ppm), which generates within the model an increase in the fresh water forcing which is perturbed stochastically. This run results in the eventual collapse of the THC as shown in Figure 3.4.

In Figure 3.4(a) the graph (corresponding approximately to the schematic diagram of Figure 3.3) is fairly linear over much of the timescale: there is no adequate early prediction of the fold bifurcation in terms of path curvature. The graph of Figure 3.4(b) shows the variation of the *first-order autoregressive coefficient*

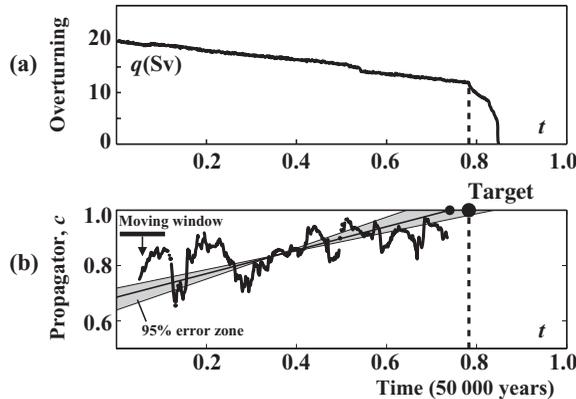


Figure 3.4 Results of Held and Kleinen (2004) which give a good prediction of the collapse of the thermohaline circulation induced by a four-fold linear increase of CO_2 over 50 000 years in a model simulation. Collapse present at $t \approx 0.8$ in (a) is predicted to occur when the propagator, $c = \text{ARC}(1)$, shown in (b), or its linear fit, reaches +1.

or *propagator*, $\text{ARC}(1)$, which is described in Section 3.4. Unlike the response diagram of $q(t)$, the time-series of $\text{ARC}(1)$, although noisy, allows a fairly good prediction of the imminent collapse using the linear fit drawn: the fairly steady rise of $\text{ARC}(1)$ towards its critical value of +1 is indeed seen over a very considerable timescale. Notice that the linear fit is surrounded by a 95 per cent zone, giving probability bounds to the collapse time. These bounds emphasise that much more precise predictions will be needed before they can be used to guide policy on whether to implement geo-engineering proposals.

3.6.2 Global glaciation and desertification of Africa

Alongside their extensive studies of past climatic events using real paleo-data, Dakos *et al.* (2008) also made some model studies as illustrated in Figure 3.5. 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 autocorrelation coefficient. Held and Kleinen (2004) and Livina and Lenton (2007) plot $\text{ARC}(1)$ at the centre of the moving window in which it has been determined. Meanwhile Dakos *et al.* (2008) plot $\text{ARC}(1)$ at the final point of this window. Here, we have redrawn the results from the latter article but shifted the $\text{ARC}(1)$ 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 the intention is to make a forward extrapolation to a*

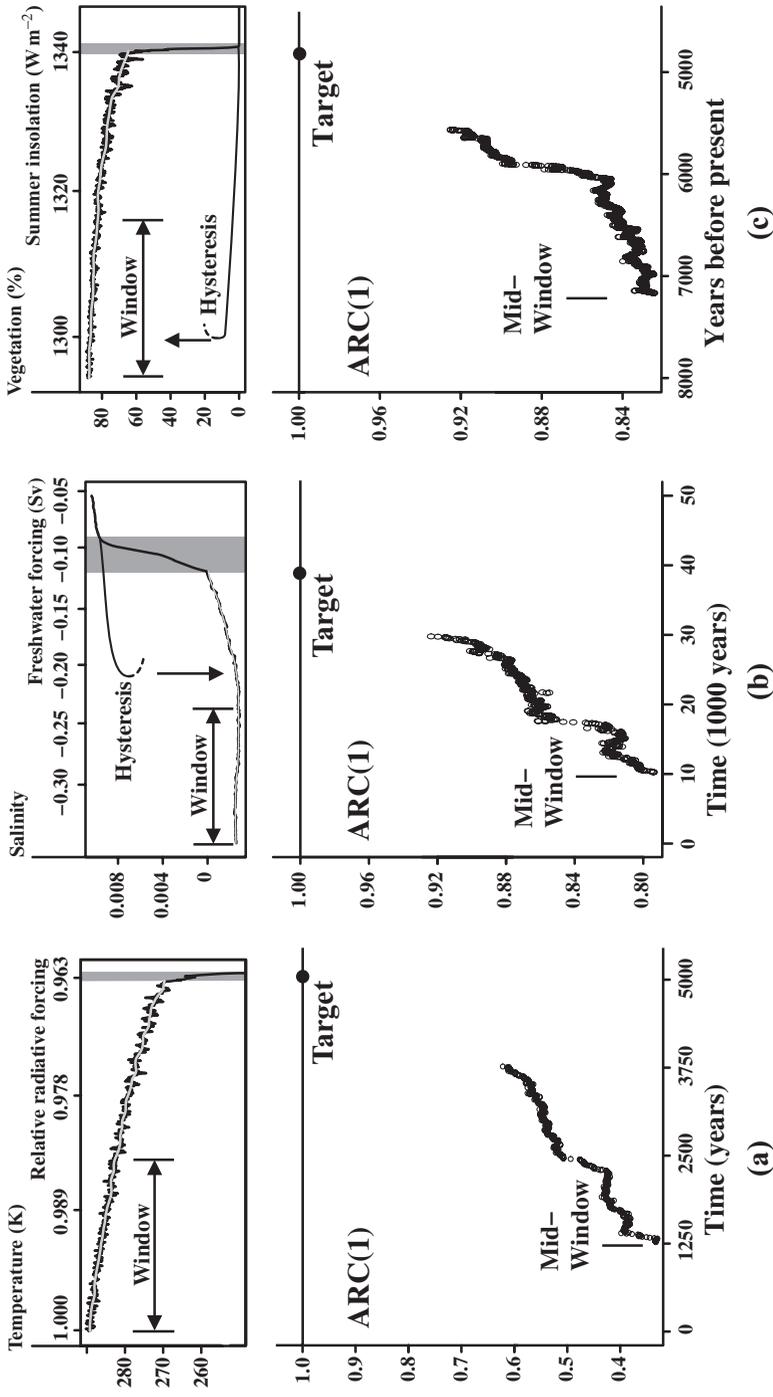


Figure 3.5 Results of Dakos *et al.* (2008) for three examples based on predictive models: (a) runaway to glaciated Earth ($N = 800$), (b) collapse of thermohaline circulation ($N = 1000$), (c) desertification of North Africa ($N = 6002$). Notice the notional hysteresis loops sketched on (b) and (c). These pictures have been redrawn as mid-window plots.

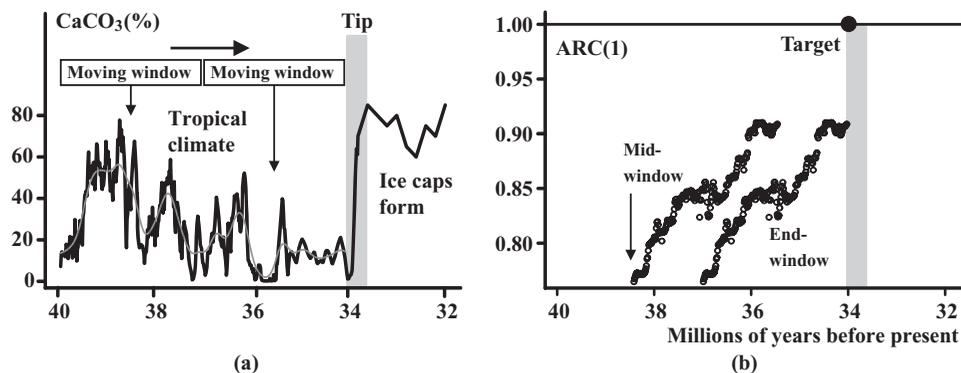


Figure 3.6 The ancient greenhouse to icehouse tipping with $N = 482$ data points. This is one of the best correlations obtained by Dakos *et al.* (2008) in their work on eight recorded tipping points. Here the sediments containing CaCO_3 were laid down 30–40 million years ago. Redrawn from Dakos *et al.* (2008), as described in the text.

target, as we are doing here (see Section 3.4.1). This forward extrapolation can be made by any appropriate method. In fact, approaching (close to) an underlying fold bifurcation, $\text{ARC}(1)$ 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.

3.7 Predictions of ancient tipplings

We have already presented the results of Livina and Lenton (2007) on the ending of the Younger Dryas event using Greenland ice-core data in Figure 3.1. 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 $\text{ARC}(1)$ propagator c extracted from the time-series of observations (as described in Section 3.4) shows a statistically significant increase, thus providing evidence that these ancient transitions indeed correspond to 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).

3.7.1 The greenhouse to icehouse tipping

We show first in Figure 3.6 their study of the greenhouse–icehouse tipping event that happened about 34 million years ago. The time-series in Figure 3.6(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

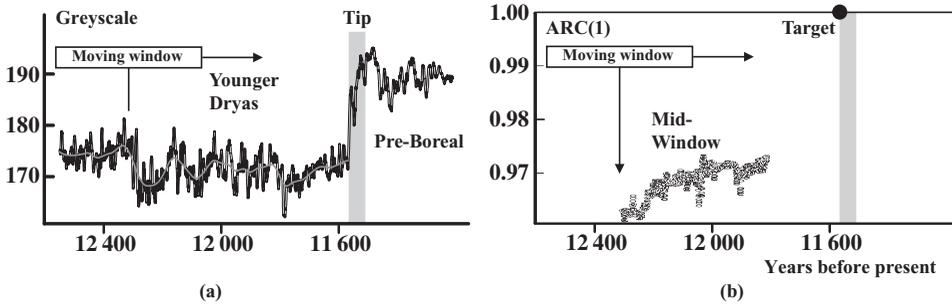


Figure 3.7 A second illustration taken from Dakos *et al.* (2008) for the end of the Younger Dryas event using the greyscale from basin sediment in Cariaco, Venezuela.

slow trends. The graph in Figure 3.6(b) shows the two plots of ARC(1) that are described in Section 3.6.2, and we notice that the mid-window projection is very close to the target, namely the known tipping point from the paleo-data.

3.7.2 End of the Younger Dryas event

To put things in perspective, Figure 3.7 shows a less well-correlated example from the Dakos paper, this one for the end of the Younger Dryas event using the greyscale 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 $\delta^{18}\text{O}$ in Greenland ice. It ended in a dramatic tipping point, about 11 500 years ago, when the Arctic warmed by 7°C in 50 years. Its behaviour 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.

In Figure 3.7(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. $ARC(1)$ is very close to its critical value $+1$ for the whole time before tipping, which suggests that the underlying deterministic system is not at an equilibrium. Note 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 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 Figure 3.7(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.

3.8 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 one may ask whether it is in fact too late to usefully apply geo-engineering because an irreversible transition is already under way.

Techniques from non-linear dynamical systems enter the modelling 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 bi-stable 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 non-linearities, currently beyond the reach of climate models, may have a strong influence on the large spatial scale. For example, Golden (2009) observes that the permeability of sea-ice to brine drainage changes drastically (from impermeable to permeable) when the brine volume fraction increases across the 5 per cent 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 non-linearities into the macroscopic picture is a challenge for future modelling 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 non-linearities, 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 3.1 involve chaos, non-linear 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 Section 3.4 for a brief description and Biggs *et al.* (2009) for a recent comparison between indicators in a fisheries model.

Acknowledgements

We are deeply indebted to many people for valuable discussions and comments. In particular, we would like to thank Professor Tim Lenton of the University of East

Anglia and his colleague Dr Valerie Livina for their continuous and detailed advice during the writing of the paper. The research group at Wageningen University in the Netherlands has also provided greatly appreciated input, notably from Professor Marten Scheffer and his research student Vasilis Dakos. Other valuable comments were received from Ian Eisenman and Eli Tziperman. Finally special thanks go to Professor Bernd Krauskopf of the non-linear dynamics group at Bristol University for his careful reading and commentary on the whole manuscript.

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. H., Dulski, P. & Nilsen, T. (2009) Rapid oceanic and atmospheric changes during the Younger Dryas cold period. *Nature Geoscience* **2**, 202–205 (doi: 10.1038/NGEO439)
- 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. (doi: 10.1073/pnas.0811729106)
- Box, G., Jenkins, G. M. & Reinsel, G. C. (1994) *Time Series Analysis, Forecasting and Control*, 3rd edition. New York: Prentice-Hall.
- Buizza, R., Petroligias, 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. *Q. J. Roy. Meteorol. Soc. B* **124**, 1935–1960.
- Carpenter, S. R. & Brock, W. A. (2006) Rising variance: a leading indicator of ecological transition. *Ecol. Lett.* **9**, 311–318.
- 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**, 14 308–14 312.
- Dijkstra, H. A. (2008) *Dynamical Oceanography*. New York: Springer.
- Dijkstra, H. A. & Weijer, W. (2003) Stability of the global ocean circulation: the connection of equilibria within a hierarchy of models. *J. Marine Res.* **61**, 725–743.
- Dijkstra, H. A. & Weijer, W. (2005) Stability of the global ocean circulation: basic bifurcation diagrams. *J. Phys. Oceanogr.* **35**, 933–948.
- Dijkstra, H. A., Raa, L. T. & Weijer, W. (2004) A systematic approach to determine thresholds of the ocean's thermohaline circulation. *Tellus Series A, Dynam. Meteorol. Oceanogr.* **56**, 362–370.
- Eisenman, I. & Wettlaufer, J. S. (2009) Nonlinear threshold behaviour 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). Lausanne and Geneva: Marcum Michaellem Bousquet.
- Golden, K. (2009) Climate change and the mathematics of transport in sea ice. *Notices of the AMS* **56**, 562–584.
- Guilyardi, E. (2006) El Niño, mean state, seasonal cycle interactions in a multi-model ensemble. *Climate Dynam.* **26**, 329–348.

- Guttal, V. & Jayaprakash, C. (2008) Changing skewness: an early warning signal of regime shifts in ecosystems. *Ecol. Lett.* **11**, 450–460. (doi:10.1111/j.1461-0248.2008.01160)
- Guttal, V. & Jayaprakash, C. (2009) Spatial variance and spatial skewness: leading indicators of regime shifts in spatial ecological systems. *Theor. Ecol.* **2**, 3–12. (doi:10.1007/s12080-008-0033-1)
- Hansen, J. E. (2005) A slippery slope: how much global warming constitutes “dangerous anthropogenic interference”? *Climatic Change* **68**, 269–279.
- Held, H. & Kleinen, T. (2004) Detection of climate system bifurcations by degenerate fingerprinting. *Geophys. Res. Lett.* **31**, L23207. (doi:10.1029/2004GL020972)
- Houghton, J. (2004) *Global Warming: The Complete Briefing*. Cambridge, UK: Cambridge University Press.
- 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. (doi:10.1029/2008GL036322)
- 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*. Cambridge, UK: Cambridge University Press.
- 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 edn. Cambridge, UK: Cambridge University Press.
- 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. *Climate Dynam.* **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 Dynam.* **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.
- 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. & Genie Team (2009) Using GENIE to study a tipping point in the climate system. *Phil. Trans. R. Soc. A* **367**, 871–884. (doi:10.1098/Rsta.2008.0171)
- 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. (doi:10.1029/2006GL028672)
- Lockwood, J. G. (2001) Abrupt and sudden climatic transitions and fluctuations: a review. *Int. J. Climatol.* **21**, 1153–1179.
- Lorenz, E. N. (1963) Deterministic nonperiodic flow. *J. Atmos. Sci.* **20**, 130–141.
- Lucht, W., Schaphoff, S., Erbrect, T., Heyder, U. & Cramer, W. (2006) Terrestrial vegetation redistribution and carbon balance under climate change. *Carbon Balance and Management* **1**, 6. (doi:10.1186/1750-0680-1-6)

- McDonald, S. W., Grebogi, C., Ott, E. & Yorke, J. A. (1985) Fractal basin boundaries. *Physica D* **17**, 125–153.
- Menocal, P. de, Ortiz, J., Guilderson, T., Adkins, J., Sarnthein, M., Baker, L. & Yarusinsky, M. (2000) Abrupt onset and termination of the African humid period: rapid climate response to gradual insolation forcing. *Quat. Sci. Rev.* **19**, 347–361.
- 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. *Climate Dynam.* **14**, 329–352.
- National Research Council (2002) *Abrupt Climate Change: Inevitable Surprises*. Washington, DC: National Academy Press.
- Nes, E. H. van & Scheffer, M. (2007) Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *Am. Nat.* **169**, 738–747.
- 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.
- 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. *Climate Dynam.* **16**, 1–17.
- Poincaré, H. (1885) Sur l'équilibre d'une masse fluide animée d'un mouvement de rotation. *Acta Math.* **7**, 259.
- Poston, T. & Stewart, I. (1978) *Catastrophe Theory and its Applications*. London: Pitman.
- Rahmstorf, S. (2000) The thermohaline ocean circulation: a system with dangerous threshold? *Climatic Change* **46**, 247–256.
- Rahmstorf, S. (2001) Abrupt climate change. In *Encyclopaedia of Ocean Sciences*, eds. Steele, J., Thorpe, S. & Turekian, K. London: Academic Press, 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. (doi:10.1029/2005GL023655)
- Rial, J. A., Pielke, R. A., Beniston, M., Claussen, M., Canadel, J., Cox, P., Held, H., de Noblet-Ducoudre, N., Prinn, R., Reynolds, J. F. & Salas, J. D. (2004) Nonlinearities, feedbacks and critical thresholds within the Earth's climate system. *Climatic Change* **65**, 11–38.
- Salinger, A. G., Bou-Rabee, N. M., Pawlowski, R. P., Wilkes, E. D., Burroughs, E. A., Lehoucq, R. Lehoucq, B. & Romero, L. A. (2002) *LOCA 1.0 Library of Continuation Algorithms: Theory and Implementation Manual*, Sandia Report SAND2002–0396. Albuquerque, NM: Sandia National Laboratories.
- Saltzman, B. (2002) *Dynamical Paleoclimatology*. London: Academic Press.
- Scheffer, M. (2009) *Critical Transitions in Nature and Society*. Princeton, NJ: Princeton University Press.
- Selten, F. M., Branstator, G. W., Dijkstra, H. W. & Kliphuis, M. (2004) Tropical origins for recent and future Northern Hemisphere climate change. *Geophys. Res. Lett.* **31**, L21205.
- Sperber, K. R., Brankovic, C., Deque, 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* **8**, 224–230.
- Takens, F. (1981), Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence*, eds. Rand, D. A. & Young, L. S. New York: Springer, p. 366.
- Thompson, J. M. T. (1982) *Instabilities and Catastrophes in Science and Engineering*. Chichester, UK: Wiley.
- Thompson, J. M. T. (1992) Global unpredictability in nonlinear dynamics: capture, dispersal and the indeterminate bifurcations. *Physica D* **58**, 260–272.
- Thompson, J. M. T. (1996) Danger of unpredictable failure due to indeterminate bifurcation. *Z. Ang. Math. Mech.* **S 4**, 199–202.
- Thompson, J. M. T. & Hunt, G. W. (1973) *A General Theory of Elastic Stability*. London: Wiley.
- Thompson, J. M. T. & Hunt, G. W. (1984) *Elastic Instability Phenomena*. Chichester, UK: Wiley.
- Thompson, J. M. T. & Stewart, H. B. (2002) *Nonlinear Dynamics and Chaos*, 2nd edn. Chichester, UK: Wiley.
- Thompson, J. M. T., Stewart, H. B. & Ueda, Y. (1994) Safe, explosive and dangerous bifurcations in dissipative dynamical systems. *Phys. Rev. E* **49**, 1019–1027.
- 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. (1997) Inherently unstable climate behaviour due to weak thermohaline ocean circulation. *Nature* **386**, 592–595.
- Tziperman, E., Toggweiler, J. R., Feliks, Y. & Bryan, K. (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.
- Winton, M. (2006) Does the Arctic sea ice have a tipping point? *Geophys. Res. Lett.* **33**, L23504. (doi:10.1029/2006GL028017)
- Zeng, N., Dickinson, R. E. & Zeng, X. (1996) Climatic impact of Amazon deforestation – a mechanistic model study. *J Climate* **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.