

Dynamics of complex systems inferred from multivariate time series



Els Weinans

Propositions

1. An increase and a decrease of complexity can both indicate the deterioration of a system.
(this thesis)
2. The traditional toolbox for the analysis of dynamical indicators of resilience is poorly equipped to deal with multivariate systems.
(this thesis)
3. The prejudice that knowledge requires exact definitions and measurements inhibits progress in qualitative modelling.
4. The career of early career scientists relies too much on personal grants at the expense of collaboratively-minded scientists.
5. Humans are bad at identifying randomness and even worse at admitting to this.
6. The systematic and unwarranted rejection of science will have catastrophic effects on public health, the economy, and the environment.

Propositions belonging to the thesis, entitled

Dynamics of complex systems inferred from multivariate timeseries

Els Weinans

Wageningen, 18 October 2021

Dynamics of complex systems inferred from multivariate time series

Els Weinans

Thesis committee

Promotor

Prof. Dr M. Scheffer

Professor of Aquatic Ecology and Water Quality Management Group
Wageningen University & Research

Co-promotors

Dr I.A. van de Leemput

Assistant professor, Aquatic Ecology and Water Quality Management Group
Wageningen University & Research

Dr E.H. van Nes

Associate Professor, Aquatic Ecology and Water Quality Management Group
Wageningen University & Research

Other members

Prof. Dr P.C. de Ruiter, Wageningen University & Research

Prof. Dr R.J. de Boer, Utrecht University

Prof. Dr T. Filatova, Delft University of Technology

Dr R.J.F. Melis, Radboud University Medical Center, Nijmegen

This research was conducted under the auspices of the Graduate School for Socio-Economic and Natural Sciences of the Environment (SENSE)

Dynamics of complex systems inferred from multivariate time series

Els Weinans

Thesis

submitted in fulfilment of the requirements for the degree of doctor
at Wageningen University

by the authority of the Rector Magnificus

Prof. Dr A.P.J. Mol,

in the presence of the

Thesis Committee appointed by the Academic Board

to be defended in public

on Monday 18 October 2021

at 4 p.m. in the Aula.

Els Weinans

Dynamics of complex systems inferred from multivariate time series

280 pages.

PhD thesis, Wageningen University, Wageningen, The Netherlands (2021)

With references, with summary in English

ISBN: 978-94-6395-888-2

DOI: <https://doi.org/10.18174/549752>

Contents

	Page
Chapter 1 General introduction	1
Chapter 2 Bridging Resilience and Complexity Indicators of Human Health	9
Chapter 3 Finding the direction of lowest resilience in multivariate complex systems	27
Chapter 4 Foreseeing the future of mutualistic communities beyond collapse	57
Chapter 5 Multivariate indicators of resilience loss	125
Chapter 6 Roots of post-truth revealed from massive language analysis	161
Chapter 7 A potential feedback loop underlying glacial-interglacial cycles	187
Chapter 8 Synthesis	219
References	229
Summary	264
Acknowledgements	268
Sense Certificate	272

Chapter 1

General introduction

Els Weinans

Complexity

Climate change, biodiversity decline and science-skepticism are three examples of challenges currently faced by science and society. These problems are not only complicated (i.e. difficult to understand), but also highly complex (i.e. composed of many interacting elements). Complex systems are characterized by interactions on different time scales and spatial scales (Simon, 1991; Arenas et al., 2006), nonlinear interactions (Golberger, 1996; Phillips, 2003), and sometimes chaotic dynamics (Lorenz, 1963; Rickles et al., 2007). Therefore traditional analysis tools have not been able to fully capture their behaviour and there is an increasing interest in tools that are designed to deal with the dynamics of complex systems (Ottino, 2003). This interest is stirred even more by the tremendous amount of high quality data that is becoming increasingly available (Yaqoob et al., 2016).

The main objective of this thesis is to explore how data sets from complex systems in different scientific fields can be used to get a fundamental understanding of their behaviour. In this first chapter, I briefly introduce some of the concepts related to complexity science, I explain why I explore these systems using time series analysis tools, and I give an overview of how each chapter in this thesis contributes to our understanding of the dynamics of complex systems.

Interactions on different scales

Complex systems are hierarchical. This means that a system is composed of interacting subsystems. These subsystems can again be broken down into smaller subsystems, until the subsystem with the lowest level is reached (Simon, 1991). Which subsystems should be taken into account to understand and predict the overall dynamics of a system, depends on the particular study objective and the knowledge of the system. For example, in physics, the atom can be taken as a low level system (and not too long ago this was the lowest level system known), but can also be considered as a complex system in itself (Simon, 1991). In biology, a cell, a leaf, a tree, a forest and a landscape are all subsystems that could be addressed. Depending on the particular question, a scale can be chosen and every scale can provide unique insights as “there is no single ‘correct’ scale on which to describe populations or ecosystems” (Levin, 1992).

An interesting observation is that various systems have similar structures that repeat on different levels. A simple example is a cauliflower, that seems to consist of smaller cauliflowers that in turn consist of even smaller cauliflowers. For these systems, a pattern is repeated on different scales, yielding a ‘scale free’ or ‘fractal’ structure (Mandelbrot, 1982). Scale free structures seem to be abundant not only in space (Mandelbrot, 1982), but also in underlying network configuration (Barabási and Bonabeau, 2003) and temporal dynamics (Zhang, 1991; Golberger, 1996).

Hierarchical structures, scale-free behaviour and fractality have not only been identified as emergent results of complex dynamics, but have also been used to quantify ‘complexity’ itself, especially in health care data. For example, the fractal dimension or the closely related Hurst exponent (Hurst, 1956), have been used to infer complexity of fMRI data (Rubin et al., 2013) and heart rate variability (Acharya et al., 2006). Furthermore, scale free properties of oscillations (i.e. indicating a signal that consists of multiple oscillations with different frequencies) are indicative of complexity for human psychology (Van Orden et al., 2011) and human gait (DiBerardino III et al., 2010).

Nonlinear interactions

Linear interactions indicate that the response of one element is proportional to the change in another driver element, as depicted in figure 1.1, left panel. Nonlinear interactions are all interactions which are not linear. A possible driver is the amount of nutrients in a lake, where the amount of algae is the response variable.

Sometimes, a linear assumption can be helpful as it greatly simplifies the analysis (Samuelson, 1947). However, for most real systems this linear assumption is only realistic for a small range of driver and response values and therefore it seems to be most useful for predictions within that range and not for mechanistic understanding (Keen, 2001). Outside of the linear regime, saturated behaviour is commonly observed. For example, increasing nutrients only enhances algae growth up to a point, after which the algae are saturated and an increase in nutrients has no effect anymore. This saturated behaviour is depicted in the middle panel of figure 1.1.

Nonlinear interactions can lead to critical transitions. For this behaviour, a system can exist in multiple states under the same conditions. A change in the driver can result in either very little change in the system, as long as it remains in the same state, or result in a very large change if the system jumps to the other state (figure 1.1, right panel). The driver value for which the system suddenly jumps from one state into the next one is called the ‘tipping point’. This behaviour can be observed in shallow lakes, where sudden shifts from clear to turbid water can be explained by the existence of a tipping point (Scheffer, 1997). These dynamics seem abundant in various fields of science, from magnetization in physics (Yang, 1952) to social-ecological systems (Biggs et al., 2018) and human psychology (Leemput et al., 2014), and have major implications for predicting and managing systems.

For a tipping point to occur, a system needs a mechanism where a small change can be amplified. One mechanism that can cause such a dynamic is a positive feedback loop. In a positive feedback loop, the interactions between variables are arranged such that an increase in one variable leads to growth of that variable itself. In the shallow lake example, this positive feedback loop can occur through vegetation. An increase in vegetation leads to clear water. Submerged vegetation, which is limited by light,




Figure 1.1: Three response curves, varying from linear to non-linear. Lines depict the stable equilibria and the dotted line represents an unstable equilibrium. Figure modified from (Scheffer et al., 2000).

benefits from this clarity and therefore their growth is enhanced, stabilizing the high vegetation/high clarity state. Similarly, when this state is disrupted, for example because an increase in nutrients has caused an increase in algae cover and therefore an decrease in water clarity, submerged vegetation will decrease, and the system will find a stable equilibrium in the low vegetation/low clarity state (Kéfi et al., 2016).

These three response curves give a first impression of the complexities that can occur in real systems that do not adhere to the linear assumption, and it has been the inspiration for many books and articles (for example Gladwell (2006)). One can imagine, that interactions in the real world do not happen in isolation, and therefore understanding one driver-response interaction seldom suffices to understand and predict a system's behaviour. An additional complicating factor is that, as multiple interactions are occurring at the same time, synergistic effects might start to play a role, where the combined effect of two drivers is not necessarily the additive effect of their individual contributions (Balduzzi and Tononi, 2008; Staal et al., 2015), leading to even more nonlinear and unpredictable dynamics.

Unpredictable dynamics

Various systems, like the weather, show highly unpredictable dynamics, making long term predictions impossible. This is partly due to stochastic fluctuations (randomness), or lack of understanding of the precise mechanisms. However, it is also possible that these systems are intrinsically highly unpredictable and 'chaotic' (Lorenz, 1963). Chaos here means that initial conditions that only differ slightly, can end up in considerably different positions as time progresses (Lorenz, 1963). Consequently, if the initial condition is not exactly known, or there is the slightest noise in the system, long term predictions are impossible.

A chaotic system may show scale-free dynamics and they can only arise from nonlinear interactions, and therefore can be a consequence of the previous mentioned two properties of complexity. Furthermore, a chaotic system is bounded, and some parts of the signal may seem to repeat themselves. This can lead to ‘attractors’ in phase space. One simple example of a chaotic attractor is the Hénon map (see figure 1.2, left panel).




Figure 1.2: The original Hénon map (one of the most simple discrete-time dynamical systems that displays chaotic behaviour, according to the equations: $x_{t+1} = 1 - 1.4x_t^2 + y_t$ and $y_{t+1} = 0.3x_t$), with its time delay embedding manifolds for X and Y. To create the time delay embedding manifolds, instead of plotting X against Y, X and Y are plotted against time-lagged versions of themselves. The three plots have different sizes and orientations, but the general shape of the Hénon map remains the same.

A remarkable finding in 1981 led to the knowledge that most properties of a chaotic attractor that consists of multiple variables, can be reconstructed using only one of the variables involved (Taken’s theorem, Takens (1981)). Even though the proof is complicated, the idea is quite simple: Different elements in a chaotic system are interacting with each other. Because of these interactions, information from one variable is contained in another variable. Therefore in principle one variable contains information of the full attractor. Indeed, for the Hénon map it is clear that both the variable X and the variable Y can reconstruct the shape of the original attractor (figure 1.2, middle and right panel).

Time series analysis

There are many ways to study complex systems. In this thesis, I focus on time series analysis tools. Three other approaches are worth mentioning. Firstly, simple models are a popular approach to study complex dynamics. One benefit of simple models are that they allow to isolate one element of the complex dynamics and therefore allow for a thorough study of one of the many complexities. Furthermore, they often are generic and therefore applicable to many different systems (Oreskes et al., 1994; Newell, 2012). The simple driver-response figures (figure 1.1) and the Hénon map (figure 1.2) are examples

of simplified models. One disadvantage is that these models can be quite theoretical in nature, and the link to the real world is not always straightforward (Evans et al., 2013b).

Secondly, computational, fully parametrized, mechanistic models have become increasingly popular with the increasing available computational power. Especially in climate science, these models have greatly enhanced our understanding of the possible future scenarios of our planet (Petersen, 2000). One disadvantage is that creating realistic models requires quite some parameters that should be chosen based on mechanistic understanding or by fitting the model to data. This problem is partly reduced by the increase of data availability that helps in the parametrization of the models (Curry and Webster, 2011). Another difficulty with realistic models, is that these models are sometimes so complex that analyzing them is as difficult as analyzing real-world data. Even when this happens, the complex models approach can provide valuable insights, especially if data collection is cumbersome (Grimm et al., 2005).

Thirdly, one approach of studying complex systems and one fundamental aspect of science is performing experiments. Experiments allow for very clean hypothesis testing, especially for inferring precise mechanisms such as causal links. However, experiments are not always feasible (e.g. the climate system) or ethical (e.g. human drug experiments).

Time series analysis tools complement these approaches. The tools that I use in this thesis are data-driven, meaning that they rely on time series data to infer a systems dynamics, and thus do not assume any prior knowledge on the system. Therefore, these tools are especially useful if precise mechanisms of underlying dynamics are unknown and high frequency data is available (Ye et al., 2015a). I focus on two groups of time series analysis tools in particular. First, I study Dynamical Indicators Of Resilience (DIORs) (sometimes referred to as ‘early warning signals’, Scheffer et al. (2009)), that are used to warn if a system is losing resilience (capacity to recover) for example because a tipping point is approaching (see figure 1.1, right panel). Secondly, I use methods based on Taken’s theorem (see figure 1.2), that use a systems attractor to infer a systems complexity (Grassberger and Procaccia, 1983b) or to detect causal links (Sugihara et al., 2012).

Simplicity

The three mentioned concepts demonstrate that complex systems cannot be assumed to adhere to the assumptions of most traditional statistical methods. Therefore, ‘complexity science’ is the relative new field of science that investigates these particular concepts and finds new ways of unraveling the mechanisms behind these complex systems. There is continuously a trade-off between diving into the complexities on one hand, but simplifying it to a graspable problem on the other hand. For the development of tools, some simplifying assumptions are often required, in order to benchmark the tools with a model

that is actually understood. However, the models should still demonstrate the properties of complex systems in order to ensure that the tools are actually complex systems analysis tools. In this thesis, I have made an attempt to balance on this line where I look for both simplicity and complexity at the same time.

Thesis overview

In **chapter 2** I give a brief overview of ‘the human’ as a complex system, with subsystems like the mood, the heart and the bones. I describe how complex system science, and in particular dynamical systems theory, has provided the tools to analyze these different subsystems. I emphasize two groups of methods in particular: Dynamical Indicators Of Resilience (DIORs), that can be used to infer the likelihood of a tipping point in a nonlinear interaction, and complexity quantification tools, that can be used to see how complex (i.e. fractal/non-linear/chaotic/irregular) the behaviour of the subsystem is. I demonstrate how both groups of methods are investigating different elements of the subsystems and how both are valuable tools when assessing human health.

Chapters 3-5 are theoretical and focus on the resilience of a system, defined as the speed in which it can recover from perturbations. In **chapter 3**, I explore the use of Min/Max Autocorrelation Factors (MAF) to detect the direction of lowest resilience in a complex system. MAF is a tool that has been developed as an alternative to PCA to reduce the dimensionality of the data by disentangling the real signal from the noise, based on the autocorrelation of the signal. I demonstrate that this tool is also surprisingly useful to find the variables that are most vulnerable to external perturbations.

Chapter 4 takes the idea that the vulnerable variables can be detected and provide a ‘direction’ in which perturbations are most dangerous, even further by exploring if this direction can also predict the future state after the system has passed a tipping point. We demonstrate that mutualistic systems have properties that allow to extrapolate the direction of these vulnerable variables to detect the direction where the system will most likely move into in the future.

During the analysis for chapters 2 and 3, I discovered some complications in applying the resilience framework to multivariate systems. Therefore, in **chapter 5**, I assess the different indicators of resilience loss that have been developed for multivariate systems and evaluate their performance in different scenarios. I demonstrate that there is not one indicator of resilience that clearly outperforms the others, but that the optimal choice depends on the scenario that the system is subject to.

In **chapters 6-7** I explore the application of these complex system analysis tools to two real world data sets. In **chapter 6** we analyze time series of the English language to look for trends and changes and see if and how they relate to societal changes. In this project we use relatively straight-forward analysis tools such as correlations and Principal

Component Analysis (PCA), and in that sense the chapter differs from the other chapters in this thesis. However, since the underlying social system can surely be considered a complex system and we investigate the temporal dynamics, it is an example of studying complexity through time series analysis. We find that the current global discontent and protests follow a shift in language where sentiment laden words become more abundant and formal language becomes less abundant.

In **chapter 7** I investigate if a causality detection method based on Taken's theorem can shed some new light on the saw-tooth shaped glacial-interglacial cycles of the past 800.000 years. Literature suggests that biological productivity, dust deposition, ocean circulation and temperature are some dominant drivers of the shape of these cycles, but the directions of the interactions (i.e. what causes what) is still a subject of debate. I demonstrate the existence of a closed loop in the data that provides a potential explanation for the saw-tooth shaped oscillations.

Finally, in **chapter 8** I reflect on these findings and their interpretation, I describe some limitations to the studies presented and some difficulties that I encountered in my attempts to understand complex systems, and I discuss some open questions and ideas for future studies.

Chapter 2

Bridging Resilience and Complexity Indicators of Human Health

Els Weinans

Jerrald L. Rector

Merlijn Olthof

Marten Scheffer

Ingrid A. van de Leemput

Abstract

In the past decades, society and science have seen an increased popularity of big data, with unsurpassed volumes, resolutions and quality. Especially in medicine there is an increasing role for big data to infer a patients health, his/her functioning, and predicting the response to interventions. These data invite on the one hand for the development of new analysis tools such as machine learning and Artificial Intelligence, but on the other hand for re-investigation of analysis tools that previously were limited by data availability. In this study we investigate two of these formerly proposed analysis tools that have a basis in complex systems dynamics: Dynamical Indicators Of Resilience (DIORs) and complexity quantification tools. We explore their similarities and differences and demonstrate how they both have specific roles when studying the human as a complex dynamical system.

2.1 Introduction

The human body is constantly regulating its own behaviour in response to both internal and external stressors, such as temperature fluctuations, viruses and perceived danger (Ramsay and Woods, 2014). For example, when the human encounters a potential danger, such as a barking dog, heart rate and respiration rate increase and the body prepares to react (fight-or-flight response). Common perturbations away from this healthy situation will be regulated, and the system will move back to the healthy situation (Ramsay and Woods, 2014). One objective of human health care is to predict how well the human body can handle perturbations (Olde Rikkert et al., 2016). For example, before a surgery an evaluation should be made whether or not the patient can survive the anesthesia and the operation and if the patient has the means to recover after the surgery (Royse et al., 2010; Scheier et al., 1989). Also, it would be useful to know in advance which subsystem might lose functionality if something goes wrong, i.e. which subsystem (such as an organ) is most likely to fail (McNicol et al., 2007).

One line of thinking that has increased in popularity during recent years, is that dynamical behaviour of various physiological processes is representative of overall ‘resilience’ (here: capacity to recover) of the patient (Scheffer et al., 2018). For example, grip strength tests have been proven to not only assess muscle fatigue (Bautmans and Mets, 2005), but also self-perceived fatigue, physical functioning and circulating IL-6 (a molecule related to inflammation) in elderly people (Bautmans et al., 2007), and even general mortality risk (Celis-Morales et al., 2018). It seems like this one sub-system (the muscles used to grip), contains information of processes on a larger scale (the whole human). Coming from a dynamical systems perspective, this idea -that one variable can be representative of full system functioning- is not surprising. In 1981, Takens wrote his influential paper where he proved that lagged states of one variable could reconstruct the dynamics of a deterministic chaotic system (Takens, 1981). In other words: If you have multiple differential equations describing the dynamics of a system, then one of the variables contains enough information to retrieve some global properties of the system such as the lyapunov exponent. Thus, it is likely that knowing one variable (such as grip strength) in a coupled system provides information of other variables as well (such as fatigue).

Methods that infer a systems property based on the dynamics of a system, such as overall health, or resilience of a patient, require fast amounts of data. This becomes especially important when the redundancy in the data cannot be extracted with linear measures, but nonlinear methods (such as the ones based on Taken’s theorem) are required. The currently unsurpassed amount of data, with increasingly high quality and resolution (Yaqoob et al., 2016), encourages re-investigation of analysis tools that were limited by data quantity and quality in the past.

In this study, we investigate how the temporal behaviour of a system can be analyzed

to infer a persons health. We focus on two popular groups of measures that aim to quantify overall health or resilience of a person: 1) Dynamical Indicators Of Resilience (DIORs), and 2) Complexity measures. Both indicators stem from a dynamical systems perspective, and find their origins in the 60s with work about nonlinear dynamics, chaos, synergetics and complexity (Haken, 1987; Prigogine et al., 1993; Thom, 2018), but despite their common origin they make different assumptions about the underlying systems and therefore cannot be used interchangeably. Before elaborating on these measures, we first provide some background information about what we want these measures to do when describing a persons health.

2.2 Background: Health & Functionality

Health has many definitions. According to the World Health Organization it is defined as “a state of complete physical, mental and social well-being and not merely the absence of disease or infirmity”. A more pragmatic definition is “having adequate physical and mental independence in activities of daily living” (Rattan, 2020). This definition is closely linked to evolutionary perspectives of health, in particular ‘physiological adaptation’, which is “concerned with maintaining energy efficiency and finding the best possible answer to lifetime encounters” (Baffy and Loscalzo, 2014). In this study we approach ‘the human’ as a complex system that is comprised of several subsystems. We consider a subsystem as well-functioning or healthy if it helps the full system (the human) in its aim for health and physiological adaptation (Baffy and Loscalzo, 2014). Both the functioning of each subsystem itself as well as the interactions between different subsystems determine overall health of the human system as a whole (Buchman, 2002).

Various subsystems in the human body aim for homeostasis. This means they work as negative feedbacks, i.e. if deviation away from a certain value are detected, the human body reverses these effects in order to return to the baseline (Ramsay and Woods, 2014). For example, when oxygen concentration in the blood drops, the human body detects this and may increase respiration in order to increase oxygen concentration back to healthy levels (Ramsay and Woods, 2014).

The concept of homeostasis is not useful to describe the behaviour of subsystems that do not aim to stay within baseline values. Therefore, in 1994 the term ‘homeodynamics’ was proposed to account for the continuously changing internal milieu of living beings (Yates, 1994). Alternatively, around that same time and for the same purpose, the term ‘allostasis’ was developed, describing “stability through change” (Sterling, 2012), where the ‘allostatic load’ are the costs related to maintaining allostasis (e.g. unrepaired molecular damage) (Rattan, 2007). Homeodynamics and allostasis can be relevant concepts when talking about systems that do not seek an equilibrium, but instead should change and adapt depending on their surrounding environment and therefore should be highly responsive.

Variables that fit in the homeostatic framework function well when they remain around a fixed value, and therefore are called regulated variables. For these variables a high variability relates to poor health (i.e. blood pressure, Parati et al. (2013)). The variables that fit in the homeodynamic/allostatic framework function well when they are able to regulate these regulated variables, and are called effector variables accordingly (Fossion et al., 2018; Rector et al., 2021). For these variables, a high variability is linked to good health (i.e. heart rate variability, Golberger (1996))

2.3 Dynamic indicators of physiological health

One way of identifying not optimally functioning subsystems, is by investigating stimulus-response patterns, such as a glucose and insulin response after an oral glucose intake (Varadhan et al., 2008). However, these experiments are not always feasible for practical or ethical reasons. Instead, non-invasive time series are becoming more abundant (Gijzel et al., 2019b). Inferring system functioning based on time series can be challenging, since mechanisms and causal links are hard to detect without experiments. Dynamical systems theory has led to several analysis tools that attempt to infer a systems functionality based on time series. In the following section we will elucidate on two concepts that have been particularly popular for medical applications and we will discuss their similarities and differences.

2.3.1 Resilience indicators

It has been suggested that various illnesses manifest themselves as a qualitatively different state of the healthy state of the same dynamical system (Olde Rikkert et al., 2016). This can happen when there are alternative stable states in the system (Strogatz, 2014). Alternative states in dynamical systems are defined as two states that can be reached under the same external conditions (Scheffer and Jeppesen, 1998). Depending on the environmental situation, states can appear or disappear. The environmental situation where a state appears or disappears, is called a tipping point or bifurcation point (Strogatz, 2014). Even when there are no alternative stable states, these tipping points can exist where very small changes in the environment can lead to very large (qualitatively different) dynamics in the system under study (Strogatz, 2014).

The idea that the loss of health of a patient reflects the shift of the dynamical system to an alternative state has major implications for diagnosis (Scheffer et al., 2018). In the last decade, several indicators of resilience have been found that signal when a system is approaching such a 'tipping point' (i.e. the point where change to the alternative state is inevitable) (Scheffer et al., 2009). These indicators are labelled 'Early Warning Signals' (or EWS, Scheffer et al. (2009)) or 'Dynamical Indicators of Resilience' (or DIORs, Gijzel et al. (2019b,a)). They rely on the fact that dynamical systems can have stable and unstable equilibria. Stable equilibria are defined as values where the system is attracted

to. If the system is perturbed, it will return back to the stable equilibrium, in line with the concept of homeostasis. Unstable equilibria are points that repel the system. If a system is perturbed slightly away from an unstable equilibria, it will travel away from it (Strogatz, 2014).

For various tipping points, such as the fold bifurcation or transcritical bifurcation, the appearance or disappearance of equilibria (or states), happens when a stable and an unstable equilibria meet each other (as in figure 2.1). Therefore, at the tipping point the system is neither attracted nor repelled by the equilibrium (Strogatz, 2014). Consequently, the recovery rate of the system to slight perturbations is exactly zero at the tipping point and approaches zero as the system approaches the tipping point. This slowing down of recovery rate in the vicinity of a tipping point (critical point) is also called ‘critical slowing down’ (CSD) (Strogatz, 2014; Scheffer et al., 2009).

most popular DIORs rely on natural fluctuations of a system around its equilibrium (Scheffer et al., 2009). All natural systems are continuously subject to small perturbations. If the recovery rate is high (i.e. far from a tipping point), each small perturbation is immediately taken care of and by the next moment in time the perturbation is ‘forgotten’. Close to the tipping point however, the recovery rate is small and therefore perturbations decay at a slower rate. This leads to a higher memory in the time series which can be uncovered by calculating the lag-1 autocorrelation of the time series (Ives, 1995; Held and Kleinen, 2004). Alternatively, since perturbations in the vicinity of the tipping point decay slowly, the system may be perturbed again before the perturbation has decayed, causing the system to move even further away from its equilibrium. Therefore, also the variance of the time series is a popular DIOR (Carpenter and Brock, 2006). Figure 2.1 demonstrates how these dynamics close and far from a tipping point can differ. Many variations to these DIORs exist that all have a basis in CSD (see Dakos et al. (2012a) for an overview of some univariate indicators or **chapter 5** for an overview of some multivariate indicators).

In order to use DIORs to infer a patient’s resilience and determine the possible risk of a perturbation such as an operation, the assumption is made that the ‘danger’ is a zero-eigenvalue bifurcation. Thus, DIORS have been developed to indicate an approaching zero-eigenvalue bifurcation. In essence, they indicate that a system is becoming slower, whether that is the result of an approaching bifurcation point, or some other reason.

Obtaining accurate DIORs is not always possible. Especially for high dimensional systems it is not always straight-forward which variables to choose, which pre-processing steps are required and to determine whether the data is of sufficient quality (See chapter 5). In many fields of science, the application of DIORs has proven difficult. However, some successful stories exist, for example in ecological systems (Van Belzen et al., 2017; Wilkinson et al., 2018), the climate (Dakos et al., 2008; Hennekam et al., 2020), financial systems (Zheng et al., 2012; Quax et al., 2013b), and also several human subsystems, of which we describe




Figure 2.1: Fluctuations of a system close and far from a tipping point. Top panel visualizes a fold bifurcation, where the driver can be any environmental variable such as temperature, and the response can be any human subsystem such as the heart. Lines reflect the stable states, and the dotted line is an unstable state. For certain values of the driver, two response states are possible. The tipping points are the points where one of the states disappears. If the driver slowly moves past a tipping point, the system will quickly jump from one state into the other, this is a critical transition. Close to the tipping point (bottom left panel), the temporal dynamics of the system are characterized by high memory (the near future is similar to the present), which can be captured by calculating the lag-1 autocorrelation. Far from a tipping point (bottom right panel), every moment in time is a random fluctuation, and the near future does not depend on the present. It is also clearly visible that close to the tipping point, the response shows a higher variance than far from the tipping point.




Figure 2.2: An intuitive definition of complexity. Purely ordered dynamics, such as a perfect sine, have a low complexity (left side). Purely random signals, such as white noise, have a low complexity as well (right side). A complexity quantification method should yield a maximum in between the two extreme cases. Figure adapted from Huberman and Hogg (1986).

four examples in the ‘examples’ section below.

2.3.2 Complexity indicators

Physiological processes are highly complex. Complexity has many definitions. One intuitive definition is that complexity finds a maximum somewhere between completely ordered data and completely disordered data (Huberman and Hogg, 1986), as we illustrated in figure 2.2. For example, if one starts with a blank sheet of paper and randomly splashes some black ink on it, the image is not considered complex, nor is it when half of the blank sheet is coloured black. However, if the ink is positioned in such a way that it creates an object, i.e. a drawing of a face, the image has a high complexity.

Quantifying this complexity can be done in several ways. First, complex dynamics are characterized by fractal behaviour. Fractal patterns are patterns that are replicated on different scales. In space, this leads to figures such as the Mandelbrot set (Mandelbrot, 1982), but this property can also be applied to time series. Self-replicating dynamics automatically lead to dynamics on different scales. Therefore ‘pink noise’ properties (i.e. noise with a power spectrum between white noise and red noise) (Van Orden et al., 2011), or ‘scale free behaviour’ (Barabási, 2005; Datta and Raut, 2006) are closely linked to fractality and have also been proposed as indicators of a system’s complexity.

It should be noted that one popular definition of complexity, namely the Kolmogorov complexity, is defined as the amount of bits needed to reproduce a string of n bits. In that sense, a completely random signal can only be reproduced by n bits, since it cannot be predicted. Therefore completely disordered (random) signals indicate a high Kolmogorov complexity instead of a low complexity. This is not concurrent with our previous definition as used in Huberman and Hogg (1986). One reason why Kolmogorov complexity based

indicators of complexity are worth using, is that chaotic signals, which can have fractal properties and could be considered as ‘complex’, are sometimes indistinguishable from randomness. Thus, one may choose to measure ‘irregularity’ (by means of a Kolmogorov entropy based quantification tool) instead of complexity according to our previous intuitive definition (Grassberger, 2012).

One last popular line of complexity quantification tools are based on correlation integrals (Grassberger and Procaccia, 1983a). The correlation integral C_d is defined as the mean probability that two states at two different times are close together. It can be calculated as

$$C_d(r) = \lim_{N \rightarrow \infty} \frac{1}{N^2} \sum_{i,j=1}^N \theta(r - |X_i - X_j|), \quad (2.1)$$

where N is the length of the timeseries, θ is the heaviside function and X_i is a vector of consecutive values of the time series $X_i, X_{i+1}, \dots, X_{i+d}$ where d is the ‘embedding dimension’. The parameter r determines what constitutes points that are ‘close together’, i.e. the distance in which you search for neighbors. $|X_i - X_j|$ is the distance between the vectors X_i and X_j . $C_d(r)$ follows a power law for small r where

$$C_d(r) \sim r^v \quad (2.2)$$

where v is closely related to the ‘fractal dimension’ of the time series (Grassberger and Procaccia, 1983a). In figure 2.3 we plotted the correlation integrals for time series of the Hénon map, a classical example of a chaotic system (see **chapter 1**), for different dimensions ranging from 1 to 5. For the Hénon map, it is known that the true dimension of the attractor is two. Figure 2.3 shows that the lines for all correlation integrals with a dimension of 2 or higher, have the same slope (the same v). This is no coincidence, if d is higher than or equal to the true dimension of the chaotic attractor, the lines will be parallel (Grassberger and Procaccia, 1983a).

The use of correlation integrals has inspired several new complexity quantification tools, of which Approximate Entropy (ApEn, Pincus (1991)), Sample Entropy (SampEn, Richman and Moorman (2000)), and Multiscale Sample Entropy (MSE, Costa et al. (2002)) are some popular examples. ApEn and SampEn are both created as approximations of the Kolmogorov complexity, based on the distance between two lines in the $C_d(r)$ plot, such as in figure 2.3. These are variations to what has been previously labeled ‘K2’ (Grassberger and Procaccia, 1983b), as an approximation of the Kolmogorov entropy ‘K’. Thus, these quantification tools, like the Kolmogorov complexity, measure irregularity as a proxy for complexity. To account for a decrease in the statistic for disordered system, MSE was developed (Costa et al., 2002). MSE calculates the SampEn for different scales of the




Figure 2.3: Correlation integrals calculated for an increasing r for the Hénon map, a classical example of a two-dimensional chaotic attractor, for different values of the dimension parameter d . When d is higher than the true dimension (2 for the Hénon map), the slope of the line does not change.

data by coarse graining the data, i.e. taking averages of windows with a varying size. It is a promising idea, but it requires choosing a number of parameters for which no rules exist, but that largely affect the results (for an illustration, see figure 2.4).

Regardless of the precise quantification tool, it has been recognized that for various physiological processes, high complexity is linked to high functionality of the subsystem, whereas low complexity is indicative of a loss of functionality and can result in illness. Four examples are discussed in the ‘examples’ section below.

2.4 Examples

Heart

From a resilience perspective, it is well known that the recovery time of the heart rate back to its baseline is an indicator of human health (Shephard, 1967; Olde Rikkert et al., 2016). Furthermore, preliminary analysis of heart rate data demonstrated that atrial fibrillation is preceded by indicators of critical slowing down such as autocorrelation (Nannes et al., 2020) and flickering (Lan et al., 2020), suggesting atrial fibrillation might be a result of a shift over a critical transition of the heart system. Furthermore, flickering might give a warning before the onset of entricular tachycardia (a type of regular, fast heart rate) (Wai, 2019).




Figure 2.4: The area under the MSE curve has been proposed as a complexity quantification tool that should be at a maximum for systems intermediary of order and disorder. We calculate the MSE for systems on this gradient by applying them to MIX models, which are models that start with a perfect sine, where a fraction of the datapoints is replaced by numbers drawn from a random distribution with zero mean and standard deviation of 1. MIX(0.3) indicates a model where 30% of the data points are replaced by random numbers. This figure demonstrates, that depending on the maximum scaling factor τ , the area under the MSE curve behaves differently, indicating that with no prior knowledge of the system, it is difficult (if even possible) to determine meaningful parameter values.

One common output measure for studying the heart is Heart Rate Variability (HRV). HRV describes the dynamical behaviour or the period between consecutive heartbeats. Heart rate variability is largely controlled by the autonomic nervous system (ANS). The ANS responds to external situations such as a good nights sleep, an encounter with a barking dog or a pleasant meal and all this information is reflected in the HRV signal (Acharya et al., 2006). The heart is a very well-known example of a subsystem that seems to be highly responsive, and whose functioning can be inferred with complexity quantification tools. For example, a high variability in heart rate is considered healthy, whereas a low variability is related to anxiety, depression and an increased risk of cardiovascular disease (Acharya et al., 2006). HRV data varies significantly for different age groups, with a higher complexity found in young healthy people. This has been tested with numerous nonlinear statistics such as Approximate Entropy, Largest Lyapunov Exponent, and Detrended Fluctuation Analysis (Acharya et al., 2006). Furthermore, complexity indicators such as approximate entropy and sample entropy have been used as an early detection of bacterial sepsis at the neonatal intensive care unit (Pincus and Viscarello, 1992; Lake et al., 2002).

Bone

Bone remodeling describes the process where old bone structure is replaced by new bone tissue. The bone adapts itself to the mechanical load (i.e. external stressors), by either adapting the density or the geometry of the new bone tissue (Weinans et al., 1992; Huiskes et al., 2000). Bone has, to our knowledge, not yet been studied using indicators of critical slowing down, however, several studies could indicate that fractures can be considered as tipping points. Bone homeostasis is a delicate balance between removal of old and micro-damaged bone parts and subsequent replacement of new strong calcified bone packages. This balance is driven by mechanical signals in a feedback system where high loading signals favors a net bone gain and low signals favors a net loss (Safadi et al., 2009). This process mimics an ‘optimal’ bone architecture in which the bone structure adapts to loads of daily life with bone struts more or less aligned with the loading pattern through the bone (Tanck et al., 2001; Huiskes et al., 2000). It has been shown that bone fracture risk at older age (osteoporosis) can be the result of loss of resilience as ongoing alignment of the bone structure makes it prone for fracture upon an infrequent “error” load that is somewhat rare but occasionally happens (Homminga et al., 2004).

Former work showed that realistic simulations models mimic the behaviour of real bone when the parameters are tuned in such a way that the system is close to a critical point (Weinans et al., 1992). The critical parameter in this model relates to information transfer in the system, where a low information transfer makes long range interaction impossible but high information transfer causes the system to homogenize, resulting in a structure of low complexity that does not resemble the complex architecture of bone (Huiskes et al., 2000). This is in line with a simple Ising spin model, where high complexity is found in the vicinity of the critical point, where the bifurcation parameter (temperature for the Ising spin model) relates to information transfer (see figure 2.5). Furthermore, complexity quantification tools such as Hurst Exponent indicate that the complexity of bone structure decreases with age (Lespessailles et al., 2002) and with diseases such as osteoporosis and inflammatory arthritis (Caldwell et al., 1998).

Mood

It is long known that clinical changes in depression do not always follow smooth trajectories but instead may be abrupt (Gaynor et al., 2003). Recent advances in human psychology demonstrate that the symptoms of psychiatric disorders, such as depression, can reinforce one another, causing feedbacks that can self-amplify or dampen the experience of emotions (Borsboom, 2017). This may lead to two stable states in the experience of emotions: one is the healthy state, the other is a depressed state. As such, this example seems to adhere perfectly to the prerequisites of critical slowing down indicators, that can signal when one states loses resilience and a shift to the alternative state becomes more likely. Indeed, critical slowing down and critical fluctuations have been found to correctly

warn before a patient ‘tips’ to a depressed state (Leemput et al., 2014; Wichers et al., 2016; Olthof et al., 2020b).

Alternatively, human emotions have been suggested to have evolved as a communication to oneself to signal the occurrence of events which are relevant to important goals (Oatley and Jenkins, 1992). As such, a healthy mood system should adapt to its environment instead of seeking to maintain an equilibrium value at all costs (i.e. emotions are always both internally and externally affected according to ‘appraisal theory’ (Ellsworth, 2013)). Therefore not only critical slowing down measures, but also a loss of complexity could indicate a loss of functioning. Indeed, recent work demonstrated that both self rated emotions and time series of self rated self-esteem are characterized by complex dynamics (Olthof et al., 2020a; Delignières et al., 2004). There is more evidence that a healthy mood system functions in a ‘complex’ way. For example, for children with anxiety, their communication pattern became less repetitive over time as their anxiety (maternal-rated) decreased (Lichtwarck-Aschoff and van Rooij, 2019) and patients suffering from bipolar disease have significantly more regular patterns in their mood record (Gottschalk et al., 1995).

Balance

The human balance system is a system that seeks an equilibrium. One test where humans were asked to use body movement to track oscillations of a videographic target, indicates a larger recovery time after a task with a higher frequency oscillations, hinting at a signal of critical slowing down (Bardy et al., 2002). Other tests rely on balance plates, where the centre of pressure of people can be analyzed during quiet standing, a test particular popular in geriatric research (Kang et al., 2009). A balance experiment between frail and healthy elderly was used to correctly distinguish between the two groups, based on their resilience. The healthier group was characterized by a fast recovery from fluctuations, indicating a high resilience, whereas the unhealthy group showed a slower response (Gijzel et al., 2019a).

Balance time series have been analyzed with several complexity measures, but the results seem to vary widely (Manor et al., 2010). A loss of complexity can be indicative of frailty. For example, decreased complexity in horizontal sway has been found to predict probability of falling in the next 48 months (Zhou et al., 2017). Furthermore, Volleyball players are taught to adapt their posture and their balance very quickly in response to game situations (Agostini et al., 2013). They are found to have larger ellipses than non-volleyball players in a balance test, in an open eyes test situation (the result disappeared in a closed eyes scenario) (Agostini et al., 2013). This suggests that high responsiveness relates to large fluctuation in horizontal sway. However, other studies showed that even though young people had a higher sample entropy in the vertical direction, elderly people had a higher standard deviation of the time series in both horizontal and vertical direction

(Borg and Laxåback, 2010), suggesting that large fluctuation relate to aging and perhaps a deterioration of the balance system. Additionally, complexity quantification analysis on balance data led to the surprising result that elderly people had a higher complexity than adults (Duarte and Sternad, 2008). This could indicate that the balance system is not aiming for high complexity (or high responsiveness) in normal situations (with the exception of professional athletes such as volleyball players), but instead aims for a high resilience.

2.5 Indicators of resilience & complexity: two sides of the same coin

So far we have made several distinctions: homeostasis vs allostasis, regulated vs effector variables and DIORs vs complexity quantification tools. Based on currently available literature, their relationship is hypothesized as follows: variables seeking for homeostasis need to be regulated. Large fluctuation around their equilibrium value often indicate a failure of maintaining homeostasis. If the system is architected in such a way that tipping points exist, a failure at maintaining homeostasis could be measured with DIORs (Gijzel et al., 2019b). The variables that have as a prime task to keep the regulated variables within their healthy range are called effector variables (Fossion et al., 2018; Rector et al., 2021). These variables are well-functioning when they are highly responsive, and therefore non-linear behaviour is highly beneficial for these variables. Their healthy dynamics are characterized by a high complexity and therefore complexity quantification tools can be used to approximate their functioning.

This explanation explicitly provides two distinctive tasks for the two dynamical indicators of physiological health. Furthermore, it hints at an explanation why nonlinear behaviour and complex dynamics are so abundant in nature: high responsiveness is beneficial for the full ‘human system’. Since high responsiveness is found near critical points, evolution is pushing systems to the non-linear regime in order to deal with the events that are continuously encountered in a complex world (Solé et al., 1996; Chialvo, 2008; Van Orden et al., 2011). For many subsystems of the human system this nonlinearity in behavior is quite intuitive. This becomes especially clear in our sensory inputs, where small signals such as soft sounds should be detected, but large signals such as loud sounds should not be experienced with the same sensitivity (see for example Nachtigall and Supin (2013) for an explanation of the mechanism in killer whales). Furthermore, the nonlinearity in sensory output allows for the filtering of noise, where part of a signal can be ignored in order to allow for the other part of the signal to become clear (for example when conversing in a crowded room) (Woods and Wilson, 2013).

DIORs have been developed for systems that are approaching a critical transition and are used to predict when such a tipping point is coming up (hence the label ‘Early Warning

Signals’). There are examples where DIORs have been used to distinguish in resilience between groups (for example Leemput et al. (2014); Gijzel et al. (2019a)), but comparison between individuals can be problematic (Bos and De Jonge, 2014). The reason for this, is that the trajectories towards illnesses show high variations between individuals. Furthermore, findings per individual do not always generalize to the group level (Van Orden et al., 2011; Fisher et al., 2018).

Complexity quantification tools have been developed as a way to characterize healthy from unhealthy individuals after the illness has developed, and are thus not used to indicate a loss of functioning before the onset of an illness. It is possible that diseases exist that are preceded by a ‘decomplexification’ in the dynamics, but so far this has received only little attention (Dakos and Soler-Toscano, 2017).

For a system that moves towards a bifurcation point that represents a tipping point to an illness, an increase in ‘complexity’ would indicate an increased danger of becoming ill. Indeed, elevated nonlinearity has not only been described as a complexity quantification tool, but also as a resilience indicator (Dakos et al., 2017). Furthermore, DIORs can signal when the Ising spin model (see figure 2.5) is pushed towards its critical temperature, which is also the location where the spatial structures are considered most complex (Morales et al., 2015). This suggests that a loss of resilience is often accompanied by an increase in complexity.

2.6 Discussion & Prospect

We have found that many subsystems can be analyzed both with DIORs and with complexity measures. Both methods measure one particular element of the dynamics. DIORs measure a slowing down of the dynamics, which could indicate an upcoming tipping point. Complexity quantification tools measure a loss of complexity, which could indicate a loss of responsiveness in the system and therefore a loss of functionality.

We hypothesize that there might be two processes at play here. On one hand, evolution has pushed systems towards a nonlinear regime, since in the nonlinear regime systems are able to be responsive and adaptive to their surroundings (Chialvo, 2008; Van Orden et al., 2011). On the other hand, this nonlinear regime may be situated next to tipping points to states representing diseases (Scheffer et al., 2018; Olde Rikkert et al., 2016). Therefore a shift towards the tipping point can be considered as dangerous, as the work on resilience indicators has demonstrated, but a shift away from the tipping point can be indicative of a loss of responsiveness and is therefore also unwanted. The work on resilience indicators and complexity quantification tools have, quite independently, described these two processes.

The subsystems we described can be analyzed with both measures. For example the heart as a subsystem has been intensively studied with both approaches. This subsystem might




Figure 2.5: A 2D Ising model on a 50x50 grid. In the Ising spin model, every cell can be either 1 (black) or -1 (white). The total energy of the grid is calculated as the negative of the product of each cell with its 4 direct neighbors, i.e. if two neighbors have the same spin the energy is low, if they have different spins their energy is high. The model is run in time, where spins are randomly flipped. If a flip results in a decrease of the energy, it is kept. If it results in an increase in the energy, it can be either kept or reversed. The temperature parameter relates to the probability of a cell to switch spin even if this increases the energy of the grid. If the temperature is low, flips are almost never accepted unless they decrease the grids energy, resulting in a grid where either all spins are 1 or all spins are -1 (see top left panel). If the temperature is high, almost all flips are accepted independent of their effect on the energy of the grid, resulting in a random grid (top right panel). The magnetization is calculated as the average spin of the grid. The magnetization shows a pitchfork bifurcation (a type of zero-eigenvalue bifurcation) at a critical temperature. For a temperature around that value, the grid shows a ‘complex’ pattern (top middle panel).

be subject to tipping points where it moves into atrial fibrillation. This shift is preceded by an increase in autocorrelation as a DIOR (Nannes et al., 2020). The subsystem can also move away from the critical regime. This shift is related to aging and illnesses such as congestive heart failure (Costa et al., 2002). For some subsystems, it is even possible that depending on the task, it moves towards or away from the responsive regime, as the work in balance suggests. In the balance system, elderly people show an increased sway, possibly indicating a loss of resilience (Gijzel et al., 2019a). However, volleyball players too have a larger sway, but this is most likely not related to a loss of resilience but instead to a skill they have developed in order to be highly responsive (Agostini et al., 2013). Furthermore, it seems like for some actions related to balance, such as walking, a loss of complexity in elderly can be restored by synchronizing them with young individuals (Almurad et al., 2018). Therefore the complexity and the resilience of a subsystem of an individual is not a static characteristic but a dynamic property by itself.

It can be expected that a decrease in complexity towards a more random regime (described as a bad thing, Golberger (1996)) relates to a decrease in autocorrelation (described as a good thing, given that the sudden situation is a healthy state, Scheffer et al. (2009)) and therefore complexity quantification tools and DIORs might contradict each other if they are used on the same data with the same questions (Rector et al., 2021). However, using system knowledge and carefully investigating specific illnesses, both measures can be used to explore the wealth of information that is contained in the dynamics of a signal. Therefore we propose that the first and most important step in investigating human time series, like many other scientific endeavors, is more about asking the right questions than asking for the right tools (Cross et al., 2010; Dodgson, 2020).

Even though the two separate processes are well-documented, the linkage between a loss of resilience and a loss of complexity has to our knowledge not yet been identified in one (sub)system. Future studies are required to explore how these two processes might work exactly. For example, simple models describing behaviour of a regulated and an effector variable can elucidate whether the mechanisms as we describe actually can give rise to the expected signals. Secondly, more realistic models can be used to test whether DIORs and complexity measures are visible in the data, and to benchmark the optimal data length and resolution for the particular subsystems. Last, we are excited to see what both DIORs and complexity measures can tell us about the human system based on empirical data. This last test will assess how these measures translate from theory to practice and might teach us new ways to look at the human as a complex system.

Chapter 3

Finding the direction of lowest resilience in multivariate complex systems

Els Weinans

J. Jelle Lever

Sebastian Bathiany

Rick Quax

Jordi Bascompte

Egbert H. van Nes

Marten Scheffer

Ingrid A. van de Leemput

This chapter is based on:

Els Weinans, J Jelle Lever, Sebastian Bathiany, Rick Quax, Jordi Bascompte, Egbert H Van Nes, Marten Scheffer, and Ingrid A Van De Leemput. Finding the direction of lowest resilience in multivariate complex systems. *Journal of the Royal Society Interface*, 16 (159):20190629, 2019

Matlab code that reproduces all figures from this chapter can be download from https://github.com/elsweinans/Direction_lowest_resilience

Abstract

The dynamics of complex systems, such as ecosystems, financial markets, and the human brain, emerge from the interactions of numerous components. We often lack the knowledge to build reliable models for the behaviour of such network systems. This makes it difficult to predict potential instabilities. We show that one could use the natural fluctuations in multivariate time series to reveal network regions with particularly slow dynamics. The multidimensional slowness points to the direction of minimal resilience, in the sense that simultaneous perturbations on this set of nodes will take longest to recover. We compare an autocorrelation-based method with a variance-based method for different time series lengths, data resolution and different noise regimes. We show that the autocorrelation-based method is less robust for short time series or time series with a low resolution but more robust for varying noise levels. This novel approach may help to identify unstable regions of multivariate systems or to distinguish safe from unsafe perturbations.

3.1 Introduction

Many complex systems are managed or structured such that they are relatively stable, in the sense that they can maintain the same functions. Examples include the human body (Olde Rikkert et al., 2016), financial systems (Zheng et al., 2012), ecosystems, or social systems (Walker et al., 2004). All of these systems can be represented as networks (Barabási et al., 2016) with multiple interacting entities, such as organs, banks or companies, species and abiotic factors, or individual human beings (Castellano et al., 2009). All network entities are continuously disturbed by external events that bring the full system somewhat out of balance. For instance, climatic extremes, diseases, or human interference may result in a temporary increase or decrease in abundance of one or more species (Scheffer et al., 2001). Environmental fluctuations and disturbances affect different species in different ways (Bender et al., 1984; Elmqvist et al., 2003), and particular compounded perturbations may have much larger impacts than when such perturbations occur in isolation (Paine et al., 1998). It is intuitively straightforward that for each system there is a particular type of perturbation (in the sense that a certain set of network entities is disturbed simultaneously in a particular way) to which the system is the most sensitive (Scheffer, 2009). This raises the question of whether we might be able to deduce such ‘weak spots’ in the myriad of possible combinations of pressures on the system.

In this study we are thus interested in finding the particular combinations of pressures from which a system will recover the slowest. In other words, we aim to identify network regions with low resilience, where resilience is defined as the rate at which a system recovers after a perturbation, also often called engineering resilience (Holling, 1996). The underlying configuration of the network and the interactions between elements is often unknown, making it hard to rely on models to simulate dynamics and find such weak spots in the system. Another approach is to use observed time series to search for combinations of variables with low resilience. If we assume a homogeneous network where all nodes and connections are similar, methods exist that may find universal patterns of resilience (Gao et al., 2016). However, for heterogeneous networks other approaches are required. One line of research has looked at the rate of recovery from perturbations as an indicator of resilience. For example, if a system is intrinsically slow, this should be reflected in the observed time series by a high autocorrelation (Ives, 1995; Dakos et al., 2008; Held and Kleinen, 2004) and a high variance (Carpenter and Brock, 2006; Wiesenfeld, 1985). While most literature on resilience indicators (also often called ‘early warning signals’) focuses on univariate data, recently the first steps towards resilience indicators based on multivariate time series (i.e. of network-type systems) have been taken. Suggested metrics to indicate the overall resilience of the system include the autocorrelation of the projection of data on the first principal component (PC) using Principal Components Analysis (PCA) (Held and Kleinen, 2004), combinations of cross-correlations between system elements and variance of individual elements (Chen et al., 2012), mean autocorrelation and variance




Figure 3.1: Using autocorrelation in a two-dimensional system to predict directions of fast and slow recovery. (A) Stability landscape of two interacting species (X and Y), showing that the speed of recovery depends on the direction of a disturbance. The speed is indicated by the slope of the stability landscape. (B) Autocorrelation along different directions in the system’s phase space. The scatterplot shows part of the time series. Red lines indicate the slowest direction (dashed line), i.e., with highest autocorrelation, and the fastest direction (solid line), i.e., with lowest autocorrelation. The colour circle indicates the autocorrelation in every direction. (C-D) Projected time series on the slowest (C) and the fastest (D) direction.

(Dakos and Bascompte, 2014) and the maximum value of the covariance matrix (Suweis and D’Odorico, 2014). However, so far, these studies have mostly focused on finding a scalar indicator of resilience, and not so much on identifying the combination of variables involved.

Here we propose that one could use observed natural fluctuations to map the multivariate pattern of indicators of slowness such as temporal autocorrelation (Fig. 1 B). The basic idea is most easily illustrated from a stability landscape illustration of a hypothetical two-dimensional system describing the dynamics of two interacting species X and Y (Fig. 3.1 A). From the shape of the stability landscape, it is intuitively clear that a disturbance resulting in an increase or decrease of both species X and Y will return to equilibrium relatively slowly. In contrast, the system will recover much quicker from a disturbance of the same strength resulting in an increase in X combined with a decrease in Y, or vice versa. Now, if we assume this system to be continuously perturbed in random directions, we can use the observed timeseries of X and Y to find the direction of slowest recovery simply by computing temporal autocorrelation or variance projected on all possible axes (Fig. 3.1.B-D). In the two-dimensional case finding this slow direction can be done by brute computational force. However, as the number of dimensions increases it becomes impossible to scan all directions. We will show how novel ways of utilizing known tools based on autocorrelation or variance allow scanning for the direction of lowest resilience even in highly complex networks.

We assess the suggested methods by applying them to synthetic data where we know the underlying mechanisms. Since in multivariate systems the link between a high autocorrelation or variance and a slow recovery is not as straightforward as in univariate systems,

we also assess what we can expect from these resilience indicators in our multivariate examples. We pick three example models with varying degrees of complexity that allow us to compare the predictions with the actual direction of slowest recovery. Furthermore, we evaluate the robustness of both autocorrelation- and variance based approaches for the length and resolution of the data and for different noise regimes. We introduce a test to assess if a particular real-world multivariate time series is suitable for the proposed analyses and discuss which method one should preferably use in which case.

3.2 Methods

Finding the direction of slow recovery

In order to find the slowest direction in a multivariate timeseries, we detect the direction of highest autocorrelation by using the Min/Max Autocorrelation Factors (MAF) analysis (Switzer and Green, 1984), which we explain below. Additionally, we detect the direction of highest variance by using the well-known PCA. We use simulated multivariate time series with equal temporal spacing between data points to investigate the general applicability and performance of both methods.

The MAF algorithm detects the direction of the highest variance of the first difference (difference between consecutive time points) of the time series. In a timeseries with high autocorrelation, the similarity between consecutive time points is high, which relates to a low variance in the first difference. Similarly, low autocorrelation relates to high variance in the first difference. The MAF algorithm detects the direction of maximum autocorrelation in a four step process:

1. We transform the data to ensure that they have an identity matrix as the covariance matrix. In line with (Haugen et al., 2015) we use an ‘SDS transform’ (spectral decomposition sphering):

$$X_{SDS} = X * U * D^{-\frac{1}{2}} * U',$$

where X is the original dataset, X_{SDS} is the transformed data, U is the eigenvector matrix of the covariance matrix of the data and D is a diagonal matrix with eigenvalues of the covariance matrix.

2. We calculate the first differences of X_{SDS} , resulting in $[X_{SDS}(t) - X_{SDS}(t + 1)]$.
3. We calculate the eigenvector matrix V and the eigenvalues E of the covariance matrix of $[X_{SDS}(t) - X_{SDS}(t + 1)]$. These eigenvalues can be used to determine how different the variances of the different eigenvectors are.
4. We calculate the MAFs:

$$W_{MAF} = U * D^{-\frac{1}{2}} * U' * V.$$

More details about the procedure can be found in (Switzer and Green, 1984; Desbarats and Dimitrakopoulos, 2000; Haugen et al., 2015). The output of the MAF analysis is a set of components called the MAFs, which are ordered from high to low autocorrelation. These can be compared with the PCs of a PCA, which are ordered from high to low variance. So like the PCs in PCA, we can project the data on the MAFs or summarize the data using only a number of MAFs to reduce the dimensionality. In contrast to PCs, the MAFs do not have to be orthogonal to each other. Since a high autocorrelation is linked to low resilience, the MAFs order the directions of the system from low to high resilience.

To be able to compare the MAFs, we use the MAF eigenvalues (E) belonging to the eigenvectors of the covariance matrix of $[X_{SDS}(t) - X_{SDS}(t + 1)]$ that we calculated in step 3. Similar to the explained variance in PCA, the MAF eigenvalues provide a weight to the autocorrelations projected on each MAF. In contrast to PCA, a MAF with a low eigenvalue indicates that the autocorrelation of the projected timeseries is higher than all other directions, whereas a high eigenvalue indicates a low autocorrelation.

Models

To test and compare the potential methods to detect the direction of lowest resilience based on multivariate time series, we apply them to time series generated by three different models. The models have a deterministic part and a stochastic part. For the stochastic models we use an Euler-Maruyama integration. For the deterministic models an Euler integration is used. To generate the time series, we used Grind for Matlab (Nes, 2017).

Metapopulation model

First, as an example of a gradient non-reactive system, we use a classical ecological model that is known for having alternative stable states (a bistable model) (May, 1977). Alternative stable states are multiple states that are stable under the exact same parameter settings. The model describes the abundance of a logistically growing species that is being harvested following a Holling's type III functional response. The modelled species could for instance represent a plant that competes for space and is being grazed by herbivores. The grazing efficiency of the herbivores may increase with plant abundance until a certain biomass is reached, at which point the herbivores become saturated. For this study, we simulate a metapopulation with three patches and assume that the modelled species can migrate between the patches,

$$dN_i = \left[N_i \left(1 - \frac{N_i}{K_i} \right) - \frac{c_i N_i^2}{1 + N_i^2} + \sum_{j \neq i} d_{ij} (N_j - N_i) \right] dt + \sigma_{N_i} dW_i \quad (3.1)$$

where N_i is the abundance of the species in location i , K_i is the carrying capacity at location i , c_i is the maximum harvesting rate and d_{ij} is a symmetric matrix describing migration between patch i and j . Finally, each patch is affected by noise, with dW_i representing a Wiener process with mean 0 and variance σ that is uncorrelated for the different variables. Default parameter settings are: $K_1=10$, $K_2=13$, $K_3=8$, $c_1=3$, $c_2=2$, $c_3=2.3$, $d_{21}=d_{12}=0.2$, $d_{31}=d_{13}=d_{32}=d_{23}=0.08$, and $\sigma_{N_i}=0.02$. A time step of 0.01 was used for integration. The parameters were chosen such that some asymmetries occur in the resilience in different directions.

It should be noted that this model is extremely simplified and the parameters are not based on observations. This first model is chosen because it is well-known and can easily be used for visualizations and for explaining how to interpret the MAF results.

Sahara model

Second, as an example of a non-gradient non-reactive system we use a simple climate model describing vegetation-precipitation interactions in four regions of the Sahara. This model has been used to explain the shift from a vegetated state to a desert state in the Sahara region. The model was developed by Brovkin et al. (1998) and made spatially explicit by Bathiany et al. (2013). The model describes the growth of the vegetation as a function of the current vegetation and the equilibrium vegetation cover, which depends on the precipitation in that location,

$$dV_i = \left[\frac{V^*(P_i) - V_i}{\tau} \right] dt + \sigma dW_i, \quad (3.2)$$

where V is the vegetation cover, $V^*(P_i)$ is the equilibrium vegetation cover as a function of the precipitation at location i and τ is the characteristic time scale.

The vegetation equilibrium is described by:

$$V^*(P_i) = 1.03 - \frac{1.03}{1 + \alpha \left(\frac{P_i - P_1}{\exp(\gamma \delta)} \right)}. \quad (3.3)$$

where δ stands for the growing degree days (minus 900 K). The dependency of vegetation on temperature in the Sahara is, however, rather unimportant compared with rainfall. The parameter γ determines how steep the $V^*(P_i)$ curve is, i.e. the sensitivity to rainfall. Precipitation reacts much faster than vegetation cover and is therefore assumed to be in its equilibrium (quasi-steady-state assumption), which depends on V ,

$$P_i(V) = P_{0i} + s_i B + \sum_{j=1}^N \mathbf{k}_{ij} V_j, \quad (3.4)$$

where $P_{0_i} + s_i B$ is the amount of precipitation if no vegetation existed and \mathbf{k}_{ij} is the sensitivity of the precipitation in location i to the vegetation in location j . Therefore \mathbf{k} is the parameter that couples the locations to each other. Default parameter settings are chosen in line with (Bathiany et al., 2013) as $N=4$, $\tau = 1$, $\sigma_{V_i} = 0.02$, $\alpha = 0.0011$, $\beta = 28$, $\delta = 9100$, $P_1 = 60.6855$, $P_0 = [-50, 40, 210, 40]$, $s = [1.7, 0.8, 0.2, 0.9]$, $B = 100$ and $k = \begin{bmatrix} 243 & 30 & 50 & 50 \\ 135 & 24 & 15 & 15 \\ 72 & 12 & 75 & 10 \\ 18 & 18 & 10 & 25 \end{bmatrix}$. A time step of 0.01 was used for integration.

This Sahara model is fitted to observations and is therefore slightly more realistic than the meta-population model.

Gene regulatory network

Third, as an example of a non-gradient reactive system we use a simple network of gene regulations among five genes, described by Chen et al. (2002). This model describes the concentration of five molecules (e.g. gene or protein expressions),

$$\begin{aligned} dz_1 &= \left[90|P| - 1236 + \frac{240 - 120|P|}{1 + z_3} + \frac{1488z_4}{1 + z_4} - 30|P|z_1 \right] dt + \sigma dW_1, \\ dz_2 &= \left[75|P| - 150 + \frac{60 - 30|P|}{1 + z_1} + \frac{(240 - 120|P|)z_3}{1 + z_3} - 60z_2 \right] dt + \sigma dW_2, \\ dz_3 &= \left[-1056 + \frac{1488z_4}{1 + z_4} - 60z_3 \right] dt + \sigma dW_3, \\ dz_4 &= \left[-600 + \frac{1350z_5}{1 + z_5} - 100z_4 \right] dt + \sigma dW_4, \\ dz_5 &= \left[108 + \frac{160}{1 + z_1} + \frac{40}{1 + z_2} + \frac{1488}{1 + z_4} - 300z_5 \right] dt + \sigma dW_5, \end{aligned} \tag{3.5}$$

where z_i is the concentration of molecule i and P is a scalar control parameter. The gene regulation growth rates are described by the Michaelis-Menten equation and the degradation rates are proportional to the concentration of the genes. There is a stable equilibrium at $\bar{Z} = (\bar{z}_1, \bar{z}_2, \bar{z}_3, \bar{z}_4, \bar{z}_5) = (1, 0, 1, 3, 2)$ and a tipping point at $P = 0$. For our simulations we use $P = 0.35$ and $\sigma = 0.2$. A time step of 0.001 was used for integration.

This model is not based on observations, but it is tuned to display dynamics not unlike real biomarker dynamics (Chen et al., 2012). Furthermore, it is a more complex model than the other two models. In this way, our models have different levels of realism and different levels of complexity.

Perturbation experiments

To verify whether the direction with the highest autocorrelation is also the direction in which perturbations recover slowest, we performed perturbation experiments in the direction of the different MAFs, using the deterministic models. We expect that the speed of recovery of perturbations in the direction of the MAF will be ordered according to the order of the MAFs. This should be true for systems that return to their equilibrium in a relatively linear way. However, when strong spiraling dynamics occur, the system can move away from the direction of the MAF after the perturbation and recovery rates may become different. Therefore we expect that the initial recovery rate is well ordered according to the MAFs, and the later recovery rates are only well ordered when the system recovers in a linear way. We capture the initial recovery time by looking at the moment when the perturbation is at 10% of its recovery. In real-life applications, there is often an interest in more than 10% recovery. For example an ecological system is normally not labeled as “recovered” until it is indistinguishable from the situation prior to the perturbation. Therefore recovery times are also calculated for 50% and 90% recovery. For all perturbation experiments, perturbation size is three times the standard deviation of the Gaussian white noise process used for the simulations.

Last, to check if the first MAF really provides the direction of slowest recovery, we did 1000 random perturbations for every model and calculated the recovery times for every one of them to assess if the perturbation on the first MAF was really the perturbation that would lead to the longest recovery.

Performance of MAF vs PCA

We evaluated the effect of data length and resolution on the performance of MAF and PCA. To test whether the time series is of sufficient length, we performed a block bootstrap with increasing block size. We started with a block size of 0.1% of the data size, and then we randomly picked 100 blocks in the data. The blocks could overlap. For every block we calculated the first MAF and first PC, resulting in a distribution based on 100 blocks of which we calculated the median and the 90% confidence interval (5% and 95% boundary). Next, we increased the block size and repeated the analysis. We repeated this until the block size was 10% of the data size. If the confidence interval converges to a small value, we conclude that the data are of sufficient length. This procedure can be done with any available dataset to evaluate whether it is of sufficient length and quality for our analysis.

To examine the effect of data resolution, we again performed a block bootstrap with blocks of size 1000 for different distances between data points. We chose the blocks by starting at a random point in the time series and then taking every n th point until the box was full at 1000 points. We let n range from 1 to 1000. Again, we used a 100 boxes per n and we calculated the median and the 90% confidence interval.

Furthermore, we tested the performance of MAF and PCA in the case that noise is unequally distributed over the variables. For this, we used the metapopulation model, with different noise levels for each variable. We simulated all combinations of noise levels, keeping the sum of the noise ($\sum_{i=1}^n \sigma_{N_i}$) at a constant of 0.2. For every noise regime, we calculated the similarity of the obtained MAF and PCA with the true direction of slowest recovery (see below).

Metric for comparing directions

In order to evaluate the performance of MAF and PCA, we compared both of them with the true direction of slowest recovery. We calculate the latter with the deterministic version of the model (using 50% recovery). Then, we calculate the angle between the first (MAF or PCA) component and the vector in which the system shows slowest recovery when perturbed along that vector with the formula

$$\theta = \cos^{-1} \frac{C \bullet V}{\|C\| \|V\|} \quad (3.6)$$

Where C is the calculated direction (first MAF or first PC) and V is the real direction of slowest recovery. The \bullet operator indicates the dot product. Next, we use the following probability density function that calculates the probability of finding an angle θ when comparing two random vectors with each other:

$$h(\theta) = \frac{1}{\sqrt{\pi}} \frac{\Gamma(\frac{d}{2})}{\Gamma(\frac{d-1}{2})} \cdot (\sin\theta)^{d-2} \quad (3.7)$$

where Γ is the Gamma function (a factorial function that can handle non-integer numbers) and d is the dimension of the input vectors (Cai et al., 2013).

We calculate the vector similarity as $1-p$, where p is the probability of finding two random vectors that have an angle that is equal to or smaller than the angle between the two vectors (using equation 3.7). This depends on both the angle of the two vectors and the dimensionality of the space (Cai et al., 2013).

3.3 Results

Interpreting MAF analysis

In this section, we will apply the MAF analysis to the three models and discuss the interpretation of the results of a MAF analysis. We start with the three-dimensional meta-population model, since the low dimensions allow for clear visualization.

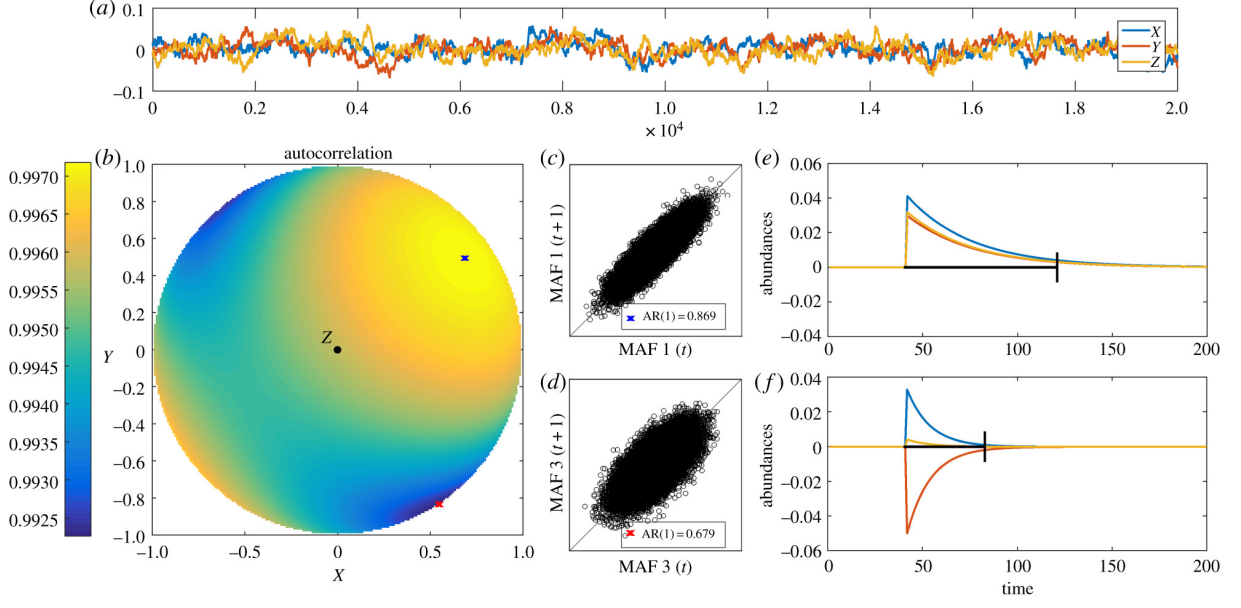


Figure 3.2: Direction of slow and fast recovery as detected by MAF. (A) Time series of the metapopulation model with three patches. (B) Autocorrelation for all possible directions in the three-dimensional plane, for $Z > 0$. The blue X indicates the direction of the first MAF (slowest direction), the red X indicates the direction of the last MAF (fastest direction). (C-D) Lag-1 autocorrelation for the projected time series on MAF 1 (C) and MAF 3 (D). (E-F) Perturbation experiments on the first and last MAF, showing that a perturbation on the first MAF results, as expected, in a slower recovery than a perturbation on the last MAF. The black lines indicate 90% recovery.

First, we calculate the autocorrelation in all possible directions. Just like the two-dimensional example in Figure 3.1, we depict the autocorrelation for the different directions with a colour gradient. In the two-dimensional example we plotted it on a circle, but in this three-dimensional case we need a sphere to visualize all directions (Figure 3.2B). It is important to note that, just like the circle in Figure 3.1B, only half of the sphere is needed since the circle is symmetrical (e.g. autocorrelation in direction $[1 \ 1 \ 1]$ is the same as autocorrelation in direction $[-1 \ -1 \ -1]$). Therefore we can look at the sphere from any side. We choose to look at the side where $Z > 0$ (Figure 3.2 B), but any other angle would give exactly the same result. We show how the MAF analysis accurately captures the direction of highest (blue X) and lowest (red X) autocorrelation in this case (Figure 3.2B). Figures 3.2 C and D indicate how the autocorrelation differs in the two directions. Perturbations in the direction of the first and last MAF show strong differences in recovery time (Figure 3.2 E and F), where a perturbation on the first MAF (Figure 3.2 E) lasts longer than a perturbation on the last MAF (Figure 3.2 F).

Next, we perform the MAF analysis for the other two models. After obtaining the MAFs, we perturb the system on the different MAFs. The expectation is that the perturbation on the first MAF, the one with the highest autocorrelation, will take the longest to recover

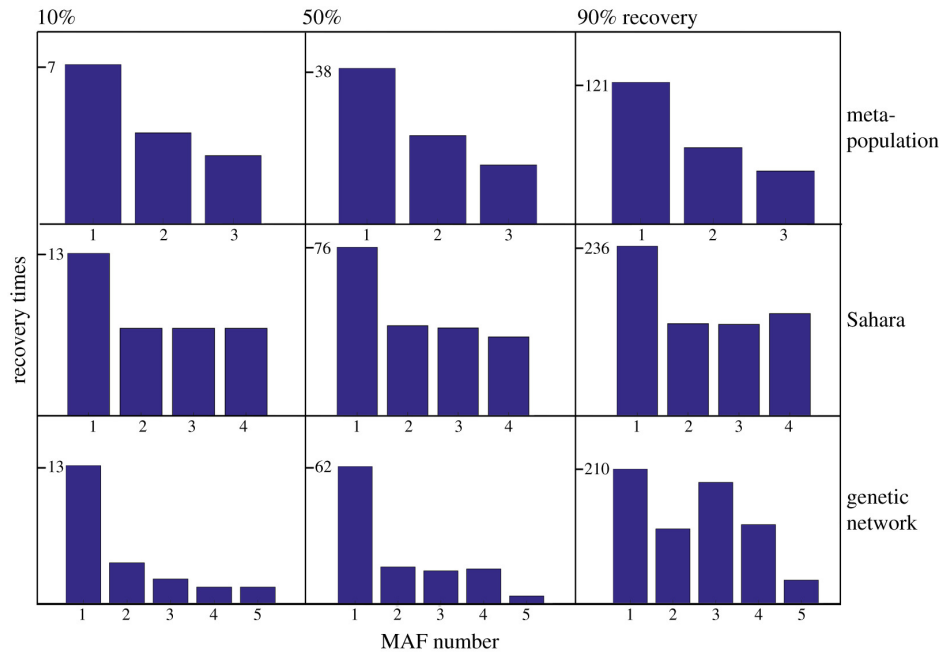


Figure 3.3: Recovery times for 10% (left), 50% (middle) or 90% (right) recovery. The initial recovery (10%, left three barplots) is well ordered by MAF, with low MAF numbers corresponding to high recovery time, 50% recovery is well ordered for the meta-population and the Sahara model and 90% recovery is only well ordered for the meta-population model which responds in a gradient way. However, in all cases the recovery time is highest for the first MAF, even for the non-linear genetic network.

and the recovery time will increase as the MAF number increases, where the shortest recovery time will be found for a perturbation on the last MAF (Figure 3.3). We see that for 10 % recovery the MAFs are indeed ordered to the recovery time of a perturbation in their direction. For 50% recovery, this is true for the meta-population and the Sahara model but not for the genetic network; and for 90% recovery it is only true for the meta-population model and not for the Sahara model or the genetic network.

The time trajectories of the perturbations are plotted in supplementary Figures S1 and S2. Here we see that for the Sahara model the recovery happens in a gradual way, just as in Figure 3.2 E and F in the meta-population model. However, in the genetic network we see some fluctuations before recovery occurs, a result of the complex eigenvalues of the model which explains why directional autocorrelation does not reflect recovery times well.

Apart from the recovery times we also calculate the MAF eigenvalues that indicate how different the autocorrelations on the different MAFs are from each other. Figure 3.4 shows the MAF eigenvalues for every MAF for the meta-population model (A), the Sahara model (B) and the genetic network (C). For the meta-population model and the genetic network there is a clear increase in the MAF eigenvalue for increasing MAF number, indicating a




Figure 3.4: Difference in autocorrelation between MAFs as indicated by their eigenvalues E . MAFs are ordered by their eigenvalues E (as described in step 3 of the MAF procedure) in the meta-population model (A), the Sahara model (B) and the genetic network (C). Low eigenvalues indicate a high autocorrelation. If two eigenvalues are similar, this indicates that there is only a little difference in the recovery times of the MAFs, for example with MAF 2 and 3 in the Sahara-model (B), which can be verified in Figure 3.3.

clear difference in autocorrelation for the different directions. For the Sahara model there is hardly any difference in autocorrelation for MAF 2, 3, and 4. This is also reflected in the recovery times of perturbations on MAFs 2, 3, and 4 (Figure 3.3 and S1).

Last, we perturbed the system in 1000 random directions and calculated the recovery time for all of them. Here we see that for the non-spiralling systems (the meta-populated and the Sahara model) the first MAF was the direction of slowest recovery. For the spiralling genetic network however, even though a perturbation on the first MAF yielded a slower recovery than a perturbation on the other MAFs, it was not the slowest direction of the system (supplementary Figures S3-S5). This shows that, for this model the direction of maximum autocorrelation is not representative for the direction of slowest recovery. This model is a reactive model (Neubert and Caswell, 1997), where perturbations exist that first grow in amplitude before they return to their equilibrium. These directions affect the MAF analysis. The other two models are not reactive (supplementary page 10).

Effect of timeseries length

We evaluate the effect of the length of the time series on the robustness of the results by performing our data suitability test, which consists of a block bootstrap with increasing block size. We find that for all our models there is clear convergence for the first MAF and the first PCA, indicating that the data are suitable for the analysis (Figure 3.5 A-B for the meta-population model and supplementary Figures S6-S7 for the other two models). MAF and PCA both need about 60 000 time points for the meta-population and Sahara model and 20 000 time points for the genetic network before convergence of the 90% confidence interval is reached.




Figure 3.5: Top panels show the effect of time series length on MAF (A) and PCA (B) for the meta-population model. MAF and PCA are calculated using a block bootstrap with 100 blocks of increasing size (horizontal axis). Solid lines indicate the median of the 100 blocks for every variable. Shaded areas show the 5% and 95% borders. Both MAF and PCA need blocks with a size of about 60.000 to reach an accurate result. Bottom panels show the effect of data resolutions on MAF (C) and PCA (D) for the meta-population model. Again, solid lines show the median of 100 blocks with an increasing distance between consecutive data points (horizontal axis). Shaded areas capture the 5% and 95% borders of the 100 blocks. The Figures clearly show that data resolution has a strong effect on the MAF analysis whereas it does not affect PCA.

Effect of data resolution

To evaluate the effect of data resolution, we perform the block bootstrap for different data resolutions (Figure 3.5 C-D for the meta-population model and supplementary Figures S6-S7 for the other two models). The first striking observation is that a sampling distance of 1 does not yield the smallest confidence interval, indicating that for both methods it is possible that the data are over-sampled, in which case reducing the amount of data could improve the result. Second, for increasing distance between points, MAF results become inaccurate, whereas data resolution does not affect PCA.

Effect of noise distribution over variables

In our previous analysis, we used Gaussian additive white noise, which is the same for all variables. To evaluate the effect of different noise types, we experiment with differently

distributed noise over the different variables. For all analyses we keep the sum of the noise at 0.2 ($\sum_i \delta_{N_i} = 0.2$).

Figure 3.6 shows the performance of MAF and PCA for different noise regimes. A location in the plot represents the noise distribution over the three variables and the color scale indicates the performance. Performance is measured as the similarity between the MAF/PCA direction and the direction of slowest recovery. For instance, a similarity of 0.8 means that the probability of finding two random vectors that have an angle that is smaller than the angle between the two vectors is 0.2. For this model, the true direction of slowest recovery is on the vector $[0.68 \ 0.52 \ 0.52]$. If the result of PCA and MAF point in the direction of only two variables, such as $[0.7 \ 0.7 \ 0]$, our similarity measure yields a score of 0.92. Therefore similarity values lower than 0.95 are not very meaningful. We consider performance of the method to be 'reasonable' when the similarity between the two vectors is higher than 0.95 and 'good' when similarity is higher than 0.99 (see contours).

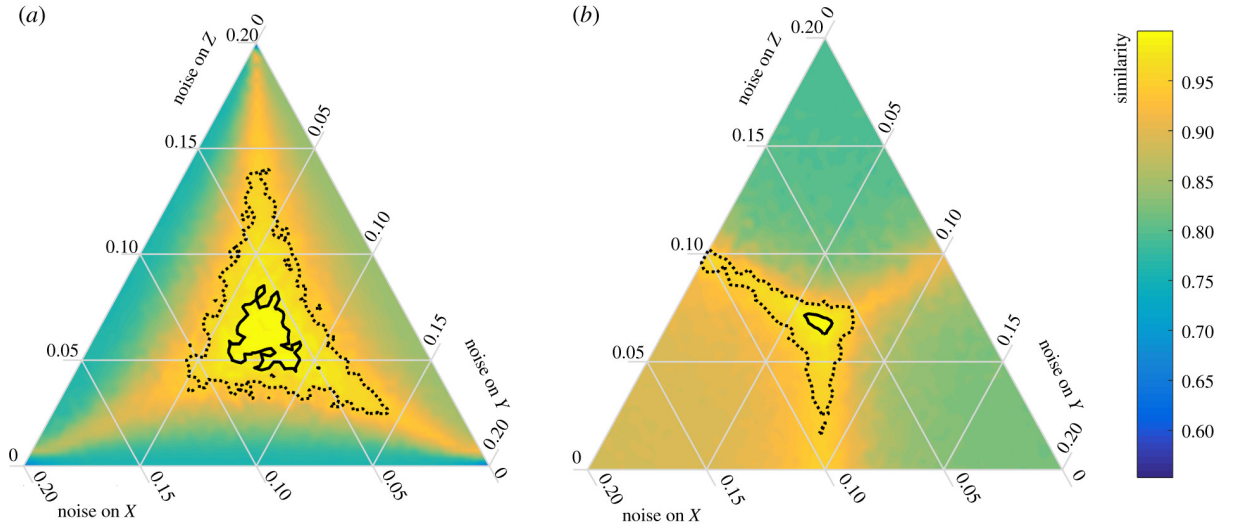


Figure 3.6: Performance of (A) MAF and (B) PCA for different noise regimes for the meta-population model with three patches. A location in the plot shows how the noise is distributed over the three variables. In the middle, every variable gets the same amount of noise. In every location, the sum of the three noise levels is 0.2. The performance is calculated by the similarity of the MAF or PCA result to the slowest direction. For noise that is the same for all variables, MAF and PCA give the same (correct) result, as indicated with the high similarity index in the middle of the panels (bright yellow colour). If one or more of the variables receive little or no noise, MAF does not perform well. MAF outperforms PCA in most other cases. The area inside the dotted lines is the area where the similarity is higher than 0.95 (reasonable performance), the area within the black lines is the area where similarity is higher than 0.99 (good performance).

Overall, under most noise regimes, MAF performs better than PCA as indicated by a larger area within the solid and the dotted black lines. Only when noise is low on one variable, and relatively high in the two other variables does PCA outperform MAF. The reason that PCA works better in that case is that the first PC will point in the direction of the two variables with noise, and this will yield a high similarity score. The same happens for MAF when there is noise on only one variable, in which case the first MAF points to the two variables without noise. The Figures show that PCA is only truly meaningful when the noise level is the same for all the involved variables. MAF is a bit more robust and, even when noise levels vary slightly, the method maintains its high accuracy. However, both methods fail to obtain the direction of slowest recovery when there is a large difference in noise levels for the different variables.

If the noise becomes larger, the results are not affected (supplementary Figure S8), assuming that the system remains in the area around its equilibrium. For shorter time series, the accuracy of both methods (MAF and PCA) is reduced (Supplementary Figure S9). The performance of MAF is more affected by data size than the performance of PCA.

3.4 Discussion & Conclusion

Our work reveals new ways in which multivariate time series may be mined to detect the direction of lowest resilience in complex systems. Since we are living in a time when more and more high-density data are becoming available (Porter et al., 2005), new methods to use these data to their full potential are a welcome expansion of the toolbox to analyse complex systems. Our method makes use of the temporal behaviour of the system on small time scales, providing information that is hard to extract from the data by more traditional statistical methods. This also means that the input data have to be sampled at a time interval that is sufficiently small. What exactly is “sufficiently small” depends on the time scale of the dynamics of the system. For instance, brain activity should be measured at much smaller time steps than tree cover. It will typically be difficult to decide a-priori what sampling frequency and time-series length are appropriate. However, a simple way to test whether or not a particular time series is suitable for the proposed analysis is to run the analyses for different time-series lengths (See Methods, figure 3.5 and supplementary figures S6-S7). If convergence is reached, and the confidence interval is small, the time series can be considered to be of sufficient length for the proposed analysis.

We showed that both high autocorrelation and high variance in a particular direction in multivariate time series can act as a pointer to the dominant slow direction of a system, provided that the system is not highly reactive and has no strong oscillating dynamics. Importantly, both methods have advantages and disadvantages, so it depends on the available data which method is expected to be most reliable. In one dimensional systems,

autocorrelation is found to be a more robust indicator of resilience than variance (Dakos et al., 2012b). Also in multiple dimensions, we show that autocorrelation outperforms variance when noise levels vary for different variables (figure 3.6). Intuitively this makes sense, since all variance based measures such as PCA, covariance and standard deviation are heavily influenced by noise levels. Still, MAF may also lose accuracy when noise only affects a subset of the variables (figure 3.6). Furthermore, if there is no noise in the slowest direction (i.e. the dominant eigenvector), resilience indicators can miss signals of slow dynamics (Boerlijst et al., 2013). Thus, in general, autocorrelation seems more robust than variance. However, the MAF analysis requires a high data resolution to capture the slow dynamics. Resolution is not an issue for PCA, which does not take the timing of the data into account. In conclusion, if the measured variables are known to be subject to different noise levels, MAF should always be preferred. If, however, data is too sparse to get a reliable estimate of a direction with high autocorrelation, PCA might be a good alternative.

There are several caveats when it comes to interpreting the results of our method. First of all, the information we obtain depends on how large the natural fluctuations are (or the noise is). We can only reliably estimate the speed of the system for the part of the state space that is visited by the system. We show that, under some conditions, the local information about slow and fast recovery may be extrapolated somewhat outside this range. However, in real systems it will typically be impossible to know whether or not this works as we lack complete insight into the properties that shape the dynamics throughout the state space. Another fundamental limitation is the assumption that the system has a stable point attractor. For systems that show oscillating, reactive or chaotic behaviour the method is not applicable, and more generally the same is true for systems that are far from equilibrium. Also, non-linear systems or reactive systems often display spiralling dynamics, even if the attractors are stable points (e.g. our gene regulatory network). For these types of systems, PCA will still find a direction of high variance and MAF will still find a direction of high autocorrelation, but these directions do not necessarily correspond to the direction of longest recovery and thus the engineering resilience of the system in that direction. Whether or not a system is expected to fall into this category can be tested based on the time series of the system with an estimation of the “worst case reactivity” of a system (Ives et al., 2003). Also, for systems that have instabilities and that could leave their equilibrium, the direction of MAF or PCA might still point to the direction in which the system will lose its stability. We have deliberately limited ourselves to detecting the mix of perturbations from which the system recovers most slowly. However, the direction of lowest resilience may in some systems also be the direction in which compound perturbations may most easily trigger a critical transition into a new state (Lever et al., 2020).

Despite these limitations, MAF and PCA offer exciting opportunities to start probing the resilience of multivariate complex systems in novel ways. Our approach builds on the

influential work on detecting instabilities based on the phenomenon of critical slowing down in the vicinity of tipping points. Clearly, the phenomena we describe are just the tip of the iceberg when it comes to probing resilience in real systems. Our results show that creative use of known computational tools allows to make the theory of resilience indicators applicable for multivariate systems. The patterns we find suggest ways to move forward to produce theoretical frameworks that help unravel resilience in the wide range of high-dimensional systems on which humanity depends.

Supplementary materials

S3.1 A note on stability and reactivity

Gradient systems are systems that have a potential. This means that if you perturb that system in one direction, it will always follow the same path. No spiralling dynamics will occur (i.e. the eigenvalues are real). Fully gradient systems are not common in the real world. However, various systems may behave as a gradient system locally around a stable equilibrium. If a system is described by $\frac{dx_i}{dt} = f_i(\vec{x})$, one can evaluate whether or not it is a locally gradient system by checking the equality:

$$\frac{\delta f_i}{\delta x_j} = \frac{\delta f_j}{\delta x_i} : i \neq j \quad (\text{S3.1})$$

This means that the Jacobian is symmetric for all possible values for \vec{x} . If the jacobian is only symmetric at a stable equilibrium, we say that the system locally has a potential. in this case, we expect simple dynamics and autocorrelation should correspond to recovery time. For these types of systems a method such as MAF is useful. Autocorrelation still corresponds to recovery time for more complex systems. If the system does not have a potential (i.e. it is not a gradient system), it can still recover from perturbations in a smooth way. However, for certain types of systems there exist directions where a perturbation will amplify before returning to the equilibrium, even though the equilibrium is stable. These systems are called reactive systems (Neubert and Caswell, 1997). Whether or not a system is reactive can be tested by comparing the dominant eigenvalue (highest eigenvalue) of the Jacobian to the dominant eigenvalue of the corresponding Hermitian. The Hermitian is calculated as $H = \frac{J+J'}{2}$, where H is the Hermitian, J is the Jacobian, and J' is the transpose of the Jacobian. If the eigenvalues for the Jacobian are all negative this indicates a stable equilibrium. However, if the Hermitian has at least one positive eigenvalue, this means that the system is reactive. For all deterministic models we calculated the Jacobian, the dominant eigenvalue of the Jacobian and the dominant eigenvalue of the corresponding Hermitian in order to evaluate whether or not the system is locally gradient and/or reactive.

The Jacobian of the threepatch model is symmetric, indicating that this system behaves like a gradient system locally (actually, this system is a completely gradient system, but we are only interested in local dynamics). Consequently, the dominant eigenvalue of the Jacobian (-0.2851) is the same as the dominant eigenvalue of the Hermitian (-0.2851) indicating no reactivity. Therefore this is a system where autocorrelation is expected to correspond to recovery time.

The Jacobian of the Sahara model is not symmetric and therefore this is not a gradient system. The dominant eigenvalue of the Jacobian (-0.5223) is different from the dominant

eigenvalue of the Hermitian (-0.4621), but since both eigenvalues are negative this should not lead to reactive behaviour and autocorrelation is expected to correspond to recovery time.

The Jacobian of the genetic network is not symmetric and therefore this is not a gradient system. The dominant eigenvalue of the Jacobian (-10.50) differs considerably from the eigenvalue of the Hermitian (10.99). Furthermore, even though the dominant eigenvalue of the Jacobian is negative, corresponding to a stable equilibrium, the dominant eigenvalue of the Hermitian is positive, which means that it is possible that some perturbations grow initially even though the equilibrium is stable. In these directions, perturbations will take a long time to recover even though the autocorrelation in this direction is not necessarily high. Consequently, autocorrelation in these systems is not a good indicator for recovery time.

S3.2 supplementary figures




Figure S3.1: Perturbations on the four MAFs for the Sahara model. 10%, 50% and 90% recovery are indicated with lines, dashed lines and dotted lines respectively. 10% and 50% are well ordered, with no or hardly any difference between the second, third and fourth MAF. 90% Recovery is not well-ordered since the perturbation on MAF 4 results in a higher recovery time. This result is summarized in figure 3.3.




Figure S3.2: Perturbations on the five MAFs for the genetic network. 10%, 50% and 90% recovery are indicated with lines, dashed lines and dotted lines respectively. 10% recovery is well ordered, 50% and 90% recovery are not well-ordered. The figure clearly shows that the perturbation on MAF 3 ends up on MAF 1 (high deviation for blue and red in the same direction and a small deviation for yellow in the opposite direction). Also, the perturbations on MAF4 and MAF5 clearly show spiralling dynamics, therefore only the initial recovery rate (10% recovery) is well ordered. This result is summarized in figure 3.3.




Figure S3.3: frequency distribution of the recovery times of 1000 perturbations in random directions for 10% (up), 50% (middle) and 90% (bottom) recovery for the meta-population model. Dotted lines indicate the recovery time of a perturbation on the last MAF, solid lines indicate a perturbation on the first MAF. It shows that for the meta-population model MAF finds the direction of slow and of fast recovery.




Figure S3.4: frequency distribution of the recovery times of 1000 perturbations in random directions for 10% (up), 50% (middle) and 90% (bottom) recovery for the Sahara model. Dotted lines indicate the recovery time of a perturbation on the last MAF, solid lines indicate a perturbation on the first MAF. It shows that for the Sahara model MAF finds the direction of slow and of fast recovery.




Figure S3.5: frequency distribution of the recovery times of 1000 perturbations in random directions for 10% (up), 50% (middle) and 90% (bottom) recovery for the genetic network. Dotted lines indicate the recovery time of a perturbation on the last MAF, solid lines indicate a perturbation on the first MAF. It shows that for the genetic network, MAF cannot find the direction of the longest and shortest recovery.

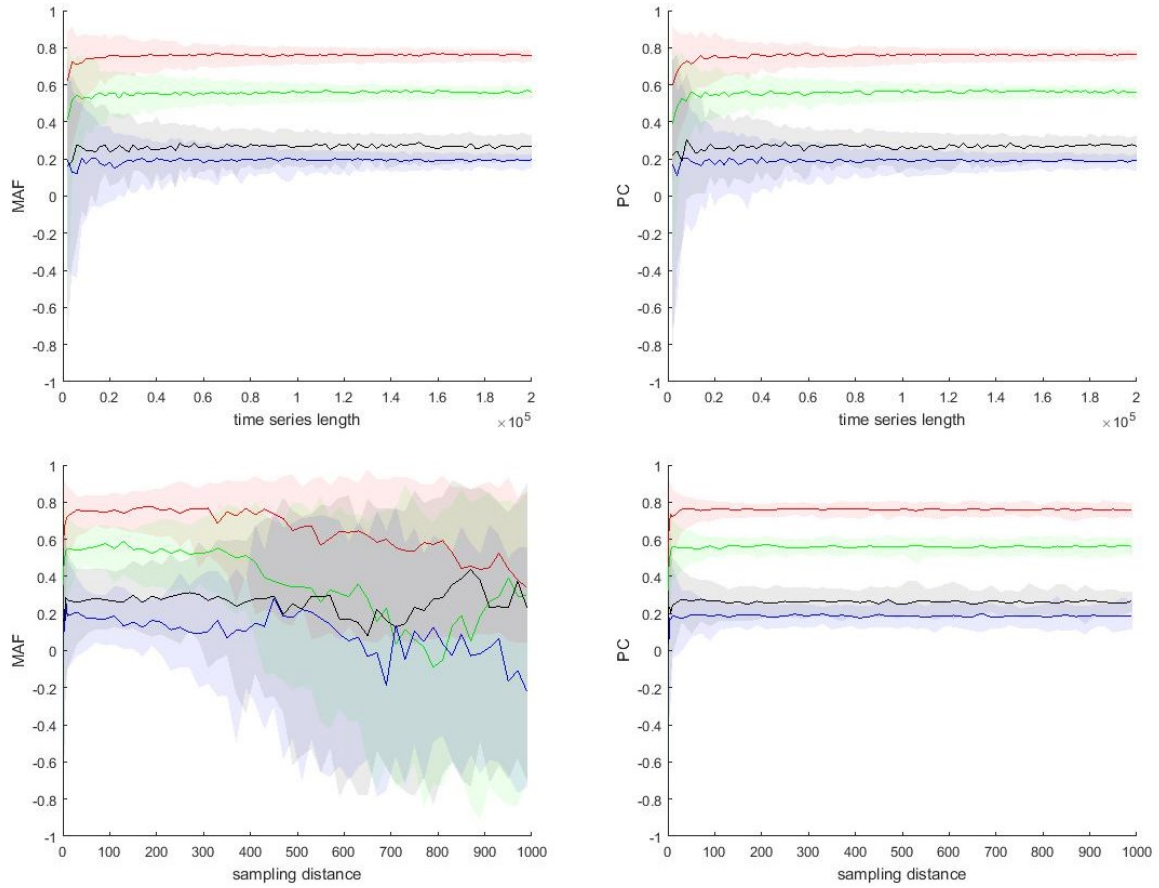


Figure S3.6: Top row: Effect of time series length on MAF (top left) and PCA (top right) for the Sahara model. Solid lines indicate the medians of 100 blocks in a block bootstrap, shaded areas show the 90% confidence interval. The figures show that both methods improve for increasing data length, but converge at around 100,000 datapoints. Bottom row: Effect of data resolution on MAF (bottom left) and PCA (bottom right) showing that if the distance between datapoints becomes too large, MAF results do not converge anymore, whereas this distance does not affect PCA.




Figure S3.7: Top row: Effect of time series length on MAF (top left) and PCA (top right) for the genetic network. Solid lines indicate the medians of 100 blocks in a block bootstrap, shaded areas show the 90% confidence interval. The figures show that both methods improve for increasing data length, but converge at around 20,000 datapoints. Furthermore, note that when converged, MAF and PCA yield slightly different directions. Bottom row: Effect of data resolution on MAF (bottom left) and PCA (bottom right) showing that if the distance between datapoints becomes too large, MAF results do not converge anymore, whereas this distance does not affect PCA.




Figure S3.8: Performance of (a) MAF and (b) PCA for different noise regimes for the meta-population model with three patches. A location in the plot shows how the noise is distributed over the three variables. In the middle, every variable gets the same amount of noise. In every location, the sum of the three noise levels is 0.4, which is twice as high as figure 3.5. The performance is calculated by the similarity of the MAF or PCA result to the slowest direction. For noise that is the same for all variables, MAF and PCA give the same (correct) result, as indicated with the high similarity index in the middle of the panels (bright yellow colour). If one or more of the variables receive little or no noise, MAF does not perform well. MAF outperforms PCA in most other cases. The area inside the dotted lines is the area where the similarity is higher than 0.95 (reasonable performance), the area within the black lines is the area where similarity is higher than 0.99 (good performance). Comparison with figure 5 indicates that for this model the results remain unchanged for higher noise levels, provided that the noise does not force the system into a new regime. This may be due to the smooth stability landscape of this gradient system.




Figure S3.9: Same as figure S3.8 but with only 300 data points, as opposed to 4000 as in figure 3.5 and S3.8. If the noise levels are equal over the three variables, PCA results are more reliable than MAF results, since there is a clear black circle (0.99 similarity to true direction of slowest recovery) around the middle of panel B. However, when noise levels vary over variables, both methods become inaccurate, but MAF has a larger region where the performance is at least “reasonable” (0.95 similarity to true direction of slowest recovery, area within dotted lines).

Chapter 4

Foreseeing the future of mutualistic communities beyond collapse

J. Jelle Lever

Ingrid A. van de Leemput

Els Weinans

Rick Quax

Vasilis Dakos

Egbert H. van Nes

Jordi Bascompte

Marten Scheffer

This chapter is based on:

J Jelle Lever, Ingrid A van de Leemput, Els Weinans, Rick Quax, Vasilis Dakos, Egbert H van Nes, Jordi Bascompte, and Marten Scheffer. Foreseeing the future of mutualistic communities beyond collapse. *Ecology letters*, 23(1):2–15, 2020

Abstract

Changing conditions may lead to sudden shifts in the state of ecosystems when critical thresholds are passed. Some well-studied drivers of such transitions lead to predictable outcomes such as a turbid lake or a degraded landscape. Many ecosystems are, however, complex systems of many interacting species. While detecting upcoming transitions in such systems is challenging, predicting what comes after a critical transition is *terra incognita* altogether. The problem is that complex ecosystems may shift to many different, alternative states. Whether an impending transition has minor, positive or catastrophic effects is thus unclear. Some systems may, however, behave more predictably than others. The dynamics of mutualistic communities can be expected to be relatively simple, because delayed negative feedbacks leading to oscillatory or other complex dynamics are weak. Here, we address the question of whether this relative simplicity allows us to foresee a community's future state. As a case study, we use a model of a bipartite mutualistic network and show that a network's post-transition state is indicated by the way in which a system recovers from minor disturbances. Similar results obtained with a unipartite model of facilitation suggest that our results are of relevance to a wide range of mutualistic systems.

4.1 Introduction

Empirical studies of lakes, arid ecosystems, coral reefs, and tropical forests suggest that remarkably sudden transitions to alternative stable states may occur when changing environmental conditions pass a critical value (Scheffer et al., 1993; Rietkerk and Van de Koppel, 1997; Scheffer et al., 2001; Hirota et al., 2011). While the outcome of such transitions is relatively predictable when a few leading species or species groups determine the state of an ecosystem, this may not be the case when ecosystem dynamics are determined by many interacting species. Species traits as well as their sensitivity to changing conditions are known to be highly diverse, and many drivers of environmental change are known to have multiple simultaneous effects on species communities. A change in climate may, for example, affect the distribution, phenology, physiology, behavior, and relative abundances of species, and these changes may, in turn, affect the strengths of interactions between species (Kareiva et al., 1993; Memmott et al., 2007; Suttle et al., 2007; Tylianakis et al., 2008; Burkle et al., 2013; Høye et al., 2013; Usinowicz and Levine, 2018). The specific ways in which interactions are arranged in complex ecological networks are known to be crucial for the stability of ecosystems (Kareiva et al., 1993; De Ruiter et al., 1995; McCann, 2000; Solé and Montoya, 2001; Neutel et al., 2002; Montoya et al., 2006; Bastolla et al., 2009; Rohr et al., 2014). Gradual changes in these patterns and other complex simultaneous effects of changing environmental conditions may therefore lead to regime shifts of which the outcomes are highly unpredictable (Scheffer et al., 2012).

The response of ecosystems to a change in environmental conditions is determined by the relative strengths of positive and negative feedback loops in the networks of interactions between species or between species and their environment. Immediate negative feedbacks, e.g. due to intraspecific competition, have stabilizing effects, while positive or ‘reinforcing’ feedbacks are destabilizing and a necessary condition for the existence of alternative stable states (Thomas, 1981; Snoussi, 1998; Gouzé, 1998). Critical transitions towards such states may occur when changing conditions alter a system’s feedbacks such that destabilizing, positive feedbacks gain in strength relative to stabilizing, immediate negative feedbacks. A classic example is found in shallow lakes where an increase in algae leads to an increased turbidity and the suppression of aquatic plants. As a consequence, more nutrients become available to algae which enhances algae growth. A clear-water, plant-dominated state may therefore switch to a turbid, algae-dominated state when gradually increasing nutrient levels pass a critical value. Recovery from such transitions requires a relatively large reduction in nutrient availability, a phenomenon called ‘hysteresis’ (Scheffer et al., 1993)). Other examples of such switching behavior are found in coral reefs, woodlands, deserts, and oceans (May, 1977; Wilson and Agnew, 1992; Scheffer et al., 2001), as well as in many other systems such as the climate (Hare and Mantua, 2000; Scheffer et al., 2001; Alley et al., 2003; Lenton et al., 2008), the economy (Diamond and Dybvig, 1983; Arthur, 1989; Easley and Kleinberg, 2010), and human cells (Hasty et al.,

2002; Ferrell Jr, 2002; Lee et al., 2002; Tyson et al., 2003; Angeli et al., 2004).

Mutually beneficial interactions are, perhaps, the most intuitive examples of positive feedback loops in complex ecological networks, metapopulations, or other complex environmental systems. Previous studies have emphasized the importance of such interactions in communities of flowering plants and animal pollinators or seed dispersers (Jordano, 1987; Bascompte et al., 2003). Mutually beneficial interactions between zooxanthellae, coral species and invertebrates occur in coral reefs where a diversity of coral species provide food, shelter and reproduction sites for other organisms (Moberg and Folke, 1999; Wilson et al., 2006; Stella et al., 2011). Nutrient exchange with mycorrhizal fungi and nitrogen-fixing bacteria is fundamental for plant communities (Kiers et al., 2011), and mutualistic interactions are of importance for microbial communities where multiple species are involved in the degradation of organic substrates (Schink, 2002; Stolyar et al., 2007). Indirect facilitation may occur between plant species when modifying harsh environments (Wilson and Agnew, 1992; Callaway, 1995; Holmgren et al., 1997; Rietkerk et al., 2004), and the exchange of individuals between habitat patches may be fundamental for metapopulations (Hanski, 1998). Previous work suggested that critical transitions may occur due to the positive feedback resulting from such mutually beneficial relationships in plant-pollinator communities because a decline in pollinator abundances may negatively affect plant abundances, which in turn is bad for pollinators. Similar transitions may occur in metapopulations due to a ‘rescue effect’ (Hanski, 1998) and in facilitative communities due to an ‘Allee effect’ (Rietkerk et al., 2004; Courchamp et al., 1999; Stephens et al., 1999)). The observation that the relative strength of facilitative interactions tends to increase with environmental stress (Bertness and Callaway, 1994; Maestre et al., 2009; Tur et al., 2016), suggests that competitive communities may become increasingly mutualistic as conditions change. The aforementioned positive feedbacks and associated critical transitions may thus also occur in communities where mutually beneficial interactions were not particularly strong under more advantageous conditions.

Here, we propose a new class of indicators that may allow us to detect the specific way in which species are affected by an increase in the relative strength of a positive feedback prior to a critical transition. The essence of our approach is that we seek the direction in a system’s phase space, i.e. a multidimensional space in which each axis corresponds to the abundance of a species, in which a system becomes increasingly sensitive to small subcritical disturbances. Earlier studies have shown that an increasingly slow recovery from small disturbances may be indicative of a loss of resilience prior to critical transitions (Wissel, 1984; Nes and Scheffer, 2007). Various indicators of this phenomenon known as ‘critical slowing down’ may therefore serve to detect an increase in the likelihood of critical transitions (Scheffer et al., 2009; Dakos et al., 2012a). Here, we take advantage of the fact that resilience is not lost equally in all directions. Disturbances have a size (i.e. the total amount of change) and a direction (i.e. the relative amount of change in each species). The more similar a disturbance’s direction to the direction in which increasingly small

perturbations may cause critical transitions, the stronger the effect of critical slowing down. Provided that there are no oscillatory, chaotic or other complex dynamics, a system's future state will most likely lie in the same approximate direction.

To get an intuitive understanding of the principle behind our approach, consider a small plant-pollinator community of which the dynamics can be represented by a landscape of valleys, hills and ridges (Fig. 4.1.A and Appendix S4.2 in Supporting Information). In this landscape, every possible combination of pollinator abundances is represented by a unique point, while the speed and direction in which abundances change corresponds roughly to the slope of the landscape. The lowest points of the landscape's valleys or 'attraction basins' represent alternative stable states. As conditions change, the shape of the landscape changes and new basins appear. When a threshold comes close to the network's initial state, a small perturbation in the right direction can invoke a transition into another attraction basin. Eventually, the basin around the network's initial state disappears altogether and the system inevitably shifts into one of the alternative basins. The question we ask is whether we may know beforehand to which of the alternative attractors a system will most likely shift. The clue is that the slope of the initial state's attraction basin changes in a characteristic way before the transition occurs. A 'mountain pass' towards the system's future state is formed, marked by a 'saddle point' in the landscape. The initial state's attraction basin becomes increasingly shallow in the direction of this pass and the recovery from perturbations increasingly slow (Fig. 4.1.B-C and Fig. S4.2). This direction is what we refer to as the 'direction of critical slowing down' and is indicative of the relative gain or loss in abundance of each species after an impending critical transition.

To explore whether the direction of critical slowing down might be indicative of the future state of mutualistic communities, we use a model of a bipartite mutualistic network in which critical transitions are known to occur (Dakos and Bascompte, 2014; Jiang et al., 2018). This model was originally developed to describe the interactions between flowering plants and animal pollinators or seed dispersers (Bastolla et al., 2009), but may describe any system characterized by competition within and cooperation between species groups. Previous work has shown that indirect facilitation occurs between pollinators when they interact with the same plant species (Moeller, 2004; Ghazoul, 2006; Bastolla et al., 2009). This indirect facilitation makes a network more resilient, i.e. the minimum size of perturbations or the amount of change in environmental conditions needed to cause a critical transition is larger. When pollinators continue to facilitate each other under increasingly harsh environmental conditions they may, however, also collapse simultaneously because they depend on each other for survival.

We generate time series in which the resilience of a network's initial state is gradually undermined by altering the relative strength of mutualistic interactions. Oscillatory, or other complex dynamics occurring after a threshold is passed may negatively affect the

performance of the here proposed class of indicators but are unlikely in purely mutualistic systems, i.e. systems in which all interspecific interactions are positive, because they require at least one delayed negative feedback, i.e. a negative feedback with a time lag, usually occurring as the result of an uneven number of negative interactions in feedback loops of two or more species (Levins, 1974; Thomas, 1981; Puccia and Levins, 1985; Goldbeter, 1996; Hastings and Powell, 1991; Snoussi, 1998; Gouzé, 1998; McCann et al., 1998; Dambacher et al., 2003). Few real ecosystems can, however, be expected to be purely mutualistic. Different scenarios are therefore explored, varying from a scenario where positive feedbacks are the only cause of instability, i.e. in purely mutualistic systems, to scenarios in which the destabilizing effects of delayed negative feedbacks are stronger, i.e. in mixed systems with mutualistic and competitive interactions. To determine the direction of critical slowing down, we study changes in the fluctuations around the species mean abundances and determine whether they can be used to predict a network's post-transition state. To explore whether the results obtained with this model may hold for a wider class of mutualistic systems, we investigate whether similar results are obtained with a more general, unipartite model of competition and facilitation between species.

4.2 Community model

We use a dynamic model describing the interactions between two types of species: plants (P) and pollinators (A). As in Bastolla et al. (2009), species of the same type compete with each other, while species belonging to a different type interact mutualistically. The dynamics of species i belonging to a group of $S^{(A)}$ pollinator species are as follows:

$$\frac{dN_i^{(A)}}{dt} = \frac{R_i(N^{(P)})}{1 + h_i R_i(N^{(P)})} N_i^{(A)} - \sum_{j=1}^{S^{(A)}} c_{ij} N_j^{(A)} N_i^{(A)} - d_i N_i^{(A)} + \epsilon_i. \quad (4.1)$$

Plant dynamics are described by a similar formula, which can be found by exchanging indices A and P . Unless stated otherwise, this procedure can be applied to all formulas in this chapter.

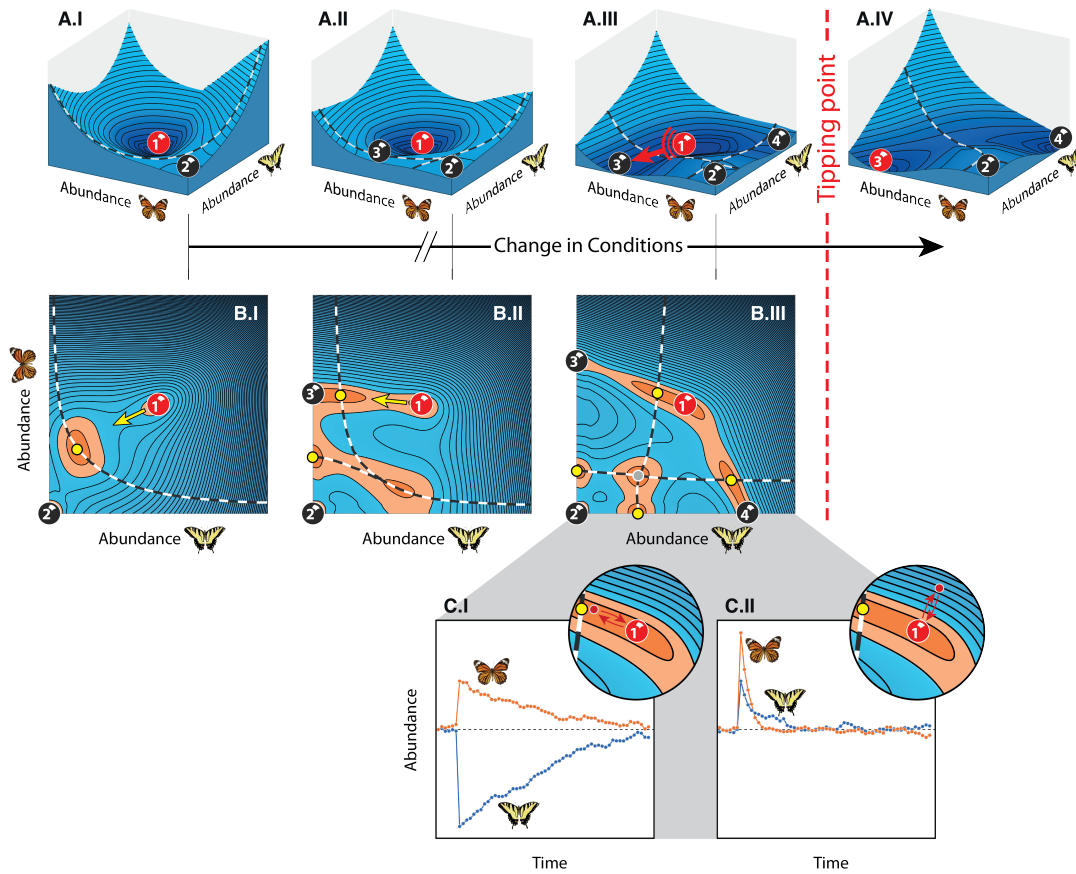


Figure 4.1: Stability properties for a small network of two pollinators (shown) and two plants (not shown). **(A)** Attraction basins (valleys) of alternative stable states (balls) are separated by thresholds (dashed curves). Initially, the only alternative to pristine state 1 is fully collapsed state 2 (A.I). When conditions change, two additional, partially collapsed states appear (states 3 and 4). The initial, pristine state loses resilience after state 3 appears (A.II-A.III). Eventually, the threshold towards state 3 approaches the pristine state so closely that a critical transition towards this state becomes inevitable (A.III-A.IV). **(B)** Alternative stable states, saddle points (yellow dots), and hilltops (grey dots) are surrounded by areas in which the landscape's slope, and thus the rate at which abundances change, is nearly zero (indicated in orange). Higher speeds are found further away from these points. The direction of slowest recovery changes substantially before future state 3 appears (yellow arrow, B.I-B.II). After state 3 appears, the system slows down in the direction of the saddle point on the approaching threshold (B.II-B.III). **(C)** Slow recovery from a perturbation towards the saddle point (C.I) as opposed to the much faster recovery from an equally large perturbation in another direction (C.II).

Species i has abundance N_i , which may increase due to mutualistic interactions with members of the other species type. The rate at which the abundance of species i increases depends on the total amount of resources provided by mutualistic partners, $R_i(N^{(P)})$, i.e. nectar for pollinators and pollen for plants. As in Okuyama and Holland (2008) and Bastolla et al. (2009), we assume that species are limited in their capacity to process resources and become saturated when the amount of resources provided is high. The rate at which species become saturated is determined by saturation term h_i . The total mutualistic benefit, $R_i(N^{(P)})$, depends on the abundance of mutualistic partners as follows:

$$R_i(N^{(P)}) = \sum_{k=1}^{S^{(P)}} \gamma_{ik} N_k^{(P)}, \quad (4.2)$$

in which γ_{ik} is the mutualistic interaction strength, i.e. the rate at which resources become available to species i , due to its interaction with species k .

Species of the same type compete directly amongst each other, e.g. plants for soil nutrients and pollinators for nesting sites. Intraspecific competition, c_{ii} , is assumed to be substantially stronger than interspecific competition, c_{ij} , such that species do not easily outcompete each other. Independent of mutualistic and competitive interactions, several processes may simultaneously enhance or reduce population growth. We assume that the combined effect of these processes is negative, which is incorporated by mortality rate d_i .

Species experience small stochastic perturbations incorporated through noise term ϵ_i :

$$\epsilon_i = \delta_i \frac{dW}{dt}. \quad (4.3)$$

ϵ_i fluctuates in time due to a Wiener process, W , with mean zero and standard deviation δ_i . The Wiener process is a continuous-time stochastic process generating white noise. To prevent noise leading to negative abundances, we assume that $dN/dt = 0$ when $N < 0.001$.

Coexistence and relative mutualistic benefits

As the number of species and/or the strength of interspecific competition increases, it becomes increasingly difficult to assign parameters such that all species may stably coexist. In previous work, a trade-off was assumed between the number and the strength of mutualistic interactions which prevented species with many interactions from becoming

overly abundant and outcompeting other species (Bastolla et al., 2009; Dakos and Bascompte, 2014; Jiang et al., 2018). Here, we assume mutualistic interaction strengths to vary continuously, i.e. pollinators may interact with all plant species and vice versa, which allows us to explore gradual changes in interaction structure beyond the fixed structure of a predefined mutualistic network. A different kind of balancing relationship is therefore required, and mutualistic interaction strengths, γ_{ik} , are determined as follows:

$$\gamma_{ik} = \frac{\theta_{ik} R_i(\hat{N}^{(P)})}{\hat{N}_k^{(P)}}, \quad (4.4)$$

in which the relative mutualistic benefit, θ_{ik} , corresponds to the fraction of the total amount of resources provided by species k , and $R_i(\hat{N}^{(P)})$ to the total amount of resources received by species i at the system's nontrivial equilibrium, i.e. the equilibrium point at which all species have a non-zero abundance. There are different costs and benefits associated to different feeding strategies, e.g. being a specialist or a generalist or interacting with specialists or generalists (Morales and Traveset, 2008; Tur et al., 2016). This way of assigning mutualistic interaction strengths makes sure that a species' total amount of resources received is independent from a species' relative feeding preferences, i.e. we assume the sum of these costs and benefits to be approximately the same for each strategy. The sum of a species' relative mutualistic benefits, $\sum_{k=1}^{S^{(P)}} \theta_{ik}$, is one. A change in relative mutualistic benefits does not affect the equilibrium abundances of species, because the total amount of resources provided to each species remains the same (see Appendix S4.5).

Changing environmental conditions and the direction in which resilience is lost

To test whether the direction of critical slowing down is indicative of a system's future state, we study our ability to predict a system's future state when changing conditions lead to substantial changes in the strength of positive feedbacks and the direction in which they have destabilizing effects. Such changes may occur when changing conditions fundamentally alter the ways in which species relate to each other.

Positive feedbacks and the direction in which resilience is lost can be studied when determining the elements of the Jacobian matrix at a system's nontrivial equilibrium. Each element in this matrix describes how a change in the abundance of species i affects the growth of species j , dN_j/dt . At a tipping point, the dominant eigenvalue of the Jacobian matrix is zero and the slope of the direction in which a system recovers slowest from perturbations is indicated by the eigenvector corresponding to this eigenvalue. The strength of the positive feedback between pollinator i and plant j can be determined by multiplying the Jacobian's off-diagonal elements; $\alpha_{ij} * \alpha_{ji}$. In a two-species system, a

tipping point is reached when the strength of this feedback is equal to the multiplication of the two direct negative feedbacks; $\alpha_{ii} * \alpha_{jj}$. Similar relationships can be obtained when studying larger systems (Levins 1974; Thomas 1981; Puccia and Levins 1985; Goldbeter 1996; Snoussi 1998; Gouzé 1998; Dambacher et al. 2003; De Ruiter et al. 1995; Neutel et al. 2002; Neutel and Thorne 2014).

Some species contribute more to the instability caused by positive feedbacks than others. The effect of a temporary change in the abundance of mutualistic partners, as described by the Jacobian matrix, for example, is small when species are highly saturated, i.e. $R_i(\hat{N}^{(P)})$ and/or h_i is large. Positive feedbacks are therefore weak and the resilience of the here studied networks is high when relative mutualistic benefits, θ_{ik} , are distributed such that most resources are obtained from the same, highly saturated species (see Appendix S4.1 and Fig. S4.1). In more complex communities such a distribution resembles a nested structure as is commonly observed in mutualistic networks, as in those networks species tend to obtain resources from the same mutualistic partners as well (Bascompte et al. 2003 and Fig. S4.6.A). The interrelationships between saturated and non-saturated species are asymmetrical as in Bascompte et al. (2006).

As a starting point for further research, we explore a scenario in which a change in the aforementioned distribution of relative mutualistic benefits, θ_{ik} , undermines the resilience of the mutualistic networks while keeping all other properties, e.g. nontrivial equilibrium abundances and the negative effects of inter- and intraspecific competition, constant (see Appendix S4.5). Increasingly strong positive feedbacks emerge when two or more non-saturated species start to interact increasingly strongly with each other. Eventually, this will lead to a full or partial network collapse depending on the specific way in which relative mutualistic benefits are changed. Conditions, M , affect relative mutualistic benefits as follows:

$$\theta_{ik}^* = \theta_{0,ik} + (\theta_{final,ik} - \theta_{0,ik})M, \quad (4.5)$$

in which $\theta_{0,ik}$ is the initial, $\theta_{final,ik}$ the final, and θ_{ik}^* the actual relative mutualistic benefit. Conditions, M , change from zero to one over time, t , such that $dM/dt = 1/T$, in which T is the total simulation time. Mutualistic interaction strengths, γ_{ik} , are updated as described in equation 4.4. The species and interactions involved in the positive feedback leading to a critical transition, the direction in which this feedback amplifies change, and the nature of a system's future state, are determined by the specific way in which interactions are altered.

In addition to the scenario in which only the relative mutualistic benefits change, we explore scenarios in which the nontrivial equilibrium abundances of species change as well due to a change in the total amount of resources received from mutualistic partners (see Appendix S4.5).

Determining the direction of critical slowing down

Although measuring the recovery rate from experimental perturbations is the most direct way to determine the direction of critical slowing down, an experimental approach may be impractical or even impossible when studying complex networks. The development of alternative methods to determine the direction of critical slowing down is therefore of importance. Previous studies suggested that small changes in the statistical properties of time series, e.g. an increase in variance, autocorrelation, skewness, and spatial correlation, may be used as an indicator of a change in the proximity to a tipping point (Scheffer et al. 2009; Dakos et al. 2012a). Here, we explore whether changes in the statistical properties of time series may be used to predict the future state of mutualistic communities.

When assuming a continuous regime of random perturbations, a system will spend most time away from its equilibrium state in the direction in which it recovers slowest from perturbations (see Appendix S4.2). When approaching a tipping point, the distribution of natural fluctuations around the species' mean abundances should therefore become increasingly elongated in the direction in which a system slows down (Fig. S4.3). To detect such change, we analyze our model-generated times series by determining the direction and magnitude of such asymmetry in a rolling window. This window has a fixed size and is moved along the time series as new data become available. To determine the direction in which abundances are distributed asymmetrically, we use a principal component analysis of which the first principal component corresponds to the line in the network's phase space along which variance is highest (see Held and Kleinen 2004; Chen et al. 2012; Suweis and D'Odorico 2014; Dakos 2018 and Chen et al. 2019 for related approaches). Abundances are distributed asymmetrically either in an up- or downward direction along this component. To determine the direction of our indicator, we orthogonally project the time series on the first principal component and determine the direction in which the projected time points are skewed (Fig. S4.4.A-E). The magnitude of the indicator is determined by the fraction of the total variance explained by the first principal component. This direction and magnitude together form a vector which is our indicator of a network's future state (Fig. S4.4.F).

A network's phase space has as many axes as there are nodes in a network. Our indicator thus has multiple components; one for each species (Fig. S4.4.F). Each component, or 'score on the indicator', gives an indication of the extent and direction in which the abundance of each individual species is distributed asymmetrically. The indicator accurately points towards the future state when its components, or 'scores', are directly proportional to the difference in abundance between a network's initial and future state. Species with a negative score are expected to decrease, while species with a positive score are expected to increase. Species with a relatively large score are expected to change more in abundance than species with a comparably smaller score. An increase in the indicator's magnitude is reflected by more extreme (positive or negative) scores.

To assess the quality of the prediction, we determine the angle between the indicator's slope, as determined by the first principal component, and the direction of the observed shift in abundance. As a measure of similarity, we take one minus the probability that the angle between two unrelated, random vectors is smaller (see Appendix S4.3). We consider the indicator's slope to be accurate when this measure of similarity is above 0.99. When time points are also skewed towards a network's future state, we consider the prediction to be fully accurate.

Simulations and parameter settings

We analyze several data sets consisting of 1000 model-generated time series in which the above described mutualistic networks approach a tipping point. For each time series, we compute the change in direction and magnitude of the indicator on the pollinator abundances (see Appendix S4.4). The distribution from which interspecific competitive interaction strengths are sampled, the number of plant and pollinator species, and the way in which changing conditions affect nontrivial equilibrium abundances differ among data sets (see Appendix S4.5). The resilience of mutualistic networks is, in all cases, undermined by a change in the distribution of mutualistic benefits leading to a substantial increase in the relative strength of positive feedbacks or delayed negative feedbacks. Declining abundances may have an additional negative effect on resilience.

To explore the effects of oscillatory, chaotic or other complex dynamics, we analyze data sets of which the strength and variability in interspecific competitive interaction strengths, c_{ij} , varies. Delayed negative feedbacks become stronger as the strength and variability of interspecific competition increases. To provide a clue as to how (un)likely it is to find transitions to oscillatory, chaotic or other complex dynamics, we determine for each time series whether the system approaches a Hopf or a saddle-node bifurcation.

Networks were discarded from a data set when they were unstable at initial conditions, $M = 0$. We determined the frequency at which this occurred as a measure of how difficult it is to find a stable solution. The final distribution of relative mutualistic benefits, $\theta_{final,ik}$, was redrawn either when a network would become unstable within the range of conditions $M = (0, 0.5)$, or when a network would still be stable at $M = 1$.

A more general, unipartite model of competition and facilitation

To explore whether the indicator may work for a wider class of systems, we investigate whether similar results are obtained with a more general model of competition and facilitation. The positive feedback between plants and pollinators in the previously described communities can be seen as an Allee effect, i.e. a positive relationship between the growth and density of populations (Courchamp et al. 1999; Stephens et al. 1999). The indirect facilitation occurring between pollinators when interacting with the same plant species

is not fundamentally different from the facilitation occurring between plant species when ameliorating the same harsh environment, or other forms of interspecific facilitation occurring in ecosystems. The most essential properties of a group of pollinator species may therefore be captured as follows:

$$\frac{dN_i}{dt} = r_i N_i \left(\frac{\sum_{j=1}^S \gamma_{ij} N_j}{A_i} - 1 \right) \left(1 - \frac{\sum_{j=1}^S c_{ij} N_j}{K_i} \right) - d_i N_i + \epsilon_i, \quad (4.6)$$

in which N_i is the abundance of species i . When the abundances of other species and mortality rates, d_i , are zero, species may grow in abundance until they reach carrying capacity K_i , or collapse to extinction when abundances are below critical abundance A_i . The speed at which species abundances change is determined by growth rate r_i . Facilitation is mediated by facilitation rate γ_{ij} . Strong interspecific facilitation allows species to recover from large disturbances, i.e. below critical abundance A_i . Species with a high critical abundance A_i depend strongly on this facilitation, and a community's overall resilience is highest when such species are facilitated relatively strongly by species with a low A_i . The relative strength of interspecific competition is determined by c_{ij} . Other causes of abundance loss are incorporated through mortality rate d_i . Species are assumed to experience small stochastic perturbations, as in the bipartite mutualistic network model, through noise term ϵ_i .

The main difference between the here presented model and the previously described plant-pollinator model is that it is a unipartite model, i.e. it describes one set of interacting species. The means by which facilitation occurs are, in contrast to the above described plant-pollinator model, not explicitly described. Parameter settings and results can be found in Appendix S4.6 and S4.7.

4.3 Results

We found that, when interspecific competitive interaction strengths are weak, instability nearly always arises from the positive feedback between plants and pollinators or from the Allee effect in the above described mutualistic or facilitative communities. Instability is caused by a saddle point approaching the communities' initial state and at least one species will collapse to extinction when a tipping point is passed. Other species may either gain or lose in abundance depending on the communities' initial properties and the way in which they are affected by changing environmental conditions (Fig. 4.2.A). Critical transitions were nearly always preceded by a period in which the indicator's magnitude would increase significantly, indicating that the distribution of fluctuating species abundances becomes increasingly asymmetric (see Appendix S4.7, Fig. 4.2.B-D and Fig. S4.7-S4.9). As with the small mutualistic network in Fig. 4.1, the indicated direction typically shifts towards a system's future state at the beginning of this period. The indicator thus consistently

pointed towards a community's future state while increasing in magnitude prior to a critical transition, when interspecific competitive interactions were weak.

A notable exception to this general pattern occurred when competitive interaction strengths were taken from a low to intermediate range, e.g. $\sim U(0.02, 0.08)$. We found that, for such a range, full network collapses were not always indicated accurately. Transitions would lead either to the collapse of relatively few species or to a collapse of the entire network (Fig. S4.9). Both the inaccurate prediction of full network collapses and the absence of intermediate-size, partial network collapses may occur because critical transitions lead to a series of cascading, partial network collapses. The likelihood of an additional collapse increases as more species collapse (Solé and Montoya 2001; Memmott et al. 2004; Rezende et al. 2007). The most likely outcome of a series of cascading, partial network collapses is therefore a collapse of the entire network. In such a scenario, the indicator will accurately indicate the initial regime shift but will not foresee the cascade of partial network collapses immediately following it (Fig. 4.3). In some time series, we observed that regime shifts consisted of several consecutive collapses (Fig. S4.10.A-B). The amount of time in between two consecutive collapses can, however, be extremely small. Also when cascades were not clearly visible, we suspect therefore that the inaccurate prediction of a full network collapse is caused by a cascading collapse.

Cascading, full network collapses were uncommon when interspecific competitive interaction strengths were drawn from other ranges (Fig. S4.9). When there is no competition between species, full network collapses are very common, well indicated and do not show signs of being caused by a cascade of partial network collapses (as in Fig. S4.10.C). When competitive interactions are strong, few species tend to collapse to extinction, while most or all other species gain in abundance from a transition. Apart from the specific range from which competitive interaction strengths were drawn, cascading collapses were found to become increasingly common when the noise level increases suggesting that they, in part, result from a low resilience of a system's future state (Fig. S4.11-S4.12). A relatively large number of species was usually indicated to lose in abundance when a, likely, cascading collapse occurred (e.g. 7 out of 10 on average, Fig. 4.3.E). As an alternative indicator of the likelihood of a cascading, full network collapse we propose therefore to use the number of species indicated to lose in abundance.

As the strength and variability of interspecific competition increases, Hopf bifurcations, leading to oscillatory, chaotic or other complex dynamics, become increasingly common. After such transitions, the system remains highly sensitive to small-scale stochastic perturbations and may end up in any of several potential future states (Fig. 4.4.A-B, and Fig. S4.13-S4.15). To which of these states a system will shift is determined by chance and thus hard to predict. For the highest competition level we tested, we found that such hard-to-predict regime shifts made up about 60% of a data set. Higher levels were not tested because, as the strength of competition increases, it becomes increasingly difficult

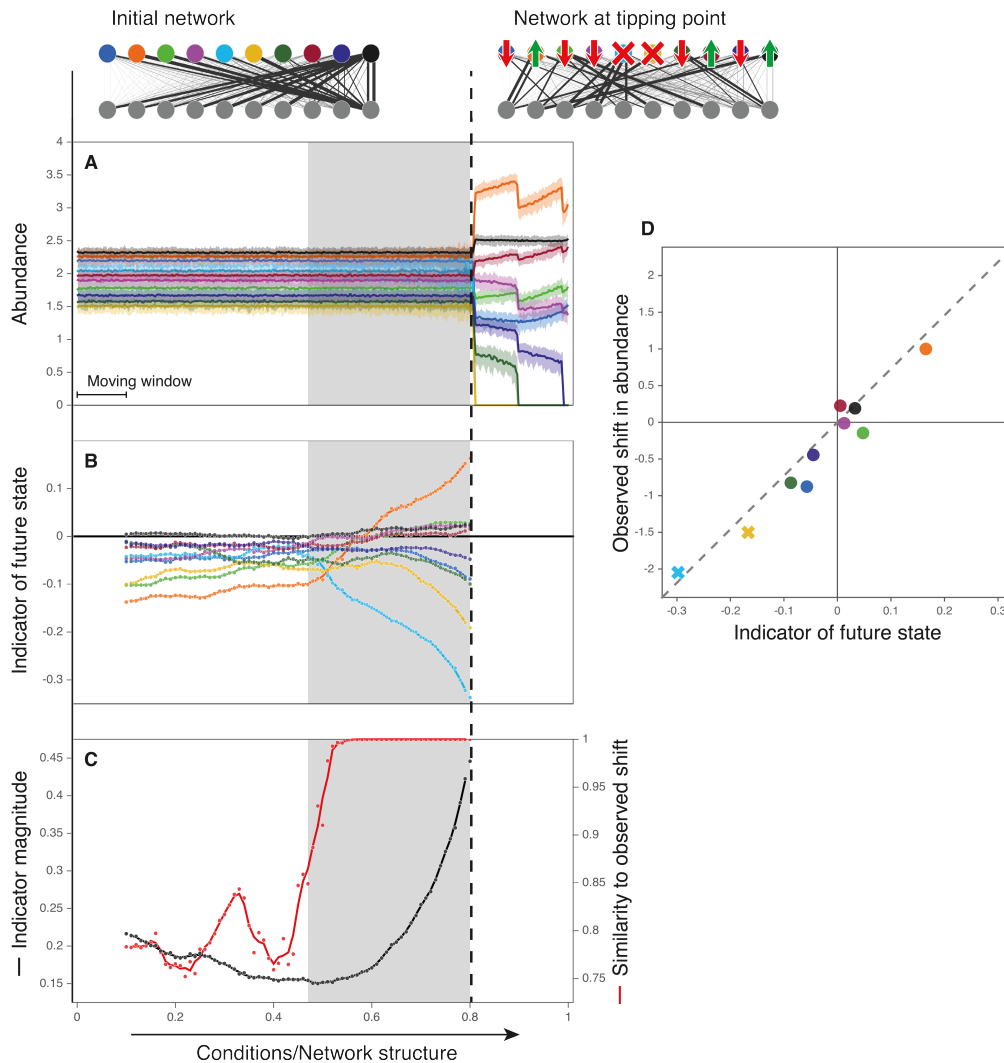


Figure 4.2: Directional slowing down in a mutualistic network as detected by our indicator. (A) Time series of species belonging to one part of a bipartite mutualistic network, i.e. the pollinators. At the tipping point two species collapse to extinction (light blue and yellow). (B) The indicator of the future state measuring the direction in which fluctuations are distributed asymmetrically. Scores on the indicator indicate the relative predicted gain or loss of each node. (C) The magnitude of the indicator, reflecting the extent to which fluctuations are distributed asymmetrically, plotted together with the accuracy measured as the similarity between its direction and the observed shift in abundance. Grey bands indicate the period in which the indicator's magnitude increases significantly. This period likely corresponds to the period in which the network rapidly loses resilience (as in Fig. 4.1.A.II-III). The accuracy increases rapidly at the beginning of this period. (D) The observed changes in abundance versus the scores on the indicator just before the tipping point. Extinct species are indicated with crosses. The observed shift is nearly proportional to the scores on the indicator as points are close to a straight line through the origin.

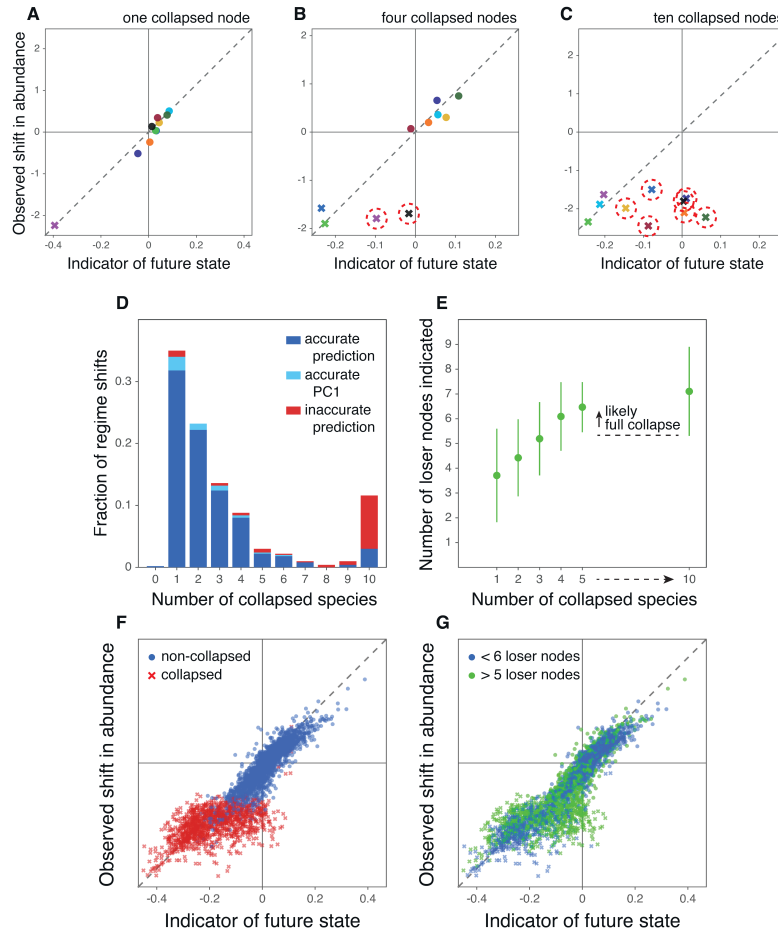


Figure 4.3: Cascading collapses and the indicator's performance when predicting the future state of mutualistic networks. **(A-C)** Examples of the relationship between the scores on the indicator and the observed shifts in abundance when a single, when four, and when all pollinator species collapse to extinction. The change in abundance of winners, losers and two or three collapsed species was almost always accurately indicated. The loss in abundance of additional species collapsing (red circles) was underestimated. **(D)** The fraction of regime shifts after which a certain number species collapsed to extinction. The fraction for which the change in abundance was not accurately indicated is shown in red. Inaccurate predictions (as in panel C) usually occurred prior to a full network collapse. **(E)** Relationship between the number of species collapsing and the number of species with a negative score on the indicator (mean and SD). When the number of species indicated to lose in abundance was high, we were often dealing with a full network collapse. **(F-G)** Combined plots of the 900 best indicated transitions in a data set of 1000 regime shifts. Species remaining after a regime shift (blue dots, panel F) are indicated more accurately than collapsing species (red crosses, panel F). Species of which the loss in abundance prior to a collapse was underestimated usually belonged to networks of which 5 or more species were indicated to lose in abundance (green dots and crosses, panel G). Competitive interaction strengths were taken from a low to intermediate range (i.e. 0.02-0.08).

to generate networks of which the initial, nontrivial state is stable. More specifically, we found that the probability of a network to be stable at initial conditions, $M = 0$, is nearly one when interspecific competitive interaction strengths were taken from the aforementioned lower ranges and below 0.01 when they were taken from the highest here reported range (Fig. S4.16). The indicator accurately indicated about 50% of the regime shifts in this ‘worst-case scenario’ (some of the hard-to-predict regime shifts were indicated accurately). When there is no competition between species, this percentage was nearly 100% (Fig. 4.4.C-D).

Qualitatively similar results were found when, in addition to a change in relative mutualistic benefits, the species’ nontrivial equilibrium abundances changed as well (see Appendix S4.7 and Fig. S4.17). Full network collapses are more frequent when abundances tend to decrease and the period in which the indicator’s magnitude increases prior to a critical transition tends to be somewhat shorter when abundances change over time. The examples in Fig. S4.18-S4.20 suggest that the direction of the first principal component is initially determined by the way in which abundances change over time. It may, therefore, take longer before the direction in which abundances are distributed asymmetrically is determined by the direction of critical slowing down. The application of a detrending method may prolong this period when trends are strong.

Qualitatively similar results were also found when analyzing data sets of communities with different numbers of species (see Appendix S4.7). Full network collapses became less common as the number of species increased, and Hopf bifurcations leading to oscillatory, chaotic or other complex dynamics became more frequent (Fig. S4.21-S4.22). These changes occurred, most likely, due to a change in the balance between intra- and interspecific competition. Interaction strengths were assigned such that the relative difference between intra- and interspecific competitive interaction strengths remained approximately the same (see Appendix S4.5). The number of interspecific competitive interactions, however, increases as the number of species increases. The combined effect of all interspecific competitive interactions is therefore larger. Systems with many species may, due to the way in which we assigned competitive interaction strengths, therefore be comparable with smaller networks in which interspecific competition is relatively strong.

Simulations with the more general, unipartite model of facilitation between species gave roughly the same qualitative results as the bipartite plant-pollinator model (see Appendix S4.7). The resilience of communities of 10, 20 and 40 species was generally a bit lower than the resilience of plant-pollinator networks with the same number of plant and pollinator species. To prevent networks from collapsing almost immediately, we chose a lower noise level with standard deviation $\delta_i = 0.05$. A relatively low resilience may also explain the relatively high frequency of likely cascading collapses in facilitative communities of 10 species (Fig. S4.23). A different way of assigning critical abundances, A_i , could have increased the resilience of the here studied facilitative communities.

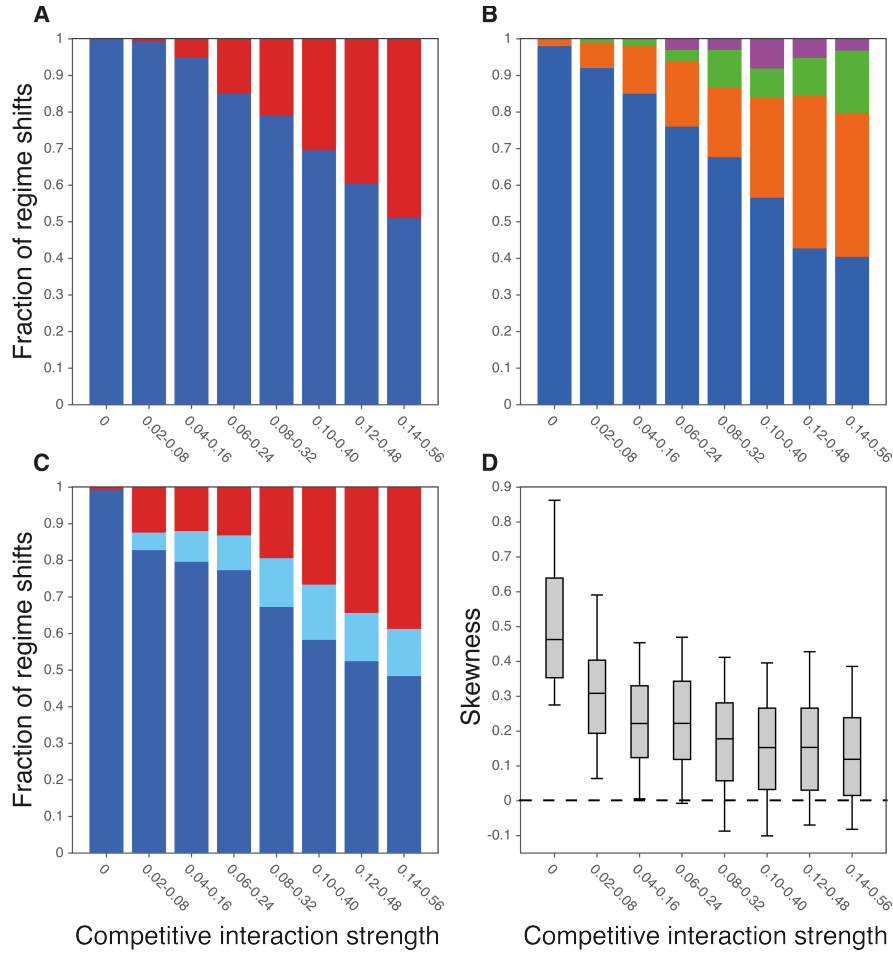


Figure 4.4: Hopf bifurcations and the predictability of a network's future state when sampling competitive interaction strengths from different parameter ranges (ranges are indicated on the x-axis). As the strength and variability of competition increases, Hopf bifurcations become increasingly frequent as well as the number of networks of which the future state is determined by chance. **(A)** The frequency of saddle-node (blue) and Hopf bifurcations (red) for different data sets. A high frequency of Hopf bifurcations indicates that transitions towards oscillatory, chaotic or other complex dynamics are common. **(B)** The fraction of cases in which, after five simulations in which a network's resilience was undermined in the exact same way, a network would always shift to the same state (blue), to one out of two states (orange), to one out of three states (green), or to one of four or more potential future states (purple). **(C)** The fraction of accurately indicated regime shifts (dark blue), the fraction accurately indicated by the first principal component, i.e. the slope of the indicator is accurate, but not by the direction in which time points are skewed (light blue), and the fraction of inaccurately indicated regime shifts (red). **(D)** The skewness of time points projected on the first principal component. A positive skewness means that time points are skewed in the direction of a network's future state. The skewness is shown for regime shifts that were accurately indicated by the first principal component.

4.4 Discussion

Human activities alter the Earth's climate and its ecosystems at unprecedented rates (Vitousek et al. 1997; Millenium Ecosystem Assessment 2005; Rockström et al. 2009; Intergovernmental Panel on Climate Change 2014; Steffen et al. 2015). These changes may jumble the patterns in the networks of interactions between species that hold complex species communities together (Kareiva et al. 1993; McCann 2007; Tylianakis et al. 2008). Monitoring and forecasting the effects of such changes thus requires a systems approach, i.e. an approach that explicitly studies the properties emerging from the complex and often unknown ways in which species relate to each other. Here, we try to make a further step towards developing such an approach by determining the direction in which destabilizing positive feedbacks undermine resilience. With model-generated time series we show that this direction is indicative of the future state of mutualistic communities, potentially providing us with a tool to assess the impact of impending critical transitions in natural communities and other complex systems.

Ecologists have emphasized the importance of improving our ability to predict the future state of ecosystems previously, and predicting future developments in complex systems is common practice in various fields of research, e.g. economics, engineering, and climatology (Clark et al. 2001; Sutherland 2006; Coreau et al. 2009; Beckage et al. 2011; Novak et al. 2011; Evans et al. 2012, 2013a; Purves et al. 2013; Petchey et al. 2015). Concerns about the forecastability of ecosystems and the limits to our capacity to predict the future state of ecosystems have however also been strong (Coreau et al. 2009; Beckage et al. 2011). Some of these concerns stem from a misunderstanding of why predictions are made. Making predictions is fundamentally different from describing a scientific law. Predictions are made when a limited amount of knowledge is available, and people rely on predictions even when they are known to often be inaccurate simply because better predictions are not available. Predictions may also be made when evaluating the risks associated with different ecological scenarios. In this spirit, we also see the indicator we propose here; as an indication of where a system's future state might lay. There is no absolute certainty as complex dynamics may occur after a critical threshold is passed.

Some general properties may, however, give a clue about the predictability of ecosystem dynamics. We found that, as the strength and variability of interspecific competition increases, dynamics change from a situation where positive feedbacks are the main cause of instability, to a mixed, intermediate situation, and, eventually, to a situation in which delayed negative feedbacks govern ecosystem dynamics. Our results suggest that the indicator performs well at predicting a system's future state when positive feedbacks are strong. Performance was reasonably good and transitions caused by positive feedbacks remained quite common in the aforementioned mixed situation, i.e. more than 50% accurate predictions. When dynamics were governed by delayed negative feedbacks, we found that the initial pristine state of the here studied systems was unlikely to be stable,

i.e. the probability of a system's nontrivial equilibrium to be stable was below 0.01. The indicator could not be applied and the interplay between several delayed negative feedbacks was likely to lead to chaotic dynamics.

Ecosystems exhibit positive feedbacks when species have direct or indirect positive effects on themselves, i.e. in loops with an even number of negative interactions, and do not only occur as the result of mutually beneficial interactions. Positive feedbacks may, for example, also occur when species positively affect themselves by suppressing other species, e.g. between a pair of competing species and in three-species omnivore loops in food webs (Nes and Scheffer, 2004; Neutel and Thorne, 2014). Despite a longstanding interest in the occurrence of complex ecosystem dynamics (May 1974; Hastings and Powell 1991; Huisman and Weissing 1999), no real classification of where and when to expect unpredictable, complex dynamics exists. As a first speculative proposal, we suggest that all the various types of mutualistic communities are likely to exhibit relatively strong positive feedbacks and predictable dynamics. Terrestrial foodwebs, where the top-down effects of herbivory are relatively small (Cyr and Face 1993), may fall in the aforementioned mixed category, while aquatic food webs are more likely to exhibit chaotic dynamics (e.g. Benincà et al. 2008). Complex dynamics are likely to occur in competitive communities when competitive interaction strengths are variable and asymmetrical. When pairs of interacting species have similar competitive effects on each other, positive feedbacks between some pairs of species are more likely to be strong and dynamics may be fairly predictable (e.g. Nes and Scheffer 2004). Further research into where and when to expect complex dynamics will greatly improve our capacity to evaluate the performance of the here proposed indicator and the predictability of ecosystem dynamics in general. Such research may, for example, involve a further investigation of the interrelationship between the structural properties of ecological networks and the occurrence of different types of critical transitions and may include transitions that are not preceded by critical slowing down (Grebogi et al., 1983; Hastings and Wysham, 2010).

Earlier studies explored different ways in which changing environmental conditions may lead to critical transitions in mutualistic networks, for example by increasing pollinator mortality rates (Jiang et al., 2018) or by declining mutualistic interaction strengths (Dakos and Bascompte 2014). In this work, assumptions were made that make the effects of these changes fairly simple from a dynamical perspective, e.g. the assumption that the intrinsic properties of species and the effects of changing environmental conditions are similar for all species, and the assumption that the structure of whom interacts with whom remains unchanged. As a consequence, there is little change in the direction of slowest recovery and the nature of the systems' alternative stable states. Here, we chose to study a more complex dynamical scenario because we wanted to test whether the direction of critical slowing down is indicative of a community's future state even when the direction of slowest recovery changes substantially prior to the period in which resilience is lost. There is no reason to assume that the indicator would perform worse at predicting a

system's future state when changing conditions affect a group of similar species in one of the aforementioned more simple ways.

The here proposed indicator has a number of advantages compared to previous methods to predict the future state of ecosystems such as extrapolation and the use of mechanistic models. Extrapolation is risky, because it assumes trends to continue outside of the range in conditions for which data are collected, and the behavior of mechanistic models, e.g. aiming to simulate feeding, reproduction, death, and other rates with as much accuracy as possible, often depends on many unknown parameters, in particular when these rates depend on environmental conditions and species abundances. Using the direction of critical slowing down as an indicator of a system's future state has the advantage that it directly relates to an emerging property of complex ecosystems, i.e. the direction in which resilience is lost. As such, it avoids the often difficult process of parameter estimation needed to develop mechanistic models, and it specifically aims to predict a system's future state when abrupt shifts away from existing trends, i.e. critical transitions, occur.

The above described results consider scenarios in which plenty of data are available. When time series are short, i.e. contain few data points, or when the rolling window used to analyze time series contains few data points, predictions become less accurate (Fig. S4.24-S4.29). This brings us to the question of how we may determine the data requirements in practice. For this, it is important to consider the two different aspects of our analysis: 'critical slowing down' and 'the direction of slowest recovery'. Critical slowing down can only be detected over a longer time periods, i.e. in which conditions change, while the direction of slowest recovery can be determined for a given set of conditions, i.e. over a short period of time. When determining critical slowing down it is not necessary to monitor the abundances of all species per se, while this is important when determining the direction of slowest recovery. A more economical approach could thus be to monitor only few species for indicators of critical slowing down (Scheffer et al., 2009; Dakos et al., 2012a), and to determine the direction of slowest recovery only once these indicators suggest that the system approaches a tipping point. In some cases, one may even consider to skip monitoring of critical slowing down indicators altogether and focus on determining the direction of slowest recovery in systems that are known to be under stress.

Two aspects could cause our approach to be less data-hungry than expected. First, we are only interested in the slope indicated by the first principal component and require, therefore, fewer data when compared to analysis in which also the higher-order components are of importance. Secondly, we expect the distribution of abundances to become highly asymmetric when a system approaches a tipping point. Dynamics become similar to a low-dimensional system and the number of observations needed to accurately determine the direction of slowest recovery becomes smaller when a system approaches a tipping point (Fig. S4.30). It remains, however, difficult to determine a priori what the data demands are.

Previous studies have proposed rules of thumb that give an indication of the minimum sample size required to perform principal component analysis, i.e. the method used to determine the slope of the indicator. Such rules are often a function of the number of variables, e.g. species abundances, and suggest that the minimum sample size required to perform a principal component analysis should be at least n , e.g. 2, 10 or 20, times more than the number of variables. Velicer and Fava (1998) and MacCallum et al. (1999) showed, however, that such rules of thumb are invalid and that the required sample size depends on the underlying correlation structure. A better approach to determine the minimum sample size is therefore to draw subsets from the data and compare results for the subset with those for the full set (Barrett and Kline 1981; Arrindell and Van der Ende 1985). When subsets give similar results to the full set, enough data is likely obtained. Methods to determine the effect of a change in sample size may vary from a simple comparison of the direction indicated (as in Fig. S4.30) to more advanced bootstrapping techniques (as in Shaukat et al. 2016).

In this study, we chose to use time-series analysis because it links closely with previous work on early warning signals (Scheffer et al. 2009; Dakos et al. 2012a), and because data collection efforts have, traditionally, focused on species abundances. For some ecosystems it may, however, be easier to monitor changes in the structural properties of ecological networks rather than in the specific way in which a system recovers from small perturbations. When such monitoring efforts could be used to estimate (changes in) the effective relationships between species as described by the different elements of the Jacobian matrix, we may be able to obtain a more direct measure of (changes in) the relative strengths of feedback loops in ecosystems, their proximity to a tipping point, and their likely future states. Our analysis suggests, for example, that the extent to which species are saturated and the relative benefits received from mutualistic partners play a crucial role in determining the resilience and future state of mutualistic communities. These properties might be measured in more direct ways, for example by determining the time spent by pollinators on handling and searching for nectar and their relative visitation rates to different plant species. Other theoretically-informed measures for other types of ecosystems may likely provide us with other potential indicators of the direction of critical slowing down.

In a time when humanity's biggest challenges and opportunities depend upon our capacity to manage complex natural systems, new tools to foresee the risks and opportunities associated with critical transitions are of increasing importance. Such tools may not only be useful when addressing the question of what a system's future state might be like, but may also help to address questions such as to what extent individual species or interactions are contributing to network resilience and which deliberate human interventions could prevent or alter the outcome of impending critical transitions. Such approaches are becoming increasingly useful as the availability of data on natural and other complex systems is rapidly increasing.

Supplementary materials

S4.1 Example: Undermining the resilience of a 3-species network

To illustrate how differences in the intrinsic properties of species and the arrangement of interactions between them may affect the overall resilience of mutualistic networks, we use a model in which one pollinator species interacts mutualistically with two plant species. The system's overall resilience is highest when this pollinator species obtains most resources from the more saturated plant species.

As conditions change from a situation in which pollinators obtain most resources from highly saturated plant species P_1 , i.e. with high saturation term h_1 , to a situation in which they obtain most resources from less saturated plant species P_2 , the network becomes increasingly sensitive to small-scale stochastic perturbations. Eventually, a critical transition occurs away from the initial pristine state of the network towards a fully collapsed network state in which both plant species and the pollinator species are extinct.

For illustrative purposes, we assume plants to be in steady-state and determine how changing conditions affect the relationship between the net growth of the pollinator species, $dN^{(A)}/dt$, and the abundance of the pollinator species, $N^{(A)}$ (Fig. S4.1). The net growth of the pollinator species is negative at low abundances. As a result, there are two alternative stable states; a pristine state in which the pollinator species has a positive abundance and a collapsed state in which the abundance of the pollinator species is zero. These two alternative stable states can be visualized more intuitively by a stability landscape of which the slope corresponds to the rate at which the abundance of the pollinator species changes, $dN^{(A)}/dt$, valleys to the attraction basins of the alternative stable states, and hilltops to the threshold between the two attraction basins. As conditions change, the attraction basin of the initial pristine state of the network becomes increasingly small and a small perturbation becomes sufficient to cross the threshold and cause a critical transition towards the alternative fully collapsed state of the network.

Parameter settings: $\hat{N}_i = 2$, $c_{ii} = 0.4$, $c_{ij} = 0.1$, $d_i = 0.2$, $h^{(A)} = 0.3$, $h_1^{(P)} = 0.3$, $h_2^{(P)} = 0.1$, and $\epsilon_i = 0.01$. Initial interaction strengths: ($M = 0$): $\theta_{0,11}^A = 1$, $\theta_{0,12}^A = 0$, $\theta_{0,11}^P = 1$, and $\theta_{0,21}^P = 1$. Final interaction strengths: ($M = 1$): $\theta_{final,11}^A = 0$, $\theta_{final,12}^A = 1$, $\theta_{final,11}^P = 1$, and $\theta_{final,21}^P = 1$.

S4.2 Example: Critical Slowing Down in a 4-species network

To illustrate the direction in which a community slows down prior to a critical transition and how this might be used to predict a community's future state, we use a model in which two pollinator species interact mutualistically with two plant species. As described in the main text, changing conditions undermine the resilience of this small network by altering relative mutualistic benefits, θ . As was the case with the earlier studied 3-species network (see Appendix S4.1), regime shifts occur in the here studied 4-species community because the community's initial pristine state is approached by a threshold (i.e. a boundary between two attraction basins, Fig. 4.1). As conditions change, the minimum size needed for perturbations to push the system over the approaching threshold becomes smaller. The likelihood of a transition caused by the small-scale stochastic perturbations incorporated in our model therefore increases and, eventually, a regime shift towards an alternative state becomes inevitable.

The outcome of a transition depends on the way in which changing conditions undermine a community's resilience. One, some or all species may collapse to extinction and remaining species may either gain or lose in abundance from a regime shift. Multiple thresholds separating the community's initial pristine state from different alternative stable states, or 'potential future states', may exist prior to a regime shift. Changing conditions may alter the number and nature of these alternative stable states, and the thresholds towards them may or may not approach the network's initial pristine state. Which alternative state eventually becomes the community's future state depends on which threshold towards which future state eventually approaches a community's initial state.

For illustrative purposes, we assume plants to be in steady state and determine how changing conditions affect the dynamics of the network. These dynamics can be visualized intuitively by a stability landscape of which the slope corresponds approximately to the rate at which the abundances of pollinator species change, $dN^{(A)}/dt$ (see methods below). Every possible combination of pollinator abundances is represented by a unique point in the stability landscape and alternative stable states are at the lowest point of the landscapes valleys or 'attraction basins'. Thresholds between attraction basins are represented by ridges in the stability landscape. These thresholds are not equally high at all places and have local maxima at hilltops and local minima at saddle points in the network's stability landscape. Attraction basins are shallow in between alternative stable states and the saddle points on the thresholds that separate them. When approached by a threshold, the attraction basin of the initial pristine state becomes increasingly shallow and the network increasingly slow when recovering from perturbations in the direction of the saddle point on the approaching threshold.

For the here studied 4-species network (Fig. 4.1) we found that the network's pristine state is initially accompanied only by a fully collapsed state, i.e. a state in which the abundance of all species is zero. The pristine state's distance from the threshold towards

this state, however, remains large even when conditions change. A regime shift towards a fully collapsed state remains, therefore, unlikely. Changing conditions start to rapidly undermine the network's resilience only after the appearance of the first of two additional alternative stable states. These states correspond to partially collapsed network states in which the abundance of some but not all species is zero. Both thresholds towards both partially collapsed states approach the network's pristine state. One threshold, however, approaches the initial pristine state more closely than the other and eventually a regime shift, caused by the small-scale stochastic perturbations to which the network is permanently subjected, towards the partially collapsed state in the attraction basin behind this threshold becomes inevitable.

As conditions change there are two decisive moments which are both preceded by a particular change in the network's dynamics. The first is the moment at which the future state of the network comes into existence as an alternative stable state in the network's stability landscape, and the second is the moment at which the regime shift towards this alternative stable state actually occurs. The direction in which the network recovers slowly from perturbations changes substantially before the future state of the network comes into existence from a direction that roughly indicates a full collapse to a direction that indicates the future partially collapsed state of the network. The speed at which the network recovers from perturbations, however, remains approximately the same. After the future state of the network comes into existence, the network slows down dramatically when recovering from perturbations in approximately the same direction (Fig. S4.2).

Methods: To determine the rate at which pollinator abundances change as illustrated in Fig. 1.B, we analytically determined this rate, $v^{(A)}$, for different pollinator abundances at 200 by 200 grid points in the network's phase plane as follows:

$$v^{(A)} = \left(\sum_{i=1}^{S^{(A)}} \left(\frac{dN_i^{(A)}}{dt} \right)^2 \right)^{0.5}, \quad (\text{S4.1})$$

in which $N_i^{(A)}$ is the abundance and $dN_i^{(A)}/dt$ the net growth rate of pollinator species i . At the same grid points we determined the height of the stability landscape with an algorithm that keeps updating the height of the landscape until all slopes in between these points are within a certain margin of error from the pollinators net growth rate. This allows us to intuitively show the position of alternative stable states, which are found at the bottom of the landscapes valleys or ‘attraction basins’, and the thresholds between them, which correspond to hills or ridges in the landscape. The stability landscape produced with this algorithm, is a useful tool to intuitively illustrate the idea behind our method. As our system is non-gradient, it is not a way to determine the potential energy of the system.

Parameter settings: $\hat{N}_i = 2$, $c_{ii} = 0.4$, $c_{ij} = 0.1$, $d_i = 0.2$, $h_1^{(A)} = 0.1$, $h_2^{(A)} = 0.3$, $h_i^{(P)} = 0.3$, and $\epsilon_i = 0.04$. Initial interaction strengths: ($M = 0$): $\theta_{11}^A = 0.7$, $\theta_{12}^A = 0.3$, $\theta_{21}^A = 0.5$, $\theta_{22}^A = 0.5$, $\theta_{11}^P = 0.5$, $\theta_{12}^P = 0.5$, $\theta_{21}^P = 0.3$, $\theta_{22}^P = 0.7$. Final interaction strengths: ($M = 1$): $\theta_{11}^A \approx 0.83$, $\theta_{12}^A \approx 0.17$, $\theta_{21}^A \approx 0.10$, $\theta_{22}^A \approx 0.90$, $\theta_{11}^P \approx 0.90$, $\theta_{12}^P \approx 0.10$, $\theta_{21}^P \approx 0.17$, and $\theta_{22}^P \approx 0.83$

Conditions analyzed for Fig. 1 in the main text: $M = 0.31$, $M = 0.66$, and $M = 0.87$.

S4.3 Similarity between the indicated and the observed shift

As explained in the main text, the slope of the indicator is determined by the first principal component (Fig. S4.4.C), while the eventual (up- or downward) direction of the indicator along the first principal component is determined by the direction in which time points are skewed (Fig. S4.4.D-E). To assess the performance of our indicator, we evaluate the performance of the first principal component and the skewness of the projected time points independently. An accurate slope, means that the indicator performs well at predicting the relative gain or loss of species and which species shift in opposite directions (i.e. an ‘accurate PC1’). The indicated direction is, however, only fully ‘accurate’ when the actual winners and losers are also indicated correctly. This depends on the direction along the first principal component in which time points are skewed.

To evaluate the performance of the first principal component, we determine the difference between the slope of our indicator and the direction of the observed shift in abundance. We do this by determining the angle, θ , between the direction of the indicator and the observed shift as follows:

$$\theta = \cos^{-1} \frac{I \cdot \Delta N^{(A)}}{|I| |\Delta N^{(A)}|}, \quad (\text{S4.2})$$

in which I is the indicator of a network’s future state and $\Delta N^{(A)}$ the observed shift in pollinator abundances. $I \cdot \Delta N^{(A)}$ indicates that we take the dot product between these two vectors. To determine $\Delta N^{(A)}$, we take the mean abundances over 200 time steps at 500 steps before the tipping point and subtract it from the mean abundances 500 steps after the tipping point was found. Because we want to evaluate the accuracy of the first principal component, and not whether points are also skewed in the right direction, we take $-I$ as the input for the formula above when we find an angle $> \pi/2$ (i.e. > 90 degrees). Both I and $\Delta N^{(A)}$ are vectors of which the number of dimensions is equal to the number of species analyzed. The smaller the angle, the more similar the direction of the two vectors.

Two random vectors in a ten-dimensional space are more likely to be orthogonal than two random vectors in a three-dimensional space. More extreme small or large angles become less likely as the number of dimensions increases (Fig. S4.5). How ‘special’ it is to find a certain angle between the indicated and the observed shift thus depends on the number of dimensions in a system. As a measure of how different the indicated direction is from the observed regime shift, we determine for the observed angle, θ , the likelihood that two unrelated random vectors have an equal or smaller angle. As a measure of similarity, we take one minus this probability, and we consider the indicator’s slope to be accurate when this measure of similarity is above 0.99.

To determine the aforementioned probability, we use the following probability density

function:

$$h(\theta) = \frac{1}{\sqrt{\pi}} \frac{\Gamma(\frac{S^{(A)}}{2})}{\Gamma(\frac{S^{(A)}-1}{2})} \cdot (\sin \theta)^{S^{(A)}-2}, \quad (\text{S4.3})$$

in which $S^{(A)}$ is the number of dimensions and $h(\theta)$ the probability density for a certain angle θ (ref. Cai et al. (2013)). Our method may be interpreted as a test whether the null hypothesis that I and N are two random vectors is true. This hypothesis is rejected when angle is found to be significantly smaller than the expected angle between two random vectors, when the one-sided p-value is smaller than 0.01 (i.e. similarity > 0.99).

To evaluate the tendency of time points to be skewed in the direction of a network's future state, we determine the skewness of the time points projected on the first principal component. When points are skewed in the direction of the network's future state, we report a positive skewness. When points are skewed in the opposite direction, we report a negative skewness. We consider a positive skewness as accurate and a negative skewness as inaccurate. A strong positive or negative skewness is considered more accurate or inaccurate than a weak positive or negative skewness.

S4.4 Time series analysis

Unless stated otherwise, we determine the dominant direction of fluctuations in a rolling window of 10% of the entire time series (e.g., 2000 out of 20.000 time points) to detect changes in the direction and extent in which time points are distributed asymmetrically. The choice of this window size is to some extent arbitrary. A too small window leads to irregular trends, while a too large window smooths out the trends. To test whether the size of the window chosen influences our results, we make additional analysis in which we use a window size of 0.005, 0.1, 0.5, 0.1, 5, 10, 20 and of 50% of the time series. The rolling window is moved along the time series with steps of 1% of the time series, independent of the window size. As time passes by, the direction and magnitude of the indicator is thus computed every 200 time steps in a window containing the last 2000 time steps when using a window size of 10% of a time series with a length of 20.000 time points.

Far from a tipping point, time points may be skewed only weakly. When this is the case, sudden shifts of nearly 180 degrees may occur in the direction of the indicator when time points are skewed in a different direction along the first principal component. Clearly, such large shifts in direction do not occur because the network's future state has changed. We, therefore, correct previously found indicator values such that there is no change larger than 90 degrees between two consecutive points at which the indicator's direction was determined. We assume the last direction in which time points were found to be skewed to be the accurate one.

To determine whether there is a significant increase in the indicator's magnitude, we determine the Kendall rank correlation coefficient, τ , for the last ten points at which the indicator's magnitude was computed. We consider the increase significant when this coefficient was positive and its p-value < 0.05 . Once a significant increase was found, we tested whether the increase remained significant by determining Kendall's correlation for the last eleven points the next time the indicator's magnitude is determined, for twelve points the time after that, and so on until the tipping point is reached. We would again look at the last ten points when the increase was found to not be significant anymore. By doing this, we could determine the range in conditions in which the indicator's magnitude increased significantly.

As a measure of a 'regime shift' we determined whether there was a change in abundance of more than 1.5 over a period of 1% of the entire time series (200 time steps). We did this by taking the mean abundances over a period of 200 time steps before this period and 200 time steps after this period and determining Euclidean distance between these two mean abundances. To make sure that this large shift in abundances was not a temporal large deviation from the species' mean abundances, we added as a second criterion that the abundance of at least one species should be near extinction, i.e. below 0.1.

We did not apply any preprocessing to handle trends in the time series. We expect

the indicator to be relatively robust against such trends, because trends only alter the direction of the first principal component when their effect on this direction is stronger than the effect of critical slowing down. Not applying any preprocessing is a good way to test this robustness. When using the indicator as part of a different study it may, however, be worth considering to apply a preprocessing method (see ref. Dakos et al. (2012a)). It may improve the performance of the indicator, especially when trends are strong.

S4.5 Additional information mutualistic networks

Nontrivial equilibrium abundances, \hat{N} , competitive interaction strengths, c , mortality rates, d , and saturation terms, h , are randomly sampled from predefined probability distributions, and the total amount of resources received by species i at the system's nontrivial equilibrium, $R_i(\hat{N}^{(P)})$, are assigned such that the rate at which abundances change at the system's nontrivial equilibrium, $d\hat{N}^{(P)}/dt$, is zero:

$$R_i(\hat{N}^{(P)}) = \frac{\sum_{j=1}^{S(A)} c_{ij} \hat{N}_j^{(A)} + d_i}{1 - h_i(\sum_{j=1}^{S(A)} c_{ij} \hat{N}_j^{(A)} + d_i)}. \quad (\text{S4.4})$$

The total amount of resources provided at the system's nontrivial equilibrium, $R_i(\hat{N}^{(P)})$, is thus approximately the same for highly specialized and more generalist species, provided that their losses due to competition, c , and mortality rates, d , and their nontrivial equilibrium abundances, \hat{N} , are similar.

The extent to which species are saturated is determined by the total amount of resources provided, $R_i(\hat{N}^{(P)})$, and the rate at which species become saturated as determined by saturation term h_i . In our simulations, we assume nontrivial equilibrium abundances, \hat{N} , and inter- and intraspecific competition, c_{ij} and c_{ii} , to be similar for all species. Highly saturated species are, therefore, the ones with a high h_i . Species are saturated relatively quickly, and, according to equation S4.4, the total amount of resources provided at the system's nontrivial equilibrium is high when species have a high h_i .

Parameters are assigned such that there are substantial differences in the extent in which species are saturated by drawing saturation terms, h_i , from a scaled beta distribution with range $\sim (0.05, 0.35)$ and shape parameters $\alpha = 1$ and $\beta = 5$. Due to this distribution, there are few highly saturated species, i.e. h_i close to 0.35, and many non-saturated species, i.e. h_i close to 0.05. Strong mutualistic interactions between non-saturated species lead to strong positive feedbacks. Non-saturated species thus need to obtain a relatively large share of resources from a few, highly saturated species for the network to be stable. Relative mutualistic benefits at initial conditions, $\theta_{0,ik}$, are therefore ordered such that larger benefits are obtained from the more saturated species. To make sure that the sum of all relative benefits is one, we take relative mutualistic benefits, $\theta_{0,ik}$, from a symmetric Dirichlet distribution. The distribution's concentration parameter, α , determines the extent in which species are specialized and is, for each species, taken from a uniform distribution between zero and one.

To explore how transitions towards oscillating, chaotic or other complex dynamics caused by delayed negative feedbacks may influence the performance of the indicator, we analyze several data sets of which the strength and variability in interspecific competitive interaction strengths, c_{ij} , varies. The tested parameter ranges are: $c_{ij} = 0$, $c_{ij} \sim U(0.02, 0.08)$,

$c_{ij} \sim U(0.04, 0.16)$, $c_{ij} \sim U(0.06, 0.24)$, $c_{ij} \sim U(0.08, 0.32)$, $c_{ij} \sim U(0.10, 0.40)$, $c_{ij} \sim U(0.12, 0.48)$, and $c_{ij} \sim U(0.14, 0.56)$. Intraspecific competition strengths, c_{ii} , are taken from $\sim U(0.9, 1.1)$. Delayed negative feedbacks become stronger as the strength and variability of interspecific competition increases. Simulations are made for communities of 10 plant and 10 pollinator species. Initial equilibrium abundances, $\hat{N}_{0,i}$, and mortality rates, d_i , are taken from $\hat{N}_{0,i} \sim U(1.5, 2.5)$ and $d_i \sim U(0.15, 0.25)$. Initial and final nontrivial equilibrium abundances are assumed to be equal, $\hat{N}_{final,i} = \hat{N}_{0,i}$.

Changing environmental conditions, M , lead to an increase in the relative mutualistic benefits received from some, and a decrease in the relative benefits received from other species. We assume the distribution of interaction strengths of the final network, at $M = 1$, to be quite heterogeneous (Fig. S4.6). We select, therefore, with a probability of 0.75, interactions of which the interaction strength goes to zero, $\theta_{final,ik} = 0$. To the remaining interactions, relative interaction strengths are assigned by taking them from a uniform Dirichlet distribution ($\alpha = 1$). The ‘diet breath’ of plants and pollinators thus tends to become more narrow as could be the case under various scenarios of global environmental change (Memmott et al. 2007; Burkle et al. 2013).

As conditions change, either a single eigenvalue or a pair of complex conjugate eigenvalues goes to zero. In the first case we are dealing with a saddle-point approaching the network’s initial state, caused by a positive feedback. In the second case, we are dealing with a Hopf bifurcation caused by a delayed negative feedback.

Data sets consist of 100 initial networks. For each network, 10 final distributions of relative mutualistic benefits, $\theta_{final,ik}$, were drawn, allowing us to determine the extent in which a community’s future state depends on the specific way in which relative mutualistic benefits are changed. Parameters were assigned such that this dependency is high. Networks were discarded from a data set when they were unstable at initial conditions, $M = 0$. We determined the frequency at which this occurred as a measure of how difficult it is to find a stable solution for the initial networks of a given data set. The final distribution of relative mutualistic benefits was redrawn either when the network would become unstable within the range of conditions $M = (0, 0.5)$, or when a network would still be stable at $M = 1$.

To test whether the indicator also works when equilibrium abundances change, we analyzed networks of 10 plant and 10 pollinator species of which the final equilibrium abundances are different. We do this by changing the nontrivial equilibrium abundances of species as follows:

$$\hat{N}_i^* = \hat{N}_{0,i} + (\hat{N}_{final,i} - \hat{N}_{0,i})M, \quad (\text{S4.5})$$

in which $\hat{N}_{0,i}$ is the initial, $\hat{N}_{final,i}$ the final, and \hat{N}_i^* the actual nontrivial equilibrium

abundance of species i . The total amount of resources provided at the system's nontrivial equilibrium, and the strengths of mutualistic interactions are determined by equations 4.4 and S4.4. We tested three scenarios. One in which the nontrivial equilibrium abundances of species tend to increase, $\hat{N}_{final,i} \sim U(2, 3)$, one in which they stay the same on average $\hat{N}_{final,i} \sim U(1.5, 2.5)$, and one in which they tend to decrease $\hat{N}_{final,i} \sim U(1, 2)$. Competitive interaction strengths were taken from the following distributions: $c_{ii} \sim U(0.9, 1.1)$ and $c_{ij} \sim U(0.02, 0.08)$. Changing abundances affect all relationships as described by the Jacobian matrix. The main effect of a decline in abundance is, however, a reduction of the direct negative effects of species on themselves which undermines resilience. Increasing abundances tend to promote resilience.

To test whether the indicator may accurately indicate the future state of larger networks, we analyzed networks of 10 and 20, 10 and 40, 20 and 10, 20 and 20, 20 and 40, 40 and 10, 40 and 20, and 40 and 40 plant and pollinator species. We assigned competitive interaction strengths such that the rate at which species lose in abundance due to competition, $\sum_{j=1}^{S^{(A)}} c_{ij} N_j^{(A)} N_i^{(A)}$, is approximately the same for different numbers of species, as well as the relative difference between intra- and interspecific competition, c_{ij}/c_{ii} . When a species group consisted of 10 species we assumed $c_{ii} \sim U(0.9, 1.1)$ and $c_{ij} \sim U(0.02, 0.08)$. When a group consisted of 20 species $c_{ii} \sim U(0.67, 0.82)$ and $c_{ij} \sim U(0.015, 0.06)$, and when a group consisted of 40 species $c_{ii} \sim U(0.44, 0.54)$ and $c_{ij} \sim U(0.01, 0.039)$. Initial and final equilibrium abundances were assumed to be equal, $\hat{N}_{final,i} = \hat{N}_{0,i}$.

The amount of noise, determined by standard deviation δ , is assumed to be equal for all species. Unless stated otherwise, we assume standard deviation $\delta = 0.1$. Additional simulations were made with lower and higher noise levels, $\delta = 0.01$, $\delta = 0.05$, $\delta = 0.15$, and $\delta = 0.2$ to make sure that this does not qualitatively alter the results. Higher noise levels were not tested because they would lead to an almost immediate collapse. Unless stated otherwise, model generated time series had a length, T , of 20.000 time steps. Additional simulations were made in which time series had a length of 100, 200, 1.000, 2.000, 10.000, and 100.000.

S4.6 Additional information unipartite model of facilitation

Nontrivial equilibrium abundances, \hat{N} , interspecific facilitation rates, γ_{ij} , critical abundances A_i , interspecific competitive interaction strengths, c_{ij} , carrying capacities, K , and mortality rates, d , are randomly sampled from predefined probability distributions. Intraspecific facilitation rates, γ_{ii} , and intraspecific competition rates, c_{ii} , are one. To make sure that the rate at which abundances change at the nontrivial equilibrium, $d\hat{N}_i/dt$, is zero, we assign the intrinsic growth rates, r , as follows:

$$r_i = \frac{d_i \hat{N}_i A_i K_i}{(\sum_{j=1}^S \gamma_{ij} \hat{N}_j - A_i)(K_i - \sum_{j=1}^S c_{ij} \hat{N}_j) \hat{N}_i}. \quad (\text{S4.6})$$

The contribution of species to the overall resilience of a network is determined by critical abundance A_i . Species with a high critical abundance, A_i , collapse more easily and the overall resilience of the community is highest when such species are facilitated by species with a low critical abundance. A change from such a distribution to a more random distribution of facilitative interaction strengths will undermine resilience. To generate time series in which the resilience of the here described facilitative communities is undermined, we assume that conditions, M , affect facilitative interactions as follows:

$$\gamma_{ij}^* = \gamma_{0,ij} + (\gamma_{final,ij} - \gamma_{0,ij})M, \quad (\text{S4.7})$$

in which $\gamma_{0,ik}$ is the initial, $\gamma_{final,ik}$ the final, and γ_{ij}^* the actual facilitative interaction strength. Conditions, M , change from zero to one over time. We assume that the total amount of facilitation received, $\sum_{j=1}^S \gamma_{ij} \hat{N}_j$, remains equal as conditions change. We therefore determine the final facilitative interaction strength as follows:

$$\gamma_{final,ij} = \frac{\theta_{ij} \sum_{k=1}^S \gamma_{ik} \hat{N}_k}{\hat{N}_j}, \quad (\text{S4.8})$$

in which θ_{ij} is the fraction of the total facilitation received by species i from species j .

We assign parameters such that there are substantial differences in the critical abundances of species by drawing critical abundances, A_i , from a scaled beta distribution with $\alpha = 5$ and $\beta = 1$ and range $\sim (0, 1.5)$. Due to the beta distribution, there are few highly vigorous species (i.e. A_i close to 0) and many non-vigorous species (i.e. A_i close to 1.5). The initial facilitative interaction strengths are taken from the following uniform distribution: $\gamma_{0,ij} \sim U(0.2, 1.8)$. Initial facilitative interaction strengths are ordered such that species receive most facilitation, i.e. highest $\gamma_{0,ij}$, from species with the lowest A_i . We assume

that as conditions change, the strength of some facilitative interactions increases strongly while others approach zero. Final relative facilitative benefits, $\theta_{final,ik}$, are therefore selected with a probability of 0.75 and set to zero. To the remaining interactions, relative benefits are assigned by taking them from a uniform Dirichlet distribution ($\alpha = 1$). As with the model of mutualistically interacting species, we chose for this distribution of critical abundances, A_i , and facilitative interaction strengths γ_{ij} , because it leads to a high variety in potential future states to which a network may shift. Other parameters and equilibrium abundances are taken from the following uniform distributions: $\hat{N}_i \sim U(1.5, 2.5)$, $c_{ij} \sim U(0.04, 0.16)$, $d_i \sim U(0.15, 0.25)$.

Simulations were made with networks of 10, 20 and 40 species. As for the bipartite model of mutualistically interacting species, we assign parameters such that the rate at which abundance is lost due to competition, $\sum_{j=1}^S c_{ij}N_j/K_i$, remains approximately the same for different species numbers, as well as the relative difference between intra- and interpecific competition (see main text). Carrying capacities, K_i , were therefore taken from respectively $K_i \sim U(5, 6)$, $K_i \sim U(7.63, 9.15)$, and $K_i \sim U(12.89, 15.47)$, depending on the number of species.

The amount of noise, determined by standard deviation δ , is assumed to be equal for all species. For the results shown in this document we assume standard deviation $\delta = 0.05$. As with the model of mutualistically interacting species time series had a length, T , of 20.000 time steps.

S4.7 Supplementary results

Independent of the parameter ranges chosen, we found that regime shifts were preceded by a substantial period in which the indicator’s magnitude increases significantly, i.e. the ‘critical range’. Our indicator would, provided that the future state is indicated accurately, point towards a network’s future state during a substantial part of this period (Fig. S4.7). In Fig. S4.8 we provide information about the critical ranges as observed in a single data set ($c_{ij} \sim U(0.02, 0.08)$). These results are exemplary for the other data sets and show that our indicator consistently indicates a network’s future state during the period in which the network slows down.

Cascading collapses occur at an intermediate range of competitive interaction strengths most likely due to the nature of effective relationships between species, i.e. the combined effect of all direct and indirect interactions (Fig. S4.9). When there is no competition, effective relationships are positive and species collapse as one group. When competition is strong, most effective relationships are negative and species collapse independently. Cascading collapses are only likely when effective relationships are a mix of positive and negative relationships. When interspecific competitive interaction strengths, c_{ij} , were taken from $\sim U(0.02, 0.08)$, we found that such likely cascading, full network collapses took up a bit more than 12% of the data set. For specific parameter ranges not tested by us, this percentage may be higher.

In Fig. S4.10 we provide examples of two cascading collapses and one immediate network collapse. Species that collapsed a bit later, were also the ones for which the indicated loss in abundance was smallest, suggesting that the indicator indicates the initial regime shift accurately. The amount of time in between two consecutive partial network collapses can be extremely small. Also when cascades are not clearly visible, we suspect therefore that the inaccurate prediction of a full network collapse is caused by the occurrence of a cascading collapse.

In Fig. S4.13 we provide an example of a network for which the future state is hard to predict because it may shift to several alternative future states. When making five simulations in which relative mutualistic benefits, θ_{ik} , are changed in the exact same way by changing conditions, M , we found that the network shifted to four different future states. The future state of this network is determined by the only stochastic element in our model; the small-scale perturbations to which the network is permanently subjected. Our indicator accurately indicates two of the future states to which the network may shift, but does not indicate the other future states. A likely explanation for the several future states to which this system may shift is the fact that this system is approaching a Hopf bifurcation, leading to oscillating (Fig. S4.14), chaotic or other complex dynamics (Fig. S4.15). Such dynamics may explain a high sensitivity to perturbations in more than one direction.

In Fig. S4.18-S4.20, we show examples of time series in which not only the relative benefits, θ_{ij} , change over time. The nontrivial equilibrium abundances, \hat{N}_i , and thus the total gain from mutualistic interactions, $R_i(\hat{N}_i)$, changes as well. We found that a change in abundance over time does not have a strong effect on the performance of the indicator (Fig. S4.17). In comparison to data sets in which abundances stay (on average) the same, full network collapses are much less frequent when abundances increase and much more frequent when abundances decrease. Quite a large fraction of full network collapses is indicated accurately when abundances decrease. Cascading collapses may occur less frequently because all species experience a similar loss in resilience as a consequence of a decline in abundance. Another difference is that the length of the critical range tends to be a bit shorter when abundances in- or decrease.

In Fig. S4.21 and S4.22, we show that the indicator performs well, also when we apply our method to networks with different numbers of plant and pollinator species. Full network collapses become less common as the number of species increases, as well as the occurrence of cascading network collapses. An explanation for this effect of an increase in species number is that the loss in abundance due to competition with other species, $\sum_{j=1}^{S^{(A)}} c_{ij} N_j^{(A)} N_i^{(A)} - c_{ii} N_i^{(A)} N_i^{(A)}$, increases substantially as the number of species increases. Systems with many species may, therefore, be comparable with smaller networks in which interspecific competition is relatively strong. In those networks we also observed that full network collapses were less frequent. Increasing numbers of species did not have clear effect on the length of the critical range, nor on the fraction of the critical range in which the future state was indicated accurately by the slope of the indicator (Fig. S4.22). We did, however, found some effect on the skewness of time points projected on the first principal component. The frequency at which we found that points were skewed in the wrong direction increased as the number of species increases.

In Fig. S4.23, we show results for a more general model of competition and facilitation (see main text). The general behavior and performance of the indicator is similar to the results obtained with the mutualistic network model. The overall resilience of the networks tested seems a bit lower than the resilience of the mutualistic networks (this depends on parameter settings). To prevent networks from collapsing almost immediately, at $M \approx 0$, we chose a lower noise level of $\delta = 0.05$. This relatively low resilience may also explain the relatively high frequency of cascading collapses in networks of 10 species.

S4.8 Supplementary figures

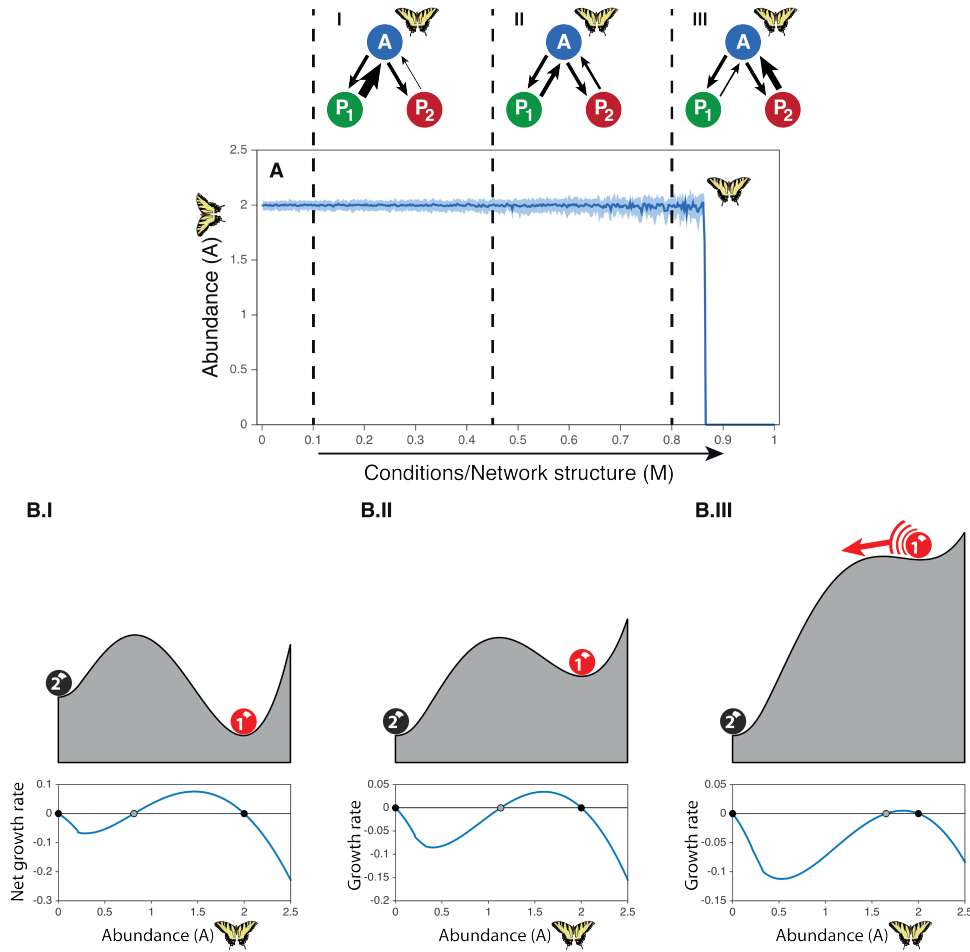


Figure S4.1: Changing conditions undermining the overall resilience of a small mutualistic network. The network consists out of one pollinator, A , and two plants species of which plant species P_1 is more saturated than plant species P_2 . For illustrative purposes, we assume plants to be in steady-state. **(A)** Time series of the pollinator species and the network at different conditions (I, II, and III). As indicated by the thickness of the network's arrows, changing conditions alter the relative mutualistic benefits, θ , such that the pollinator species becomes increasingly dependent on non-saturated plant species P_2 . This undermines the overall resilience of the network and leads to a full collapse of the network at which both plant species (not shown) and the pollinator species (shown) collapse to zero. **(B)** The net growth rate, dA/dt , and the stability landscape of the pollinator species at conditions I, II and III. As conditions change, the initial pristine state of the network, 1, is approached by a threshold, i.e. a hilltop in the stability landscape, and a small perturbation becomes sufficient to cause a regime shift towards fully collapsed state 2.




Figure S4.2: The slope of the small mutualistic network's stability landscape reflecting the speed at which pollinator abundances change, v_a , at different conditions, M . As in Fig. 2 of the main text, alternative stable states (balls), saddle points (yellow dots), and hilltops (grey dots) are surrounded by areas in which the landscape's slope, and thus the rate at which abundances change, is nearly zero (indicated in orange). Higher speeds (blue) are found further away from these points. The network recovers slowest from perturbations in the direction of the saddle point on the nearest threshold and slows down in the direction of the saddle point on the threshold approaching the network's initial pristine state. Changing conditions alter the shape of the network's stability landscape in a non-linear way. After a period in which there is almost no change ($M = [0, 0.31]$), the direction in which the network recovers slowest from perturbations (see yellow arrow) changes substantially from a direction that roughly indicates a full collapse to a direction indicating the future partially collapsed state of the network ($M = [0.31, 0.59]$). After the network's future state comes into existence, the network slows down dramatically in approximately the same direction ($M = [0.59, 0.87]$).

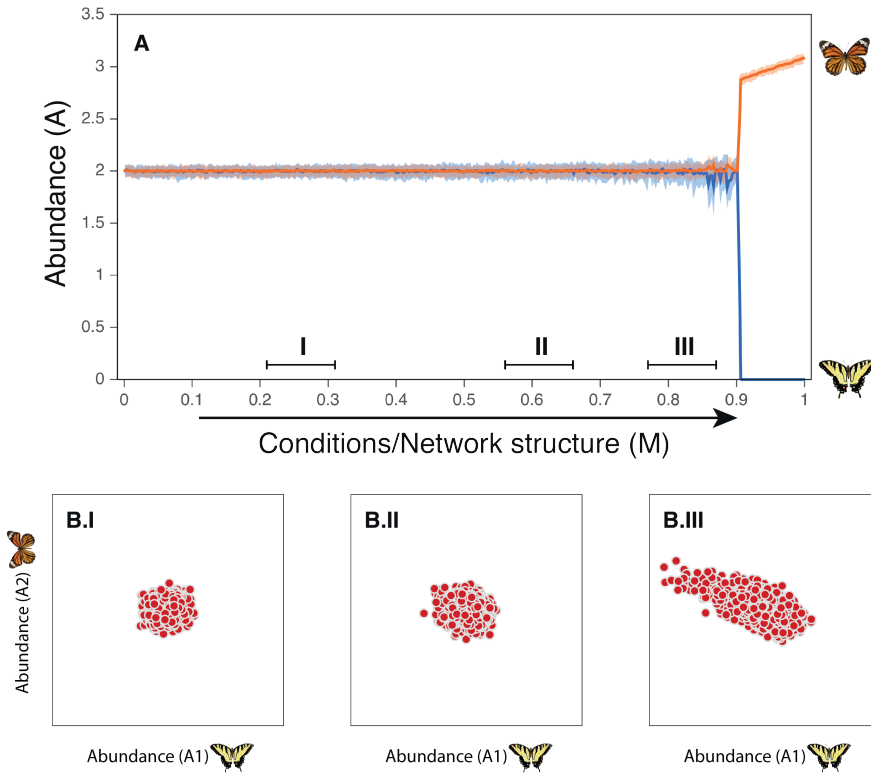


Figure S4.3: Example of a time series in which the small mutualistic network in Appendix S4.2 approaches a tipping point. Conditions at window I, II and III correspond to the conditions for which stability landscapes are shown in Fig. 1 of the main text. **(A)** At the tipping point ($M \approx 0.9$) one pollinator species collapses to extinction, while the other gains in abundance. **(B)** The distribution of points in the network's phase plane representing the abundances of species at different moments in time for time window I, II and III (see A). Far from the tipping point, in window I and II, deviations from the species' mean abundances are relatively small. Close to the tipping point, in window III, the distribution of points in the network's phase space is highly asymmetrical. Deviations from the mean abundances in time window III usually involve a simultaneous increase in the abundance of species A1 and a relatively larger decrease in the abundance of species A2, suggesting that this will also be the direction in which the network will shift once a threshold is passed.

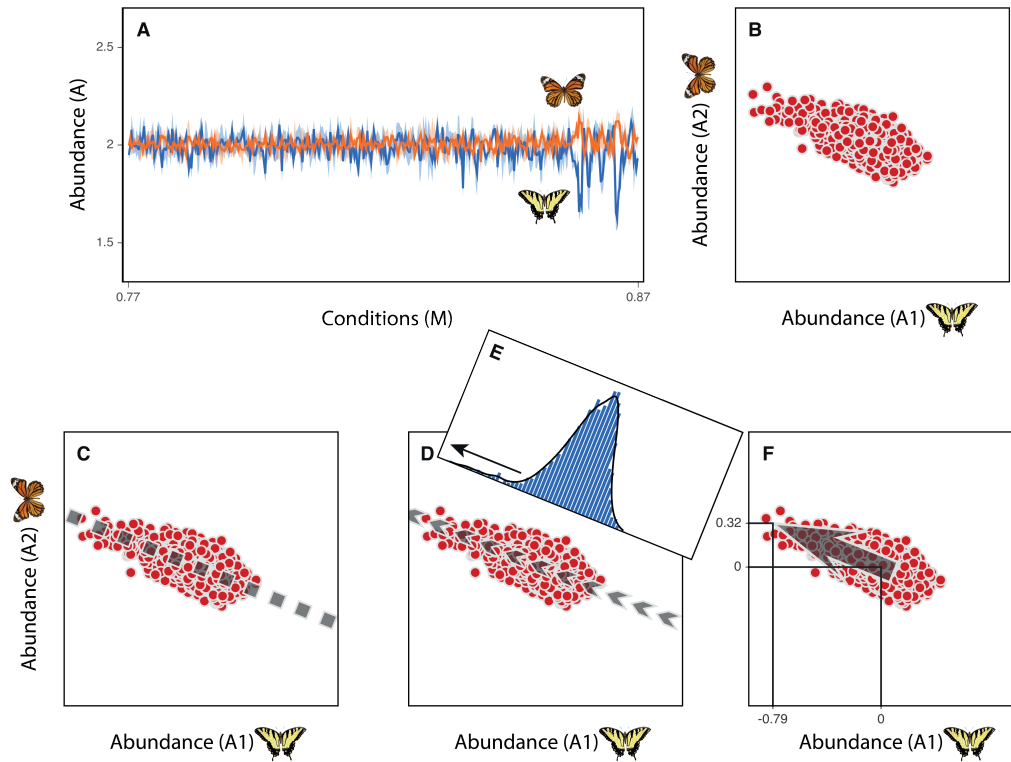


Figure S4.4: The measures of asymmetry together forming our indicator as they were determined for window III in Fig. S4.3. **(A)** Time series of the two pollinator species in the moving window. **(B)** Time points, representing species abundances at different moments in time, in the phase plane of the network. **(C)** The first principal component (grey dotted line) corresponding to the line in the phase plane along which variance is highest. **(D)** Direction along the first principal component (grey arrows) in which time points deviate the most from the species' mean abundance, i.e. the direction in which time points projected on the first principal component are skewed. **(E)** Distribution of the projected time points. **(F)** The indicator, corresponding to a vector in the phase plane of the network (grey arrow). The two components of this vector correspond to the species 'scores on the indicator'. In this example, we found a large negative score (-0.79) indicating a relatively large decline in abundance for the pollinator on the x-axis and a relatively smaller positive score (0.32) indicating a relatively smaller increase in abundance for the pollinator on the y-axis. The length of the indicator corresponds to the amount of variance explained by the first principal component.




Figure S4.5: Cumulative distribution function of the angle between two random vectors. **(A)** Cumulative distribution function when these vectors have three dimensions. **(B)** Cumulative distribution function when these vectors have ten dimensions. As can be seen from the distributions, the probability of finding an angle of, for example, 40 degrees or less is much smaller in a high dimensional system. Cumulative distribution functions are determined with the help of the probability density function in ref. Cai et al. (2013).




Figure S4.6: Example of **(A)** a highly resilient mutualistic network and **(B)** a network with a low resilience. Plant (circles) and pollinator species (squares) are ordered from highly saturated (green/left) to non-saturated (red/right). The thickness of the lines between nodes indicates relative mutualistic benefit θ_{ij} . In the highly resilient network species receive most of their resources from highly saturated species, while this is not the case in the network with a low overall resilience. The resilience of a network is undermined when relative benefits are changed from the situation in **A** to the situation in **B**.




Figure S4.7: The critical range (grey band) in which the indicator's magnitude increases significantly and the fraction of this range in which the indicator's similarity to the observed shift in abundance is larger than 0.99. In the here shown example, the length of the critical period is $0.8 - 0.46 = 0.34$. The slope of the indicator accurately indicates the future state, i.e. similarity is > 0.99 , during a fraction of $0.29 / 0.34 = 0.85$ of this period. The full time series is shown in Fig. 2.

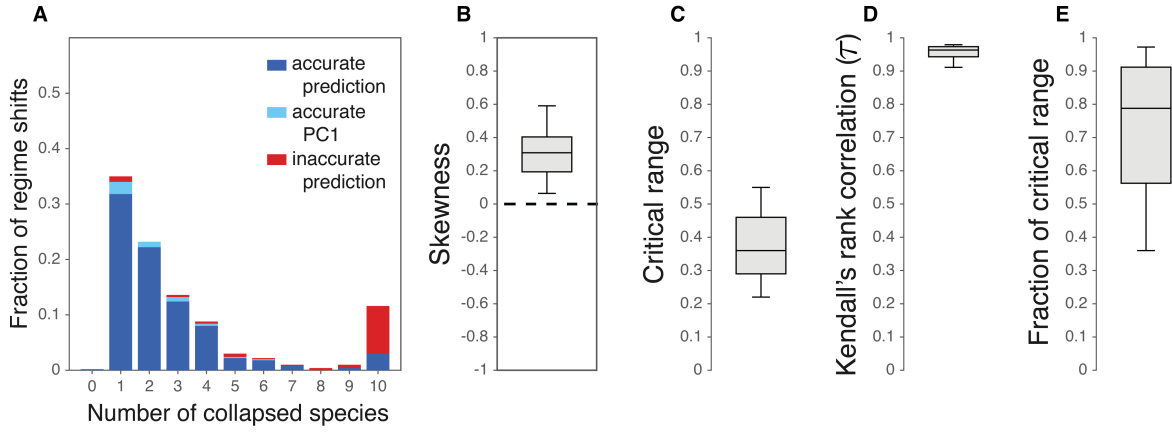


Figure S4.8: Overall statistics on the performance of the indicator when competitive interaction strengths, c_{ij} are taken from $\sim U(0.02, 0.08)$. **(A)** The performance of the indicator for different numbers of collapsed species. The fraction of regime shifts for which the change in abundance was not well indicated is shown in red. The fraction accurately indicated by the first principal component, but not by the direction in which time points are skewed is shown in light blue. Fully accurate predictions are indicated in dark blue. **(B)** The skewness of time points projected on the first principal component. A positive skewness means that time points were skewed in the direction of the network's future state. **(C)** The length of the critical range in which the indicator's magnitude increases significantly. **(D)** Kendall's rank correlation, τ , as determined for the critical range. **(E)** The fraction of the critical range in which the slope of the indicator accurately indicates the future state, i.e. in which the similarity between the first principal component and the observed shift in abundance is > 0.99 . Results in panels (B-E) are shown for regime shifts that were accurately indicated by the first principal component. Box plots show the median and the upper and lower quartiles. Whiskers correspond to the 9th and the 91st percentile.

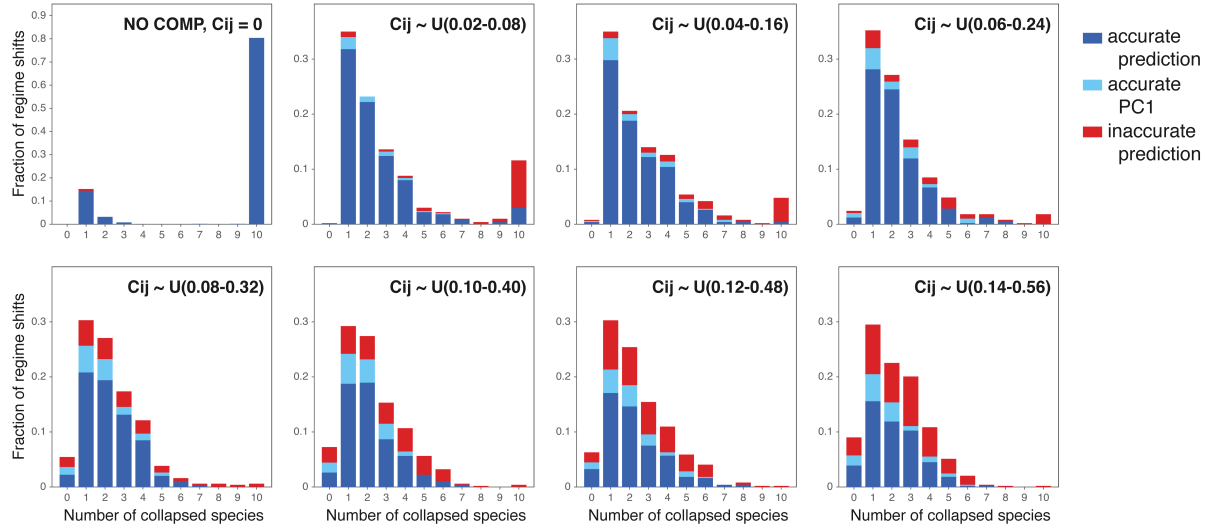


Figure S4.9: The number of pollinator species collapsing to extinction as observed in data sets of 1000 regime shifts. Each panel shows results when sampling competitive interaction strengths from a different parameter range (see ranges indicated). In the extreme case where there was no competition (top left panel), we found almost exclusively full network collapses (i.e. all ten pollinator species collapsed to extinction). As the strength of competition increases, full network collapses become less frequent. Partial network collapses tend to be small independent of the strength of competition, i.e. the most common partial collapse led to the extinction of only one single pollinator species. The fraction of regime shifts for which the change in abundance was not well indicated is shown in red. The fraction accurately indicated by the first principal component, i.e. the slope of the indicator is accurate, but not by the direction in which time points are skewed is shown in light blue. Fully accurate predictions are indicated in dark blue.

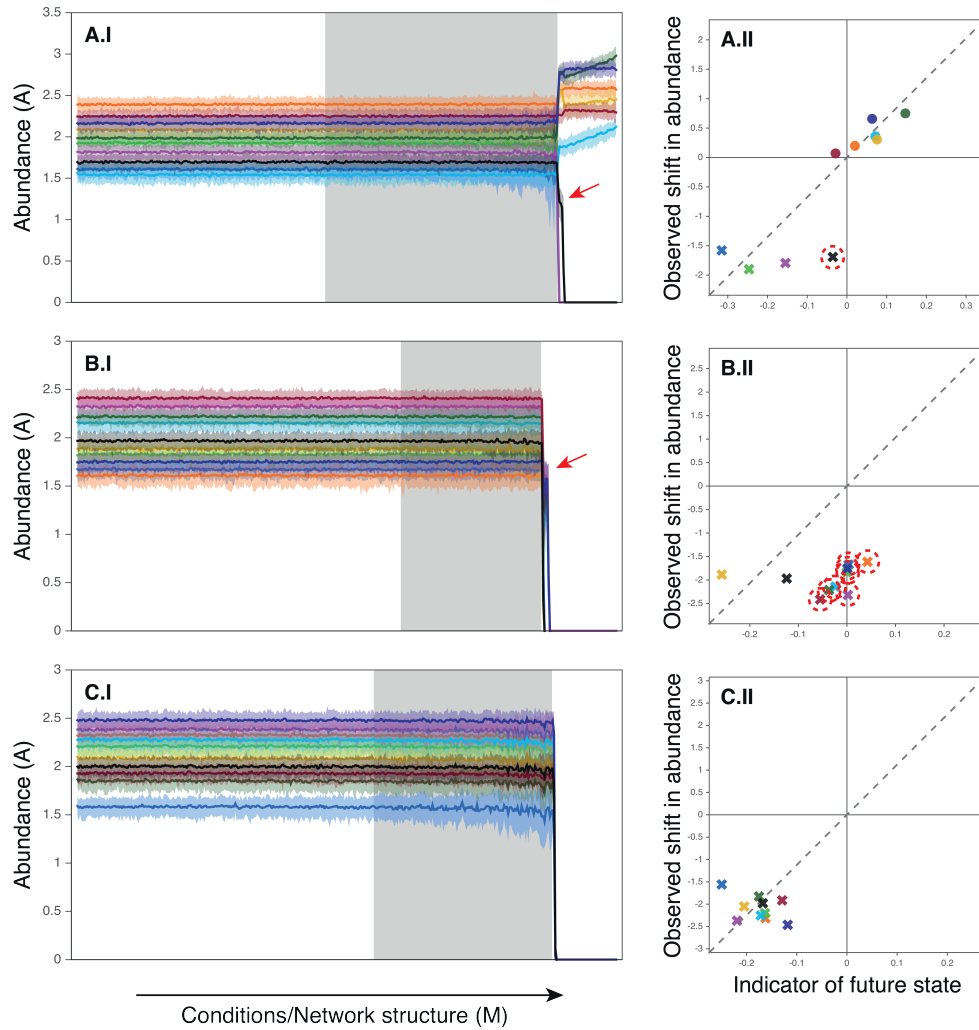


Figure S4.10: Two cascading collapses and one immediate collapse. **(A)** Example of a cascading collapse that eventually leads to the collapse of four pollinator species. Three species (blue, green and purple) collapse to extinction rapidly. A fourth (black) species collapses as well, but remains for a short while at a lower abundance before collapsing to extinction (red arrow, A.I). Out of the four species that collapse to extinction, the black species is also the one for which the indicated loss in abundance is smallest (red circle, A.II). **(B)** Example of a cascading collapse that eventually leads to a full collapse of the network (i.e. the most common outcome of a cascading collapse). Two species (black and yellow) collapse to extinction rapidly. The other species collapse as well, but remain for a short while at a lower abundance before collapsing to extinction (red arrow, b.I). The indicated loss in abundance of the rapidly collapsing species is much bigger than the loss indicated for the species that collapse a bit later (red circles, B.II). **(C)** Example of a full network collapse that was accurately indicated. All species collapse at approximately the same time (C.I). All species were indicated to lose in abundance (C.II).

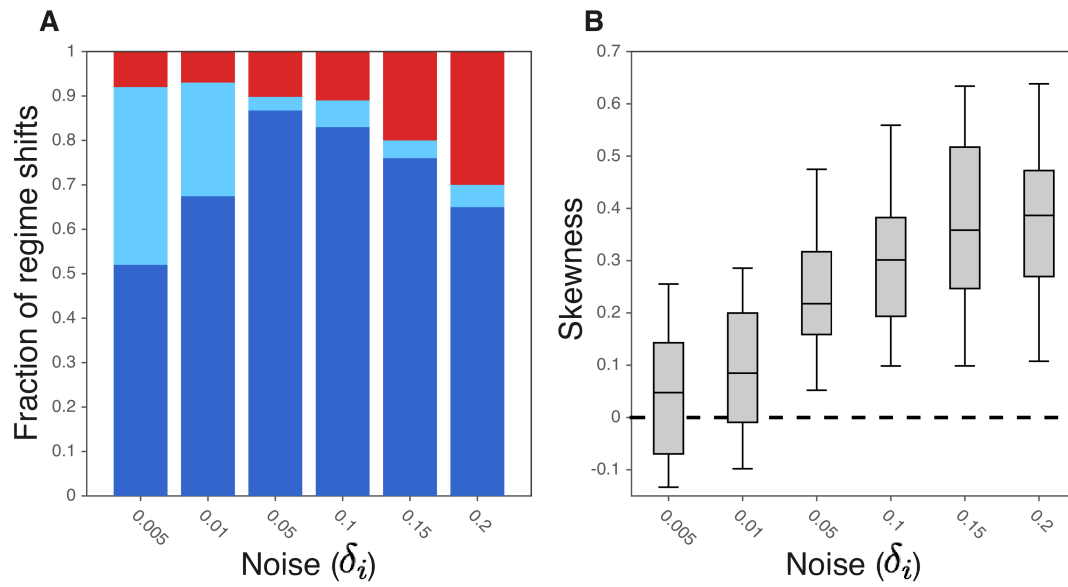


Figure S4.11: Performance of the indicator for different noise levels (noise levels, ϵ_i , are indicated on the x-axis). **(A)** The fraction of accurately indicated regime shifts (dark blue), the fraction accurately indicated by the first principal component, i.e. the slope of the indicator is accurate, but not by the direction in which time points are skewed (light blue), and the fraction of inaccurately indicated regime shifts (red). **(B)** The skewness of time points projected on the first principal component. A positive skewness means that time points are skewed in the direction of a network's future state. The skewness is shown for regime shifts that were accurately indicated by the first principal component.

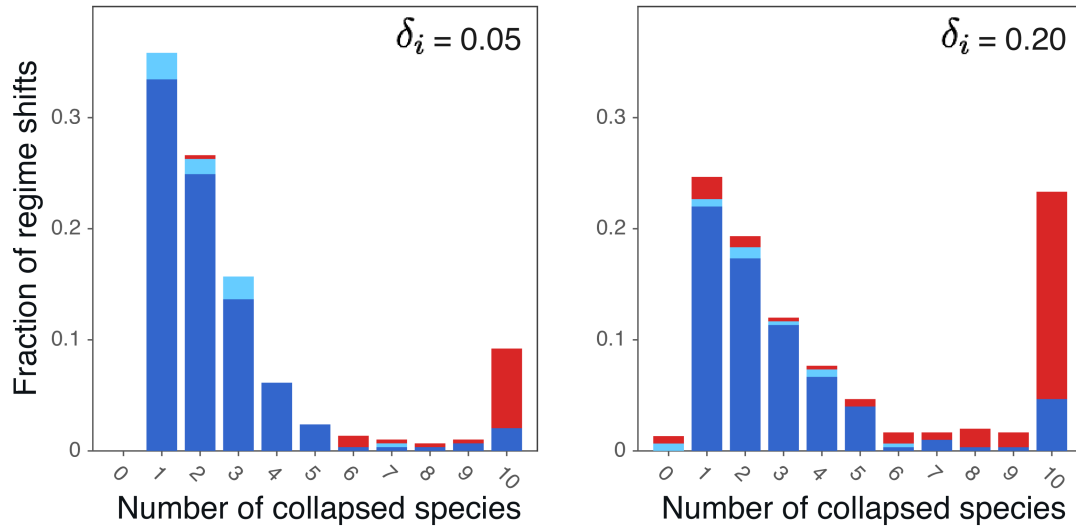


Figure S4.12: The number of pollinator species collapsing to extinction as observed in data sets of 1000 regime shifts when noise levels are low (left panel, $\epsilon_i = 0.05$) and when noise levels are high (left panel, $\epsilon_i = 0.2$). Full network collapses were found to occur more frequently when noise levels are high. The fraction of regime shifts for which the change in abundance was not well indicated is shown in red. The fraction accurately indicated by the first principal component, i.e. the slope of the indicator is accurate, but not by the direction in which time points are skewed is shown in light blue. Fully accurate predictions are indicated in dark blue.

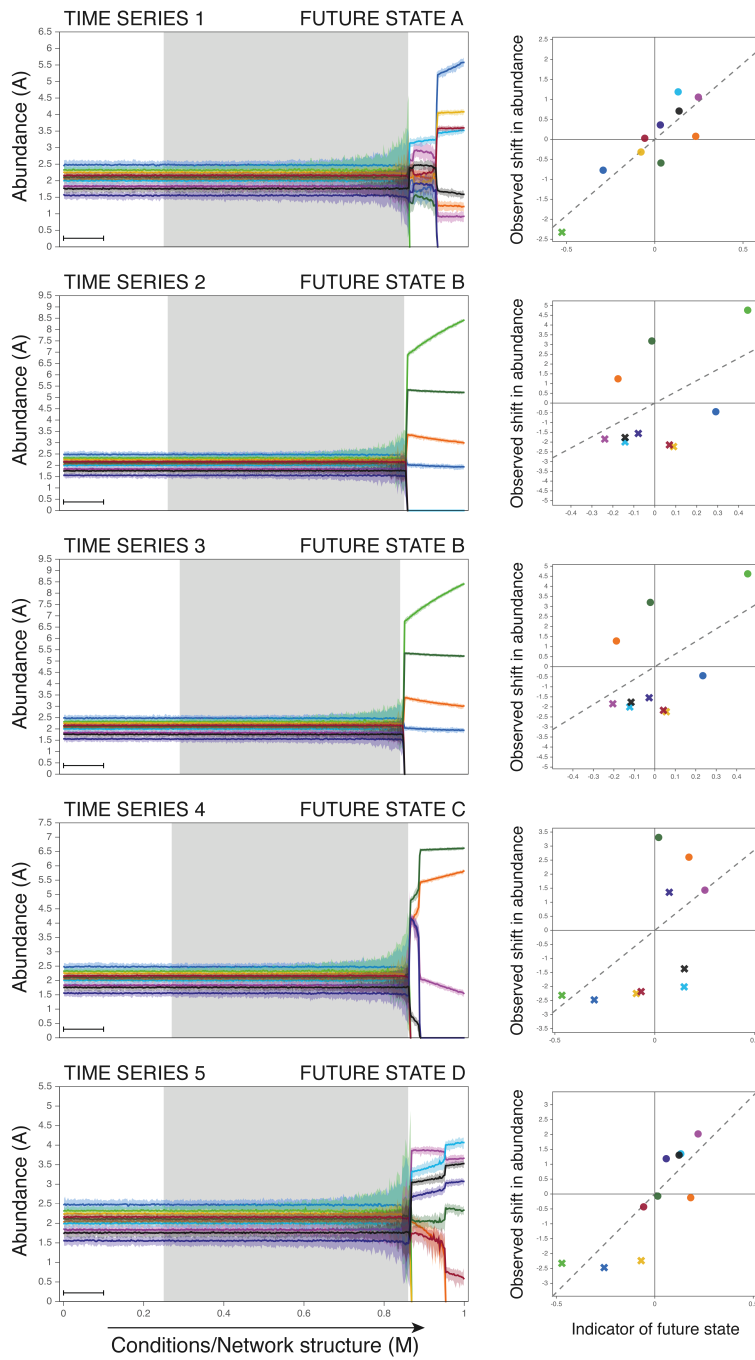


Figure S4.13: Five time series of a network that shows ‘unpredictable’ behavior. Even though the resilience of the network is undermined in the exact same way, the network may shift to several alternative future states. The future state of the network is determined by the only stochastic element in our model; the small-scale perturbations to which the network is permanently subjected. We found that this network may shift to (at least) four different future states (i.e. Euclidean distance between future states > 1.5). Of these future states, future state A and D are well indicated by the indicator (i.e. similarity > 0.99). The future state of the network is the same only in time series 2 and 3.

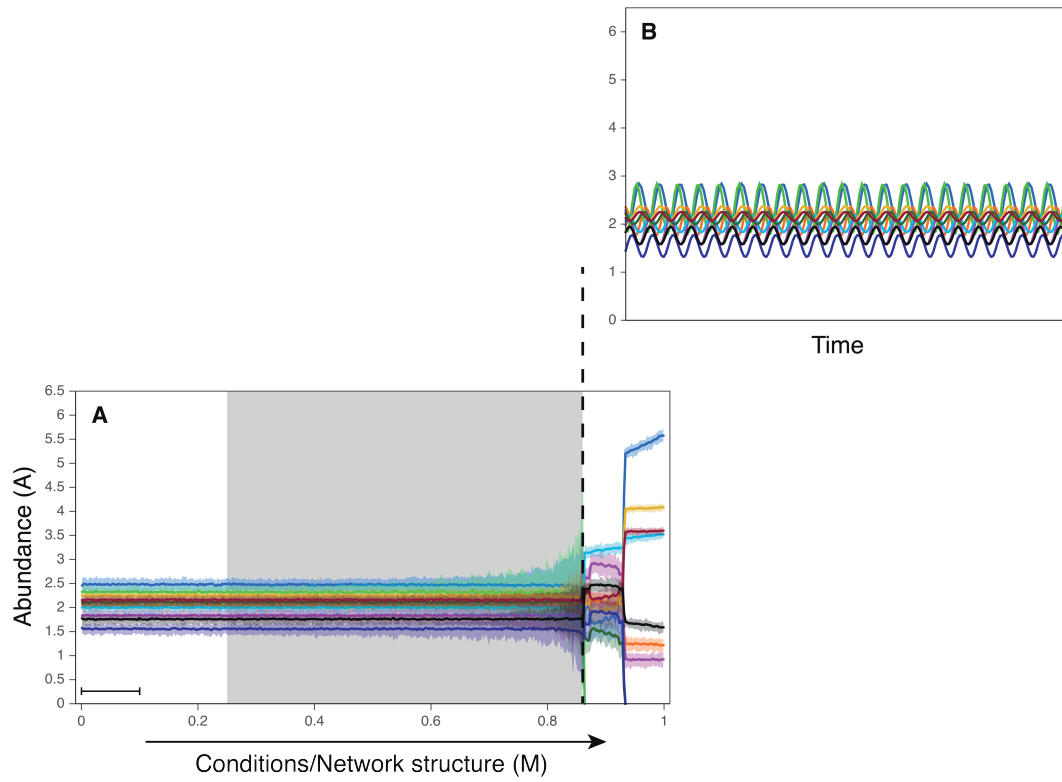


Figure S4.14: Example of a network approaching a supercritical Hopf bifurcation. **(A)** Time series of the network as conditions change. **(B)** Time series at fixed conditions just after the bifurcation point when assuming there are no external perturbations ($\epsilon = 0$). As can be seen from the dynamics we are dealing with a limit cycle. In the presence of external perturbations, the fluctuations caused by these dynamics are amplified and lead to a partial collapse of the network.

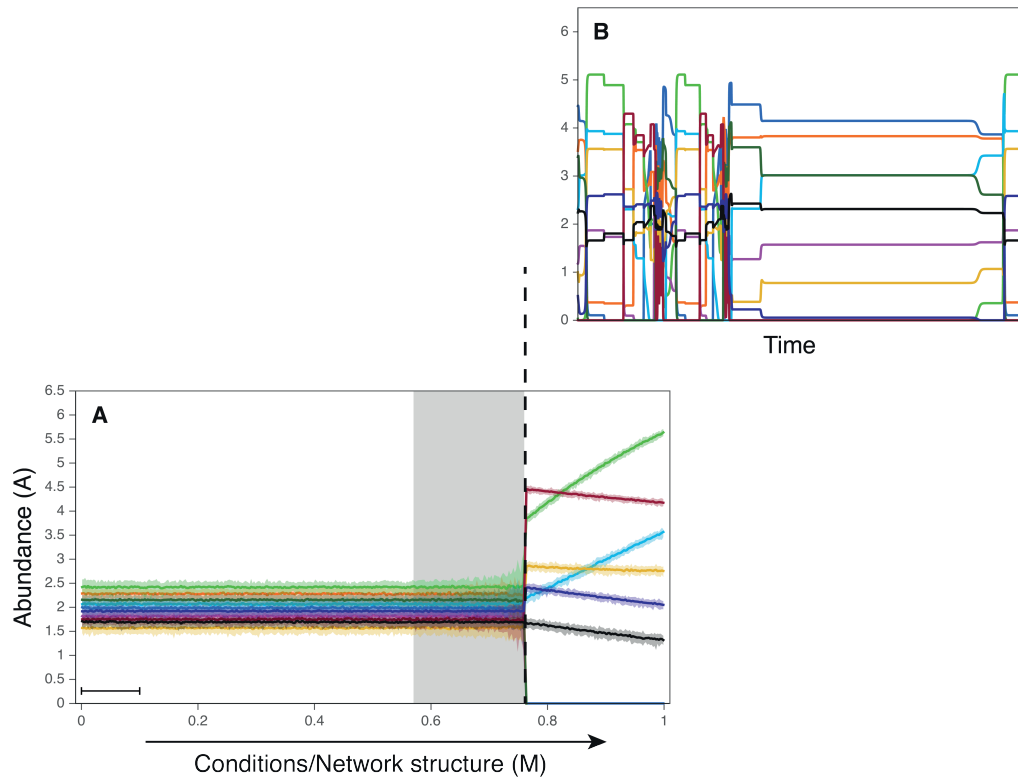


Figure S4.15: Example of a network approaching a subcritical Hopf bifurcation. **(A)** Time series of the network as conditions change. **(B)** Time series at fixed conditions just after the bifurcation point when assuming there are no external perturbations ($\epsilon = 0$) and when excluding the condition that populations of a size smaller than 0.001 have a zero growth rate ($dN/dt=0$). As can be seen from the dynamics we are dealing with chaotic/heteroclinic dynamics. The condition that populations of a size smaller than 0.001 have a zero growth rate leads to a partial collapse of the network. Which species are the first to cross this threshold is strongly influenced by the stochastic perturbations that are constantly disturbing the network.




Figure S4.16: The probability of finding a stable solution at initial conditions, $M = 0$, when sampling competitive interaction strengths from different parameter ranges (ranges are indicated on the x-axis). As the strength of competition increases, it becomes increasingly difficult to find a stable solution. When there is no competition between species, the probability of finding a stable solution is nearly one. For the highest competition level we tested, i.e. (0.14,0.56), this probability was below 0.01. Results are shown for networks of 10 plants and 10 pollinators as described in Appendix S4.5.

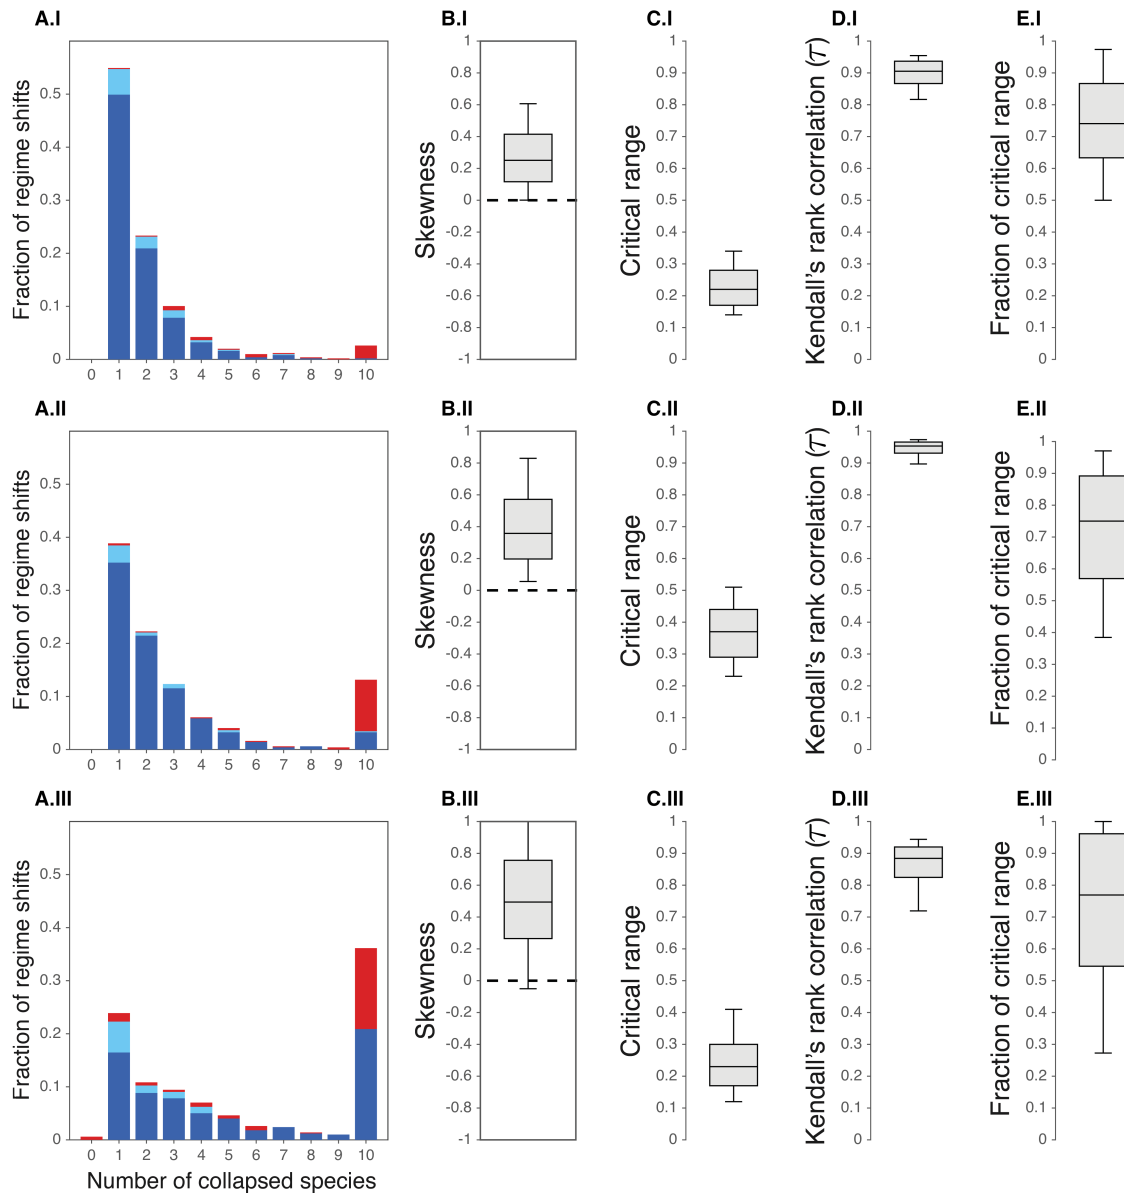


Figure S4.17: Overall statistics on the performance of the indicator when nontrivial equilibrium tend to increase (I), change but stay the same on average (II), and when abundances tend to decrease (III). Results are shown for data sets of 1000 regime shifts. As in Fig. S4.8 we show: (A) the performance of the indicator for different numbers of collapsed species, (B) the skewness of time points projected on the first principal component, (C) the length of the critical range in which the indicator's magnitude increases significantly, (D) Kendall's rank correlation, τ , as determined for the critical range, and (E) the fraction of the critical range in which the slope of the indicator accurately indicates the future state. Box plots show the median and the upper and lower quartiles. Whiskers correspond to the 9th and the 91st percentile.

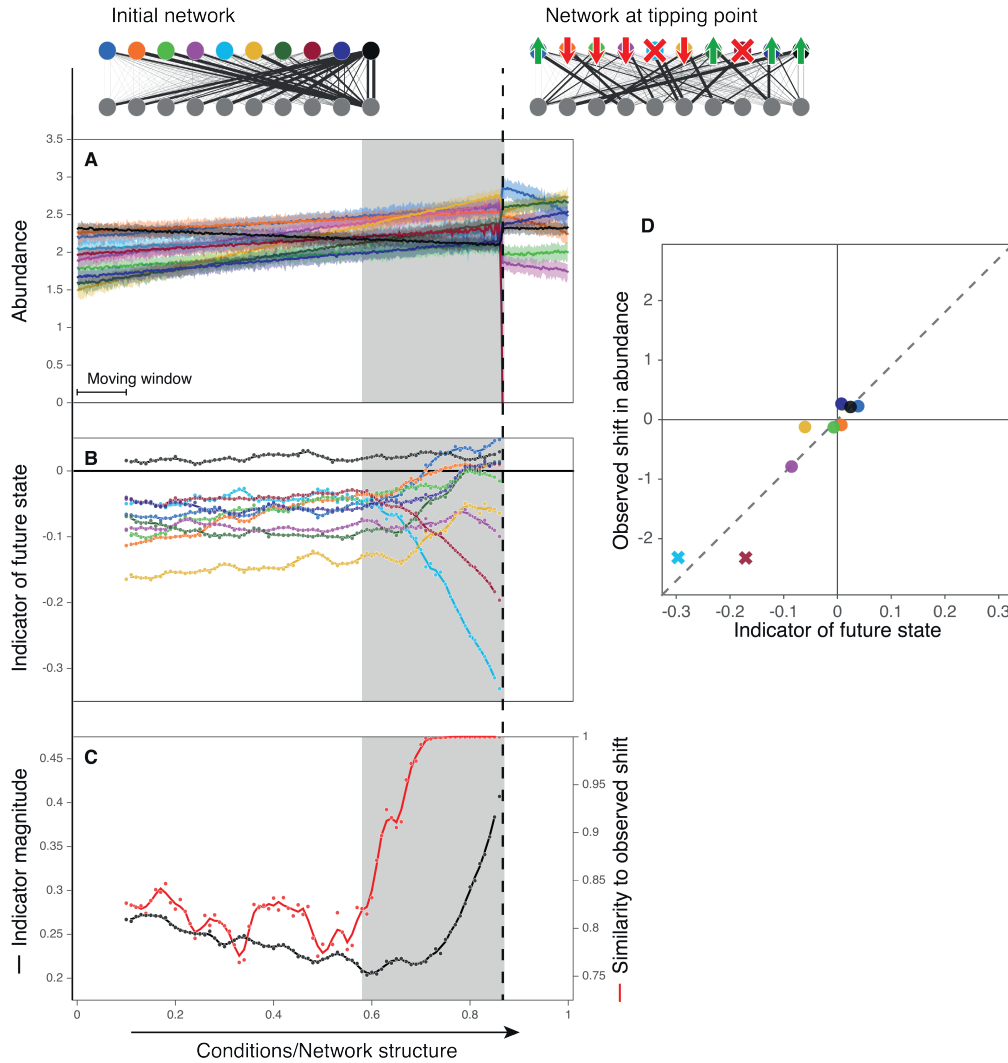


Figure S4.18: Directional slowing down when abundances tend to increase. **(A)** Time series of species belonging to one set of a bipartite mutualistic network, i.e. the pollinators. At the tipping point two species collapse to extinction (red and light blue). **(B)** The indicator of the future state measuring the direction in which fluctuations are distributed asymmetrically. **(C)** The magnitude of the indicator, reflecting the extent in which fluctuations are distributed asymmetrically, plotted together with the accuracy measured as the similarity between its direction and the observed shift in abundance. Grey bands indicate the period in which the indicator's magnitude increases significantly. **(D)** The observed changes in abundance versus the scores on the indicator just before the tipping point. Extinct species are indicated with crosses. The initial network, at $M=0$, is the same as in Fig. 4.2.

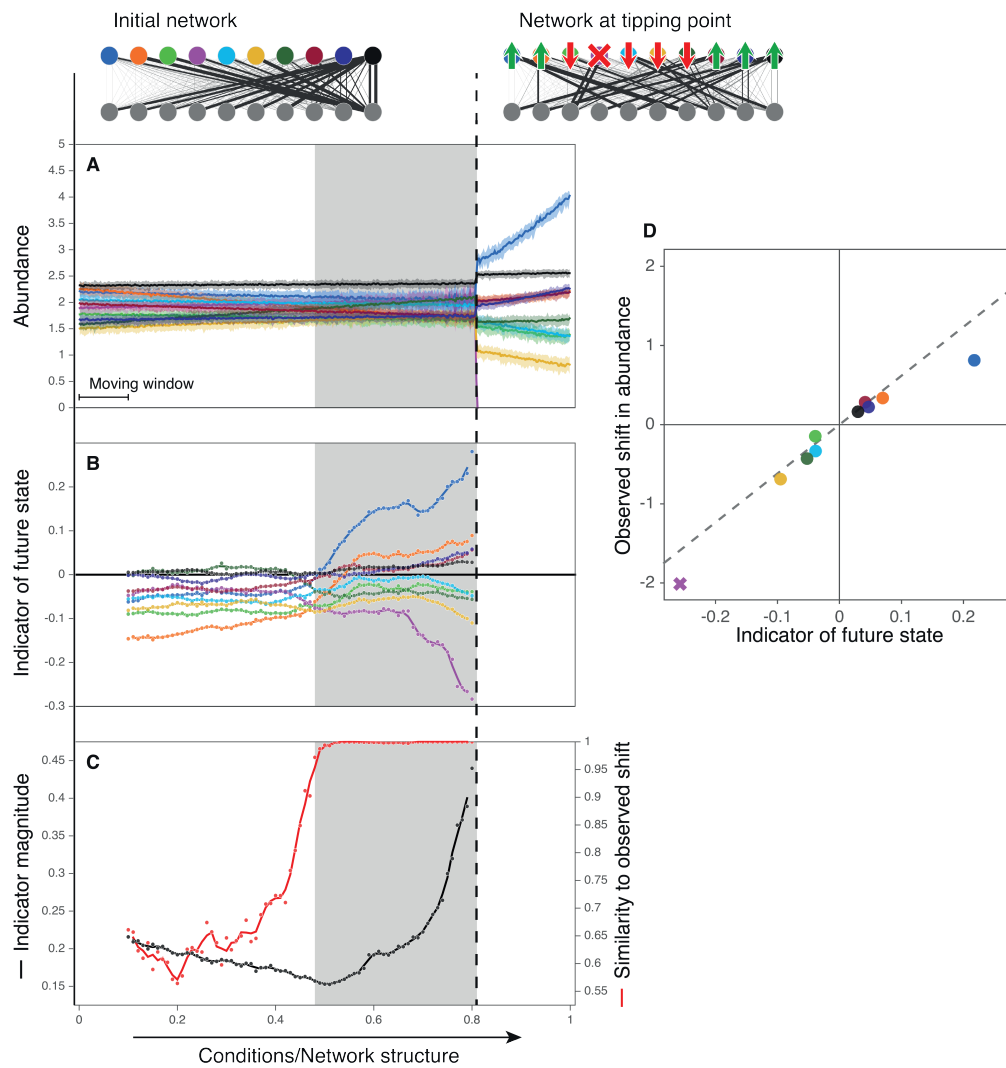


Figure S4.19: Directional slowing down when abundances change, but stay the same on average. **(A)** Time series of species belonging to one set of a bipartite mutualistic network, i.e. the pollinators. At the tipping point a single species collapses to extinction (purple). **(B)** The indicator of the future state measuring the direction in which fluctuations are distributed asymmetrically. **(C)** The magnitude of the indicator, reflecting the extent in which fluctuations are distributed asymmetrically, plotted together with the accuracy measured as the similarity between its direction and the observed shift in abundance. Grey bands indicate the period in which the indicator's magnitude increases significantly. **(D)** The observed changes in abundance versus the scores on the indicator just before the tipping point. Extinct species are indicated with crosses. The initial network, at $M=0$, is the same as in Fig. 4.2 and Fig. S4.18.

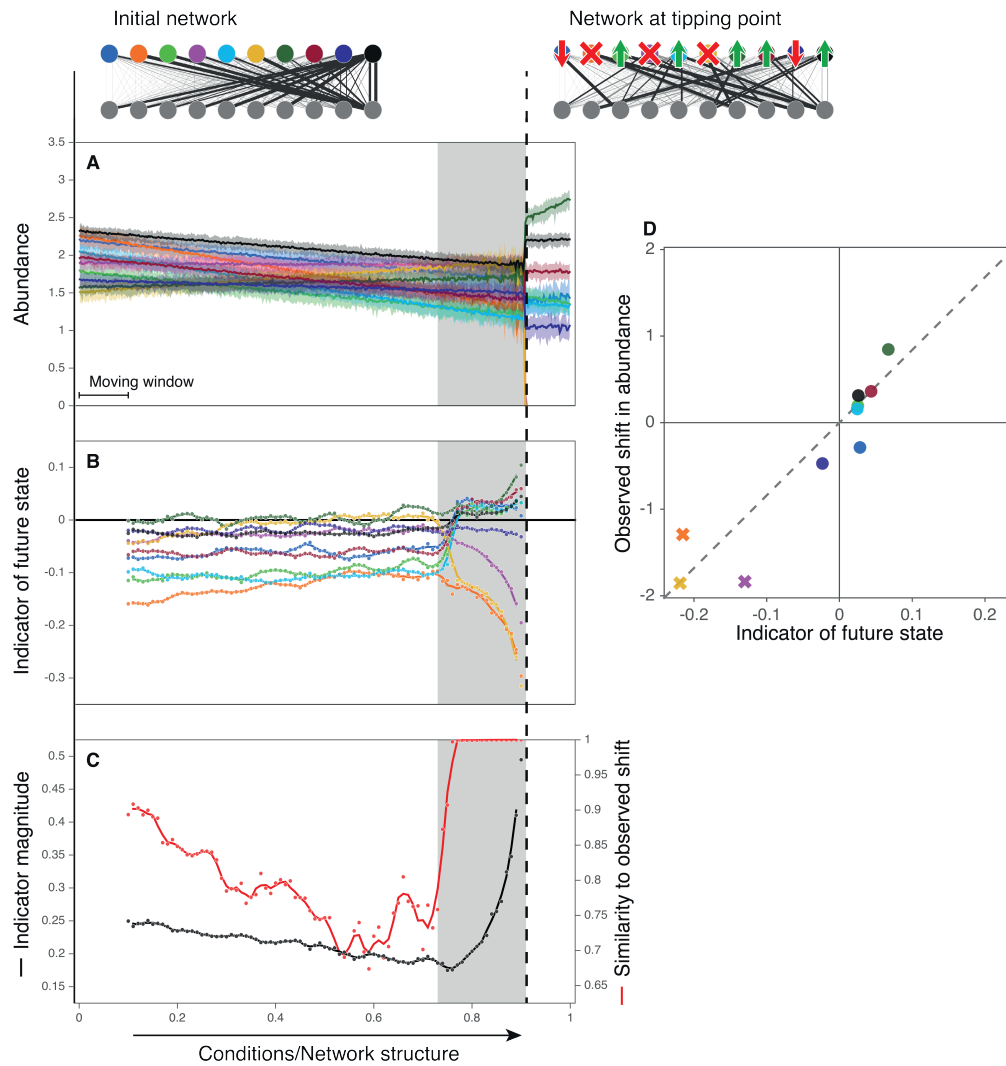


Figure S4.20: Directional slowing down when abundances tend to decrease. **(A)** Time series of species belonging to one set of a bipartite mutualistic network, i.e. the pollinators. At the tipping point three species collapse to extinction (yellow, purple and orange). **(B)** The indicator of the future state measuring the direction in which fluctuations are distributed asymmetrically. **(C)** The magnitude of the indicator, reflecting the extent in which fluctuations are distributed asymmetrically, plotted together with the accuracy measured as the similarity between its direction and the observed shift in abundance. Grey bands indicate the period in which the indicator's magnitude increases significantly. **(D)** The observed changes in abundance versus the scores on the indicator just before the tipping point. Extinct species are indicated with crosses. The initial network, at $M=0$, is the same as in Fig. 4.2, Fig. S4.19, and Fig. S4.18.

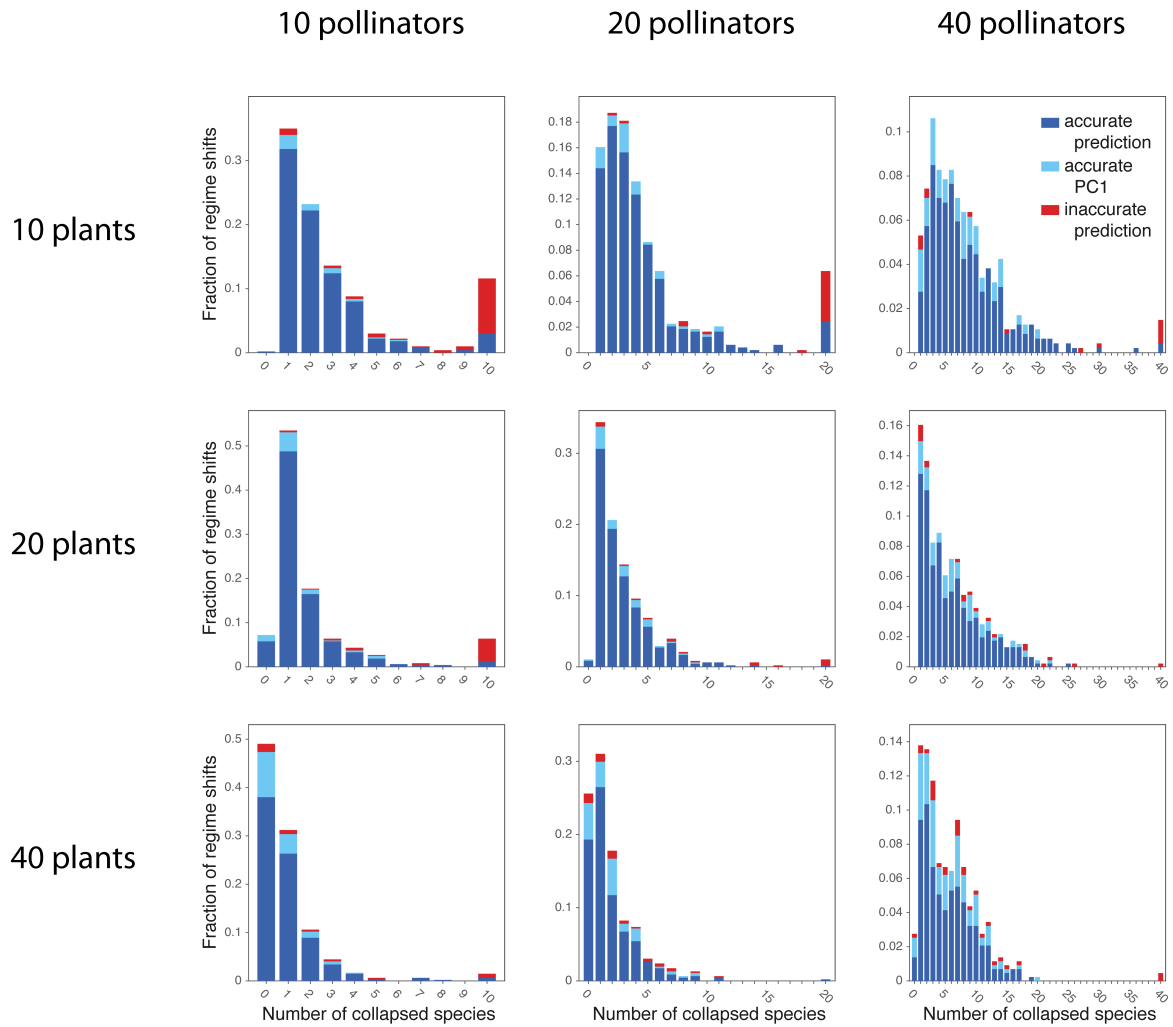


Figure S4.21: The performance of the indicator for networks of different size, i.e. for different number of plant (rows) and pollinator species (columns). Each panel shows the number of pollinator species collapsing to extinction as observed in data sets of 1000 regime shifts. The fraction of regime shifts for which the change in abundance was not well indicated is shown in red. The fraction accurately indicated by the first principal component, but not by the direction in which time points are skewed is shown in light blue. Fully accurate predictions are indicated in dark blue.




Figure S4.22: The performance of the indicator for networks of different size, i.e. for different number of plant (rows) and pollinator species (columns). Each panel shows the skewness of time points projected on the first principal component, the length of the critical range in which the indicator's magnitude increases significantly, Kendall's rank correlation, τ , as determined for the critical range, and the fraction of the critical range in which the slope of the indicator accurately indicates the future state. Results are shown for regime shifts that were accurately indicated by the first principal component. Results are shown for data sets of 1000 regime shifts. Box plots show the median and the upper and lower quartiles. Whiskers correspond to the 9th and the 91st percentile.

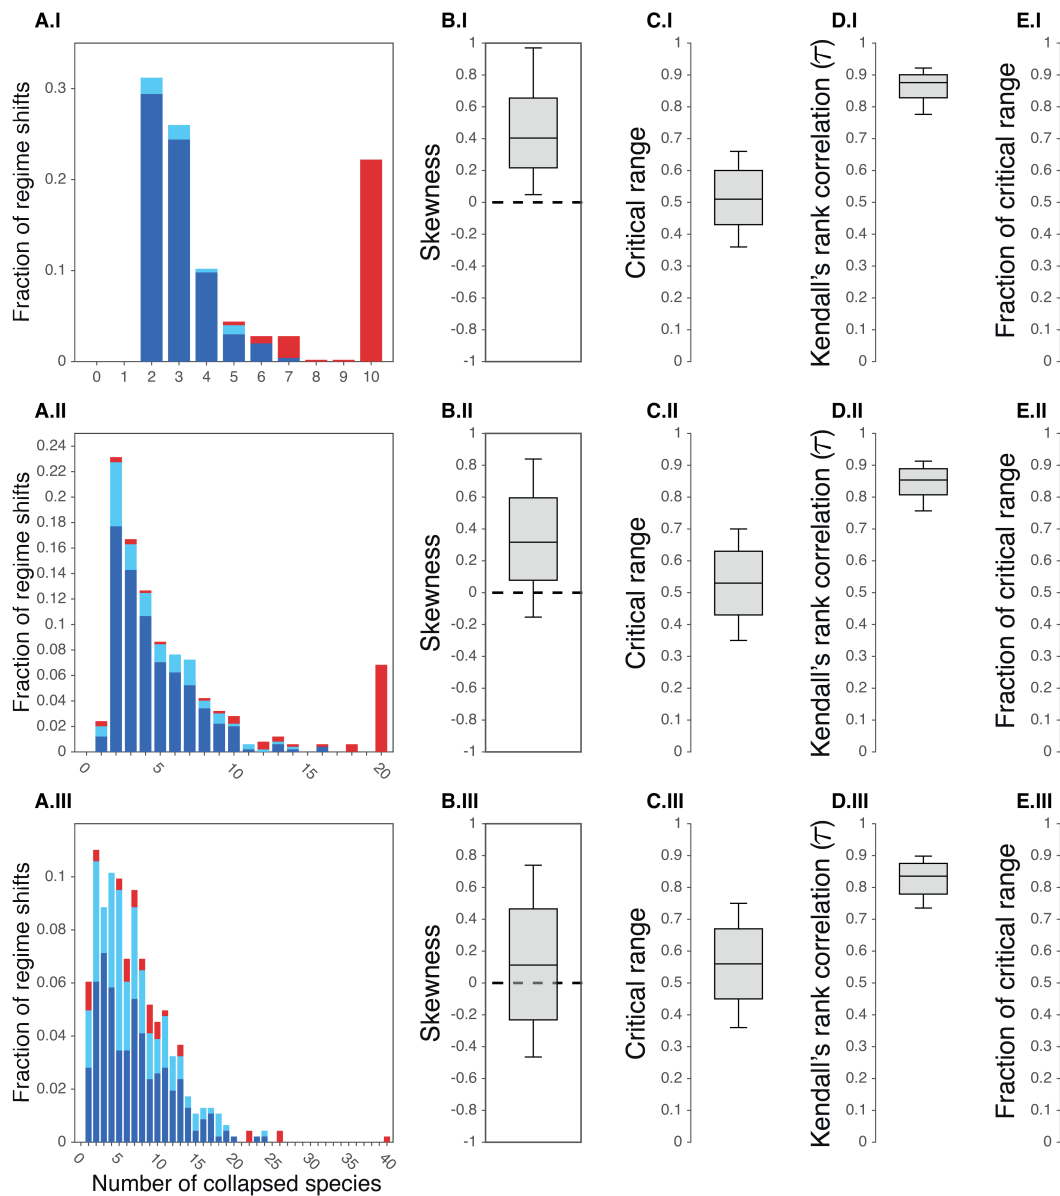


Figure S4.23: Overall statistics on the performance of the indicator when predicting the future state of a more general model of competition and facilitation. Results are shown for networks of 10 (I), 20 (II) and 40 species (III). (A) The performance of the indicator for different numbers of collapsed species. (B) The skewness of time points projected on the first principal component. (C) The length of the critical range in which the indicator's magnitude increases significantly. (D) Kendall's rank correlation, τ , as determined for the critical range. (E) The fraction of the critical range in which the slope of the indicator accurately indicates the future state, i.e. in which the similarity between the first principal component and the observed shift in abundance is > 0.99 . Results in panels (B-E) are shown for regime shifts that were accurately indicated by the first principal component. Box plots show the median and the upper and lower quartiles. Whiskers correspond to the 9th and the 91st percentile.

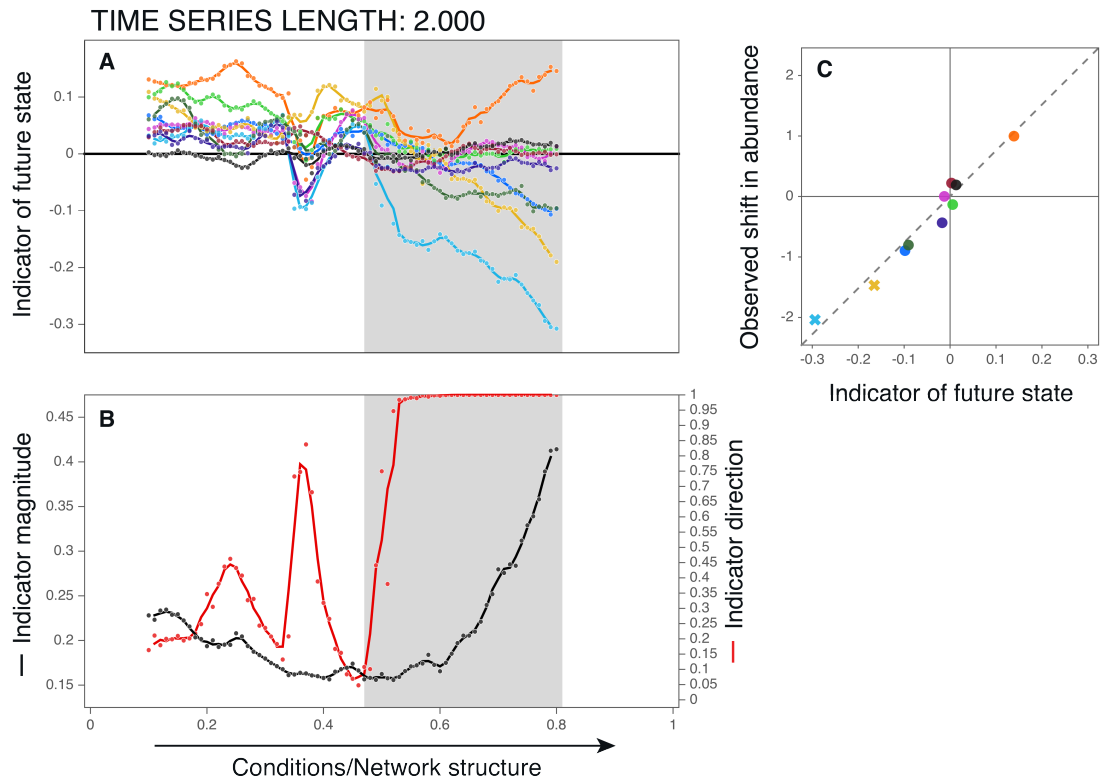


Figure S4.24: Directional slowing down as detected by the indicator when the total length of a time series is 2,000 time steps. **(A)** The indicator of the future state measuring the direction in which fluctuations are distributed asymmetrically. **(B)** The magnitude of the indicator, reflecting the extent in which fluctuations are distributed asymmetrically, plotted together with the accuracy measured as the similarity between its direction and the observed shift in abundance. Grey bands indicate the period in which the indicators magnitude increases significantly. **(C)** The observed changes in abundance versus the scores on the indicator just before the tipping point. Extinct species are indicated with crosses. The initial network, at $M = 0$, and the way in which this network is affected by changing environmental conditions, M , is the same as in Fig. 4.2. Changes in the direction and magnitude of the indicator are determined with a rolling window of 10% of the entire time series, i.e. 200 out of 2,000 time steps.

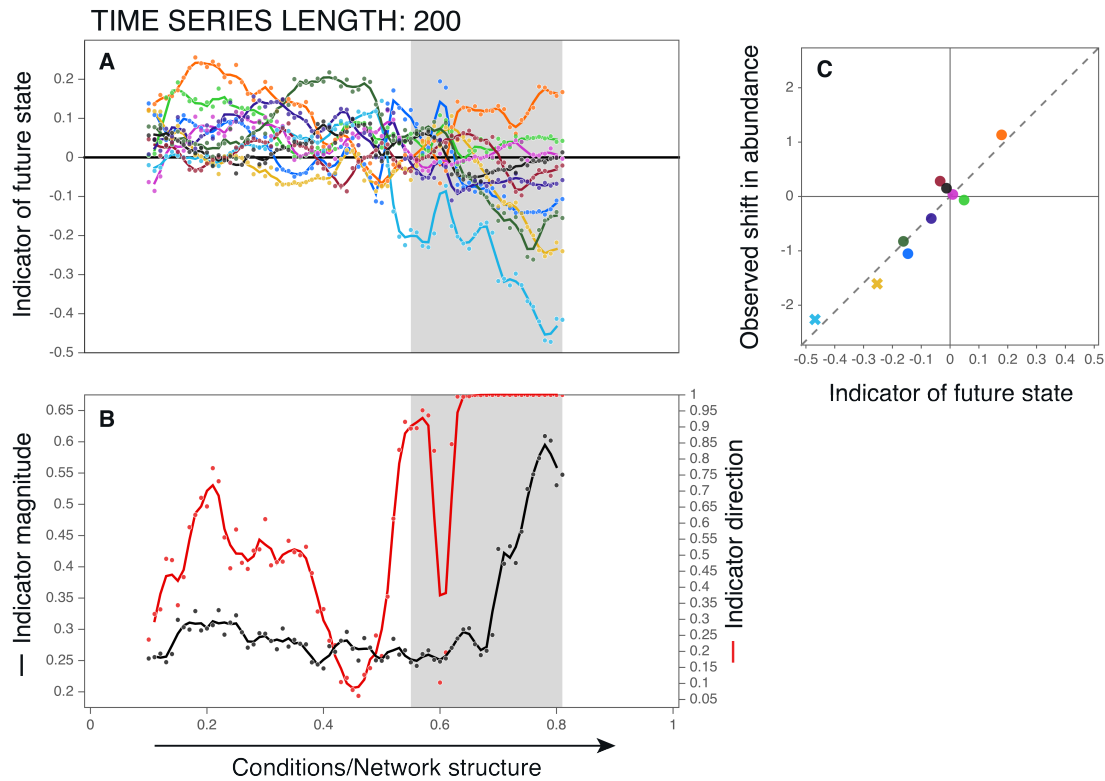


Figure S4.25: Directional slowing down as detected by the indicator when the total length of a time series is 200 time steps. **(A)** The indicator of the future state measuring the direction in which fluctuations are distributed asymmetrically. **(B)** The magnitude of the indicator, reflecting the extent in which fluctuations are distributed asymmetrically, plotted together with the accuracy measured as the similarity between its direction and the observed shift in abundance. Grey bands indicate the period in which the indicators magnitude increases significantly. **(C)** The observed changes in abundance versus the scores on the indicator just before the tipping point. Extinct species are indicated with crosses. The initial network, at $M = 0$, and the way in which this network is affected by changing environmental conditions, M , is the same as in Fig. 4.2. Changes in the direction and magnitude of the indicator are determined with a rolling window of 10% of the entire time series, i.e. 20 out of 200 time steps.

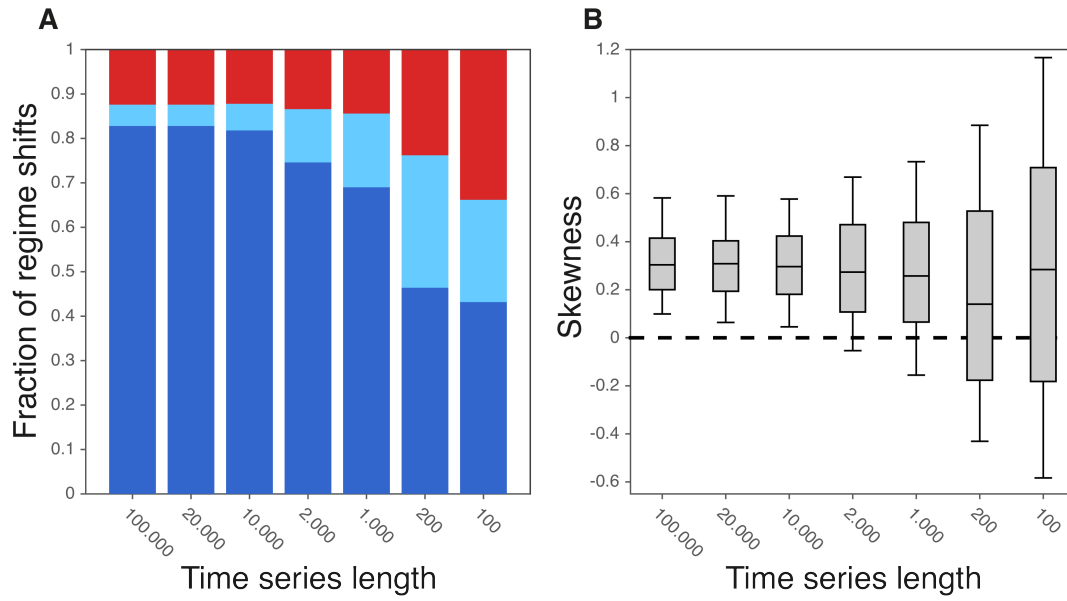


Figure S4.26: Performance of the indicator when time series have a different length (lengths are indicated on the x-axis). **(A)** The fraction of accurately indicated regime shifts (dark blue), the fraction accurately indicated by the first principal component, i.e. the slope of the indicator is accurate, but not by the direction in which time points are skewed (light blue), and the fraction of inaccurately indicated regime shifts (red). **(B)** The skewness of time points projected on the first principal component. A positive skewness means that time points are skewed in the direction of a network's future state. The skewness is shown for regime shifts that were accurately indicated by the first principal component. Changes in the direction and magnitude of the indicator are determined with a rolling window of 10% of the entire time series.

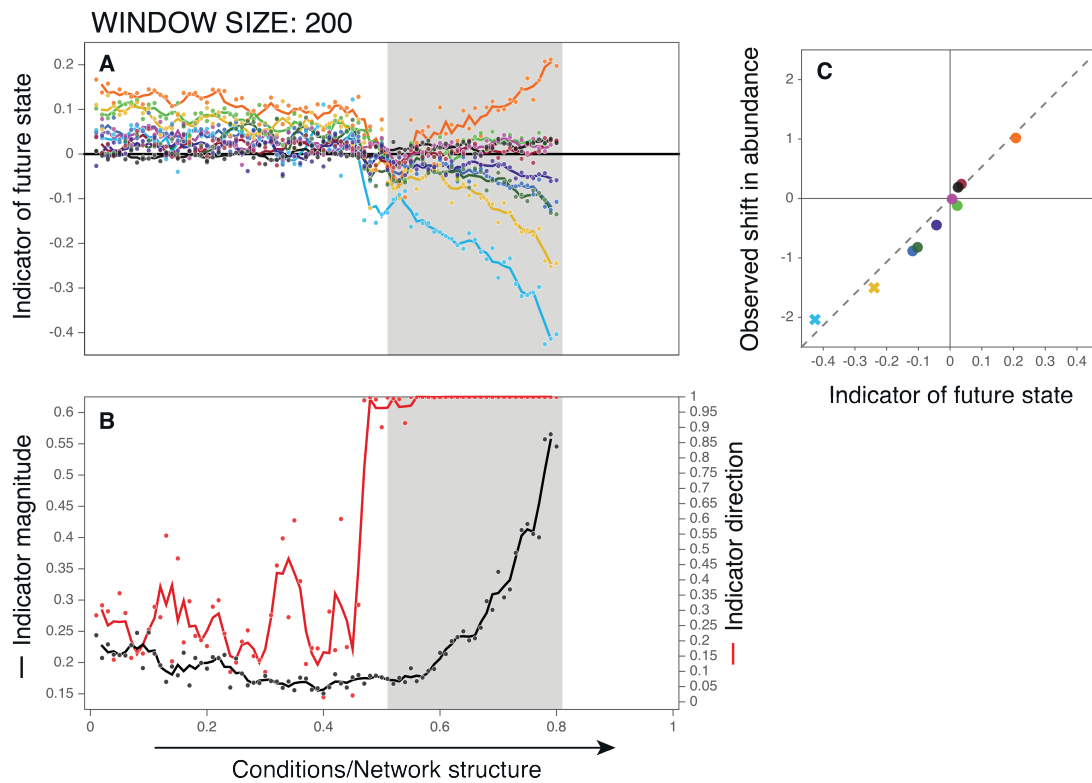


Figure S4.27: Directional slowing down as detected by the indicator when using a rolling window of 1% of the entire time series, i.e. 200 out of 20,000 time steps. **(A)** The indicator of the future state measuring the direction in which fluctuations are distributed asymmetrically. **(B)** The magnitude of the indicator, reflecting the extent in which fluctuations are distributed asymmetrically, plotted together with the accuracy measured as the similarity between its direction and the observed shift in abundance. Grey bands indicate the period in which the indicators magnitude increases significantly. **(C)** The observed changes in abundance versus the scores on the indicator just before the tipping point. Extinct species are indicated with crosses. The initial network, at $M = 0$, and the way in which this network is affected by changing environmental conditions, M , is the same as in Fig. 4.2.

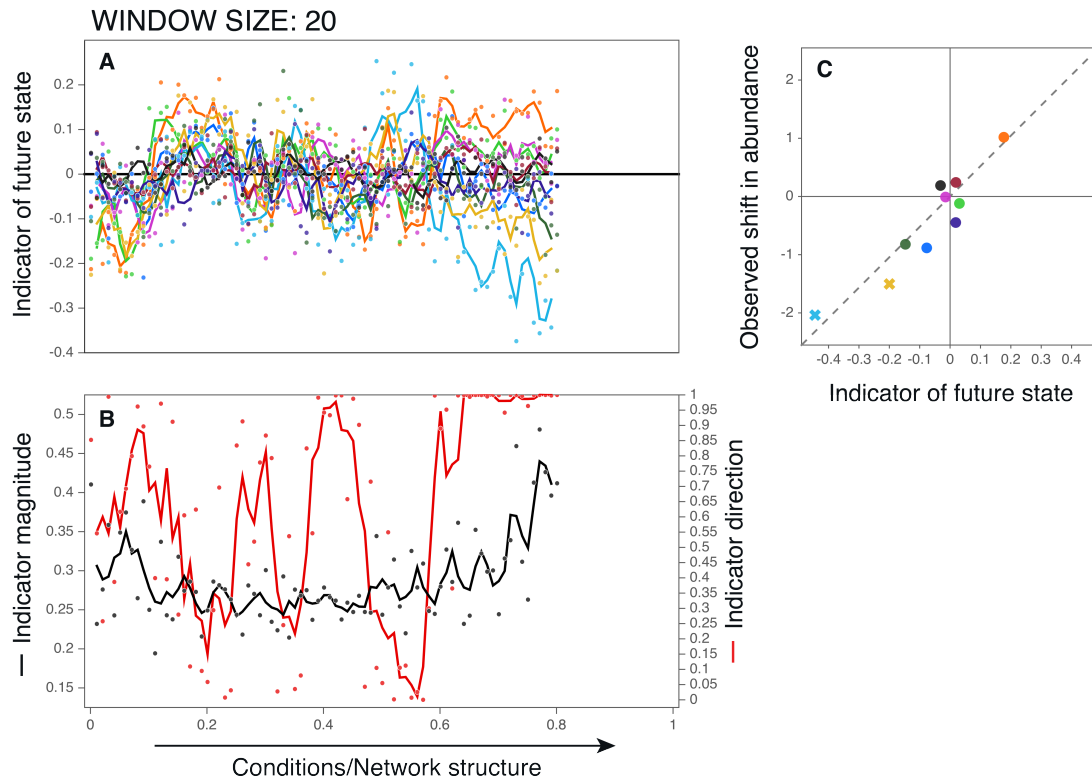


Figure S4.28: Directional slowing down as detected by the indicator when using a rolling window of 0.1% of the entire time series, i.e. 20 out of 20,000 time steps. **(A)** The indicator of the future state measuring the direction in which fluctuations are distributed asymmetrically. **(B)** The magnitude of the indicator, reflecting the extent in which fluctuations are distributed asymmetrically, plotted together with the accuracy measured as the similarity between its direction and the observed shift in abundance. No significant increase in the indicator's magnitude was detected. **(C)** The observed changes in abundance versus the scores on the indicator just before the tipping point. Extinct species are indicated with crosses. The initial network, at $M = 0$, and the way in which this network is affected by changing environmental conditions, M , is the same as in Fig. 4.2.

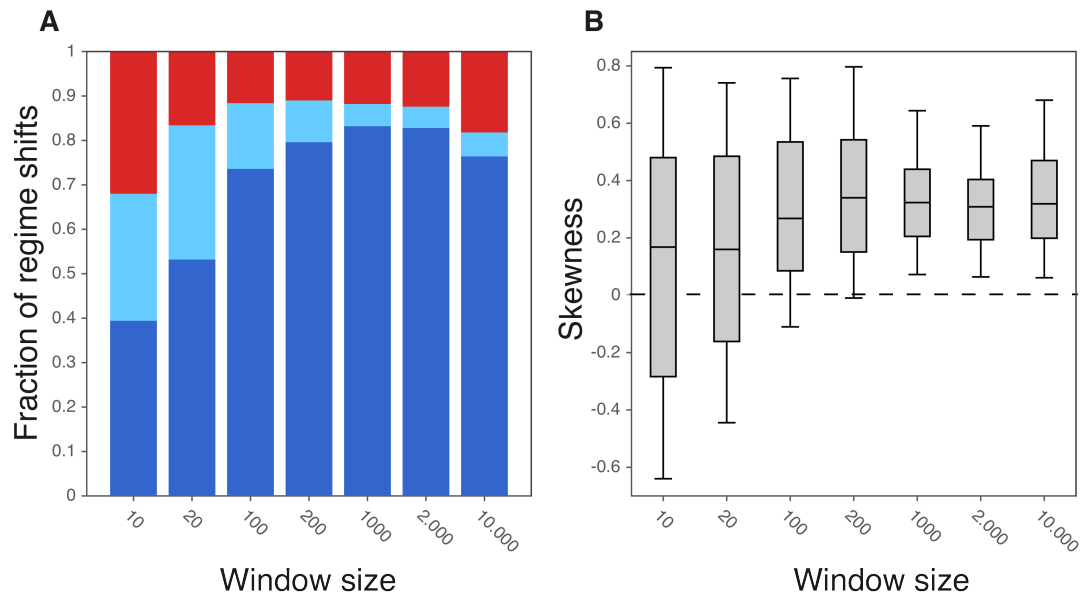


Figure S4.29: Performance of the indicator when the rolling window has a different length (lengths are indicated on the x-axis). **(A)** The fraction of accurately indicated regime shifts (dark blue), the fraction accurately indicated by the first principal component, i.e. the slope of the indicator is accurate, but not by the direction in which time points are skewed (light blue), and the fraction of inaccurately indicated regime shifts (red). **(B)** The skewness of time points projected on the first principal component. A positive skewness means that time points are skewed in the direction of a network's future state. The skewness is shown for regime shifts that were accurately indicated by the first principal component. Time series have a length of 20.000 time steps.




Figure S4.30: The extend in which the size of a rolling window affects the slope indicated by the first principal component far from a tipping point (blue) and close to a tipping point (orange). The y-axis corresponds to the difference in angle between the first principal component obtained for a window of 200 observations and for a window containing the number of observations indicated on the x-axis. The effect of an increasingly small window size on the direction of the first principal component is, in this example, much smaller close to a tipping point. Results are shown for the time series in Fig. 4.2.A at $M=0.1$ (blue) and $M=0.78$ (orange).

Chapter 5

Multivariate indicators of resilience loss

Els Weinans

Rick Quax

Egbert H. van Nes

Ingrid A. van de Leemput

This chapter is under review at *Scientific Reports*.

Abstract

Various complex systems, such as the climate, ecosystems, and physical and mental health can show large shifts in response to small changes in their environment. These ‘tipping points’ are notoriously hard to predict. However, in the past 20 years several indicators pointing to a loss of resilience have been developed. These indicators use fluctuations in time series to detect critical slowing down preceding a large shift. Most of the existing indicators are based on theories of one-dimensional systems. However, most if not all complex systems consist of multiple interacting entities. Moreover, due to technological developments and wearables, multivariate time series are becoming increasingly available in different fields of science. In order to apply the framework of resilience indicators to multivariate systems, various extensions have been proposed. Not all multivariate indicators have been tested for the same types of systems and therefore a systematic comparison between the methods is lacking. Here, we evaluate the performance of the different multivariate indicators of resilience loss in different scenarios. We show that there is not one method outperforming the others. Instead, which method is best to use depends on the type of scenario the system is subject to. We propose a set of guidelines to help future users choose which multivariate indicator of resilience is best to use for their particular system.

5.1 Introduction

Some systems may show large transitions in response to very small changes in their environment. Such nonlinear responses have been documented in systems from various seemingly unrelated fields of study, including algae coverage in shallow lakes (Scheffer, 1997), self-reported emotions of persons (Leemput et al., 2014), abundance of fish (Mangel and Levin, 2005), climatic variables such as ice cover (Crowley and Baum, 1995), or illness of animals or human beings (Scheffer et al., 2018; Liu et al., 2012). Often these observed transitions are argued to be shifts from one stable state into another one, and these shifts between two stable equilibria are the focus point of this study. Because of internal feedback mechanisms, reversing conditions to a pre-shift situation does not necessarily cause a shift back to the old state (Strogatz, 2014). Sometimes the shift back to the preferred state might not be possible at all. The likelihood of such transitions to be triggered by a perturbation, or in more technical terms the size of the stability landscape, is called the resilience of the system (Holling, 1996). Being able to indicate if a system is losing resilience is one fundamental goal of the research on critical transitions (Scheffer et al., 2009). If interactions and feedbacks of a system are well understood, fully parameterized models can help to simulate transitions. However, most of the aforementioned examples are inherently so complicated that accurate models do not exist. As an alternative to models, there are data driven methods that need time series data as input and that can provide a signal when a system is approaching a tipping point (the point where a shift to the alternative state is inevitable). The most well-known indicators of resilience loss are an increase in temporal autocorrelation (most often lag-1 autocorrelation) (Ives, 1995; Held and Kleinen, 2004) and an increase in variance (Carpenter and Brock, 2006). These indicators are based on the phenomenon of critical slowing down. If a system approaches a tipping point, it will become intrinsically slower, such that recovery rate of disturbances decreases (Scheffer et al., 2009). Indicators of critical slowing down have been applied to various lab-experiments (Veraart et al., 2012; Dai et al., 2012) and have been observed in real life systems (Dakos et al., 2008; Wilkinson et al., 2018).

One limitation of the resilience indicators is that the theoretical framework used for their development is generally based on one-dimensional systems. It is therefore not clear if it is one-to-one applicable to the complexities that may occur in multivariate systems (Brock and Carpenter, 2010). Multivariate systems, or network systems, are systems whose dynamics are described by multiple entities (Barabási et al., 2016). Examples include food webs of multiple interacting species, social networks where multiple individuals are observed, or spatial systems where the different locations in space can be viewed as different variables. Obtaining resilience indicators from time series of multivariate systems is fundamentally different from resilience indicators in univariate systems for two reasons. First, for most systems, it is not possible to obtain a quality time series of all variables, so a measurable subset should be chosen. Sometimes, the variables of inter-

est are not measurable at all and a proxy is used (i.e. self-reported levels of emotions to track a persons mood in psychological studies (Wichers, 2014), or isotope measurements of $\delta^{18}O$ in a sediment core to get an indication of past temperatures (Epstein et al., 1951)). Second, for multivariate systems, recovery trajectories depend on which nodes are perturbed (Rodríguez-Sánchez et al., 2020). In mathematical terms it is said that multivariate systems may not have smooth potentials or that they are ‘non-gradient’. It has been suggested that non-gradient behaviour gives a dramatic boost for the possibilities in a system’s dynamics that are completely overlooked by the more traditional analyses (Green et al., 2005; Hastings and Wysham, 2010). One example of such behaviour of systems without a smooth potential is reactive behaviour, in which a perturbation leads to an initial response away from a stable equilibrium and only later recovers to it’s equilibrium position (Neubert and Caswell, 1997). Reactivity has been proposed to be an intermediary step between stable systems and unstable systems and therefore might have properties of both (Tang and Allesina, 2014).

Despite this worry about the applicability, several multivariate indicators of resilience have been proposed. These can be divided into indicators based on univariate measures and multivariate indicators. A straightforward option using commonly used univariate measures is to choose one ‘representative’ variable to track over time (Dakos, 2018). However, this immediately leads to the problem of deciding which variable to measure. Another approach is to use the mean or median of commonly used univariate resilience indicators, such as autocorrelation and variance, for all variables (Bathiany et al., 2013; Dakos and Bascompte, 2014). However, the average value of all variables might be influenced by outliers, and does not exploit the full amount of information that is available in the multivariate signal. Among univariate indicators, autocorrelation is considered the more direct indicator of resilience which is more robust to noise, whereas variance is easier to measure and less sensitive to varying time intervals between consecutive data points (Dakos et al., 2012b). A nonlinear alternative to autocorrelation is mutual information (Kraskov et al., 2004).

Proposed multivariate indicators consist of a first step of dimension reduction technique, followed by the 1-D framework on the newly created 1-D data. This has led to the development of for example degenerate fingerprinting (Held and Kleinen, 2004), that calculates the autocorrelation of the data projection on the first principal component of a Principal Component Analysis (PCA). The advantage is that these techniques offer some new properties that in turn have been suggested as multivariate indicators of resilience loss, such as the explained variance of a PCA analysis (Lever et al., 2020), the eigenvalue of a Min/Max Autocorrelation Factor (MAF) analysis (Weinans et al., 2019), the maximum value of the covariance matrix (Suweis and D’Odorico, 2014; Chen et al., 2019), and the cross-correlation between individual elements (Leemput et al., 2014; Chen et al., 2012). Multivariate extensions to mutual information are for example information dissipation length (IDL) (Quax et al., 2013b) or information dissipation time (IDT) (Quax et al.,

2013a) which measures how long a signal remains in the system before the information is lost. A possible disadvantage of these multivariate metrics is that they are data-hungry and that their theoretical foundation and link to critical transitions is not as well-developed as the more simple metrics. The multivariate indicators of resilience loss that we investigate in this study are listed in table 5.1.

The set of recent indicators clearly reflects the interest in multivariate resilience indicators and the promising new opportunities. However they also pose new questions. Can we expect the multivariate indicators of resilience loss to accurately indicate an upcoming tipping point? What type of data do the different methods require? And can all the proposed indicators deal with the above mentioned multivariate data issues? Since most indicators have not yet been systematically compared (although some have (Suweis and D’Odorico, 2014)), we here evaluate the performance of the indicators from table 5.1. The model we use to generate the time series is well-known plant-pollinator interaction model (Bastolla et al., 2009; Lever et al., 2020). This model has been suggested to also be applicable to other bipartite networks where both facilitation and competition play a role. We chose to use this model, because it can be tuned to display the different types of dynamics that we explore here and it is simple enough to be representative of many systems that can undergo a critical transition such as a fold- or transcritical bifurcation. These bifurcations are part of a group of ‘zero-eigenvalue bifurcations’ where at the bifurcation point one eigenvalue of the systems Jacobian Matrix is zero.

We use the plant-pollinator interaction model to generate data where we know the true outcome which allows us to evaluate the performance of the different proposed indicators. We investigate the effect of six scenarios on the performance of the different indicators. These scenarios are meant as an illustration of some common data issues and system issues that can demonstrate the pros and cons of the list of indicators that we investigate, but obviously do not represent all issues that may be encountered when dealing with time series data. Four scenarios are associated with data acquisition: 1) limited data length, 2) limited data resolution, 3) observational noise/measurement noise, and 4) multiplicative noise as an example of a complex noise regime. The other two scenarios are associated explicitly with multivariate systems: 5) an incomplete set of observed variables and 6) reactivity as an example of non-smooth potential behaviour. The six scenarios are summarized in table 5.2. All scenarios (except reactivity) are tested on a 4-dimensional (4D) version of the model and on two 20-dimensional (20D) models, of which one undergoes a full network collapse (all pollinator species are affected by the shift, all species go extinct), and the other approaches a partial network collapse (only half of the pollinators are directly affected by the shift, half of the pollinators go extinct, the other half of the pollinators and all plants remain alive). For reactivity, we only use the 4D model, since only this model can be tuned to display reactive behaviour.

Table 5.1: Multivariate indicators of resilience loss used in this study. References are 1: Held and Kleinen (2004), 2: Weinans et al. (2019), 3: Quax et al. (2013b), 4: Dakos (2018), 5: Suweis and D'Odorico (2014), 6: Lever et al. (2020), 7: Leemput et al. (2014)

Indicator	Description	Average based	Autocorrelation based	Variance Based	Dimension reduction technique
Degenerate fingerprinting ¹	The autocorrelation of the data projection on the first principal component of a PCA.	✗	✓	✓	✓
MAF autocorrelation ²	The autocorrelation of the data projection on the first MAF. Alternatively his measure can be seen as the maximum autocorrelation in the system	✗	✓	✗	✓
MAF eigenvalue ²	The minimum eigenvalue of a MAF analysis, calculated as the eigenvalues of the covariance matrix of the first difference of an SDS-transform of the original data.	✗	✓	✗	✓
Mutual information ³	Mutual information of time series with lagged time series of itself	✗	✓	✗	✗
Average autocorrelation	Autocorrelation averaged over all variables.	✓	✓	✗	✗
Node maximum autocorrelation ⁴	The autocorrelation of the variable with the highest autocorrelation.	✗	✓	✗	✗
MAF variance ²	The variance of the data projection on the first MAF.	✗	✓	✓	✓
Node maximum variance ⁴	The variance of the variable with the highest variance.	✗	✗	✓	✗
Average variance	Variance averaged over all variables.	✓	✗	✓	✗
PCA variance ¹	The variance of the data projection of the first principal component of a PCA. Alternatively this measure can be seen as the maximum variance in the system.	✗	✗	✓	✓
Maximum value of covariance matrix ⁵	The maximum value of the covariance matrix.	✗	✗	✓	✓
Explained variance ⁶	The explained variance of a PCA based on the covariance matrix, calculated as the maximum eigenvalue of the covariance matrix divided by the sum of all eigenvalues of the covariance matrix.	✗	✗	✓	✓
Average absolute cross-correlation ⁷	The average of the absolute values of all possible cross-correlations between variables.	✓	✗	✗	✗

5.2 Methods

Model

To investigate the effect of various multivariate data issues, we use a well-known simplistic model that we can tune to display a wide range of dynamics and which can be pushed towards a fold or a transcritical bifurcation (two types of critical transitions). The model has been used to describe plant-pollinator interactions (Lever et al., 2014), but can be used to describe a wide range of phenomena where facilitation and competition both play a role (Lever et al., 2020). The model has a deterministic part and a stochastic part. The deterministic part is used to calculate the dominant eigenvalue of the Jacobian (as the ‘true’ resilience of the system) and the eigenvalue of the corresponding Hermitian (as the reactivity of the system). The model is implemented in Grind for Matlab (Nes, 2017) and integrated using an Euler-Maruyama scheme with an integration step of 0.01. The time unit is arbitrary. We sampled data points with an interval of 0.1 time point (so after 10 integration steps), unless stated otherwise.

$$dA_k = [r_k^{(A)} A_k + \frac{\sum_{i=1}^{S_P} \gamma_{ki}^{(A)} P_i}{1 + h_k \sum_{i=1}^{S_P} \gamma_{ki}^{(A)} P_i} A_k - \sum_{l=1}^{S_A} c_{kl}^{(A)} A_l A_k] dt + \sigma_{A_k} dW \quad (5.1)$$

$$dP_i = [r_i^{(P)} P_i + \frac{\sum_{k=1}^{S_A} \gamma_{ik}^{(P)} A_k}{1 + h_i \sum_{k=1}^{S_A} \gamma_{ik}^{(P)} A_k} P_i - \sum_{j=1}^{S_P} c_{ij}^{(P)} P_j P_i] dt + \sigma_{P_i} dW \quad (5.2)$$

In this model, A_k represents the abundance of pollinator species k and P_i represents the abundance of plant species i .

The parameter r describes the per capita growth rate, which in this case can be negative. The parameter γ describes the mutualistic interactions with other species of the other group (where $\gamma_{ik}^{(P)}$ stands for the positive effect that plant species i experiences from pollinator species k). In the model, a saturation is assumed for high abundances of mutualistic partners, where parameter h is the half-saturation constant. The parameter c is a competition term describing the negative effect that pollinators have on each other ($c_{kl}^{(A)}$) and on themselves ($c_{kk}^{(A)}$) and that plants have on each other ($c_{ij}^{(P)}$) and on themselves ($c_{ii}^{(P)}$). Consistent with previous work (Lever et al., 2014), we assume that species can not out-compete each other in the absence of mutualistic partners, so each species has a higher competition with itself than with the other species.

First we use this model to simulate data with two plants ($S_P = 2$) and two pollinators ($S_A = 2$) with default parameters set as $r^{(P)} = [-0.5, -0.5]$ (non-reactive) or $r^{(P)} = [2.2, 2.2]$ (reactive), $\gamma_{11} = \gamma_{22} = 1$, $\gamma_{12} = \gamma_{21} = 0.8$ and $c_{11} = c_{22} = 0.3$, $c_{12} = c_{21} = 0.1$ for both plants and pollinators. The half saturation constant h is 0.5 for all species.

In this model, the relative growth rates between the plants and the pollinators can cause the system to become reactive or pass a tipping point. We find the combinations of parameter values of $r^{(P)}$ and $r^{(A)}$ for which the system undergoes a tipping point by calculating when the Jacobian matrix is zero (Supplementary figure 3). Additionally, we find the parameter values where the system becomes reactive (Supplementary figure 3).

We find that the combination of parameter values for which the model becomes reactive and for which it passes a tipping point by setting the noise to zero and calculating the eigenvalues of the Jacobian and Hermitian matrix dependent on the values of $r^{(P)}$ and $r^{(A)}$ (Supplementary figure 1). The parameter $r^{(A)}$ is the bifurcation parameter and changes in the basic model from $[-0.3 -0.2]$ to $[-0.68 -0.58]$, in the reactive model from $[-0.91 -0.81]$ to $[-1.45 -1.35]$. We tuned this bifurcation parameter such that the change in resilience for both scenarios is exactly the same: the dominant eigenvalue moves from -0.45 to -0.15 in both scenarios. In line with previous studies, in the reactive scenario, the system is not reactive in the beginning but it becomes reactive as the system moves towards the tipping point (Tang and Allesina, 2014) (in supplementary figure 1 the dashed line at $r^{(P)} = 2.2$ crosses the blue line indicating the system becomes reactive).

For our simulations, we increase $r^{(A)}$ stepwise in 50 steps. For every value of $r^{(A)}$ we generate stochastic time series by setting σ to 0.02 for both the plants and the pollinators, unless stated otherwise.

Next, we use the same model to simulate data with ten plants ($S_P = 10$) and ten pollinators ($S_A = 10$). For this model, we could not use default parameters from other studies, so we did a random search to find parameters for which the system starts in an equilibrium where all species are present (abundance > 0.1) and that slowly moves either to a full network collapse (10 pollinators are affected by the changing parameter, all 20 species become extinct) or a partial network collapse (5 pollinators are affected by the changing parameter, those 5 species become extinct). An illustration of the behaviour of these 20D models can be found in supplementary figures 1 (full network collapse) and 2 (partial network collapse). In line with the 4D model $h = 0.5$, $c_{ii}^{(P)} = c_{kk}^{(A)} = 0.3$ and $\gamma_{ii}^{(A)} = \gamma_{kk}^{(P)} = 1$. Parameter settings for the other variables can be found in the supplementary materials section 1 and 2.

Indicator performance

We use the generated data to calculate how the indicators in Table 5.1 change for different values of the bifurcation parameter. In line with previous work (Dakos et al., 2012a; Chen et al., 2019), we calculate Kendall tau correlation between the indicator and the value of the bifurcation parameter as a measure of how well the indicator performs. If the Kendall tau correlation was lower than zero, the performance was set at zero. The Kendall tau correlation is a rank correlation and therefore distinguishes if there is a trend. It provides

no information on how this trend evolves over time (i.e. a linear trend can have the same correlation as an exponential trend).

All indicators with the exception of mutual information were implemented in Matlab. Mutual information (Kraskov et al., 2004) was calculated in Python using the NPEET package (Versteeg, 2014), using 3 neighbors for the kNN algorithm ($k=3$), a base of 2 and no bias correction ($\alpha=0$).

We test for the sensitivity of the indicators to data length by repeating the tests for different lengths of data while keeping the sampling interval constant. A data length of 1000 was chosen for the reduced data length scenario. Next, we test for the sensitivity to data resolution by sampling every 100^{rd} time point (instead of the default of 0.1). Last, we test for the sensitivity to accuracy of the data by applying artificial measurement noise (or observational noise) by adding random values from a normal distribution to the data with $\mu = 0$ and σ ranging from 0 to 0.3. A σ of 0.08 was chosen for the reduced data accuracy scenario. We create data with multiplicative noise by multiplying the Wiener process with the species abundance. We analyze the effect of not observing all variables by repeating the analysis for all possibilities of half of the variables, so for our 4D model this leads to 6 options (4 choose 2) and for the 20D model this leads to 184765 options (20 choose 10). In this scenario we determine the performance as the 5% and 95% quantile of all Kendall tau correlations and label them 'worst case' and 'best case' respectively. Using the 4D model, we test for the effect of reactivity by generating data with the reactivity parameter set to its reactive value ($r^{(P)} = 2.2$) and repeating the same tests as described before. This was only done for the 4D model as the 20D model could not be made reactive while ensuring all species were present in the system. The scenarios, including a short description, are summarized in Table 5.2.

5.3 Results


The performance of the indicators based on the 4D model are summarized in figure 5.1. The results of the 20D model with a full network collapse are summarized in figure 5.2, and of the 20D model with a partial network collapse in figure 5.3. The performance of each indicator per scenario is measured as the Kendall tau correlation between the indicator itself and the value of the bifurcation parameter $r^{(A)}$ (see methods). In the next subsections we will discuss the performance per scenario, so per column in these figures.

Performance in the basic model

The basic model scenario is unlimited by data length, resolution and accuracy, it has the most simple noise regime, all variables are observed, and the system behaves like a gradient system locally. Therefore, it adheres to all assumptions needed to apply indicators of

Table 5.2: Scenarios used in this study.

Scenario	Description
Basic model	Time series have a length of 10.000 with a resolution of 0.1 steps. Gaussian white noise is implemented as a Wiener process with a standard deviation of 0.02. There is no measurement noise. All variables are taken into account.
Data length	Time series have a length of 1.000 with a resolution of 0.1 steps. Gaussian white noise is implemented as a Wiener process with a standard deviation of 0.02. There is no measurement noise. All variables are taken into account.
Data resolution	Time series have a length of 10.000 with a resolution of 100 steps. Gaussian white noise is implemented as a Wiener process with a standard deviation of 0.02. There is no measurement noise. All variables are taken into account.
Measurement noise	Time series have a length of 10.000 with a resolution of 0.1 steps. Gaussian white noise is implemented as a Wiener process with a standard deviation of 0.02. Afterwards, a random number drawn from a normal distribution with 0 mean and a standard deviation of 0.08 is added to every datapoint in the time series. All variables are taken into account.
Multiplicative noise	Time series have a length of 10.000 with a resolution of 0.1 steps. Gaussian white noise is implemented as a Wiener process with a standard deviation of 0.02 multiplied by the value of the variable at that moment in time. There is no measurement noise. All variables are observed.
Subset variables	Time series have a length of 10.000 with a resolution of 0.1 steps. Gaussian white noise is implemented as a Wiener process with a standard deviation of 0.02. There is no measurement noise. All possible subsets of half of the variables are analyzed and their Kendall tau correlation is calculated. The 5% (worst case) and 95% (best case) percentiles of these correlations are depicted as the performance.
Reactive system	Same as basic model, but model parameters are chosen in such a way that the equilibrium in which the system resides is reactive.

**Figure 5.1:** Performance of all indicators for different situations of the four-dimensional plant-pollinator model. Performance is calculated as the Kendall tau correlation of the change in indicator as the system approaches a critical transition.




Figure 5.2: Performance of all indicators for different situations of the 20-dimensional plant-pollinator model with a full network collapse.

resilience loss. In this scenario, as expected, all indicators increase as the bifurcation parameter increases, reflected by a high Kendall tau correlation. An example of the change in the indicator values for the 4D model can be found in supplementary figure 4. In the 4D model the lowest Kendall tau correlation is found for the MAF eigenvalue and the explained variance (figure 5.1). Additional analyses indicate that the Kendall tau correlation of all indicators has quite high statistical specificity, i.e. it is able to distinguish between time series of a system that moves towards a tipping point and time series of a system that does not move towards a tipping point (supplementary figure 7).

For the 20D system, a weak Kendall tau correlation is found for the average absolute cross-correlation between variables (figure 5.2 - 5.3). In the 20D model with a partial network collapse all indicators have a lower performance than in the 4D and the 20D models where the entire network collapses. This indicates that partial network collapses, which are quite common especially in highly dimensional systems, are harder to detect than full network collapses.

Performance for reduced data length

The data length scenario shows that all indicators are negatively affected by a reduction of the data length (from 10.000 data points in the basic model to 1000 data points in this scenario) (figures 5.1-5.3, second column). An illustration of the performance for 10.000, 1000 and 100 time points is provided in supplementary figure 4-6. The effect of a gradually increasing data length is visualized in supplementary figure 8. Interestingly, for the 4D model, most indicators quickly converge when the time series length is increased from 10 to 1000 data points. For the average absolute cross-correlation, explained variance and MAF eigenvalue, convergence happens slower and these indicators still seem to not have converged for a data length of 10.000 time points (supplementary figure 8).

The best performing indicators in the reduced data length scenario are average autocorrelation and average variance, both in the 4D (figure 5.1) and the 20D model with the full network collapse (figure 5.2). These indicators are closely followed by the dimension reduction techniques degenerate fingerprinting, autocorrelation of the data projection on the first MAF (MAF autocorrelation), the variance of the data projection on the first PC (PCA variance) and the variance on the first MAF (MAF variance). In the 20D model with the partial network collapse (figure 5.3), all indicators performed quite poorly in this scenario. The best indicators are the node with the maximum autocorrelation or maximum variance, the average variance and the maximum value of the covariance matrix. The already poorer performing indicators MAF eigenvalue, explained variance, and average absolute cross-correlation are still the least effective indicators in this scenario. Another important observation is that the mutual information performs quite well for the reduced data length scenario in the 4DD model, even though this method is known to be data-hungry. As the dimensions increase in the 20D models, the mutual information suffers more from the reduced data length scenario than in the 4D situation. The mutual information approximates a joint distribution from the data, and therefore it requires more data as dimensions increases.

Performance for reduced data resolution

In the scenario of reduced data resolution, all autocorrelation-based indicators perform poorly compared to the variance based indicators (figures 5.1 - 5.3). Even the variance on the first MAF, which is also variance-based, performs poorly here, probably because the direction of the first MAF is not meaningful when data resolution is low. An illustration of the effect of a decreasing data resolution for the 4D model is provided in supplementary figure 9. Perhaps surprisingly, all variance based indicators perform even better in the low-resolution scenario than in the basic model scenario (figures 5.1 - 5.3). Furthermore, the autocorrelation-based indicators too seem to benefit from a small decrease in data resolution (supplementary figure 9). This can be explained by the fact that we fixed the data length at 10.000 points. The decrease in data resolution in this scenario thus




Figure 5.3: Performance of all indicators for different situations of the 20-dimensional plant-pollinator model with a partial network collapse.

entails an increase in simulation time. The increase in performance indicates that a higher sampling rate does not always provide more accurate results than a lower sampling rate, if the data length remains unchanged. The average variance is the indicator that outperforms the others for all three models in this scenario, which is especially clear in the 20-dimension model with a partial network collapse (figure 5.3).

Performance for reduced data accuracy

A reduction in data accuracy affects the performance of all indicators negatively. The effect of a gradually increasing amount of measurement noise (or observational noise) is visualized in supplementary figure 10. The least affected indicators are the autocorrelation on the first MAF, the average autocorrelation and the average variance. The mutual information is most negatively affected by the increase in measurement noise. The mutual information is the only indicator that does not make any assumptions on the distribution of the data, and therefore might have trouble estimating the distribution when the signal becomes more noisy.

Performance for multiplicative noise

When the noise is modelled in a multiplicative way, all autocorrelation-based indicators maintain their high performance in the 4D model (figure 5.1). the mutual information and average autocorrelation are even completely unaffected by this change in noise type. In both 20D models however (figures 5.2 - 5.3), the mutual information fails to detect the upcoming critical transition. In this scenario noise depends on the abundances of the variables and therefore the total noise changes over time. This might affect the bias in the mutual information, something we did not correct for in this study. The variance-based indicators, with the exception of the explained variance, are unable to detect an upcoming critical transition in this scenario, both in the 4D and 20D models. This is an interesting observation, since the explained variance is closely related to the maximum value of the covariance matrix. The advantage that the explained variance might have here, is that it is relative to the other values in the covariance matrix and is therefore less affected than the maximum value of the covariance matrix by unevenly distributed noise.

Performance for reduced number of observed variables

The effect of only sampling a subset of the involved variables in the system depends on the type of collapse. Our reported performance levels reflect the 5% (worst case) and 95% (best case) quantile. In the case of a full network collapse we find that the MAF eigenvalue, the explained variance and the average absolute cross-correlation have the lowest performance (figures 5.1 - 5.2, column 'subset worst case'). For the 4D model, the other indicators work as good as in the basic model. For the 20D model however, all indicators are compromised, although the average autocorrelation, average variance, and mutual information are least affected by this scenario. In the case of a partial network collapse however, the performance of all indicators is highly reduced (figure 5.3 - column 'subset worst case'). The worst performance is found in the mutual information and the average-based indicators. Surprisingly, the best case scenario here even outperforms the basic model (figure 5.3 - 'subset best case'). Additional analyses of the distribution of the Kendall tau correlations show a large variation in performance (supplementary figures 11-13). The average absolute cross-correlation does not detect a signal on average, but all other indicators have a high probability of detecting the upcoming shift. The highest probability of detecting a shift in this scenario is for the node with the maximum autocorrelation, the node with the maximum variance, and the maximum value of the covariance matrix (supplementary figures 11-13). The high performances in this distribution correspond to situations where all or most variables that are affected by the bifurcation parameter are observed. The low performances appear for situations where the collapsed variables are not part of the observed subset (supplementary figure 14). Interestingly, the increased performance in the best case column compared to the basic model, suggests that all indicators are hindered by the inclusion of variables that are not taking part in the shift.

Performance for reactive systems

In the 4D version of our model, we tuned the parameters in such a way that the system is reactive. This did not substantially affect the performance of the indicators, other than that the average absolute cross-correlation performed slightly better in this scenario. Also the MAF-based indicators seemed to slightly improve their performance in this reactive system (figure 5.1).

5.4 Discussion

Our results show that in our basic model (figures 5.1 - 5.3, first columns) all proposed indicators rise preceding a critical transition (see also supplementary figure 4). Furthermore, even though non-gradient behaviour has been described as a major issue for multivariate dynamical systems (Neubert and Caswell, 1997; Hastings and Wysham, 2010), our results find no evidence that reactivity is a problem for the multivariate indicators of resilience tested for here (figure 5.1, last column). However, not all indicators perform well in every situation and our modelled scenarios help us to understand why certain resilience indicators may fail under particular circumstances.

In the scenario of reduced data length, all indicators have a lower performance than in the basic model. This can simply be explained by the fact that less data leads to weaker statistics.

In the reduced data resolution scenario all variance-based indicators and the absolute average cross-correlation remain strong, while all the autocorrelation-based indicators are weakened. If the sampling resolution is too low, the indicators directly capturing the speed of the system, i.e. based on autocorrelation, will fail to indicate critical slowing down. However, the variance-based indicators, which are essentially indirect measures of slowing down, are not affected by lower resolutions (Dakos et al., 2012b). This is true for both univariate and multivariate timeseries. Whether data resolution is problematic depends on the sampling frequency and the actual speed of the system. For instance, in systems where the activity of taking measurements might affect the dynamics of the system, such as the questionnaires about the mood or behaviour of individuals (Schwarz, 1999), obtaining data with high data length and high resolution can be challenging. Also many ecosystems are difficult to sample on a sufficiently high resolution and sufficiently long time scale, since interactions often occur on long time scales (order of magnitude of several years are not uncommon) (Hastings, 2010). However, long time series do exist in ecology (Benincà et al., 2008), and modern sensor techniques are proving their potential, in both aquatic ecology (Carpenter et al., 2020; Duarte et al., 2021) and vegetation studies (Lamchin et al., 2018). Fields of science that more commonly gather data via automated monitoring devices, such as in medical applications (e.g. wearables), high resolution data is more abundant (Lee and Yoon, 2017).

Higher measurement noise leading to reduced data accuracy has a detrimental effect on all indicators. Mutual information is most strongly affected by reduced data accuracy in comparison to the other indicators. Mutual information is closely related to autocorrelation. It has for instance been shown that for a bivariate normal distribution the lag-1 autocorrelation and mutual information are directly linked to each other (with the relation $MI = -\frac{1}{2}\log(1 - \rho^2)$, where ρ is the lag-1 autocorrelation) (Gelfand and Yaglom, 1959). Therefore, the difference in performance between the two has to be linked to data length. This links to our observation both in the noisy scenarios, as well as in the high dimensional systems, where mutual information needs an increasing amount of data to estimate the distribution of the data, since it is not assuming any predefined distribution. Autocorrelation has fewer degrees of freedom and is therefore less influenced by noisy data when data is limited. Our analysis was based on stationary time series of 10.000 time points for every value of the control parameter, which is considered quite a lot of data in most applications. Therefore mutual information might not be the most practical choice as an indicator of resilience unless the number of degrees of freedom can be significantly reduced or large amounts of data are available.

We simulated a complex noise scenario by replacing additive noise with multiplicative noise, meaning that the amount of noise depends on the variable level (here species abundance). The multiplicative noise scenario hardly affects the autocorrelation based indicators in the 4D model, but it has a detrimental effect on all variance-based indicators. In the 20D systems it also affects degenerate fingerprinting and mutual information. Degenerate fingerprinting is most likely affected because a PCA, which is a step in degenerate fingerprinting, is heavily influenced by changes in noise regimes. In our multiplicative noise scenario we multiplied the noise by the species abundances. Other ways to investigate complex noise regimes are by changing the way it is distributed over variables (Boerlijst et al., 2013; Weinans et al., 2019) or by adding the noise to modeled parameters instead of variables (Dakos et al., 2012b). Ecological systems are often most realistically modeled with noise added to a parameter (for example a simple system of bacteria where temperature affects the growth rate and in that way affects the abundance), and therefore in ecological systems autocorrelation-based indicators might be the more robust choice (Dakos et al., 2012b).

Our findings suggest that the subset of variables that is looked at, can have a major influence on the results. This problem was first described for univariate analyses of three variables in a three-dimensional system (Boerlijst et al., 2013), and still appears for the multivariate methods that we investigate here. Importantly, not all variables in a multivariate system approaching a critical transition are subject to critical slowing down, so when only a subset is observed, it depends on the subset whether a loss of resilience will be signalled. In models with a full network collapse, several indicators hardly show a loss of resilience: the MAF eigenvalue, the explained variance and the average absolute cross-correlation (figures 5.1 - 5.2). In contrast, in our 20D model where only a part of the

variables collapses, all indicators have a reduced performance. Cross-correlation is most affected by not observing all variables (supplementary figures 11-13). Some combinations of variables even give an indication of an increase in resilience as the system moves towards the tipping point (supplementary figures 7-9). In this scenario, only 5 of the 20 species are directly affected by the changing environmental condition. The variables are all interacting with each other, so in principle this information could be distributed through the network. However, as we show, the signal is not picked up by all subsets of variables. The performance essentially depends on the fraction of observed variables that are part of the transition. Performance is high if all variables that are collapsing are observed, and performance reaches a minimum when none of the observed variables are collapsing (supplementary figure 14). One important note is that in real systems it is typically unknown how many of the observed variables will be part of the shift (i.e. the location on the horizontal axis in supplementary figure 14). The prevalence of this scenario for all real-world cases therefore suggests a cautious interpretation of all univariate and multivariate resilience analyses.

Our analysis of multiple indicators of resilience loss in multivariate timeseries shows that there is not one indicator that clearly outperforms the others. It depends on the scenario which indicator is favourable, similar to what was previously found for one-dimensional systems (Dakos et al., 2012a). Some of the scenarios that we present here can be tested for, providing a user with an indication of the type of system they are dealing with. Box 1 provides some guidelines and questions to consider before choosing a multivariate indicator of resilience. Given the complexity of the matter, we are unable to propose a step-by-step flow chart that recommends one indicator based on a set of questions. However, we do believe that these considerations can guide the decision-making process.

It is important to realize that the scenarios we discussed generally do not happen in isolation. In order to fully map the performance of resilience indicators, ultimately a multidimensional coordinate system is needed where each of the scenarios that we used can be seen as one of the axes in the coordinate system (as illustrated in figure 5.4. This coordinate system has many additional axes that are not investigated here). In every region of this coordinate system some indicators might perform well, whereas others are unable to detect any change in resilience. Some regions might not be suitable for any (of the currently used) multivariate indicators of resilience loss. This figure illustrates one limitation of our study: our scenarios investigate what happens along each axis, whereas the spaces in between are left unstudied. For example in our 4D model, in every scenario we demonstrate that some indicators can correctly predict the upcoming tipping point. However, when the data has a low resolution, in combination with multiplicative noise, and only a subset of data is observed (not uncommon for empirical datasets), all indicators will likely have a low performance.

Box 1: Guidelines for choosing a multivariate indicator of resilience

Based on our findings, we recommend the following guidelines to help deciding which indicator to use.

1. Before anything else, decide if the system you are dealing with could potentially be subject to a zero-eigenvalue bifurcation. Are there known processes that could cause a positive feedback loop? Have shifts been observed before? Is the system fluctuating around an equilibrium?
2. Check the autocorrelation of the included variables. If the autocorrelation is not significantly different from zero, this suggests the resolution of the data is too low to use autocorrelation-based indicators, so use a variance-based indicator instead.
3. Use system knowledge or take multiple measurements at the same time to determine how accurate the data is. If there is a low data accuracy (high observational or measurement noise), use a dimension-reduction technique or use the average autocorrelation or variance.
4. Use system knowledge or perform experiments to get an understanding of how the noise behaves. In real systems the noise is most likely a combination of observational/measurement noise and system noise. If the system noise varies over time, such as in our multiplicative noise scenario, use an autocorrelation-based indicator.
5. Irrespective of which method is chosen, we recommend to do a data suitability test (Weinans et al., 2019), to test whether the data is of sufficient length to reliably calculate the resilience indicators.

In this study, we focus on the performance of indicators for approaching zero-eigenvalue bifurcations as classic examples of critical transitions. Obviously, there are many other types of transitions that are out of the scope of this study. Consequently, the parameters in our model were tuned in such a way that more complex bifurcations such as hopf bifurcations, or global bifurcations, are not taken into account. Also, we did not consider abrupt shifts caused by an increase in noise (Horsthemke, 1984), in which case flickering might be a more promising indicator than any of the indicators we have investigated here (Dakos et al., 2013). Also, some rapid transitions might be caused by rapid changes in the environment in which case no indicators of critical slowing down are expected (Boettiger et al., 2013; Dakos et al., 2015). Therefore our conclusions should only be seen in the

context of slowly approaching fold or transcritical bifurcations, which is why the first guideline in Box 1 suggests to always first consider the type of transition that is expected before applying any indicator of resilience to the data.

Some multivariate indicators of resilience loss have been suggested to not only predict when a systems is losing resilience, but also which variables are most affected by it (Bathiany et al., 2013; Chen et al., 2019; Weinans et al., 2019) or even where the future state of the system can be found (Lever et al., 2020). To answer these ‘follow-up questions’, MAF or PCA related indicators seem quite appropriate, since they provide a ‘direction’ (i.e. a combination of variables that are most affected by the changing conditions or that recover slowly when perturbed simultaneously) in addition to a resilience indicator. These are the indicators that fail in our ‘subset of variables’ scenario and therefore we do not recommend to use them to predict an upcoming tipping point. However, we do see a use for these indicators to answer questions such as “which variables are affected by an upcoming shift?” and “what will the future state look like?” (Weinans et al., 2019; Lever et al., 2020).

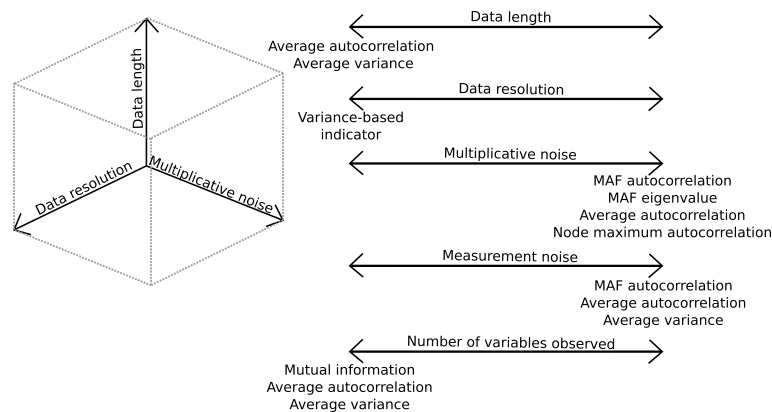


Figure 5.4: Conceptual image of the multivariate space a system can be in, of which we here visualize three axes. In reality, any issue (data issue such as data length of system issue such as reactivity) can be seen as an axis in this coordinate system, resulting in a high-dimensional space. Depending on the region in this coordinate system, some indicators are preferable over others. For example, as the system under investigation moves to the low end of the ‘data resolution axis’, variance based indicators are preferred, whereas on high resolution side there is no limitation on the choice of indicator. If the system is found on the low end of the ‘data length’ axis, average autocorrelation, average variance, or mutual information could be the best choice for an indicator. As a system moves along the ‘multiplicative noise’ axis, our analysis suggests that autocorrelation-based indicators, with the exception of mutual information, are the best choice. In our analysis we investigate the effect of moving along several axes of this conceptual coordinate system. The regions in between are yet to be explored.

There is a group of indicators based on information theory that could potentially also be useful to predict upcoming transitions (such as Fisher information (Eason and Cabezas, 2012) or entropy measures (Richman and Moorman, 2000)). These indicators also infer properties of dynamical systems, but so far have not been applied to infer a systems resilience. Therefore, we did not include them in our analysis. We do touch upon information theoretical indicators by including the mutual information indicator. This indicator accurately points to a loss of resilience in the basic model, but it seems quite sensitive to the scenarios we have investigated here. Furthermore, some indicators not directly linked to critical slowing down have been found to precede critical transitions, such as critical fluctuations (Olthof et al., 2020b). Future comparisons can take into account these, and other, extensions to the resilience indicators investigated here.

The idea that critical transitions are preceded by generic indicators that cross over multiple scientific domains is an exciting premise that has attracted justifiable widespread attention. However, even though the theoretical work has led to many successful discoveries, the application to multivariate empirical data remains challenging. Awareness of the potential pitfalls that can be encountered in real data and an understanding of their effect on the different indicators may help to interpret the growing body of resilience research.

Supplementary materials

S5.1 parametersettings 20D plant pollinator model - full network collapse

$$c^{(A)} =$$

0.3000	0.0360	0.0462	0.0166	0.0365	0.0149	0.0536	0.0368	0	0.0221
0.0220	0.3000	0.0242	0.0885	0.0292	0.0602	0.0371	0.0809	0.0229	0.0086
0.0233	0.0763	0.3000	0.0547	0.0354	0.0228	0.0528	0	0.0549	0.0618
0.0510	0.0637	0	0.3000	0.0355	0.0057	0.0087	0.0556	0.0185	0.0498
0.0577	0.0251	0.0795	0.0500	0.3000	0.0534	0.0302	0.0537	0.0325	0.0523
0.0707	0.0165	0.0128	0.0471	0.0492	0.3000	0.0224	0.0573	0.0351	0.0350
0.0371	0.0774	0.0375	0.0434	0.0245	0.0150	0.3000	0.0771	0.0413	0.0834
0.0516	0.0164	0.0183	0.0406	0.0463	0.0355	0.0307	0.3000	0.0507	0.0457
0.0289	0.0318	0.0278	0.0602	0.0326	0.0403	0.0308	0.0332	0.3000	0.0306
0.0551	0.0336	0.0267	0.0392	0.0454	0.0556	0.0248	0.0124	0.0488	0.3000

$$c^{(P)} =$$

0.3000	0.0512	0.0299	0.0454	0.0033	0.0304	0.0392	0.0284	0.0382	0.0571
0	0.3000	0.0264	0.0192	0.0393	0.0504	0.0755	0.0020	0.0708	0.0283
0.0458	0.0463	0.3000	0.0394	0.0216	0.0335	0.0404	0.0209	0.0097	0.0329
0.0400	0.0478	0.0414	0.3000	0.0348	0.0370	0.0022	0.0331	0.0321	0.0345
0.0807	0.0348	0.0521	0.0429	0.3000	0.0675	0.0318	0.0622	0.0369	0.0425
0.0516	0.0079	0.0439	0.0347	0.0460	0.3000	0.0768	0.0280	0.0779	0.0484
0.0294	0.0456	0.0646	0.0001	0.0266	0.0385	0.3000	0.0769	0.0436	0.0576
0.0522	0.0476	0.0355	0.0641	0.0255	0.0359	0.0339	0.3000	0.0523	0.0398
0.0026	0.0699	0.0547	0.0246	0.0471	0.0396	0.0630	0.0355	0.3000	0.0591
0.0320	0.0109	0.0470	0.0213	0.0489	0.0139	0.0517	0.0353	0.0298	0.3000

$$\gamma^{(A)} =$$

1.0000	0.7256	0.8165	0.8032	0.8804	0.7901	0.8946	0.7545	0.9086	0.8047
0.7243	1.0000	0.7728	0.7887	0.8128	0.8364	0.9992	0.8955	0.7779	0.8616
0.8501	0.8354	1.0000	0.8472	0.7449	0.8142	0.6800	0.6601	0.9537	0.9296
0.7165	0.8669	0.6222	1.0000	0.8427	0.8191	0.7604	0.6323	0.6300	0.7554
0.8607	0.8280	0.5856	0.7303	1.0000	0.8277	0.8152	0.7924	0.9418	0.7169
0.7751	0.8574	0.8164	1.0000	0.7621	1.0000	0.8343	0.7611	0.8105	0.7735
0.8250	0.9945	0.8621	0.8314	0.8036	0.8109	1.0000	0.8486	0.7733	0.5809
0.6731	0.7307	0.8968	0.7996	0.8268	0.6182	0.6832	1.0000	0.8913	0.5988
0.8080	0.9226	0.9267	0.9080	0.8974	0.9306	0.6748	0.7757	1.0000	0.8893
0.7584	0.7300	0.7736	0.9048	0.8394	0.8409	0.8633	0.6953	0.9092	1.0000

$$\gamma^{(P)} =$$

1.0000	0.7717	0.8184	0.7633	0.8644	0.9198	0.9089	0.8588	0.8697	0.8381
0.9262	1.0000	0.7196	0.7353	0.9588	0.9840	0.8463	0.7853	0.7648	0.7666
0.7936	0.8113	1.0000	0.7441	0.8513	0.9274	0.7042	0.8250	0.8528	0.7291
0.7977	0.8136	0.9174	1.0000	0.8501	0.7205	0.8204	0.7631	0.8086	0.8773
0.8256	0.7828	1.0000	0.7830	1.0000	0.7872	0.8519	0.7275	0.7577	0.8263
0.8329	0.7944	0.5845	0.8226	0.7242	1.0000	0.8258	0.8666	0.8125	0.6453
0.7948	0.6633	0.6590	0.7301	0.7309	0.7954	1.0000	0.6924	1.0000	0.7572
0.6132	0.8242	0.9830	0.8946	0.9385	0.8337	0.6134	1.0000	0.7439	0.9050
1.0000	0.8272	1.0000	0.7849	0.8794	0.8221	0.8042	0.8693	1.0000	0.9213
0.6377	0.8034	0.8878	0.7836	0.8224	0.8371	0.8678	0.8240	0.8049	1.0000

$$r^{(P)} = -0.1$$

The bifurcation parameter $r^{(A)}$ changes linearly in 50 steps from -0.0424 to -1.2106 for all pollinators, chosen in such a way that the dominant eigenvalue of the Jacobian matrix changes from -0.45 to -0.15 as the system approaches a critical transition (fig. S5.1).




Figure S5.1: Behaviour of the 20D model without noise with a full network collapse. Green lines indicate plant abundances and red lines indicate pollinator abundances. The dashed black line depicts the values for $r^{(A)}$ that we use for our analysis.

S5.2 parametersettings 20D plant pollinator model - partial network collapse

 $c^{(A)} =$

0.3000	0.0566	0.0525	0.0181	0.0516	0.0286	0.0510	0.0510	0.0635	0.0271
0.0517	0.3000	0.0724	0.0255	0.0508	0.0317	0.0358	0.0216	0.0444	0.0194
0.0085	0.0427	0.3000	0.0183	0.0235	0.0615	0.0009	0.0363	0.0468	0.0283
0.0418	0.0431	0.0361	0.3000	0.0572	0.0134	0.0524	0.0481	0.0603	0.0184
0.0421	0.0233	0.0480	0.0555	0.3000	0.0097	0.0220	0.0550	0.0407	0.0331
0.0381	0.0481	0.0158	0.0485	0.0583	0.3000	0.0480	0.0286	0.0388	0.0181
0	0.0608	0.0324	0.0282	0.0298	0.0629	0.3000	0.0200	0.0543	0.0451
0.0273	0.0594	0.0258	0.0364	0.0591	0.0148	0.0119	0.3000	0.0237	0.0559
0.0519	0.0439	0.0403	0.0210	0.0445	0.0274	0.0481	0.0243	0.3000	0.0070
0.0332	0.0069	0.0347	0.0470	0.0101	0.0496	0.0309	0.0264	0.0234	0.3000

 $c^{(P)} =$

0.3000	0.0723	0.0379	0.0616	0.0127	0.0531	0.0275	0.0421	0.0248	0.0378
0.0460	0.3000	0.0038	0.0612	0.0664	0.0333	0.0293	0.0449	0.0360	0.0020
0.0344	0.0320	0.3000	0.0374	0.0585	0.0582	0.0680	0.0382	0.0676	0.0128
0.0184	0.0228	0.0371	0.3000	0.0392	0.0627	0.0575	0.0208	0.0461	0.0283
0.0402	0.0470	0.0392	0.0435	0.3000	0.0495	0.0195	0.0429	0.0591	0.0340
0.0431	0.0520	0.0581	0.0718	0.0269	0.3000	0.0079	0.0379	0.0408	0.0534
0.0615	0.0129	0.0262	0.0431	0.0396	0.0186	0.3000	0.0475	0.0253	0.0425
0.0461	0.0583	0.0134	0.0490	0.0126	0.0216	0.0446	0.3000	0.0577	0.0657
0.0290	0.0306	0.0290	0.0600	0.0147	0.0692	0.0807	0.0492	0.3000	0.0211
0.0345	0.0322	0.0443	0.0193	0.0409	0.0349	0.0504	0.0506	0.0441	0.3000

 $\gamma^{(A)} =$

1.0000	0.8529	0.7304	0.9240	0.6789	0.8022	0.6468	0.5195	0.9032	0.8485
0.7804	1.0000	0.8725	0.7045	0.8580	0.7597	0.8700	0.6865	0.8380	0.6699
0.8359	0.6892	1.0000	0.7063	0.7855	0.7843	0.6484	0.9288	0.7659	0.8341
0.7470	0.7224	0.9097	1.0000	0.7419	0.6929	0.8887	0.7431	0.9261	0.8208
0.7901	0.6600	0.9390	0.8583	1.0000	0.9169	0.6562	0.8008	0.8312	0.7562
0.6107	0.8838	0.7761	0.6947	0.8547	1.0000	0.8929	0.6877	0.7465	0.7407
0.7592	0.8284	0.6102	0.7943	0.8421	0.8563	1.0000	0.8955	0.8385	0.8438
0.7883	0.8154	0.8696	0.8296	0.8414	0.8169	0.6932	1.0000	0.7579	0.8582
0.6963	0.7577	1.0000	0.7533	0.8661	0.9028	0.8552	0.7388	1.0000	0.7857
0.7951	0.7947	0.6607	0.8321	0.7693	0.8694	0.7682	0.6940	0.7994	1.0000

$$\gamma^{(P)} =$$

1.0000	1.0000	0.7609	0.7574	0.8573	0.8231	0.6850	0.7173	0.8598	0.8089
0.8550	1.0000	0.8879	0.7032	0.9633	0.8709	0.8280	0.8565	0.9481	0.7359
0.6711	0.7216	1.0000	0.7861	0.7660	0.7909	0.8975	0.7717	0.7193	0.7823
0.6764	0.7527	0.9938	1.0000	0.7275	0.6884	0.8215	0.7282	0.7369	0.6859
0.7902	0.7127	1.0000	0.9249	1.0000	0.9663	0.7082	0.8111	0.6957	0.8916
0.9422	0.7436	0.8804	0.8826	0.8622	1.0000	0.8426	0.8001	0.7651	0.8607
0.8188	0.6296	0.7034	0.8614	0.8303	0.7966	1.0000	0.7652	0.9753	0.7626
0.9492	0.7197	0.7458	0.7138	0.6965	0.9203	0.9097	1.0000	0.6572	0.6437
0.6836	0.9059	0.8163	0.7971	0.6762	0.9395	0.7365	0.7183	1.0000	0.8436
0.7899	0.6714	0.9072	0.5238	0.8976	0.7449	0.8339	0.7834	0.7596	1.0000

$$r^{(P)} = -0.1$$

The bifurcation parameter $r^{(A)}$ changes linearly in 50 steps from -0.4318 to -0.8537 for the first 5 pollinators. For the other pollinators it is set to -0.4318 and it remains there. These values are chosen in such a way that the dominant eigenvalue of the Jacobian matrix changes from -0.45 to -0.15 as the system approaches a critical transition (fig. S5.2).




Figure S5.2: Behaviour of the 20D model without noise with a partial network collapse. Green lines indicate plant abundances and red lines indicate pollinator abundances. The dashed black line depicts the values for $r^{(A)}$ that we use for our analysis.

S5.3 supplementary figures




Figure S5.3: Illustration of the combination of values for $r^{(A)}$ (here denoted ra) and $r^{(P)}$ (here rp) where the 4 dimensional plant-pollinator system becomes reactive and where the tipping point is located. Dashed lines indicate the range of $r^{(A)}$ that we use for our analysis where we push the system towards the tipping point, chosen in such a way that the dominant eigenvalue of the Jacobian ranges from -0.45 to -0.15.

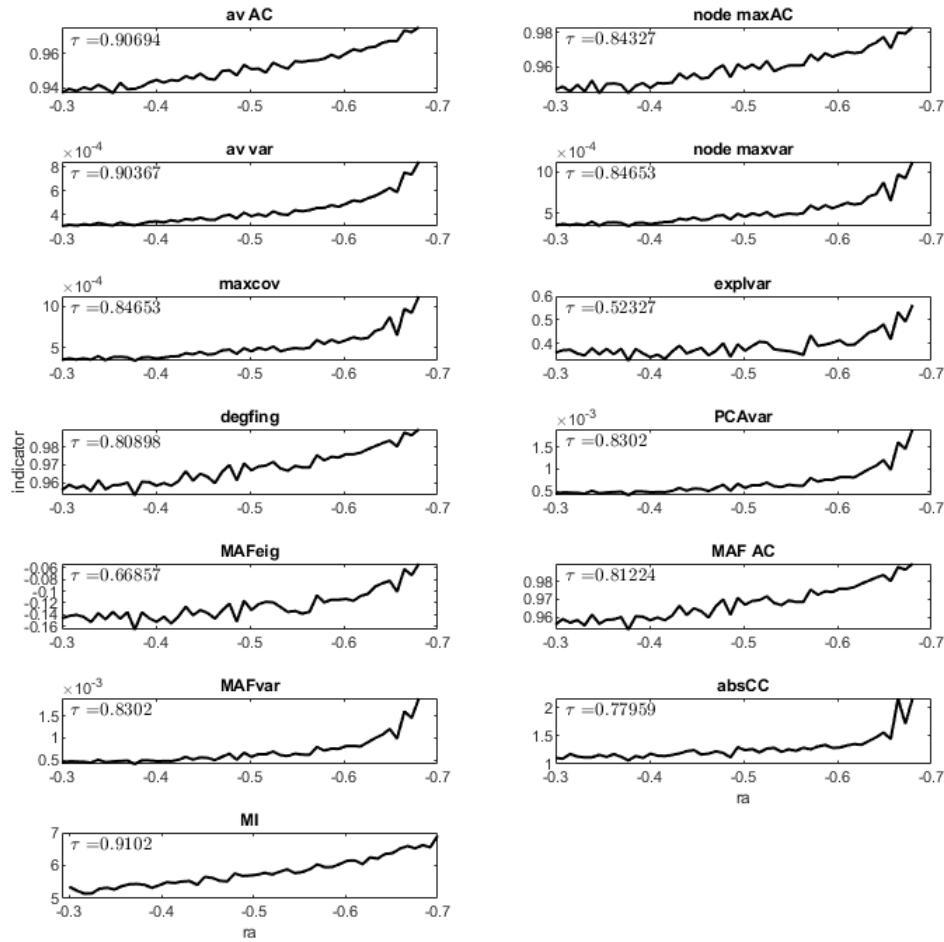


Figure S5.4: Illustration of the increase in all of the indicators as the bifurcation parameter $r^{(A)}$ (here denoted ra) reaches the critical value. Kendall tau correlation for the trend is indicated in every subplot.

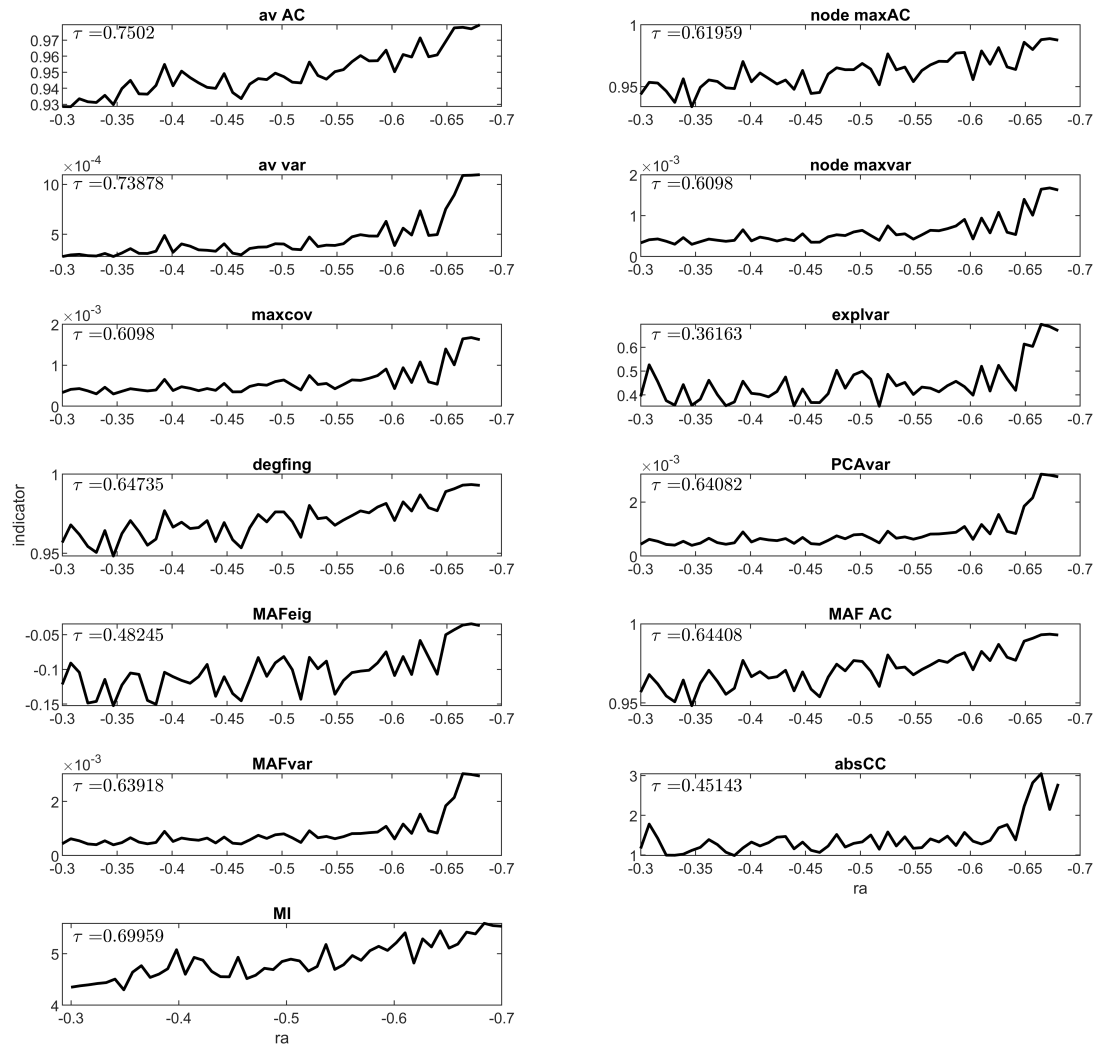


Figure S5.5: Same as figure S5.4, but with time series with a length of 1000 instead of 10,000 per value for ra, yielding a lower performance than with 10,000 points.

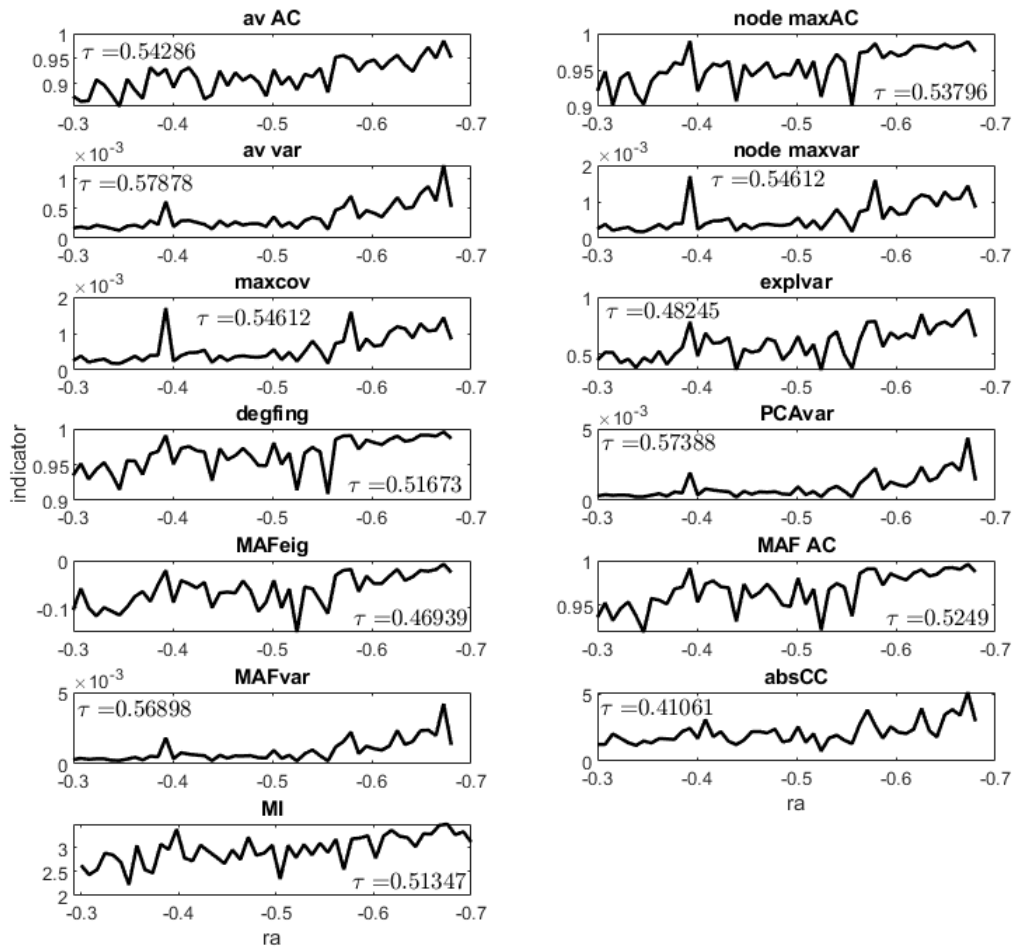


Figure S5.6: Same as figure S5.4, but with time series with a length of 100 instead of 10,000 per value for ra , yielding a lower performance than with 10,000 or 1000 points.

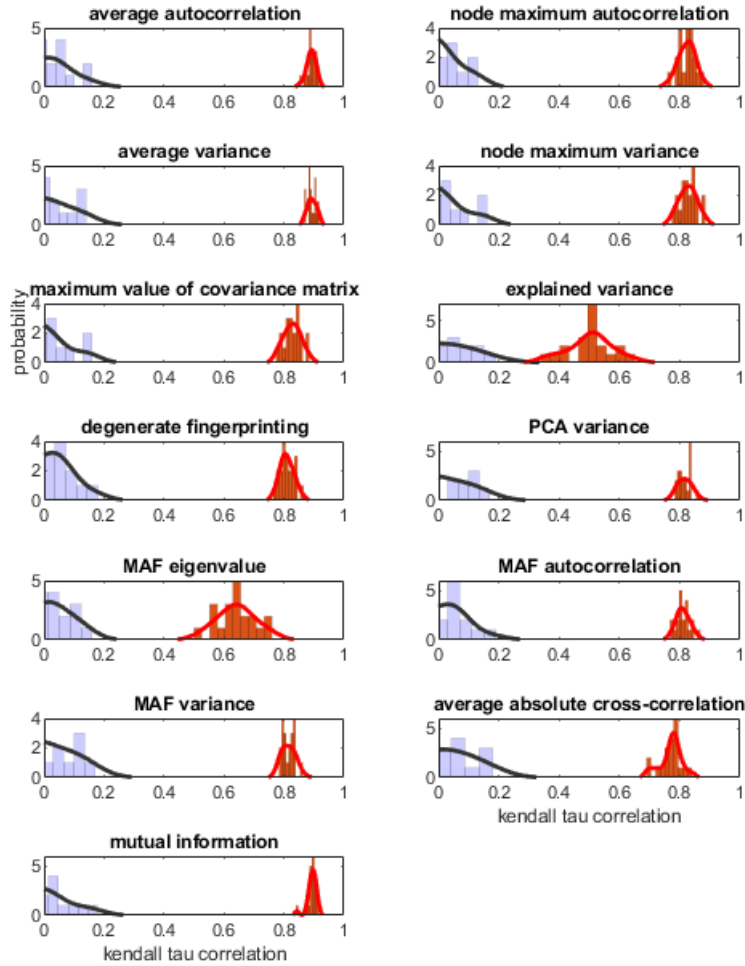


Figure S5.7: Distributions of 20 simulations for the 4-dimension model. The blue distributions are kendall tau correlations for different indicators for time series that do not change over time (the parameter $r^{(A)}$ is fixed). In this situation, no change is expected and thus the kendall tau correlation is centered around zero. The red distributions are distributions of the kendall tau correlations for data that is pushed towards a tipping point (for one instance, see figure 2). This figure shows that the distributions do not overlap and for this particular dataset the specificity of the kendall tau correlation is high, suggesting a low rate of false positives.

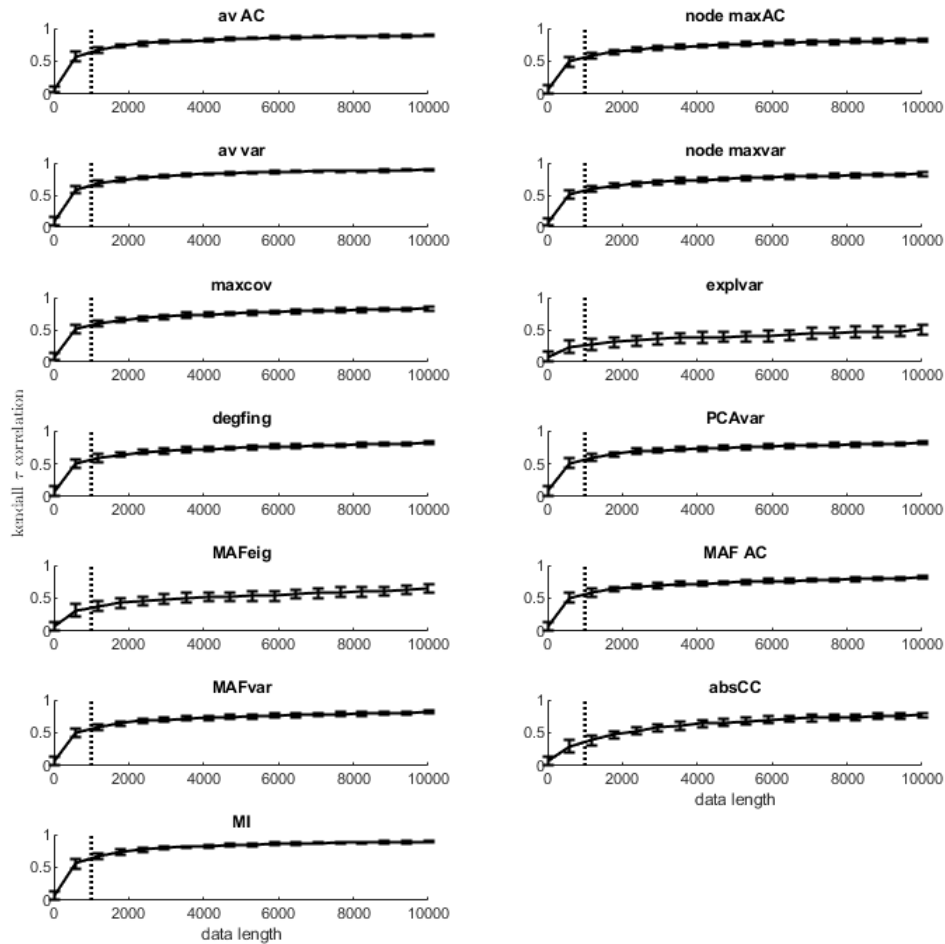


Figure S5.8: The effect of different data lengths for the analyses. Lines demonstrate the increase in kendall tau correlation as the data length increases. Error bars indicate standard deviations of 20 different simulations. The dotted line indicates the data length used for our ‘reduced data length scenario’. All other analyses were performed on the largest data size indicated here (10,000 time points).

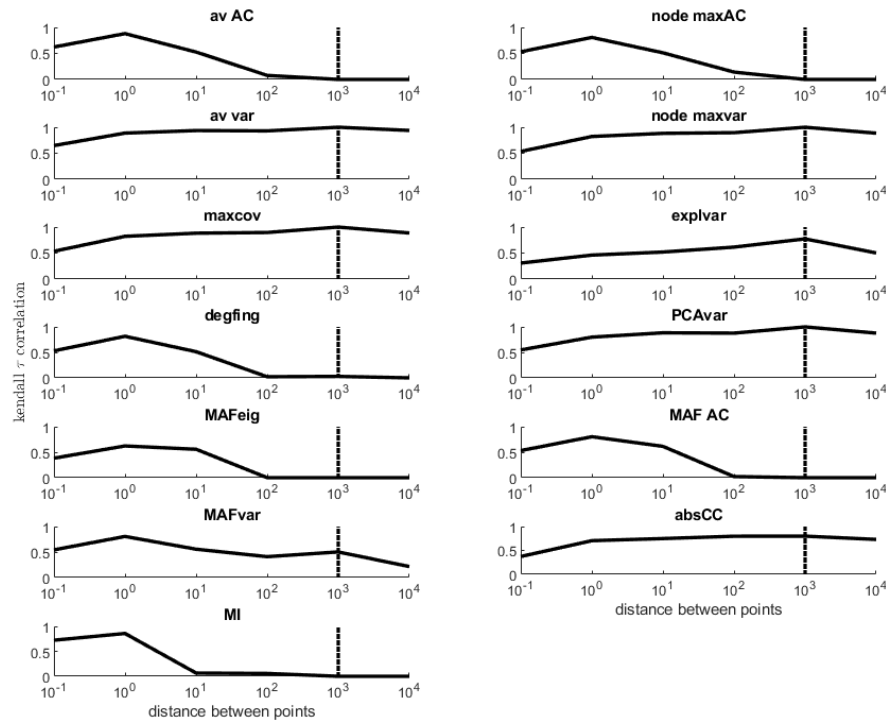


Figure S5.9: The effect of different data resolutions for the analyses with data with a length of 1000. Lines demonstrate the increase or decrease in kendall tau correlation as the data length increases. The dotted line indicates the data length used for our ‘reduced data resolution scenario’. All other analyses were performed on the highest data resolution indicated here (a distance of 10^{-1} time points).




Figure S5.10: The effect of different measurement noise or observation noise for the analyses. Lines demonstrate the decrease in kendall tau correlation as the standard deviation of the measurement noise increases. Error bars indicate standard deviations of 20 different simulations. The dotted line indicates the measurement noise used for our ‘measurement noise scenario’. All other analyses were performed without measurement noise.




Figure S5.11: Violin plots of the distributions of kendall τ correlations of all 6 possibilities of taking 2 out of the 4 variables in the 4D model. White circle indicates the median value. The lowest values are used as the indicator's performance in figure 5.1 in the main text. Kendall τ correlations lower than zero indicate a reversed effect, i.e. the indicator predicts an increase in resilience as the system moves towards the tipping point.




Figure S5.12: Violin plots of the distributions of kendall τ correlations of all 184765 possibilities of taking 10 out of the 20 variables in the 20D model with a full network collapse. White circle indicates the median value. The lowest values are used as the indicator's performance in figure 5.2 in the main text. Kendall τ correlations lower than zero indicate a reversed effect, i.e. the indicator predicts an increase in resilience as the system moves towards the tipping point.




Figure S5.13: Violin plots of the distributions of kendall τ correlations of all 184765 possibilities of taking 10 out of the 20 variables in the 20D model with a partial network collapse. White circle indicates the median value. The lowest values are used as the indicator's performance in figure 5.3 in the main text. Kendall τ correlations lower than zero indicate a reversed effect, i.e. the indicator predicts an increase in resilience as the system moves towards the tipping point.

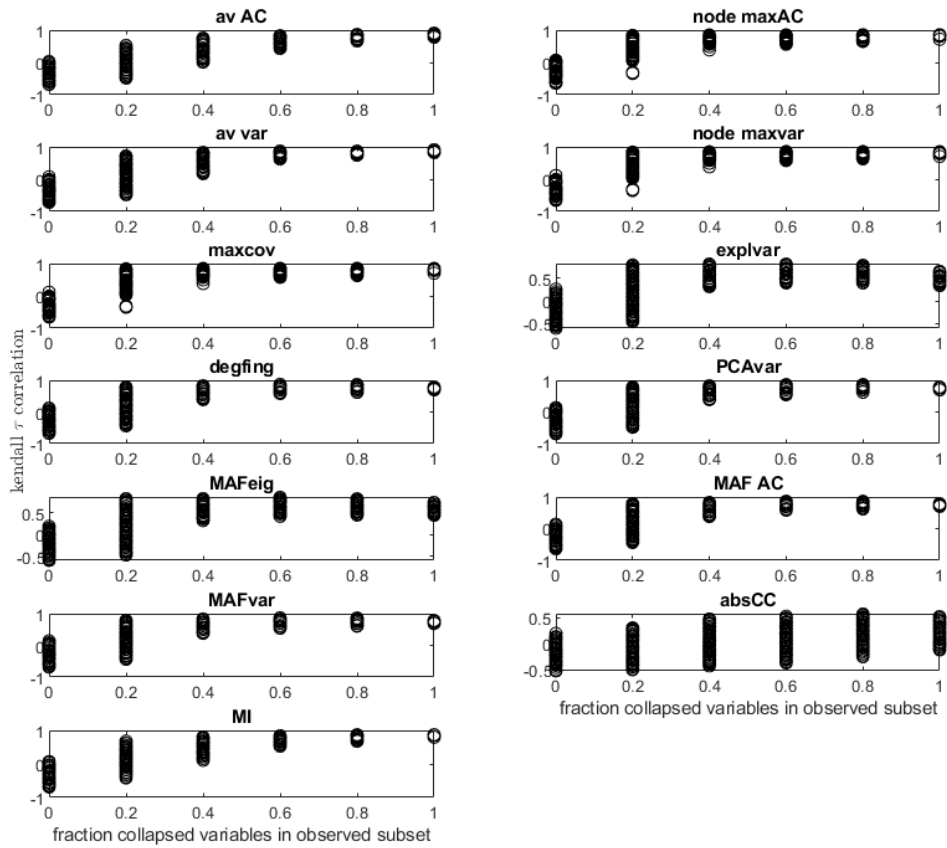


Figure S5.14: Kendall τ correlations per indicator for all simulations where a fraction of the collapsed variables are present. A fraction of 0 indicates that none of the collapsing variables are within the subset of observed variables and a fraction of 1 indicates that all collapsing variables are within the subset of observed variables. Highest performance is obtained when all collapsed variables are observed.

Chapter 6

Roots of post-truth revealed from massive language analysis

Marten Scheffer

Ingrid A. van de Leemput

Els Weinans

Johan Bollen

This chapter is submitted to *Science Advances*

Abstract

The post-truth era where ‘feelings trump facts’ has taken many by surprise after a century of societal, technological, and scientific progress. Here we demonstrate that the rise of fact-free argumentation may be understood as part of a deeper change. We analyze text from millions of books revealing that after the industrial revolution the use of sentiment-laden words declined systematically, while the use of words associated with fact-based argumentation such as ‘test’, ‘conclude’ and ‘information’ rose steadily. However, this pattern stagnated in the 1980s and reversed sharply around 2007, when across languages the frequency of fact-related words dropped suddenly while emotion laden language surged, together with words related to ‘spirituality’, ‘intuition’ and ‘feeling’. Within two decades language sentiment has bounced back to pre-war levels while interest in words related to science and societal organization dropped sharply. This reversal of the historical trend happened in close synchrony with the rise of social media.

Introduction

The surge of fact-free, post-truth political argumentation suggests that we are living a special period in history when it comes to the balance between emotion and reasoning. However, quantifying this intuitive notion remains difficult as systematic surveys of public sentiment and world-views do not have a very long history. Here we address this gap by systematically analyzing word-use in millions of books across eight languages covering the period from 1850 to 2019 (Michel et al., 2011). Reading this amount of text would take a single person millennia, but computational analyses of trends in relative word frequencies may hint at aspects of cultural change (Michel et al., 2011; Kesebir and Kesebir, 2012; Lin et al., 2012). Print culture is selective and cannot be interpreted as a faithful representation of culture in a broader sense (Pettit, 2016). Also, the popularity of particular words and phrases in a language can change for many reasons including technological context (e.g. carriage or computer), and the meaning of some words can change profoundly over time (e.g. gay). Nonetheless, across large amounts of words, patterns of change in frequencies may reflect changes in the way people feel and see the world (Michel et al., 2011; Kesebir and Kesebir, 2012; Lin et al., 2012), assuming that concepts that are more abundantly referred to in books in part represent concepts that readers at that time were more interested in. Here, we systematically analyzed long-term dynamics in the frequency of the 5000 most used words in American English, German and Spanish (Brysbaert et al., 2019) in search of indicators of changing world views. We compare patterns for selected key words in the corpuses of Chinese, French and Russian as well as British English and English fiction to gauge the robustness and generalizability of our results.

Principal components of change

As a first unbiased exploration of language change we perform a principal component analysis (PCA) of changes in word frequencies over time (supplementary material 1 and 2). This approach seeks to capture patterns of change in a large dataset without relying on prior assumptions. In each of the three languages, a dominant principal component emerges that corresponds to a monotonic trend over time, and another one that shows an asymmetric U-curve, or ‘tilted hockeystick’, declining gradually since the industrial revolution, and surging sharply on recent decades (Fig. 6.2 left hand panels). Examination of the words that score highly on opposite ends of either principal component axis (Supplementary material 8) suggests that, in each language the monotonic axis captures general trends of word popularity over time. On the low end (representing earlier times) we find more archaic terms such as favor, mere, wit, iron, liberty and crown. On the high end (corresponding to more recent times) we find words such as computer, privacy, lesbian, casino and taxi. By contrast, the tilted hockeystick axes are dominated on the high side (more recently) by words reflecting human nature related concepts such as personal pronouns, sensory concepts and emotion laden words (Table 1). By contrast, the opposite end this axis (earlier times) has words related to rational decision making, procedures and

systems (Table 2). Tellingly, in German where it carries a marked signal of the two world wars, this emotion-linked hockeystick pattern is in the first principal component, whereas it comes second in American English and Spanish.

Sentiment trends

As a next step we analyzed changes in the relative frequencies of words that have been independently assessed as indicators of different aspects of emotion (using the ANEW lexicon; Affective Norms for English Words) (Nielsen, 2011; Warriner et al., 2013) and comparable lexicons for German and Spanish (See supplementary material 3). ‘Valence’ or pleasantness associated with a word is a dominant aspect of many models of human emotion. ANEW valence values range from unhappy (e.g. for words such as torture, rape, terrorism) to happy (e.g. words such as vacation, enjoyment, free). Some models of human emotions also include an orthogonal affective dimension of arousal, which can be evoked by words going from none (e.g. dull, scene, asleep) to strong (e.g. rampage, sex, tornado). Integrating frequency-weighted valence and arousal sentiment scores over our set of 5000 words we arrive at an index of positive and negative sentiment as well as arousal for the entire content of books published each year since 1850. In all three languages we find similar patterns of these affective indicators, comparable to the hockeystick axis of the PCA (Fig. 6.2 second column of panels). Indeed, sentiment indicators are correlated to this principal component of language change in all languages (see supplementary table S1).

Correlated concepts

Although sentiment analysis allows for an objective evaluation of the affective components of language change, sentiment indicators do not necessarily capture the full essence of the marked pattern revealed by the Principal component analysis. This becomes evident if we examine the words correlating and anticorrelating most strongly to the PCA axis, to the sentiment score, and to the overall hockeystick pattern. A glance at the 1% top scoring words for American English reveals that the patterns we find correspond to a seesaw between two markedly opposing poles of concepts (Table 1). On the one hand we have words that may be broadly characterized as related to human nature (Table 1, top row). At the opposite end there are words that could be characterized at first glance as related to societal systems (Table 1 bottom row). The human nature words contain many so-called stop words which through their sheer abundance tend to dominate language analyses (Perkins, 2014; Saif et al., 2014). Removing those eliminates personal pronouns which may be considered characteristic of the human oriented nature, but also allows a richer spectrum to appear in the top 1% (supplementary table S2). A closer look at the unfiltered and filtered, short and long lists of words suggests that both the human and the society poles encompass several distinct groups of concepts. On the human nature pole we have words that relate to sentiments (tear, delicious, cold, lovely), words

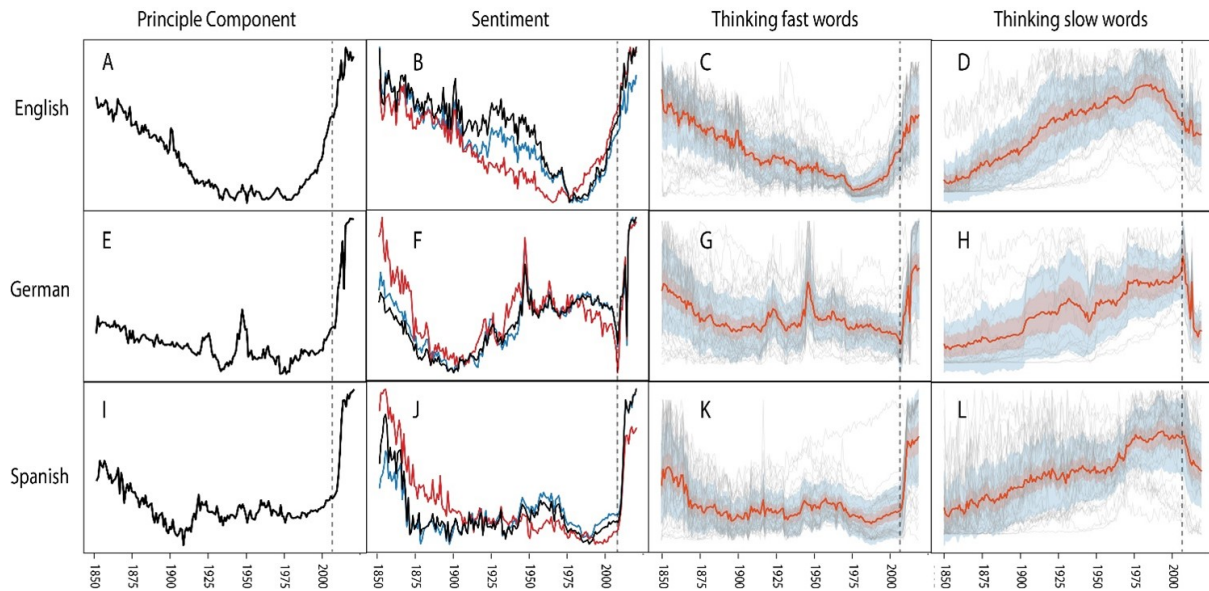


Figure 6.1: Dynamics of American English, German and Spanish book language characterized in four complementary ways. A, E and I) Principal component of change in frequency of the 5000 most used words. B, F and J) Relative level of arousal (black), positive sentiment (blue) and negative sentiment (red). C, G and K) time development of the frequencies of flag-words related to intuition, believing, spirituality, sapience: *spirit, imagine, wisdom, wise, hunch, mind, suspicion, believe, think, trust, faith, truth, true, belief, doubt, hope, fear, life, soul, heaven, eternal, mortal, holy, god, pray, mystery, sense, feel, soft, hard, cold, hot, smell, foul, taste, sweet, bitter, hear, sound, silence, loud, see, light, dark, bright, always, never, everything, nothing, surely*. All frequencies are standardized so that they range from zero to unity. The central line represents the mean, the red shaded area the 95% confidence interval of the mean, and the blue shaded area the standard deviation. D, H and L) similar but for flag-words related science, technology, quantification, economy, procedure: *science, technology, scientific, chemistry, chemicals, physics, medicine, model, method, fact, data, math, analysis, conclusion, limit, result, determine, transmission, assuming, system, size, unit, pressure, area, percent, business, company, payment, expenses, manager, profit, investment, market, employee*. Timelines for a few selected English example words are shown in Fig. 2.




Figure 6.2: Examples of American English words associated to intuition and faith-based views (System-I thinking, panel A) and rational evidenced-based views (System-II thinking, panel B).

associated to spirituality (God, mind, spirit, heaven) and relationships (father, mother, friend, brother), personal pronouns (they, us, him, them), and absolutes (never, nothing, ever). By contrast, at the opposing pole tentatively labeled as ‘societal systems’ we have words such as education, office, city, and employment, but also words related to science and technology (result, determine, medical, tested, size, physics, transmission).

To test if those tentatively discerned concepts do indeed contribute individually to the seesaw that we see in the PCA and sentiment, we examined the historical dynamics of the separate concept groups. Groups such as personal pronouns and absolutes are straightforward to delineate. To populate the other word groups we used a thesaurus algorithm available at relatedwords.org (combining search techniques such as word embedding and Concept-Net). Examining dynamics of each of these separate groups reveals a striking synchrony, confirming that within each language the shifts in interest across the concepts we identified happen very much in concert (Supplementary material 4, 5 and 6).

Arguably, the opposing poles of human versus society concepts we identified may also be interpreted in terms of how they relate to two fundamentally different cognitive modes of operation (Han and Pöppel, 2011; Allen and Thomas, 2011; Baas et al., 2008; Morewedge and Kahneman, 2010), namely system-I (‘thinking fast’, loosely intuition) vs system-II (‘thinking slow’, loosely rational). We test this idea by exploring clusters of words that we now filter to specifically reflect those opposed modes of thinking (for details see supplementary information 4). Selected, system-I flag words are rooted in the concepts spirituality, sapience, intuition, believing and dichotomization, while the system-II flag

words we used are rooted in the concepts science, technology, quantification, business and economy. Plotting the dynamics of the frequencies of words in those clusters supports the view that the PCA and sentiment patterns we revealed are closely related with a systematic change along the rationality-intuition gradient (Fig. 1 right-hand panels and Fig. 2). This pattern remains robust if smaller focused subsets are chosen (Supplementary Fig. S5).

In summary, the observed seesaw pattern consist of two trends. First, the steady decline in the use of sentiment laden words starting after the industrial revolution and continuing up to about 1980 went hand-in-hand with a decrease in the use of words related to intuitive thinking, and a rise of terms associated with argumentation based on data and objective procedures. The robust set of trends we find from 1850 to 1980 is suggestive of an ongoing development towards a rising interest in rational arguments and data-based approaches. Second, the subsequent stagnation of those trends around 1980 and the sharp reversal in most languages around 2007 suggests a societal revaluation of the role of intuition and spirituality, coupled with a declining interest in science, technology, and procedures, consistent with a development towards rejection of fact-based argumentation.

Robustness and potential biases

It is possible that these patterns are in part artifacts of the Google Books data and our choice of words. With respect to the latter, the 5000 most frequent words in any language represents an overwhelming sample of common language use, buffering against the problem that any individual word may be subject to fashions or change meaning. Our findings may however be subject to other systematic biases. For instance, our list of 5000 most common words and the emotional ranking of words was determined in recent years and therefore reflect relatively recent language use. Also, an important caveat of using book texts is that they are a biased representation of language, a bias that may change over time (Zhang, 2015; Pechenick et al., 2015). What ends up in the university libraries used for the Google n-gram data varies with trends in editorial practices, library policies, and popularity of genres. For instance, over time there has been an increase in the relative proportion of nonfiction (a trend that noted as early as the 1950's (Walters, 1953)) and scientific works (Zhang, 2015; Pechenick et al., 2015). Importantly, while such developments could logically contribute to the long-term trends we find up to 1980, they cannot explain the reversal after 1980, as the rise of nonfiction and science books continues until today (Rowe, 2018).

It is also worth noting that the link between book language and social sentiment has been validated in other studies (Hills et al., 2019), and that the long-term trend we find until 1980 is in line with what has been found in other studies including different text corpuses. For instance, a study comparing a corpus of New York Times articles from 1851 to 2015 and the Google books corpus 1800-2000 showed (Iliev et al., 2016) that in both

corpora there has been a significant downward trend in positive as well as negative words as classified through the LIWC system (Pennebaker et al., 2001). Another study, using a New York Times corpus, a corpus of Scientific American articles and Google books, revealed that over the past two centuries in all those corpora there has been a significant increase of words reflecting causal reasoning as reflected by words in the ‘cause’ category in the LIWC system (a list of 108 words such as because, since, hence, how, why, depends, and implies) (Iliev and Axelrod, 2016).

In conclusion, studies using different corpora as well as different marker words confirm the long term decline of sentiment laden (positive and negative) language and a rise of words related to causal reasoning. Meanwhile, the stagnation of this trend around 1980 and the subsequent dramatic reversal we reveal obviously cannot be explained by ongoing trends towards non-fiction. Thus, while it will be important to explore alternative word collections, sentiment classifications and text corpora it seems likely that the marked U-shaped pattern we find reflects a true dimension of change.

Potential drivers

Inferring the drivers of this stark pattern necessarily remains speculative, as language is affected by many overlapping social and cultural changes. Nonetheless, it is worthwhile reflecting on a few potential mechanisms. One possibility when it comes to the trends from 1850 to 1980 is that the rapid developments in science and technology and their socio-economic benefits drove a rise in status of the scientific approach, which gradually permeated culture, society, and its institutions ranging from education to politics. As argued early on by Max Weber, this may have led to a process of “disenchantment” as the role of spiritualism dwindled in modernized, bureaucratic, secularized societies (Jenkins, 2000; Storm, 2017). It remains difficult to pinpoint what precisely caused the observed stagnation in the long-term trend around 1980 visible in the PCA analyses of German and Spanish, but particularly clear in all indicators of change in American English. However, one could speculate that there is a connection to tensions arising from the increasing neoliberal policies which were defended on rational arguments, while the economic fruits were reaped by an increasingly small fraction of societies (Dehm, 2018; Duménil and Lévy, 2005; Razavi et al., 2020).

Much more spectacular, however, than the 1980 stagnation is the sharp change we see across languages around 2007 (Fig. 3). This change coincides with the global financial crisis which may have had an impact. However, earlier economic crises such as the great depression (1929-1939) did not leave discernable marks on our indicators of book language. This suggests we are observing a societal trend that transcends financial and economic concerns. Perhaps more significantly, 2007 was also roughly the start of a near-universal global surge of social media. This becomes obvious if we plot dynamics of the word ‘Facebook’ (or Weibo in Chinese) alongside the frequency of the word for ‘science’ as a

marker for interest in concepts related to the ‘slow thinking’ cluster in eight different languages (Fig. 3).

Various lines of evidence underpin the plausibility of an impact of social media on emotions, interests and worldviews. For instance, there are negative effects of the use of social media on subjective wellbeing (Hanley et al., 2019). This may in part be related to distortions such as the perception that your friends are more successful, have more friends and are happier (Feld, 1991; Bollen et al., 2017) and more beautiful (Fardouly and Holland, 2018) than you are. At the same time, a perception that problems are abound may have been fed by activist groups seeking to muster support (Gerbaudo and Treré, 2015) and lifestyle movements seeking to inspire alternative choices (Haenfler et al., 2012). For instance, social media catalyzed the Arab spring among other things by depicting atrocities of the regime (Breuer et al., 2015), jihadist video’s motivate terrorists by showing gruesome acts committed by US soldiers (Weimann et al., 2014), and veganism is promoted by social media campaigns highlighting appalling animal welfare issues (Haenfler et al., 2012). Many of the problems highlighted on social media will be real, and often they may have been largely hidden from the public eye in the past. However, independently of whether problems are exaggerated or merely revealed online, the popular effect of such awareness campaigns may be the perception of an unfair world entangled in a multiplicity of crises. Further down the gradient from revelation to exaggeration we find misinformation. The spread of misinformation (Southwell et al., 2018) and conspiracy theories (Douglas et al., 2019) may be efficiently facilitated by social media. Recent research shows that the online diffusion of false news is significantly broader, faster, and deeper than that of true news and efforts to debunk (Vosoughi et al., 2018). Conspiracy theories originate particularly in times of uncertainty and crisis (Van Prooijen and Douglas, 2017; Douglas et al., 2019), and generally depict established institutions as hiding the truth and sustaining an unfair situation (Van Prooijen et al., 2015). As a result, they may find fertile grounds in social media promulgating a sense of unfairness, subsequently feeding anti-system sentiments. Neither conspiracy theories, the exaggerated visibility of the successful, nor the overexposure of societal problems are new phenomena. However, social media may plausibly have acted as an amplifier of societal arousal and sentiment, potentially stimulating an anti-system backlash, including its perceived emphasis on rationality and institutions.

Outlook

It seems unlikely that we will ever be able to accurately quantify the role of different mechanisms driving language change. However, the universal and robust shift that we observe suggests a historical rearrangement of the balance between the rational and the emotional. It is consistent with a sharp reversal of a long term rise in the interest in ‘slow thinking’ and fact-based argumentation. Careful weighing of evidence and arguments, is essential to science, technology, and governance, the very institutions upon which modern




Figure 6.3: Synchrony of recent change across languages. Curves represent the frequency of words corresponding to the personal pronoun 'I' (*Ich, I, je, yo, 我, я, I, scienza, io*), and 'science' (*science, science, ciencia, 学术, научный, science, scienza*). The rising curves starting around 2007 in each panel correspond to the word 'Facebook' (for Chinese, Weibo: 微博). The grey box marks the period after 2007, characterized by the rise of social media and by a declining use of words related to science, technology, quantification, economy, procedure (see also figures 6.1 and 6.2).

societies are based. The complementary mode of ‘fast thinking’, requires not only less time, but also less cognitive energy (Kahneman, 2017). The rapid pace of social media discussions may provoke fast emotional responses. More generally, ‘fast thinking’ becomes more prominent when people lack the resources to make carefully considered judgments which may have a ranges of causes and consequences. For instance, the surge of absolutes (e.g. never, always, all) is consistent with black-and-white thinking, which correlated among other things to low scores on cognition tests, poor educational backgrounds, harsh childhood years, impulsive behaviors, and aggression (Oshio et al., 2016; Mieda et al., 2021).

Clearly, rational, fact-based approaches are essential for maintaining functional democracies and addressing global challenges such as global warming, poverty and the loss of nature. The trends we observe suggest a profound shift in societal interest away from ‘slow-thinking’ towards emotionality and subjective experiences. It may well be impossible to reverse this sea-change. Instead, societies may need to find a new balance, explicitly recognizing the importance of intuition and emotion, while at the same time making best use of the much needed power of rationality and science to deal with topics in their full complexity (Kofman, 2018).

Table 6.1: Human nature (top row) vs societal organisation related words (bottom row) emerging from correlation with Principle Component results, sentiment analysis and the hockeystick pattern. Listed are the 1% words that score highest vs lowest on the PCA axis depicted in Fig. 1, the 1% that correlating most positively vs negatively with sentiment, and the 1% words that increased most clearly (measured as Kendall tau) after 1980 while declining between 1850 and 1980 vs words that show the opposite pattern. We used positive sentiment for computing the correlations, but this is closely correlated to negative sentiment and arousal too. The human nature words contain many so-called stop words (8, 9). Removing those allows a richer spectrum to appear in the top 1%, but also eliminates personal pronouns which may be considered characteristic of the human oriented nature (supplementary table S2).

<p>Words scoring highest on surging PCA axis: I, he, his, and, to, it, her, was, you, she, that, but, had, my, not, they, a, him, their, me, them, so, with, as, what, we, who, said, when, your, all, do, have, then, god, like, if, into, could, upon, how, did, one, which, up, us, can, our, great, the</p>	<p>Words correlating most positively to sentiment: when, had, forgot, tear, wonder, he, thing, eager, rushing, come, gone, mind, lighten, they, forgotten, delicious, rest, never, happier, lost, rid, brother, forget, stood, dress, silent, night, forgetting, cold, take, again, glow, warm, perfume, here, speaking, die, lovely, quiet, dead, uncle, who, stray, into, hear, deep, gorgeous, drew, thirsty</p>	<p>Words rising after 1980 and declining before that year: thrown, perfectly, admit, ashes, perfect, every, true, believing, hearts, them, throw, prayer, ever, mistaken, place, safely, faithful, regret, distant, so, soon, supposed, understood, embrace, touching, spirit, lay, mighty, kindness, alone, nothing, but, shed, constantly, besides, fallen, sight, him, shine, bless, never, thankful, glory, ruin, harmony, owe, awful, grave, seat, before</p>
<p>Words scoring lowest on surging PCA axis: board, education, industry, water, plant, material, field, office, unit, city, method, air, employment, farm, interest, agreement, building, present, pressure, health, table, market, social, highway, radio, hospital, machine, paper, metal, month, milk, quality, item, outstanding, phase, part, avenue, failure, vehicle, truck, food, car</p>	<p>Words correlating most negatively to sentiment: limit, manual, transfer, current, section, result, determine, issue, technical, medical, request, final, phoenix, tested, separate, affected, prior, national, failure, limited, base, size, recent, private, addition, growth, physics, transmission, chemical, local, policy, agreement, library, assuming, standard, delta, notify, reliable, central, chemicals</p>	<p>Words declining after 1980 and rising before that year: program, basis, personnel, available, equipment, technical, determine, area, hearing, million, percent, indicate, final, funds, result, addition, separate, procedure, federal, basic, rate, automatic, maximum, development, budget, base, manual, included, responsible, replacement, plus, transportation, involved, staff, control, training, nuclear, billion, limited, initial</p>

Supplementary materials

S6.1 Data and preprocessing

We used the 2019 release of the Google Books n-gram data which the Google Books team made freely available (<https://storage.googleapis.com/books/ngrams/books/datasetv3.html>). The data covers a period from the 16th century to the year 2019. Time series from 1850 till 2019 of the 5000 most frequent words in US English (Brysbaert et al., 2019) were downloaded from the Google n-gram viewer at <https://books.google.com/ngrams>. To correct for the increasing volume of text towards more recent times, we scaled the time series dividing word frequencies by the frequency of the word ‘an’. In addition, the word ‘war’ was removed. The sentiment of ‘war’ is 2.08 for ANEW (on the original scale from 0 to 10) and the word has a massive increase in frequency around the two world wars. Therefore it would skew the sentiment analysis around war periods.

S6.2 Principal Component Analysis

PCA was performed on time series of all words after the preprocessing step. A PCA finds a direction in a multidimensional space where the variance is maximized, i.e. where most dynamics are observed. The multidimensional space in our case, consists of axes that describe the frequency of a word, so every axis reflects the dynamics of one word in our analysis. The obtained coefficients of the principal components reflect the weight that a word has on that particular principal component. The 5% words with the highest coefficient and 5% words with the lowest coefficient were investigated and analyzed for their sentiment. The complete time series were projected on the first two principal components in order to evaluate the dynamics of the language as a whole in the direction of these PC’s.

Words scoring strongly in the two opposite directions on either axes are given in the next sections. Examining those words suggests that the first axis in American English and Spanish and the second axis in German capture mostly the natural turnover of vocabulary while the other axis in each language is dominated on the high side by words reflecting human nature related and emotion laden concepts, while the opposite end is populated by concepts related to rational decision making, procedures and systems. This ‘human nature’ axis invariably surges in recent decades and is correlated to the sentiment value of words (Table S1).

S6.3 Computing sentiment scores

For each of the words we assigned scores for mean valence and arousal if such scores were available in the sentiment lists that we used for English (Warriner et al., 2013), German (Schmidtke et al., 2014) and Spanish (Redondo et al., 2007). Since we aimed to separate

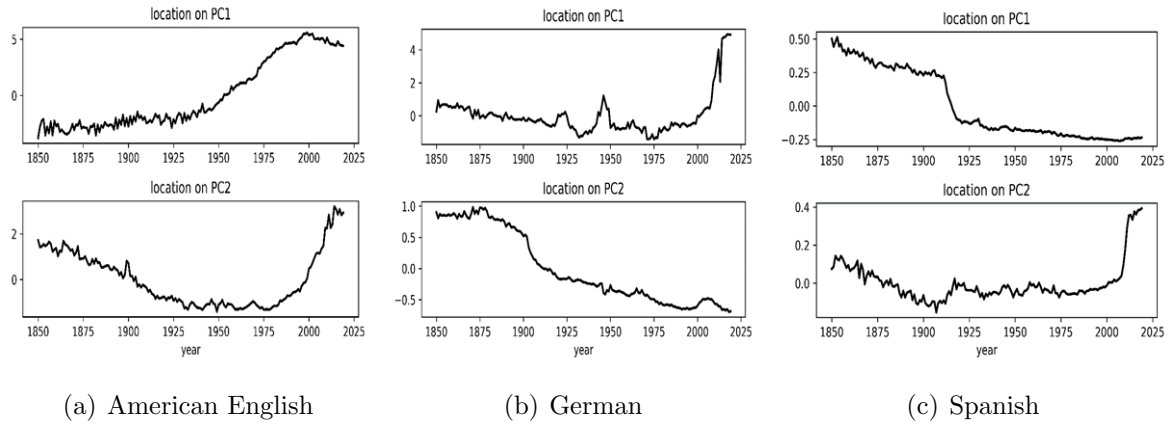


Figure S6.1: Time development of the 5000 most used words in three languages represented by the weight on their first two principal components. Note that in each language there is one axis that correlates strongly to sentiment scores and surges in recent decades. This is the axis which we examine further in the main text. Explained variance for English, German and Spanish resp. are 0.83, 0.72 and 0.77 for PC1 and 0.13, 0.17 and 0.12 for PC2.

Table 6.2: Spearman rank correlations (ρ) of positive and negative sentiment with the PC that shows an inverse U curve (this hockeystick pattern occurs in PC2 for American English and Spanish, but PC1 for German).

		Positive Sentiment			Negative sentiment		
		rho	n	p-value	rho	n	p-value
American English	PC2	0.24*	435	7e-7	0.17*	251	0.007
German	PC1	0.06	279	0.33	-0.13	144	0.12
Spanish	PC2	0.14	263	0.03	-0.01	138	0.90

negative and positive sentiment in our analysis, we first subtracted 5 from the ANEW valence scores, to obtain scores from -5 to 5 (-3 to 3 for the German list), where all scores lower than zero reflected negative sentiment whereas all scores higher than zero reflected positive sentiment. Next, we multiplied the scaled time series with their sentiment score to obtain a sentiment per word per time. The sentiments per word for each year were summed over all words, resulting in a sentiment per moment in time.

S6.4 Flag-word selection

Our exploration of words that scored strongly on the hockeystick PCA axis or correlated strongly with sentiment led us to tentatively define groups of correlated concepts that seemed to go hand-in-hand. To check for each of those concepts if they indeed follow similar dynamics, we first populated sets of flag-words related to each of the concepts using a thesaurus algorithm available at relatedwords.org (combining search techniques such as word embedding and Concept-Net). Subsequently we plotted dynamics of each group of flag-words separately. As we are using the 5000-word most frequent word collections for plotting, not all of the words resulting from the thesaurus search end up in the final selections. Also, inevitably, there is no precise 1-to-1 connection between words in the different languages.

Human related word categories

1. Spirituality, intuition, believing and sapience : spirit, imagine, wisdom, wise, hunch, mind, suspicion, believe, think, trust, faith, truth, true, belief, doubt, hope, fear, life, soul, heaven, eternal, mortal, holy, god, pray, mystery
2. Senses: sense, feel, soft, hard, cold, hot, smell, foul, taste, sweet, bitter, hear, sound, silence, loud, see, light, dark, bright
3. Relationships: father, mother, sister, brother, friend
4. Personal pronouns: I, we, you, he, she, they, mine, our, their

Society related word categories

1. Science & Technology: science, technology, scientific, chemistry, chemicals, physics, medicine, model, method, fact, data, math, analysis, conclusion, limit, result, determine, transmission, assuming, system
2. Quantification: size, unit, pressure, area, percent
3. Business and economy: business, company, payment, expenses, manager, profit, investment, market, employee
4. Social organization: commission, lawyer, government, law, nation, community, administration, education, city, agreement, health, policy, central

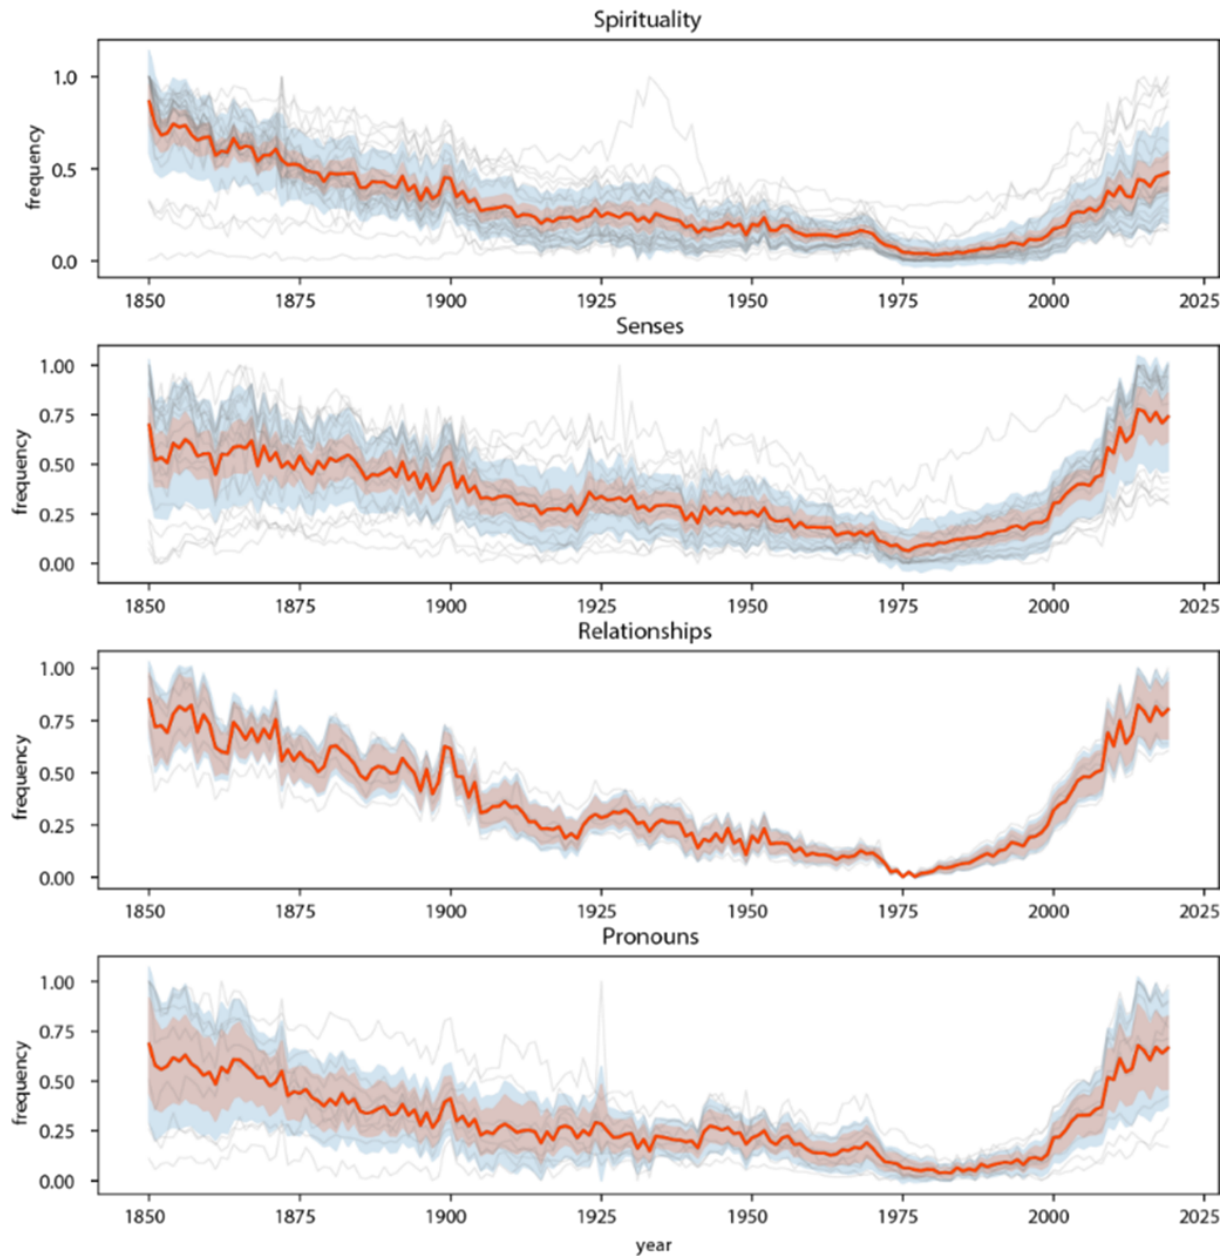


Figure S6.2: Time development of the frequencies of separate groups of human-related flag-words (see text above). All frequencies are standardized so that they range from zero to unity. The central line represents the mean, the red shaded area the 95% confidence interval of the mean, and the blue shaded area the standard deviation.

Two clusters related to cognition modes

Lastly, we explore the idea that the human versus society groups of words may also correspond to two types of thinking that have been shown to represent fundamentally different cognitive modes of operation (Han and Pöppel, 2011; Allen and Thomas, 2011; Baas et al., 2008; Morewedge and Kahneman, 2010): system-I (‘thinking fast’, loosely intuition) vs system-II (‘thinking slow’, loosely rational). To slightly tailor our word groups to this end we removed some concepts and added others and subsequently plot the behavior of the two resulting over-arching clusters as a whole. To turn the human related cluster into an System-I related cluster we removed relationships and personal pronouns, and added absolutes. To turn the society related cluster into an System-II related cluster we removed words related specifically to social organization.

System-I flag words thus become words related to

1. Spirituality, intuition, believing and sapience (as above)
2. Senses (as above)
3. Absolutes: always, never, everything, nothing, surely

Resulting in: spirit, imagine, wisdom, wise, hunch, mind, suspicion, believe, think, trust, faith, truth, true, belief, doubt, hope, fear, life, soul, heaven, eternal, mortal, holy, god, pray, mystery, sense, feel, soft, hard, cold, hot, smell, foul, taste, sweet, bitter, hear, sound, silence, loud, see, light, dark, bright, always, never, everything, nothing, surely

System-II flag words are now related to:

1. Science & Technology (as above)
2. Quantification (as above)
3. Business & Economy (as above)

resulting in: science, technology, scientific, chemistry, chemicals, physics, medicine, model, method, fact, data, math, analysis, conclusion, limit, result, determine, transmission, assuming, system, size, unit, pressure, area, percent, business, company, payment, expenses, manager, profit, investment, market, employee

Dynamics of those clusters are presented in Fig. 6.1 of the main text. The additional group of absolute words associated to dichotomous all-or-nothing thinking (Mieda et al., 2021; Oshio et al., 2016) follows similar dynamics as illustrated by Fig. S6.4.

Robustness of cognition mode results to differently-sized sets of words

In the figures presented in the main text for System-I and System-II words, some clusters are more prominent than others, because the number of words differs per cluster. We

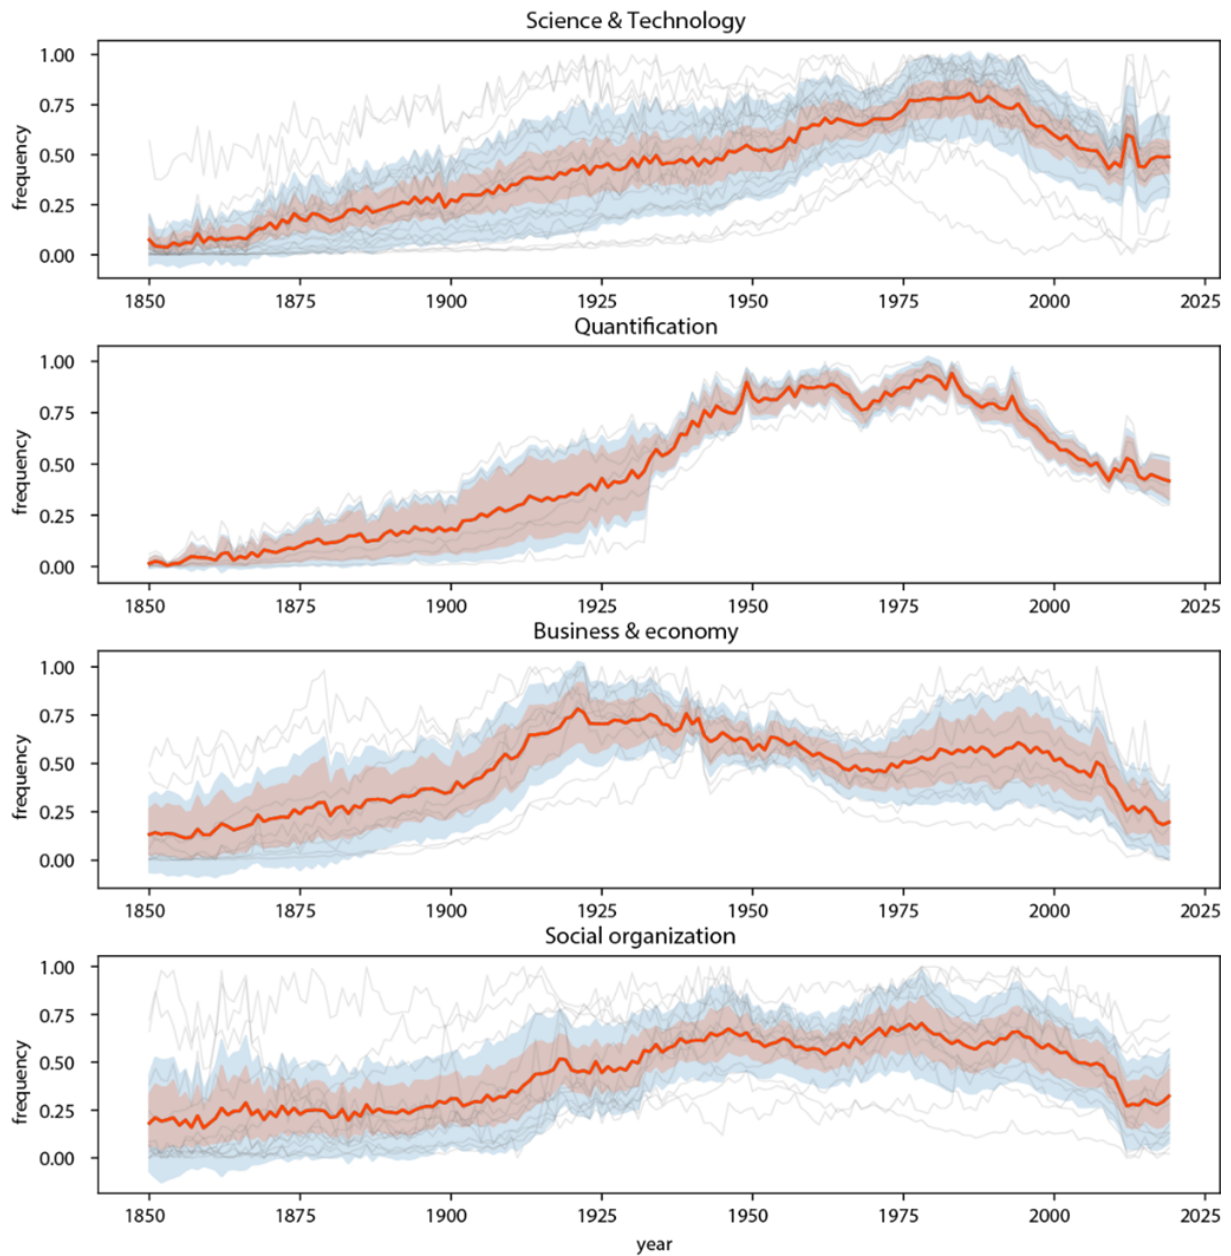


Figure S6.3: Time development of the frequencies of separate groups of society related flag-words in (see text above). Use of colors and shading as in Fig S6.2.




Figure S6.4: Time development of the frequencies of absolute words (see text above). Use of colors and shading as in Fig S6.2.

therefore checked the robustness against this effect by picking five representative words per cluster:

Subsets of system-I flag words, five words per cluster

1. Spirituality, intuition, believing and sapience : spirit, mind, trust, faith, god
2. Senses: sense, smell, taste, hear, see
3. Relationships: father, mother, sister, brother, friend
4. Personal pronouns: I, we, you, he, she

Resulting in the following list of System-I flag words: spirit, mind, trust, faith, god, sense, smell, taste, hear, see, father, mother, sister, brother, friend, I, we, you, he, she, spirit, mind, trust, faith, god, sense, smell, taste

Subsets of system-II flag words, five words per cluster:

1. Science & Technology: science, technology, method, data, analysis
2. Quantification: size, unit, pressure, area, percent
3. Business and economy: business, payment, manager, investment, market

Resulting in the following list of System-II flag words: science, technology, method, data, analysis, size, unit, pressure, area, percent, business, payment, manager, investment, market

Dynamics of these five-word clusters show very much the same pattern as found for the full sets of words, albeit with less variance (Fig. S6.5).

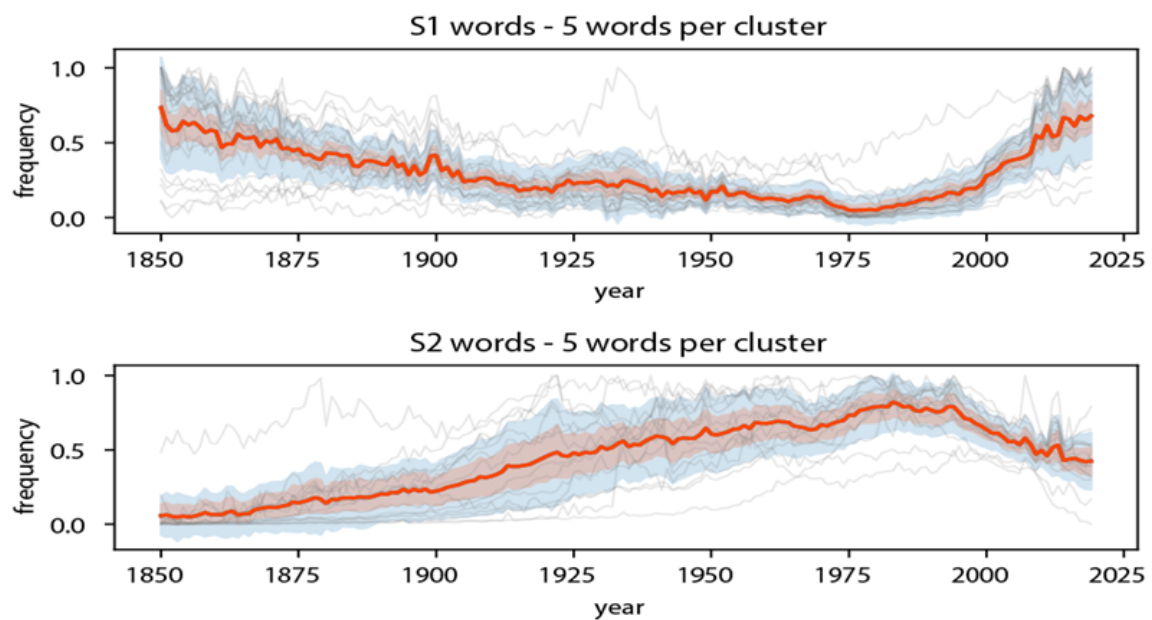


Figure S6.5: Time development of the frequencies of smaller subsets of System-I versus System-II flag words (see text above). Use of colors and shading as in Fig S6.2.

S6.5 Flag-words in German

Human related word categories

1. Spirituality, intuition, believing and sapience: Geist, Ahnung, Verstand, denke, Vertrauen, glaube, Wahrheit, Zweifel, Hoffnung, Angst, Leben, Seele, Himmel, ewig, Gott, beten, weise
2. Senses: Gefühl, weich, hart, kalt, heiß, Geruch, übel, Geschmack, süß, hören, Ton, laut, sehen, Licht, dunkel, hell
3. Relationships: Vater, Mutter, Schwester, Bruder, Freund
4. Personal pronouns: Ich, wir, du, er, sie, meine, unsere, ihre

Society related word categories

1. Science and technology: Wissenschaft, Technologie, Medizin, Modell, Tatsache, Daten, Grenze, Ergebnis, System
2. Quantification: Größe, Einheit, Druck, Prozent
3. Business and economy: Unternehmen, Manager, Arbeit, Markt, Mitarbeiter
4. Social organization: Regierung, Gesetz, Nation, Gesundheit, Politik

Absolutisms

The translation of the class of absolute words we used is: immer, nie, alles, nichts, sicher

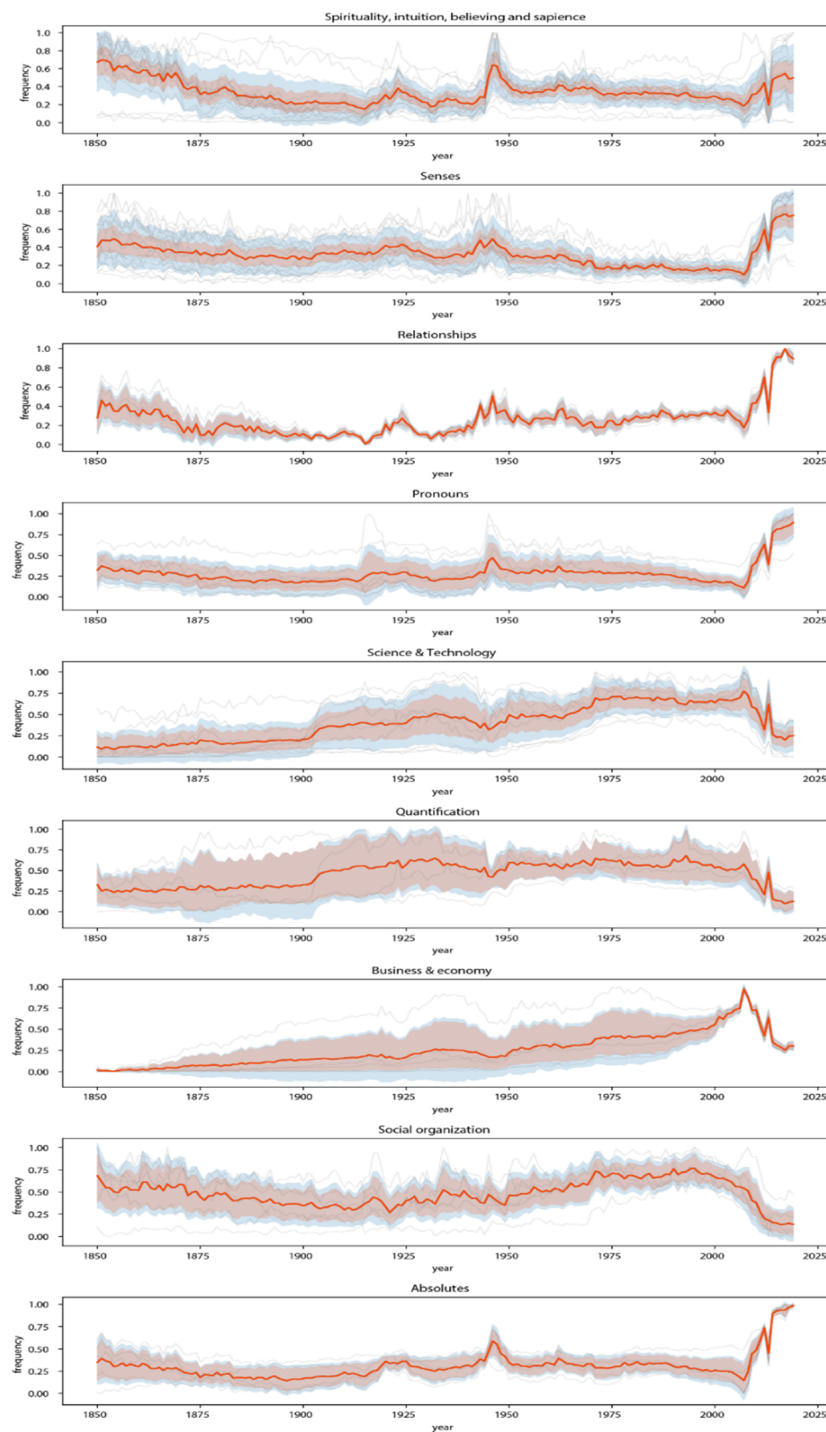


Figure S6.6: Time development of the frequencies of separate groups of flag-words in German (see text above). Use of colors and shading as in Fig S6.2.

S6.6 Flag-words in Spanish

Human related word categories

1. Spirituality, intuition, believing and sapience: espíritu, imaginar, sabiduría, mente, sospecha, creer, pensar, fe, verdad, duda, esperanza, miedo, vida, alma, cielo, santo, dios, misterio
2. Senses: sentido, sensación, sentir, suave, duro, frío, caliente, gusto, dulce, oír, silencio, fuerte, ver, mirar, oscuro, brillante
3. Relationships: padre, madre, hermana, hermano, amigo
4. Personal pronouns: yo, nosotros, tú, él, ella, ellos, mío, nuestro, su

Society related word categories

1. Science and technology: ciencia, tecnología, científico, química, productos, física, medicina, modelo, método, dato, datos, hipótesis, estadísticas, cálculo, análisis, conclusión, límite, resultado, determinar, transmisión, sistema
2. Quantification: tamaño, unidad, presión, área, densidad, porcentaje
3. Business and economy: comercio, empresa, pago, gastos, gerente, beneficio, inversión, trabajo, ingresos, mercado, trabajadores
4. Social organization: ministerio, comité, gobierno, ley, nación, comunidad, administración, municipio, educación, ciudad, acuerdo, salud, política, central

Absolutisms

The translation of the class of absolute words we used is: siempre, nunca, todo, nada, seguro

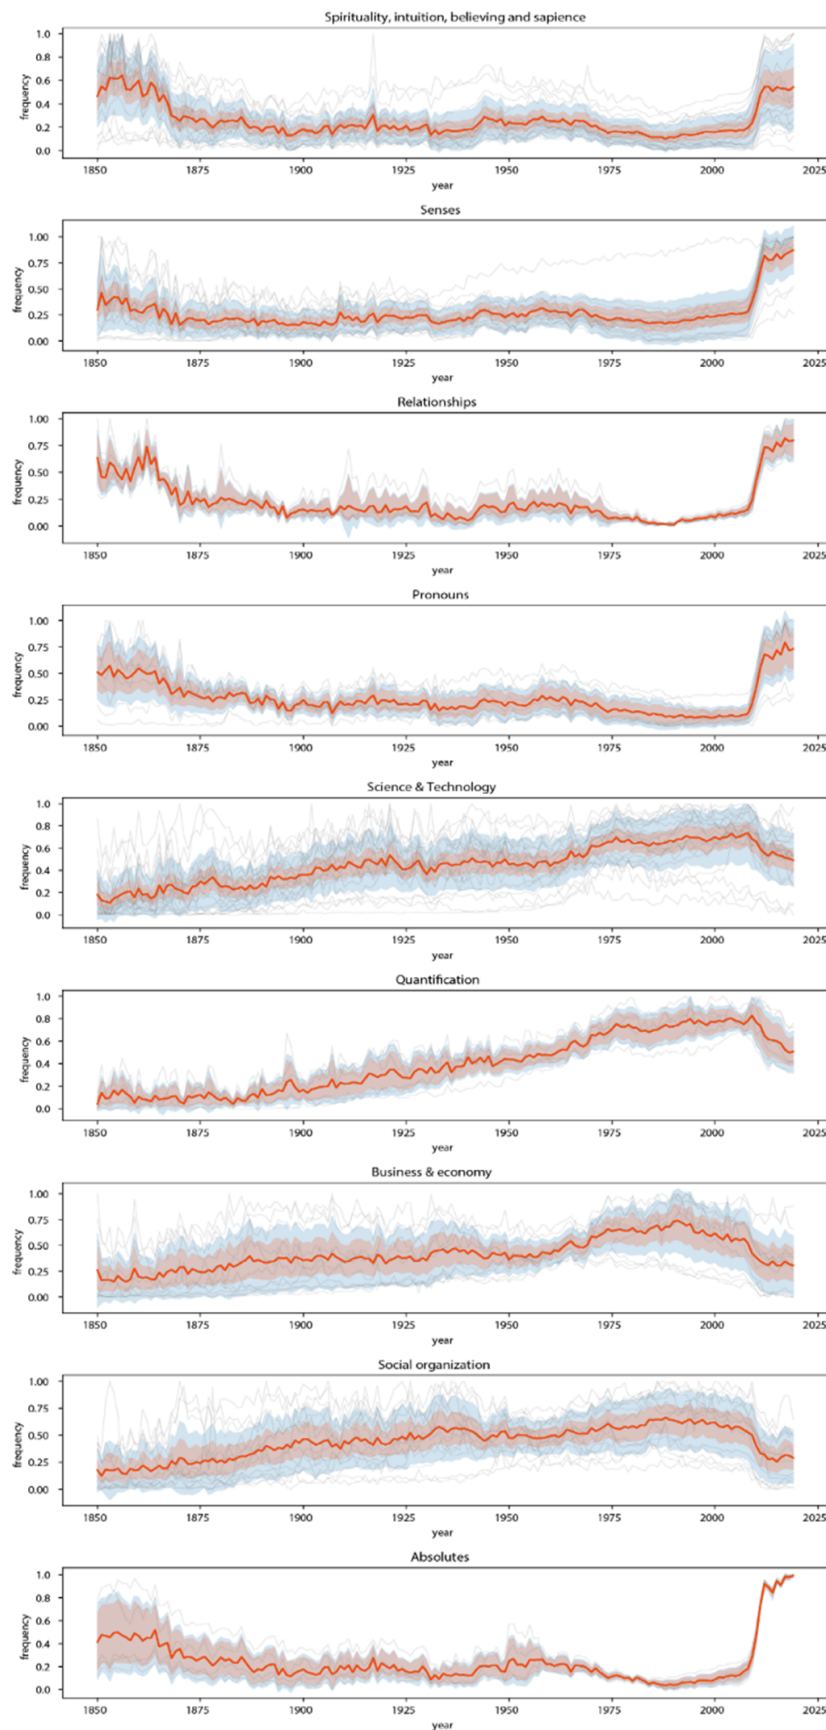


Figure S6.7: Time development of the frequencies of separate groups of flag-words in Spanish (see text above). Use of colors and shading as in Fig S6.2.

S6.7 Top 1% strongest correlated words with PCA excluding stopwords

Table 6.3: Human nature related words obtained as the 1% words that score highest on the PCA axis depicted in Fig. 1, the 1% that correlating most positively with sentiment, and the 1% words that increased most clearly (measured as Kendall tau) after 1980 while declining between 1850 and 1980. We used positive sentiment for computing the correlations, but this is closely correlated to negative sentiment and arousal too. Those lists are filtered for ‘stop words’ using the list provided in the Python Natural Language Processing Tool Kit (1). In an unfiltered version of this table (main text table 1) the list of words scoring highest on the surging PCA axis (let-hand column) is strongly dominated by personal pronouns.

words scoring highest on surging PCA axis	words correlating most positively to sentiment	words rising after 1980 and declining before that year
said, god, like, could, upon, one, us, great,man, every, never, know, back, day, yet, life, love, much, people, good, even,lord, little, still, though, thought, eyes, whole, heart, let, away, way, say, place, men, might, would, ever, hundred, go, house, nothing, left, mind, father, head, come, around, look, took	forgot, tear, wonder, thing, eager, rushing, come, gone, mind, lighten, forgotten, delicious, rest, never, happier, lost, rid, brother, forget, stood, dress, silent, night, forgetting, cold, take, glow, warm, perfume, speaking, die, lovely, quiet, dead, uncle, stray, hear, deep, gorgeous, drew, thirsty, broke, threw, devil, spoke, life, dark, together, lonely, thirsty	thrown, entertain, perfectly, admit, disturb, ashes, perfect, every, great, true, believing, hearts, throw,prayer, ever, mistaken, place, safely, worthy, faithful, regret, distant,declare, hesitate, useless, bear, soon, forbid, supposed, acquaintance, understood, embrace, touching, possessed, sooner, raised, lay, spirit, impress,mighty, pass, altogether, kindness, twelve, alone, whatever, nothing, much, shed, fatal

Chapter 7

A potential feedback loop underlying glacial-interglacial cycles

Els Weinans

Anne Willem Omta

George A. K. van Voorn

Egbert H. van Nes

This chapter is based on:

Els Weinans, Anne Willem Omta, George AK van Voorn, and Egbert H van Nes. A potential feedback loop underlying glacial-interglacial cycles. *Climate Dynamics*, pages 1–13, 2021.

Abstract

The sawtooth-patterned glacial-interglacial cycles in the Earth's atmospheric temperature are a well-known, though poorly understood phenomenon. Pinpointing the relevant mechanisms behind these cycles will not only provide insights into past climate dynamics, but also help predict possible future responses of the Earth system to changing CO₂ levels. Previous work on this phenomenon suggests that the most important underlying mechanisms are interactions between marine biological production, ocean circulation, temperature and dust. So far, interaction directions (i.e., what causes what) have remained elusive. In this paper, we apply Convergent Cross-Mapping (CCM) to analyze paleoclimatic and paleoceanographic records to elucidate which mechanisms proposed in the literature play an important role in glacial-interglacial cycles, and to test the directionality of interactions. We find causal links between ocean ventilation, biological productivity, benthic $\delta^{18}\text{O}$ and dust, consistent with some but not all of the mechanisms proposed in the literature. Most importantly, we find evidence for a potential feedback loop from ocean ventilation to biological productivity to climate back to ocean ventilation. Here, we propose the hypothesis that this feedback loop of connected mechanisms could be the main driver for the glacial-interglacial cycles.

7.1 Introduction

The past 0.9 million years have been characterized by large cycles in global Earth temperature with a periodicity of about 100 ky. During warm (interglacial) periods, ice volume was small and CO₂ concentrations were high. During cold (glacial) time periods, ice volume was large and CO₂ concentrations were on average 90 ppm lower (Ruddiman, 2001). Understanding the processes involved in these cycles can help to pinpoint the relevant processes in the carbon cycle. Given the rapid change that our climate system is undergoing today, knowledge on the existing feedbacks is of particular importance.

There is an ongoing debate about which mechanisms drive the reduction of CO₂ during glacials and which mechanisms drive the rapid increases in CO₂ during deglaciations (Menviel et al., 2018). Interactions with the ocean's carbon reservoir are an obvious first candidate for two main reasons:

1. The ocean is by far the largest carbon reservoir interacting with the atmosphere on the relevant timescale (Fasham, 2003).
2. The $\delta^{13}\text{C}$ of the ocean-atmosphere system was isotopically lighter at the Last Glacial Maximum than during the Holocene (especially in the deep ocean) (Eggelston et al., 2016; Peterson and Lisiecki, 2018). Therefore, it is unlikely that the excess carbon was stored in the terrestrial biosphere or methane hydrates, which are both isotopically light carbon reservoirs (Zeebe and Wolf-Gladrow, 2001).

It has been hypothesized that the excess carbon was stored in an isolated abyssal reservoir (Broecker and Barker, 2007). At the glacial-interglacial transition, carbon from this reservoir would have been released to the atmosphere through upwelling (Broecker and Barker, 2007; Marchitto et al., 2007). However, it has turned out to be difficult to locate a stagnant reservoir of sufficient size to account for the excess carbon (Broecker and Barker, 2007; Broecker and Clark, 2010). Another proposed mechanism that has received considerable attention is that changes in ocean circulation caused the increase in CO₂ (Siegenthaler and Wenk, 1984) by bringing carbon from the deep to the sea surface (Anderson et al., 2009). Ocean circulation or ocean ventilation could have caused the rise of CO₂ in other ways as well. For example, meltwater pulses in the North Atlantic during the last deglaciation may have led to a temporary shutdown of the Atlantic Meridional Overturning Circulation (AMOC) and an associated increase in Antarctic Bottom Water (AABW) (Broecker, 1998). According to a hypothesis by Toggweiler (1999), such a shift would in turn have led to outgassing of CO₂ from the ocean to the atmosphere due to the poor nutrient utilization subpolar Southern Ocean, where AABW is formed. Even so, modeling studies have been ambiguous about the impact of meltwater pulses on the ocean carbon cycle. In some simulations, the addition of meltwater in the North Atlantic led to carbon release from the ocean to the atmosphere (Schmittner et al., 2007; Bouttes et al., 2012; Matsumoto and Yokoyama, 2013; Schmittner and Lund, 2015), whereas it led to a

net uptake of carbon by the ocean from the atmosphere in others (Obata, 2007; Bozbiyik et al., 2011; Chikamoto et al., 2012). According to a set of model experiments by Menviel et al. (2014), the net effect depends strongly on the detailed salt budget.

Other potential mechanisms for CO₂ fluctuations during glacial-interglacial cycles are based on biological activity. For a long time, the main hypothesis was that enhanced productivity in polar regions during glacial times drove down carbon from the atmosphere into the ocean (Sarmiento and Toggweiler, 1984; Sigman and Boyle, 2000). An early hypothesis by Martin (1990) states that enhanced biological productivity during glacial times was caused by iron deposition (or dust deposition). This idea is supported by experiments that show that current densities of phytoplankton (biological productivity) are limited by iron, for example within the Southern Ocean Iron RElease Experiment (SOIREE) project (Boyd and Law, 2001). It has been suggested that changes in iron concentration can be amplified by local feedbacks, for example by increased ocean surface temperatures. These feedbacks could further increase the effect of iron fertilization on climate (Ridgwell, 2002). However, according to the iron fertilization hypothesis, primary productivity in the subpolar Southern Ocean should be higher during glacial times, which is not supported by proxy records (Kohfeld and Chase, 2011; Kohfeld et al., 2005, 2013). Furthermore, general circulation models have shown that atmospheric CO₂ did not respond as much to increased productivity as suggested by earlier box models (Archer et al., 2000). A hypothesis by Omta et al. (2013) describes another role for biological productivity: it suggests that spikes in the densities of marine calcifiers may lead to a quick reduction in sea surface alkalinity which explains the observed rapid increase in temperature and CO₂ during deglaciations. Instead of extracting CO₂ from the sea surface and thus reducing atmospheric CO₂ and temperature, marine productivity increases atmospheric CO₂ according to this explanation.

It is generally assumed that the different mechanisms are linked to each other (Crucifix et al., 2017), and many of them are not mutually exclusive. Therefore, most modeling studies use a combination of proposed mechanisms (Brovkin et al., 2007). Furthermore, many geological variables are synchronized with the Milankovitch forcing (the Summer insolation at subpolar latitudes in the Northern Hemisphere), which makes it hard to distinguish between links between the variables themselves or common links with the forcing (Daruka and Ditlevsen, 2016). What makes the topic even more complicated is the cyclic behaviour that is observed, which may be linked to positive or negative feedback loops involving different elements (Lenton et al., 2008). One example is the ice-albedo effect, where higher temperatures lead to reduced ice caps which leads to a reduction in the albedo, which in turn increases the temperature. Thus, two processes reinforcing each other can give a strong positive feedback loop. The addition of more processes may weaken the feedback, particularly if different scales are involved. This suggests that a potential feedback loop underlying glacial-interglacial cycles would likely be dominated by a few key causal links.




Figure 7.1: Timeseries of Ba/Fe, $\delta^{18}\text{O}$ and Pacific (P), North-East Atlantic (NEA) and East Atlantic (EA) $\delta^{13}\text{C}$ and dust deposition (time left to right).

It is still unclear what these few dominant causal links would be, although many different potential causal interactions have been described (see table 7.1 and figure 7.2). Here we take a statistical approach to investigate whether existing data suggest any particularly strong causal links. The recent increase in available records that go back hundreds of thousands of years and that have improved temporal resolution and quality allows for such an approach. A suitable method for unravelling causal links from nonlinear time series is Convergent Cross-Mapping (CCM) (Sugihara et al., 2012). It has been applied in ecological systems to find causal links between temperature and anchovy and sardine abundance (Sugihara et al., 2012), in physiology to distinguish between normal (healthy) and impaired cerebral autoregulation (Heskamp et al., 2014), in social systems for predicting the behaviour of users of social media (Luo et al., 2014) and in the climate system for detecting directionality in the relation between temperature and CO₂ (Van Nes et al., 2015b). In this study, we apply CCM to published marine and ice-core records spanning the past 800,000 years in order to shed light on the (directions of) interactions between climate, ocean circulation, biological productivity and dust deposition.

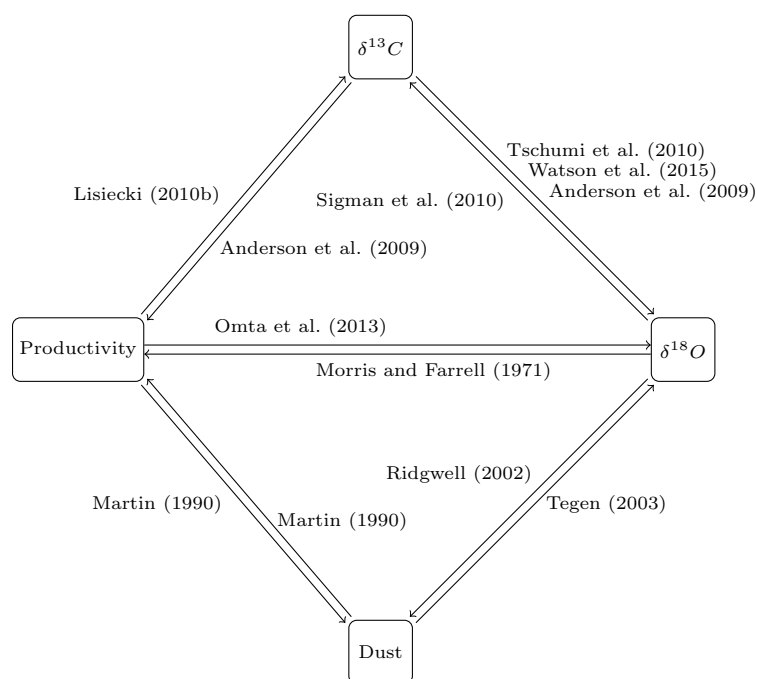


Figure 7.2: A schematic depiction of the links described in the literature between $\delta^{13}C$, biological productivity, $\delta^{18}O$ and dust, indicating that all links are potentially possible. Mechanisms are described in table 7.1.

Table 7.1: Literature describing the links between ocean circulation, biological productivity, $\delta^{18}\text{O}$ and dust. $X \rightarrow Y$ means the hypothesis that X causes Y.

Link	Paper	Mechanism
$\delta^{13}\text{C} \rightarrow \delta^{18}\text{O}$	Tschumi et al. (2010)	Break of stratification in southern ocean leading to better ventilated ocean and increased CO_2 .
	Watson et al. (2015)	Changes in meridional overturning circulation affect the amount of CO_2 outgassing.
	Anderson et al. (2009)	Upwelling events increase $\delta^{13}\text{C}$ measurements and atmospheric CO_2 .
$\delta^{18}\text{O} \rightarrow \delta^{13}\text{C}$	Sigman et al. (2010)	Increased CO_2 facilitates CaCO_3 dissolution which increases alkalinity.
$\delta^{13}\text{C} \rightarrow \text{Productivity}$	Anderson et al. (2009)	Upwelling events increase $\delta^{13}\text{C}$ measurements and increase ocean nutrients which positively affect productivity.
$\text{Productivity} \rightarrow \delta^{13}\text{C}$	Lisiecki (2010b)	Biological productivity in surface water changes $\delta^{13}\text{C}$.
$\text{Productivity} \rightarrow \delta^{18}\text{O}$	Omta et al. (2013)	Increased abundance in calcifiers decreases alkalinity, which in turn inhibits calcifier growth.
$\delta^{18}\text{O} \rightarrow \text{Productivity}$	Morris and Farrell (1971)	At lower temperatures assimilation of CO_2 happens more rapidly. Temperature has a non-linear effect on most enzyme activity.
$\text{Productivity} \rightarrow \text{Dust}$	Martin (1990)	Uptake of iron by plankton decreases iron availability in sea surface.
$\text{Dust} \rightarrow \text{Productivity}$	Martin (1990)	Increased iron in ocean leads to higher biological productivity which is iron-limited.
$\delta^{18}\text{O} \rightarrow \text{Dust}$	Ridgwell (2002)	A colder dryer glacial climate has a less vigorous hydrological cycle leading to higher dust deposition rates.
$\text{Dust} \rightarrow \delta^{18}\text{O}$	Tegen (2003)	Absorption or reflection of incoming solar radiation affects temperature which is linked to $\delta^{18}\text{O}$

7.2 Methods

7.1 Data

We use benthic $\delta^{18}\text{O}$ from the LR04 stack of 57 globally distributed sites (Lisiecki and Raymo, 2005). Benthic $\delta^{18}\text{O}$ is a function of the temperature and the $\delta^{18}\text{O}$ of the surrounding water, which is influenced by the global ice volume. Therefore, $\delta^{18}\text{O}$ is considered a proxy for temperature and ice volume.

In the Supplementary Materials, we also analyze the role of atmospheric CO_2 (as opposed to $\delta^{18}\text{O}$) to see if a climate proxy that is more related to alkalinity gives similar results. We use atmospheric CO_2 data from the European Project for Ice Coring in Antarctica (EPICA Lüthi et al. (2008)). This CO_2 record can likely also be used as a proxy for ocean alkalinity (ALK) for the following reason. $[\text{HCO}_3^-]$ and $[\text{CO}_3^{2-}]$ account for more than 95% of ocean alkalinity (ALK) and 99% of dissolved inorganic carbon (DIC) (Sarmiento and Gruber, 2006; Williams and Follows, 2011) and therefore:

$$\text{ALK} \approx [\text{HCO}_3^-] + 2[\text{CO}_3^{2-}] \quad (7.1a)$$

$$\text{DIC} \approx [\text{HCO}_3^-] + [\text{CO}_3^{2-}] \quad (7.1b)$$

Thus, $\text{ALK} \approx \text{DIC} + [\text{CO}_3^{2-}]$. Whole-ocean DIC is likely higher during glacials than during interglacials, as excess atmospheric carbon is stored in the ocean. Furthermore, both the lysocline depth and the B/Ca proxy indicate relatively small glacial-interglacial changes in deep-ocean $[\text{CO}_3^{2-}]$ (Catubig et al., 1998; Yu et al., 2010; Raitzsch et al., 2011; Yu et al., 2014). Hence, ocean alkalinity probably increases from an interglacial to a glacial period and decreases from a glacial to an interglacial.

Ocean circulation cannot be measured directly, but the $\delta^{13}\text{C}$ of benthic foraminifera is a proxy for ocean ventilation and ocean circulation (Curry and Oppo, 2005), because it decreases in “older” water (measured from the time it was last in contact with the atmosphere) due to the remineralization of organic material. We use the Lisiecki (2010a) benthic foraminifera $\delta^{13}\text{C}$ data compilation.

We use a Ba/Fe record from the Southern Ocean as a proxy for biological productivity (Jaccard et al., 2013). It has been shown that the vertical flux of marine Ba is a viable proxy for productivity (Paytan and Griffith, 2007) and that the Ba/Fe in turn reflects the vertical flux of marine Ba (Studer et al., 2015).

Lastly, we use Antarctic dust data from the EPICA dome C ice core (Lambert et al., 2008).

The data are available at resolutions of 0.2 ky (Ba/Fe), 1 ky ($\delta^{18}\text{O}$), 0.015 ky (CO_2), 2 ky ($\delta^{13}\text{C}$), and varying steps with a mean of 0.6942 ky (dust). The data dates back to the last 1557 ky (Ba/Fe), 5320 ky ($\delta^{18}\text{O}$), 798 ky (CO_2), 802 ky ($\delta^{13}\text{C}$), and 800 ky (dust). Therefore, we only use data up to 798 ky before present and we use the data at intervals of 2 ky (linear interpolation if necessary) for all interactions with $\delta^{13}\text{C}$ and with intervals of 1 ky (linear interpolation if necessary) for all other interactions.

7.2 CCM

For data analysis we use convergent cross-mapping (CCM) (Sugihara et al., 2012). The method is based on Takens' Theorem (Takens, 1981) that shows that dynamical systems involving more than one variable can be reconstructed from time-lagged series of only one of those involved variables, as long as sufficient time-lagged states are included: the so-called embedding dimension. In CCM, time-lagged time series of two variables are analysed to see if time-lagged states of those two variables can predict each other's current state. If time series of X can be predicted based on the time-lagged time series Y, it is concluded that X causes Y. The prediction is tested by calculating the CCM skill, which is the correlation between predicted values and true values. The method also tests for convergence, which means that the predictions improve with the length L of the time series that is used to predict X from Y. We calculate the 'CCM skill' for $X \rightarrow Y$ (X causes Y) with a slowly increasing length of data. We start with a library size (the number of data points used for the prediction) of 15, and then increase with steps of 15, to end with a library size of 270. More information about the CCM method can be found in the supplementary materials or in Sugihara et al. (2012).

The significance of causal links is tested by generating 100 surrogate data sets by cutting the time series at a random point and then exchanging the first and second part. In that way, the surrogate data have exactly the same characteristics as the original data in terms of power spectrum, autocorrelation and variance, but the linkage between the two data sets is broken. This method has been used previously by Van Nes et al. (2015b). To test the robustness of our results, we repeat the analysis with an alternative way of surrogate data generation by Ebisuzaki (1997).

If causal links are found in two directions, one cannot distinguish between the following options of having 1/ truly bidirectional causal links, 2/ synchronization with a third variable, or 3/ a one-way strong interaction where one variable behaves as a follower of the other. One possible method to distinguish between these three options is to perform CCM analysis on time lags of the variables by shifting the two time series relative to each other (Ye et al., 2015b). If one link has a clear optimum at a negative time lag, this means that past values from one variable can better predict future values from another variable, which is what you expect from a true causal link. If however the optimum is found at a positive time lag (in the future), this suggests that the link might not be causal but a

mere artifact from a strong coupling between the variables (Ye et al., 2015b). If one link has a negative time lag and one link has a zero time lag, we cannot distinguish between a bidirectional link and a strong unidirectional link. To be conservative and because we are not sure about the link with the time lag of zero, we will call these unidirectional links where we exclude the link with a time lag higher than or equal to zero. Lastly, it is important to note that the method cannot distinguish between positive and negative interactions.

7.3 Choice of parameters

Using the method by Cao (1997) (for details, see Supplementary Materials) for finding an optimal embedding dimension, we find a dimension of 10. The choice of parameters specific to a method is often a topic for debate. Therefore, we also calculate the sensitivity to the choices we made, and we find that a deviation from $E=10$ does not significantly affect our results (Supplementary Materials Fig. S7.6). Subsequently, we tried an alternative method from Sugihara (1994) for finding E and that also yields a maximum embedding dimension of $E=10$ (Supplementary Materials Fig. S7.2–S7.4).

A time lag $\tau = 1$ is used. Since climate time series are not over-sampled, it seems reasonable to take the time lag as small as possible. Increasing τ can cause a loss of the signal, which is to be expected for the data resolution we have (Supplementary Materials Fig. S7.7).

7.3 Results

The causal links found in our analysis are summarized in Figure 7.3. Importantly, these links include a causal loop, consisting of the link from $\delta^{13}C$ to Ba/Fe to $\delta^{18}O$ back to $\delta^{13}C$ in both the Pacific and East-Atlantic ocean. Within this loop, bidirectional significant links are found between Ba/Fe (proxy for productivity) and $\delta^{18}O$ (climate proxy) (Fig. 7.4 A). As these interactions are strong, this could either indicate truly bidirectional causal links, or a coupling that is so strong that one variable becomes a ‘slave’ of the other. The time-delayed results from convergent cross-mapping shows that the link from Ba to $\delta^{18}O$ has a clear negative time lag (Fig. 7.5 B) whereas the other link has a zero time lag (Fig. 7.5 A). This suggests that productivity (Ba/Fe) affects climate ($\delta^{18}O$) with a time lag of 4 ky whereas $\delta^{18}O$ either causes Ba/Fe as well, or behaves as a “slave” of Ba/Fe.

As discussed in the Introduction, glacial-interglacial changes in ocean productivity are thought to be caused by (i) changes in the ocean circulation (for which deep-ocean $\delta^{13}C$ is a proxy) or (ii) variations in iron input (for which dust from Antarctic ice cores is a proxy).

In case (i), we expect a causal link from deep Pacific and/or deep Atlantic $\delta^{13}C$ to Ba/Fe.




Figure 7.3: A schematic depiction of the network of causal links between deep Pacific $\delta^{13}\text{C}$ (P), deep North-East Atlantic $\delta^{13}\text{C}$ (NEA), deep East-Atlantic $\delta^{13}\text{C}$ (EA), biological productivity (Ba), $\delta^{18}\text{O}$ and dust, according to our CCM analysis. Arrows indicate causal links, i.e., $X \rightarrow Y$ means that X causes Y.

The results for the analysis with $\delta^{13}\text{C}$ show that indeed there are links from deep Pacific and deep East Atlantic $\delta^{13}\text{C}$ to and from $\delta^{18}\text{O}$ (Fig. 7.4 C, G). Furthermore, there are links from deep Pacific and deep East Atlantic $\delta^{13}\text{C}$ to productivity (Fig. 7.4 B, F). The time-lag analysis shows that all the links have an optimum for a negative time lag (Fig. 7.5 D-F, L-N), indicating causal links with a time lag. These results are consistent with the hypothesis that changes in ocean circulation could have had an effect on glacial-interglacial oscillations. Furthermore, the closed loop suggests a potential mechanism for nonlinear dynamics, which could explain the observed cycles (see Discussion). Note that in the North-East Atlantic, the pattern is different from that in the Pacific and East Atlantic: no significant causal links are found between the deep North-East Atlantic $\delta^{13}\text{C}$, Ba/Fe and CO_2 (Fig. 7.4 D-E).

In case (ii), we expect a causal link from dust to Ba/Fe. Our results are consistent with this hypothesis too, but only via a confounded route. Even though the link from dust to Ba/Fe is significant (Fig. 7.4 H), our time lags analysis shows that the optimum time lag for this interaction is positive, suggesting that this link is not a causal link but an artifact of the coupling between Ba/Fe and dust. However, we find causal links between dust and $\delta^{18}\text{O}$ (Fig. 7.4 I) and as mentioned before from $\delta^{18}\text{O}$ to $\delta^{13}\text{C}$ and from $\delta^{13}\text{C}$ to Ba/Fe. Time lags analysis suggest that these are truly causal links (Fig. 7.5 Q, R, E, M, D, L).

An alternative way of surrogate data generation to test for significance does not affect our results (Supplementary Materials Fig. 8-9). Furthermore, replacing $\delta^{18}\text{O}$ for CO_2 , as an alternative climate and temperature proxy that is closely linked to ocean alkalinity, only alters one link in the results: the link from CO_2 to Ba also has an optimum for a negative time lag, suggesting a causal link with a time lag of 4 kyr (Supplementary Materials Fig. 10-11).

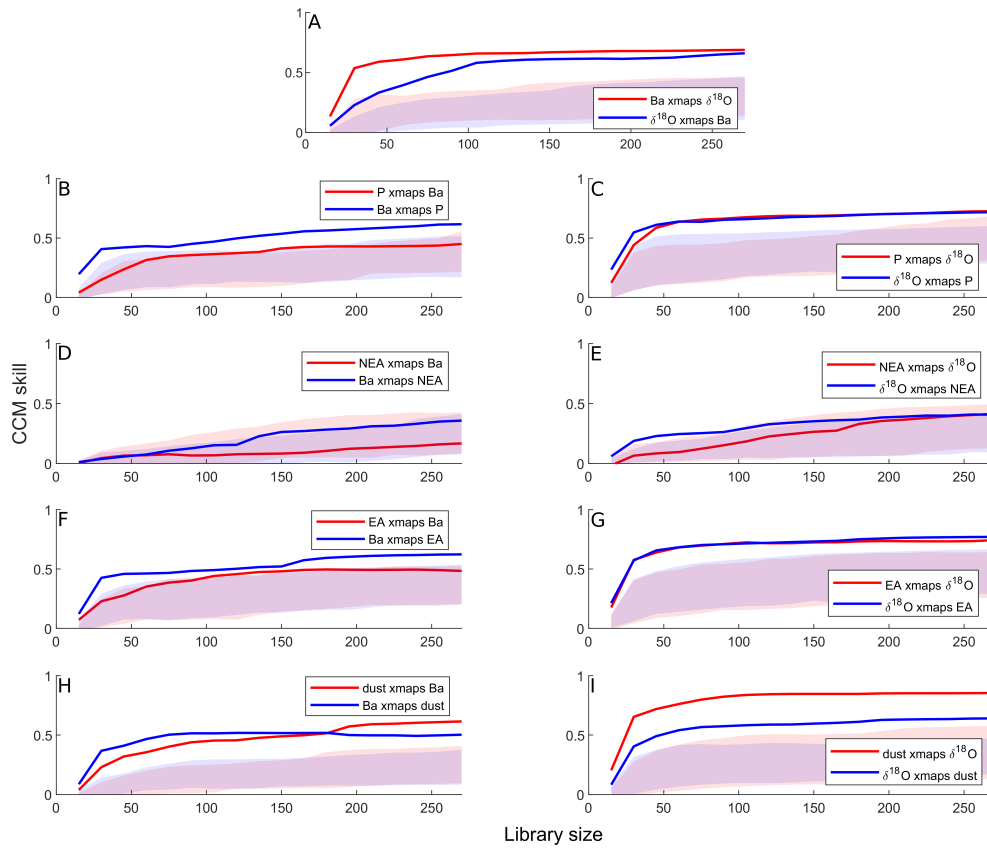


Figure 7.4: CCM analysis results, where the lines show the correlation between the predicted and true values (the CCM skill) and the shaded areas indicate a 90% confidence interval. The analysis indicates significant causal links between Ba and $\delta^{18}\text{O}$. Also, the links from Pacific (P) and East Atlantic (EA) $\delta^{13}\text{C}$ to Ba and $\delta^{18}\text{O}$ and from $\delta^{18}\text{O}$ to Pacific (P) and East Atlantic (EA) $\delta^{13}\text{C}$ are significant. The links between North-East Atlantic (NEA) $\delta^{13}\text{C}$ to and from Ba and $\delta^{18}\text{O}$ are not significant. Lastly, we find a significant bidirectional link between dust and Ba.

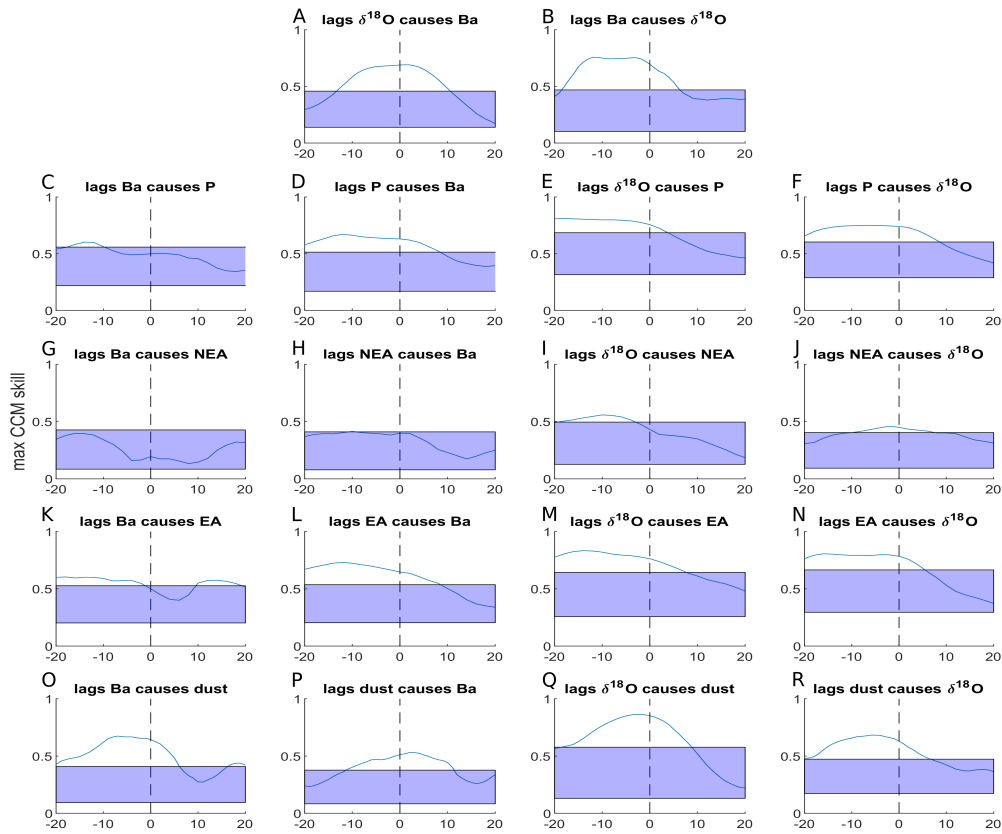


Figure 7.5: Timelags show that almost all links found in Fig. 7.4 have their optimum for a negative time lag, suggesting true causal links. The only exception is the link from $\delta^{18}\text{O}$ to Ba, so this is most likely not a truly causal link.

7.4 Discussion

For the first time, we have applied CCM to identify causal relationships between ice-core and paleoceanographic proxy records. Our results suggest the existence of several closed feedback loops. These feedback loops could provide an explanation for the dynamics underlying glacial-interglacial cycles.

The first feedback loop is suggested by the unidirectional causal link from deep Pacific and deep Eastern Atlantic $\delta^{13}\text{C}$ (Lisiecki, 2010a) (proxies for the strength of the overturning circulation) to Southern Ocean Ba/Fe (Jaccard et al., 2013) (1) and the bidirectional links between Ba/Fe (productivity) to $\delta^{18}\text{O}$ (climate) (2) and between $\delta^{18}\text{O}$ and deep-ocean $\delta^{13}\text{C}$ (3). The first bidirectional link could also be a strong unidirectional link from Ba/Fe to $\delta^{18}\text{O}$, since the lag between Ba/Fe and marine $\delta^{18}\text{O}$ is negative (~ 4 ky) whereas the lag for the link between $\delta^{18}\text{O}$ and Ba/Fe is positive (~ 2 ky). For this feedback loop, the underlying sequence of mechanisms could be:

1. A stronger Southern cell of the overturning circulation (of which $\delta^{13}\text{C}$ in the deep Pacific and deep Eastern Atlantic are proxies) is associated with enhanced upwelling off Antarctica (Downes et al., 2011; Gent, 2016; Morrison and McC. Hogg, 2013), which in turn drives a larger nutrient supply to the ocean surface. Thus, the overturning circulation may impact Southern Ocean productivity through variations in the nutrient supply rate.
2. Southern Ocean productivity may impact atmospheric global climate and $\delta^{18}\text{O}$ and CO_2 . According to a longstanding hypothesis, high $\delta^{18}\text{O}$ and low atmospheric CO_2 during glacial times are (partly) caused by enhanced Southern Ocean productivity (Siegenthaler and Wenk, 1984; Sarmiento and Toggweiler, 1984; Sigman and Boyle, 2000). However, the Southern Ocean Ba/Fe record (Jaccard et al., 2013) (Fig. 1) that we used for our analysis does not suggest particularly high productivity during times of low atmospheric CO_2 . Rather, the highest productivity appears to be associated with the glacial-interglacial transitions. Therefore, we think that a causal link between productivity and atmospheric CO_2 as described in Omta et al. (2018) would be more consistent with this record. That is, a spike in burial of organic matter leads to a decrease in deep-ocean alkalinity and the eventual outgassing of CO_2 from the ocean to the atmosphere due to carbonate compensation. This appears consistent with the recently reported finding that the variation in productivity in the Southern Ocean became larger as the amplitude of the glacial-interglacial cycles increased across the Mid-Pleistocene Transition (Fig. 3E in Hasenfratz et al. (2019)). In other words, small-amplitude glacial cycles are associated with relatively constant levels of productivity, whereas large-amplitude glacial cycles are associated with large periodic spikes in productivity.

3. $\delta^{18}\text{O}$ and CO_2 may impact the overturning circulation in various ways. For example, it has been suggested that a more northerly position of the Southern Hemisphere westerly winds during glacial periods would lead to weaker upwelling around Antarctica, which would in turn weaken the Southern cell of the overturning circulation (Toggweiler et al., 2006). Alternatively, a northward shift of the regions of upwelling and deep-water formation could be buoyancy-driven through changes in the Southern Ocean sea-ice extent (Ferrari et al., 2014; Jansen, 2017). Regardless of the precise mechanism, it appears that the structure of the overturning circulation varies significantly between glacial and interglacial times (Curry and Oppo, 2005; Lynch-Stieglitz et al., 2007; Burke et al., 2015).

Beside this three-variable loop, CCM detects two loops consisting of two variables, namely between $\delta^{18}\text{O}$ and Pacific and East-Atlantic $\delta^{13}\text{C}$ and between $\delta^{18}\text{O}$ and dust. However, we want to interpret these loops a bit more cautiously than the others, because CCM cannot always distinguish between a strong unidirectional link and a bidirectional link. Even so, these two loops can also be considered candidate potential drivers for the nonlinear behaviour observed in the glacial-interglacial dynamics. In particular, lower sea levels during glacial times led to larger areas of the continental shelf off Patagonia being exposed (Iriondo, 2000). Since most Antarctic dust seems to have a Patagonian origin (Fischer et al., 2007), changes in climate cause changes in dust deposition. Furthermore, it has been suggested that the larger equator-to-pole surface temperature gradient during glacial times led to more storms in the Southern Hemisphere, which in turn led to more dust transport (McGee et al., 2010). The link from dust to $\delta^{18}\text{O}$ may be explained from the radiative impacts of airborne dust (Tegen, 2003; Yoon et al., 2005; Rosenfeld et al., 2006) and the impact of dust deposited in Antarctica on the snow albedo (Bar-Or et al., 2008).

In addition to our standard CCM analysis, we also performed the same analysis for different shifts in the data to see if the optimum causal link was pointing in the right direction: from the past to the future. That is, a true causal link implies a negative optimum time lag. For the link from $\delta^{18}\text{O}$ to Ba/Fe, we found an optimum at a zero time lag and the link from dust to Ba/Fe had a positive time lag. To be conservative, we therefore interpret these interactions as unidirectional links from Ba/Fe to $\delta^{18}\text{O}$ and to dust. We want to point out that with the current quality of the data, the time lags analysis should not be used to interpret the precise time scale of the interactions we find. However, it can be used as an additional test to see if a link could be causal (i.e., past values can only cause the future and not the other way around). It should be noted that, like any statistical method, the absence of a significant causal link in our analysis does not prove that this link is absent in the real world. It is possible that, as longer and more accurate data become available, links will be revealed that currently remain undetected.

The unidirectional link from Ba/Fe to dust suggested by our analysis may not reflect a direct causal link but may rather be the consequence of the transportation route of the dust. From Patagonia, the dust first travels over the Southern Ocean where it may enhance productivity. Subsequently, some of the remaining dust is transported further and ultimately deposited on the Antarctic ice sheet (See Supplementary Figure 1 for a location of these sites). As a result of this chronological order, the signal measured by the productivity proxy Ba/Fe will be present in the Antarctic dust signal (and therefore we find that productivity causes dust). However, elements of the Antarctic dust signal that developed in the later stages of the dust's journey will not be visible in the Ba/Fe data (and therefore we find a positive optimal time lag for the causal link from dust to Ba/Fe).

An interesting implication of the existence of feedback loops is the possibility for nonlinear dynamics in the ocean-climate system. Feedback loops can be negative, in which case they work stabilizing or cause oscillatory behaviour if there are long time lags in the interactions (Levins, 1974). Feedback loops can also be positive, in which case they can cause self-amplifying dynamics, but time lags in positive feedback loops can either weaken or amplify the destabilizing effect. Based on our results, one cannot determine whether the identified dominant feedback loop is stabilizing or not, and under which conditions. There is no reason to believe that the interactions we investigate here are linear interactions, and therefore the loops can vary in their response under different circumstances. Furthermore, the present study cannot prove the existence of positive feedback loops, both because CCM does not provide information about the way the effect-variable responds to the cause variable (i.e., positive or negative) and because of large uncertainties in the precise time lags of the analysis. However, as one of our objectives of this study is to find an explanation for the nonlinear glacial-interglacial cycles, we believe that these loops are worth investigating.

Positive feedback in a dynamic system can generate alternative system states that reinforce themselves. Under certain conditions the system may 'flip' from one state to another - the conditions under which this occurs are generally referred to as 'tipping points'. The model hypothesized by Omta et al. (2013) involves calcifiers in a feedback loop, and in a numerical study it was shown that 1/ this model can generate the glacial-interglacial sawtooth pattern under Milankovitch forcing of 20 ky, and 2/ that the period of the sawtooth can vary under minor perturbations with the same forcing (Omta et al., 2016), effectively suggesting the possibility of a tipping point. This can explain the glacial interglacial cycles of the past, but it can also guide the predictions for current climate dynamics. Currently, there is a significant perturbation in the form of a considerable increase of atmospheric CO₂. The existence of a feedback loop as found in the data could indicate a potential change in the future from the sawtooth pattern observed in the data over the last 800,000 years to some unknown behaviour.

As we mentioned in the Introduction, various models have been formulated to explain glacial-interglacial dynamics. Crucifix (2012) reviewed many such models, based on a range of different mechanisms. It turned out to be difficult to decide between these models, since they all generate a sawtooth oscillation in temperature/ice volume and/or CO_2 . Therefore, Crucifix (2012) argued, we need more stringent criteria to operate our model selection. In Omta et al. (2016), one such criterion was formulated: models will need to reproduce the observed linear proportionality between the length and amplitude of the glacial-interglacial cycles. The causal relationships that emerged from our analysis provide further criteria. More specifically, model output could be analysed with CCM to test whether the same causal relationships emerge as from the analysis of the proxy records. However, it should be kept in mind that interpreting these links requires knowledge of the system under study (see for example for our link from productivity to dust).

Quite remarkably, our analysis suggests that mid-depth Northeast Atlantic $\delta^{13}\text{C}$ (Lisiecki, 2010a) has no significant causal relations with any of the other measured records. The locations of the cores that form the basis of this record are likely within the Northern overturning cell (North Atlantic Deep Water and Glacial North Atlantic Intermediate Water) during both glacial and interglacial times. Therefore, one possible interpretation would be that the strength of the Northern cell has no impact on, and is not impacted by, the strength of the Southern overturning and climate in general. However, ocean models generally indicate that a stronger and deeper Northern cell is associated with a weaker and shallower Southern cell and vice versa (Shin et al., 2003; Liu et al., 2005; Jansen, 2017). Another possible interpretation would be that variations in the mid-depth Northeast Atlantic $\delta^{13}\text{C}$ primarily reflect atmospheric changes, since the water at these sites is so ‘young’. Thus, our analysis would suggest that atmospheric $\delta^{13}\text{C}$ does not have a strong causal relationship either way with the overturning circulation, climate or atmospheric CO_2 . This appears in contradiction with the generally accepted view that the release of isotopically light ‘old’ carbon from the abyss causes both the observed minimum in atmospheric $\delta^{13}\text{C}$ and the rise in atmospheric CO_2 during the deglaciation (Schmitt et al., 2012; Broecker and McGee, 2013). Lastly, we show that the deep Northeast Atlantic $\delta^{13}\text{C}$ data is not clearly distinguishable from noise according to one of the tests we did (supplementary Fig. 7.5), making it hard to pick up signals with CCM. It is possible that if data of higher resolution becomes available in the future, signals might become clear that are not significant when using the currently available data sets.

The Ba/Fe record is from a single location in the Southern Ocean, whereas the other records represent either global signals or averages for oceanic regions. Nevertheless, the deglacial marine productivity maxima appear to be a global phenomenon. Deglacial maxima in productivity proxies such as opal accumulation and biogenic Ba have been found at various sites throughout the Atlantic (Kasten et al., 2001; Romero et al., 2008; Gil et al., 2009; Meckler et al., 2013), the Pacific (Crusius et al., 2004; Jaccard et al., 2005; Galbraith et al., 2007; Jaccard et al., 2010; Kohfeld and Chase, 2011; Hayes et al.,

2011), and the Southern Ocean (Anderson et al., 2009). A global compilation of marine sediment proxies reveals an increase in suboxic conditions in Oxygen Minimum Zones (OMZ) during the deglaciation (Jaccard and Galbraith, 2012) which could result from enhanced productivity. Furthermore, particularly strong deglacial maxima in $\delta^{15}\text{N}$ have been recorded in OMZ (Pride et al., 1999; Thunell and Kepple, 2004; Deutsch et al., 2004), a possible indication of enhanced denitrification due to the low oxygen conditions.

Paleoceanographic records are often dated by overlapping them with other records, for which the dating ultimately relies on orbital tuning. For example, the $\delta^{13}\text{C}$ records Lisiecki (2010a) were dated by aligning their $\delta^{18}\text{O}$ records to the LR04 benthic $\delta^{18}\text{O}$ stack (Lisiecki and Raymo, 2005), which was in turn orbitally tuned. It has been pointed out that this may cause spurious correlations with orbital cycles (Huybers and Wunsch, 2004). However, the Fourier spectra of records using a depth-derived age model and an orbitally tuned age model appear rather similar (Fig. 9 in Huybers and Wunsch (2004)). Perhaps more importantly, any shifting of the time series due to tuning should only affect the estimated lag between records and not the direction of the causality. As CCM does not primarily establish causality based on time lags but rather based on the structure of the data, we therefore believe the causalities found with the method are true causalities and not the result of artifacts.

7.5 Conclusion

Our CCM analysis on paleoceanographic and ice-core records suggests a dominant loop from ocean ventilation to biological productivity to climate and then back to ocean ventilation. While all separate mechanisms have been suggested earlier as explanations of the sawtooth-patterned glacial-interglacial dynamics, the existence of the dominant loop as a whole has to our knowledge not been identified from any data before. These loops provide possible mechanisms that explain why the climate system dynamics might not react in a linear way to perturbations in elements of this loop. The current considerable increase in atmospheric carbon dioxide could present a perturbation that results in a shift of the climate system dynamics from the glacial-interglacial cycle to new (unpredictable) behaviour. To elucidate the ramifications of feedback mechanisms in the climate system further, numerical models could be used to quantify these mechanisms and simulate possible glacial-interglacial dynamics as well as possible transitions between dynamical regimes.

Supplementary materials

S7.1 locations of data

Our data for $\delta^{18}\text{O}$ and $\delta^{13}\text{C}$ were compilations of multiple locations. For Ba/Fe and dust the data were obtained from specific locations: ODP site 1094 and EPICA Dome C, as depicted in figure S7.1.




Figure S7.1: The location for Ba/Fe (ODP site 1094) and dust (EPICA Dome C).

S7.2 CCM

For our analysis we follow the CCM method as proposed in Sugihara et al. (2012). In the following section we briefly describe the CCM algorithm. For a more extensive description including some examples we refer to Sugihara et al. (2012).

If a multivariate time series of n variables is described with $Z = Z_1, \dots, Z_L$ where every Z_i contains information for all variables, we can plot the manifold of the full system in an n -dimensional space. Takens theory states that the dynamics of this manifold of the full system are fully preserved by a lagged time series of just one of the variables (the ‘shadow manifold’, i.e., the projection of the manifold in the selected variable plane, like a shadow), provided that sufficient lags are included. CCM tests whether a manifold, build with lagged time series of one variable can be used to predict the dynamics of a shadow manifold build from lagged time series of another variable to see whether they belong to the same dynamical system. If the manifold of variable X (M_X) can predict the dynamics of the manifold of variable Y (M_Y), it can be stated that information of Y has been included in X and therefore there must be a causal link from Y to X . The cases that can occur are:

- X and Y can not be used to predict each other: The two time series are unrelated.
- X can be used to predict Y, but Y cannot be used to predict X: The systems are related with a unidirectional link from Y to X.
- Both variables can be used to predict each other: There is either a strong unidirectional link, a bi-directional link or are both driven by an external variable. (To test whether one option is more likely than another, we look at time lags of the interactions, as described in 2.2 of the main article.)

We use the manifold of time series X (M_X) to estimate values for Y(t) ($\hat{Y}(t)|M_x$) by ‘cross-mapping’ according to the following algorithm:

1. Find $\mathbf{x}(t)$ and its E+1 nearest neighbors
2. Get the time indices of the E+1 nearest neighbors: t_1, \dots, t_{E+1}
3. Estimate Y(t) from the locally weighted mean of the E+1 $Y(t_i)$ values using

$$\hat{Y}_t|M_x = \sum_{i=1}^{E+1} w_i Y(t_i) \quad (\text{S7.1})$$

where the weight w_i is calculated by

$$w_i = \frac{u_i}{\sum_{j=1}^{E+1} u_j} \quad (\text{S7.2})$$

and

$$u_i = \exp\left(\frac{-d[\mathbf{x}(t), \mathbf{x}(t_i)]}{d[\mathbf{x}(t), \mathbf{x}(t_1)]}\right) \quad (\text{S7.3})$$

where $d[\mathbf{x}(t), \mathbf{x}(t_i)]$ denotes the euclidean distance between $\mathbf{x}(t)$ and $\mathbf{x}(t_i)$.

If X and Y are coupled, the nearest neighbors on M_x should give time indices that find the nearest neighbors on Y_x . For longer time series (larger L), the manifold fills and $\hat{Y}(t)|M_x$ should converge to Y(t). If cross-mapping of Y from X does not converge, this means that Y is not causal to X (i.e. the information of Y is not contained in X).

S7.3 Choice of embedding dimension

We calculate the optimal embedding dimension using the method of Cao (1997). This method uses false nearest neighbors, with an improved stopping criterion. In line with Cao (1997) we calculate two values. ‘E1’ is a measure that sees if two points are nearest neighbors even if the embedding dimension increases. If the points are nearest neighbors for an embedding dimension i, but not anymore for embedding dimension i+1, they are labeled as “false nearest neighbors”. We increase the embedding dimension until no more false nearest neighbors are found. ‘E2’ is a measure that tests if the data is in fact

generated by a deterministic model and is not an effect of noise. The system is considered deterministic if $E2$ depends on the embedding dimension and there exist an embedding dimension for which $E2 \neq 1$. More information and several examples can be found in Cao (1997).

The analysis results in the choice of $E = 10$ (figure S7.2). Another analysis from Sugihara et al. (2012) yields the same maximum embedding dimension of E , but here there is more variation among variables (figure S7.4).

It is important to note that the $E2$ signal does not strongly indicate a deterministic process. However, we still think the data is suitable for CCM because of the following reasons:

1. The resolution of the data might be so low that the signal is missed. This is further emphasized by repeating the analysis for the high-resolution data set that is available for Barium which clearly shows a dependency of $E2$ on the embedding dimension (figure S7.3).
2. The $E2$ measure of Cao (1997) is quite strict. Even in the simulated examples in the original article by Cao (1997), the $E2$ signal is not always clearly indicative of a deterministic process. To further explore this idea, we did another analysis designed to distinguish a deterministic process from red noise based on the method of Sugihara (1994). In this analysis, we calculate local and global forecast skill by changing a parameter “theta”. the system is considered deterministic if correlations decrease for high theta (global forecast skill) compared to low theta (local forecast skill), and furthermore if the correlations are higher than the autocorrelation. Our results clearly show deterministic dynamics for Barium, CO_2 , Pacific $\delta^{13}C$, East Atlantic $\delta^{13}C$ and dust. For North-East Atlantic $\delta^{13}C$ we are unable to distinguish it clearly from red noise, since the correlations from the predictions are lower than the autocorrelation (figure S7.5).
3. We observe convergence in our CCM analysis. This should not happen if the signal was actually noise.

Clearly, based on these results it is hard to prove that the embedding dimension is accurate. So the next question we ask is how important the choice of the embedding dimension actually is. We show that the the analysis on our data is not sensitive to deviations away from the chosen E of 10 (Figure S7.6). The parameter b was always put to $E + 1$.

S7.4 Choice of τ

In theory, the choice of the time lag τ for the embedding dimension should not influence the results Cao (1997), however, in practice this parameter can play a role (Cao, 1997). For time series of discrete maps the best choice for τ is 1, but for continuous data there are




Figure S7.2: Values for E1 and E2 for different embedding dimensions for all variables, where P, NEA and EA stand for Pacific, North-East Atlantic and East Atlantic respectively. Analysis led to the choice of an embedding dimension of 10.




Figure S7.3: Values for E1 and E2 for different embedding dimensions for the high resolution dataset of Ba, indicating some changes in E1, and still converging at an embedding dimension of 10.




Figure S7.4: Simplex projection based on Sugihara (1994) to find E. The best embedding dimension is the dimension where we find the maximum correlation. Results differ for different variables, yielding 10, 3, 4 and 5. E=10 is the highest value we find and therefore a good choice.




Figure S7.5: forecast skill using simplex projection based on Sugihara (1994) to distinguish deterministic dynamics from red noise. Lines indicate the forecast skill dependent of theta, dashed lines indicate the autocorrelation. A deterministic process has a better forecast skill for low values of theta (local predictions) than for high values of theta (global predictions), as we see for all variables. However, for North-East Atlantic $\delta^{13}\text{C}$ this forecast skill is lower than the autocorrelation, which could indicate that the data is dominated by red-noise.

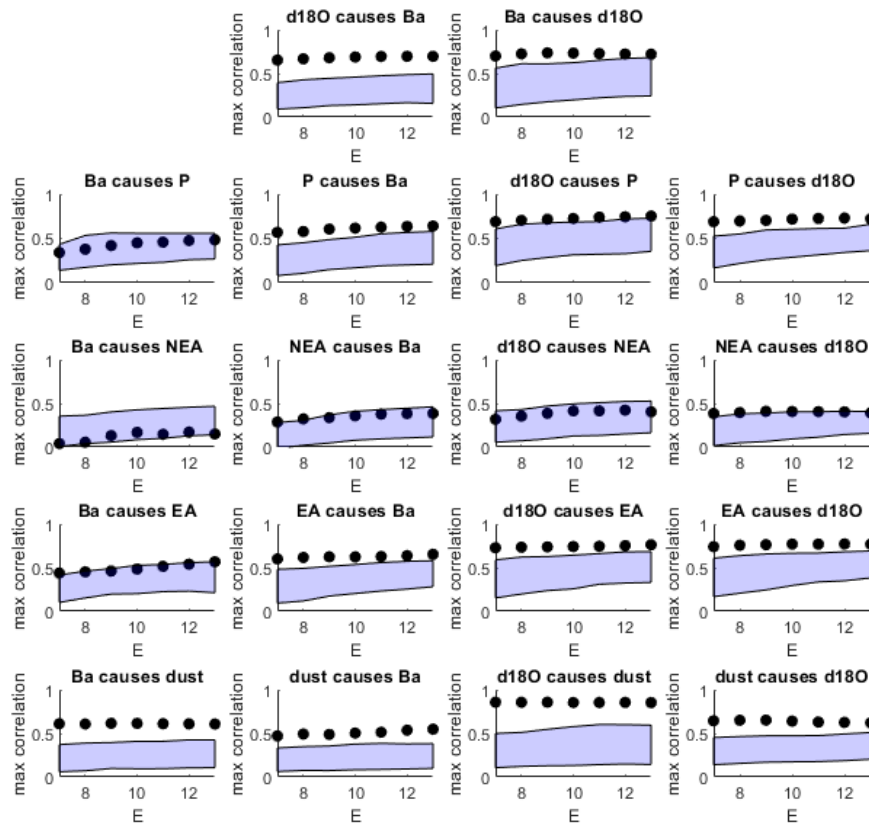


Figure S7.6: CCM result for all variables, where P, NEA and EA stand for Pacific, North-East Atlantic and East Atlantic $\delta^{13}\text{C}$ respectively, for different values of E to test for sensitivity. The black dots form mostly horizontal curves, therefore we conclude that the results are not sensitive to deviations away from our chosen embedding dimension.




Figure S7.7: CCM result for all variables, where P, NEA and EA stand for Pacific, North-East Atlantic and East Atlantic $\delta^{13}\text{C}$ respectively, for different values of τ to test for sensitivity. Results show that for a choice of τ larger than 1, often the signal is lost.

various methods to calculate or think about τ . Since our data is clearly not oversampled (with a time step of 2000 years) and the embedding dimension is quite high already (10), it seems sensible to choose τ as small as possible. Therefore we set $\tau = 1$. Also for τ we calculate the sensitivity of the method to different values. Our results show quite some fluctuation in CCM-skill dependent on the value of τ . However, overall the signal seems to be strongest for smaller values of τ , where the signal is sometimes lost for higher values. The strongest fluctuations are seen for the links from and to North-East Atlantic $\delta^{13}\text{C}$. However, since this variable is already excluded from the analysis based on figure S7.5, we do not consider this to be a problem.

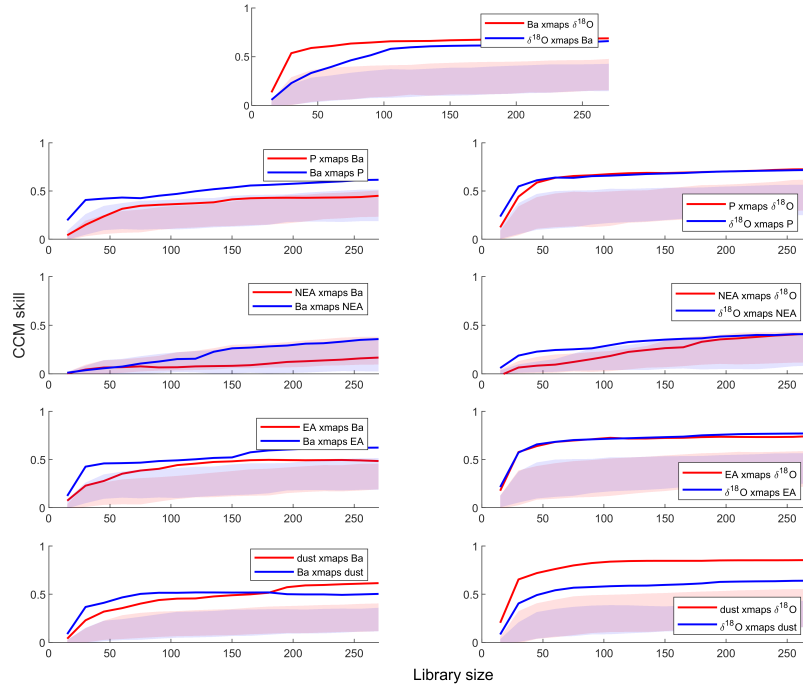


Figure S7.8: CCM result where the Ebisuzaki method of surrogate data generation is used, giving very similar results as Fig. 7.4 in the main text, indicating that these results are robust for variations in significance determination.

S7.5 Analysis with Ebisuzaki method of surrogate data generation

As our default method, we calculated our confidence intervals by creating surrogate data by cutting the data in random locations and then swapping the first and last part. In this way all characteristics of the data are kept intact, but the link between two time series is broken. An alternative way to create surrogate data is to generate random data with the same autocorrelation function and the same power spectrum as the original data. This can be done with a method proposed by Ebisuzaki (1997). We use this analysis to see if our results were robust for slight alterations in the way that we calculated the significance. Supplementary Fig. S7.8 – S7.9 show very similar results as Fig. 7.4 – 7.5 in the original text and therefore we conclude that the analysis is robust and that the results do not depend on the way that the surrogate data is generated.

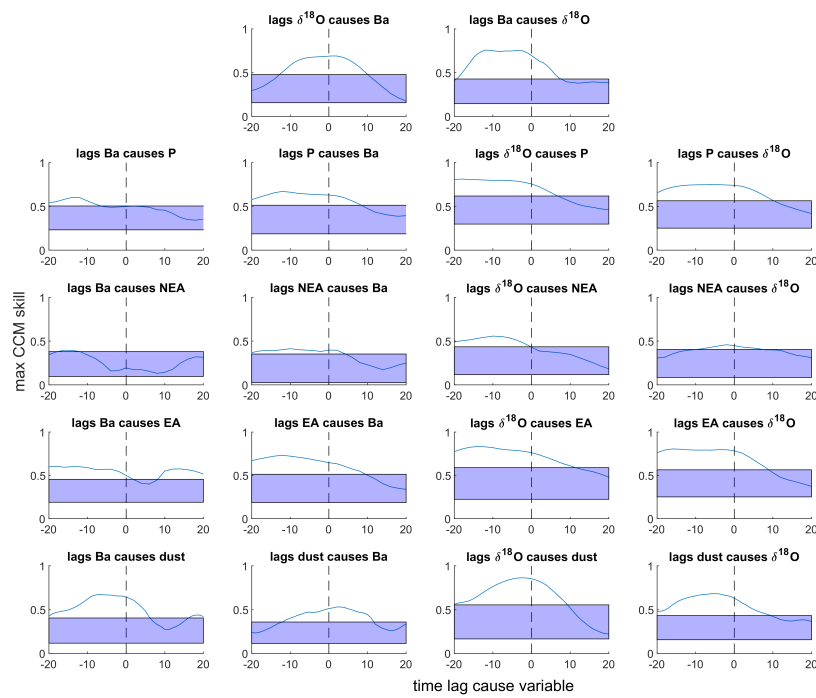


Figure S7.9: CCM time-lags result where the Ebisuzaki method of surrogate data generation is used, giving very similar results as Fig. 7.5 in the main text, indicating that these results are robust for variations in significance determination

S7.6 Analysis with CO₂ instead of $\delta^{18}\text{O}$

In our main analysis we choose $\delta^{18}\text{O}$ as a “climate proxy” since most mechanisms underlying the links use $\delta^{18}\text{O}$ in the explanation. However, an argument could be made that instead of $\delta^{18}\text{O}$, we could have used CO₂. Especially in the link from productivity to “climate”, the explained mechanisms heavily rely on ocean alkalinity which is more closely linked to CO₂ than to $\delta^{18}\text{O}$. Furthermore, since CO₂ and $\delta^{18}\text{O}$ are closely linked, the analysis with CO₂ should give similar results and is therefore a good additional evaluation about the robustness of our results.

The analysis with CO₂ yields very similar results as the analysis with $\delta^{18}\text{O}$ (Supplementary Fig. S7.10 - S7.11). One small change is that the link from CO₂ to Ba has a negative time-lag, so the link between these two variables is a bidirectional link in this situation.

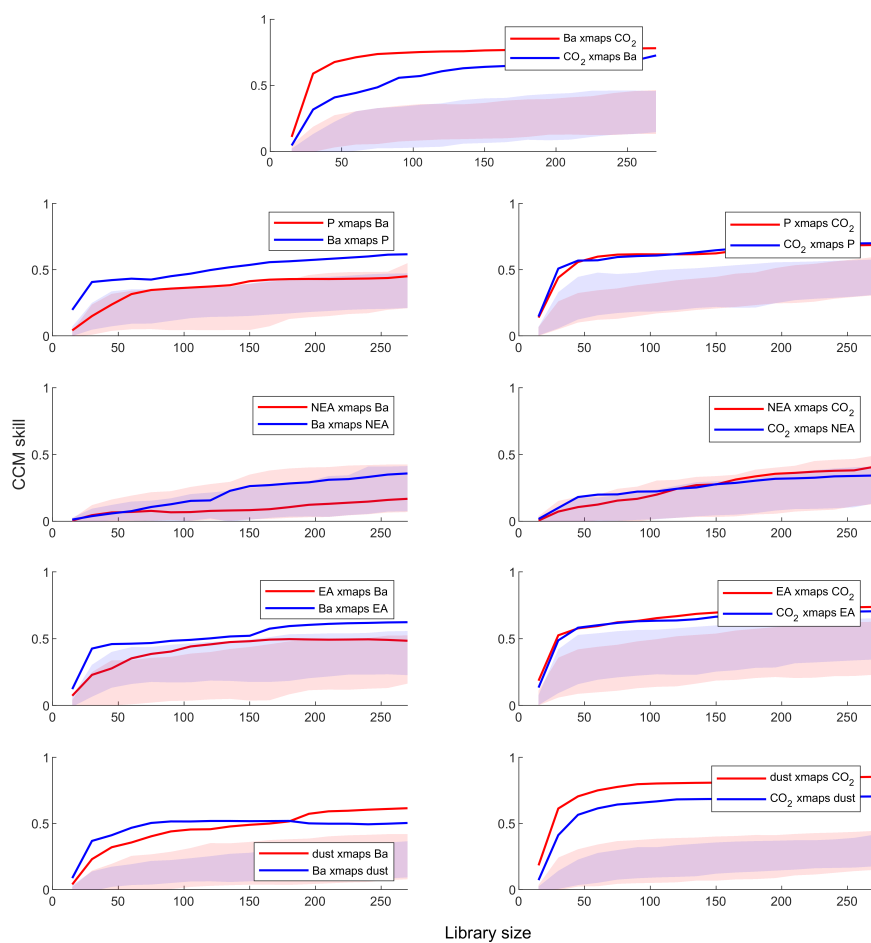


Figure S7.10: CCM result where $\delta^{18}\text{O}$ is replaced by CO_2 , giving very similar results as Fig. 7.4 in the main text.

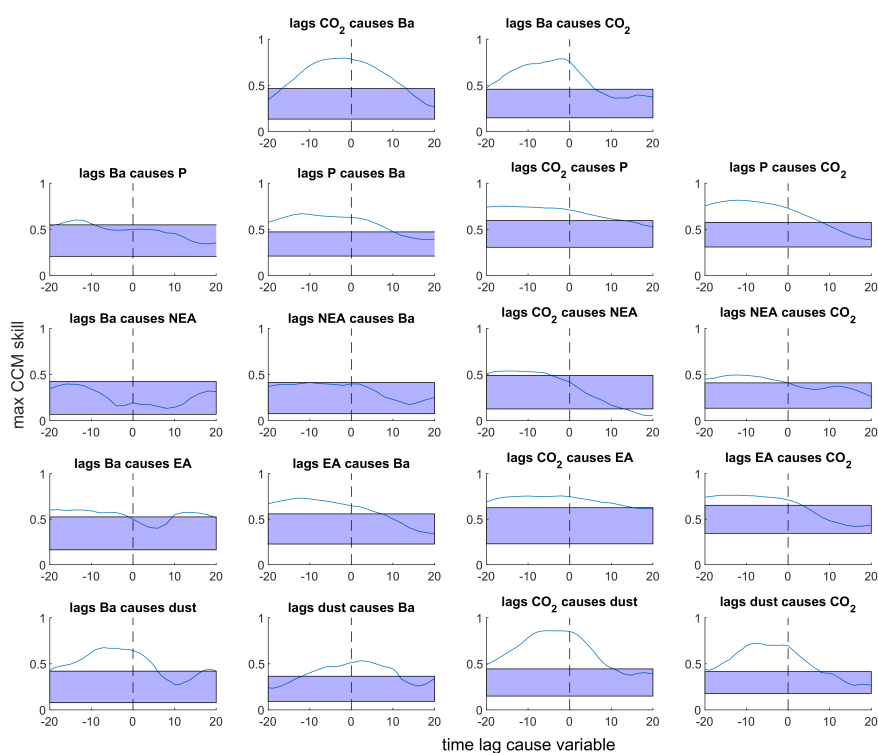


Figure S7.11: CCM time-lags result where $\delta^{18}\text{O}$ is replaced by CO_2 , giving very similar results as Fig. 7.5 in the main text.

Chapter 8

Synthesis

Els Weinans

Main findings

The main objective of this thesis was to explore how data-driven methods can be used to study the behaviour of complex systems. Here I will shortly recap how the chapters contribute to this overall goal and point to some limitations of the work presented. My main findings consist of three sections: 1) resilience in multivariate systems (**chapters 2-5**), 2) extracting information from chaotic attractors (**chapters 2, 7**), and 3) finding the best methods (**all chapters**).

Resilience in multivariate systems

Detecting slowing down of a system as an indication for a loss of capacity to recover and possibly for an upcoming critical transition, has received considerable attention (Scheffer et al., 2009). In most studies on resilience indicators, the proposed indicators are applied to one-dimensional systems, and therefore the multivariate nature of complex systems was somewhat underrepresented (Boettiger et al., 2013). Various studies have contributed to this knowledge gap (Held and Kleinen, 2004; Chen et al., 2012; Bathiany et al., 2013; Boerlijst et al., 2013; Suweis and D’Odorico, 2014; Chen et al., 2019). The first chapters in this thesis too, consider the issues related to applying resilience indicators to multivariate systems.

In **chapter 2** I question how and if resilience indicators could be applied to systems that need to respond quickly and adequately to changes in the environment, instead of aiming for a baseline value at all cost. Various subsystems of the human system, such as mood, the heart, or bones, can undergo rapid change. It has been suggested that these rapid changes can be manifestations of tipping points towards illnesses such as depression (Leemput et al., 2014), atrial fibrillation (Nannes et al., 2020) and osteoporosis (Homminga et al., 2004). However, for human subsystems a loss of functioning may not be preceded by slowing down, but instead by a lack of responsiveness, allowing for different kind of resilience indicators based on the complexity of the signal, which I will clarify in the next section on ‘extracting information from chaotic attractors’.

In **chapter 3** I extend on the idea of ‘degenerate fingerprinting’ (Held and Kleinen, 2004), where data is first projected on a ‘dominant mode’ after which the one-dimensional resilience indicators can be applied. Previous work has shown that this mode can provide information about the vulnerable variables in the system, i.e. the ‘hotspots’ (Bathiany et al., 2013). The mode in degenerate fingerprinting is found by a Principal Component Analysis (PCA). In **chapter 3** I explore if Min/Max Autocorrelation Facors (MAF) can be an alternative. MAF is a method that finds the modes (the directions in phase space) that have the highest or lowest autocorrelation. The main reason for exploring MAF, is that previous work on one-dimensional resilience indicators suggested that autocorrelation-based indicators might be more robust to noise than variance-based indicators (Dakos

et al., 2012b). Since MAF is an autocorrelation-based alternative to a PCA, I hypothesized that it could improve the robustness of methods that use a PCA. Indeed, my simulations demonstrate that MAF was more robust to PCA for one example of a ‘complex noise regime’, where I distributed the noise unevenly over the variables. However, this came at the cost of requiring data of a larger length and, as with all autocorrelation-based indicators, higher resolution.

In **chapter 4**, we explore if the properties of mutualistic systems allow for an extrapolation of the direction of the dominant mode to predict the future state of a complex system, after a tipping point has passed. The hypothesis that mutualistic interactions lead to more predictable behaviour has been proposed before, for example by Suweis and D’Odorico (2014). We found that this predictable behaviour could be used to find a pointer to the future state, provided that there are no complex dynamics such as oscillations in the system.

In **chapter 5**, I compare the multivariate indicators of resilience loss that were proposed by others and by ourselves to find which indicator is best to use in which scenario. The simulations in this chapter clearly demonstrate that different scenarios require different methods. We confirm that autocorrelation-based indicators require high resolution whereas variance-based indicators require a simple noise regime. Furthermore, we found that major problems can occur for all indicators when only a subset of the variables are measured, a problem identified earlier by Boerlijst et al. (2013). It is possible that future work will find ways to extract the relevant information from the variables that are observed. However, it is also possible that the information of the upcoming transition is simply not present in all the variables. This chapter suggests a cautious interpretation of all work on resilience indicators in multivariate systems.

Extracting information from chaotic attractors

Some systems are not situated in equilibria and might not be subject to tipping points but are complex because their dynamics are chaotic. Taken’s theorem describes how for chaotic systems, time lags of one variable can be used to reconstruct the attractor of the full dynamical system (Takens, 1981). The idea behind this theorem is that when variables interact with each other, information from one variable is embedded in the other variable, and thus in principle you don’t need all variables to determine some overall properties (such as the Hurst exponent) of the system (as explained in the introduction of this thesis).

In **chapter 2** I explore not only resilience, but also the meaning of several indicators of complexity, with applications to the human system. While studying previous work, I found out that complexity has several definitions, some are conceptual (such as the description of how complexity should have an optimum between complete order and complete disorder, i.e. figure 2 in chapter 2, as explained in Huberman and Hogg (1986)),

others are mathematical (such as a Kolmogorov complexity approximation by Grassberger and Procaccia (1983b)). Nevertheless, I found out that many of them find a basis in Taken's theorem by analysing time-lagged versions of one variable. It is now, almost 40 years after the above mentioned articles, still an exciting insight that time lagged versions of a variable contain information about full system functioning. The current data availability stimulates re-exploration of these tools. In this chapter, I demonstrate how resilience indicators and complexity indicators can be used to determine the resilience and responsiveness of human subsystems, and as such can together tell us something about the functioning of each subsystem. Even though the two processes have been extensively described, the connection between the two remains unclear. This chapter attempts to unify the theoretical framework of the two indicators by describing different functions of human subsystems (i.e. regulated or effector variables). Furthermore, it elucidates how being close to a tipping point can be dangerous because it increases the chance to tip to another state, but being far from a tipping point can be dangerous as well because this could lead to a loss of responsiveness. I shortly illustrate this idea with the Ising Spin model. I do realize that this simplistic model, although interesting, does not yet proof the usability of complexity and resilience indicators for the human system. Therefore I propose an outline where simplistic models, realistic models, and real data are all required before we will be able to successfully apply this framework. For the human system, the subsystems mentioned in **chapter 2** (the heart, bones, mood and balance) seem like a good place to start.

In 2012, Sugihara and colleagues suggested that reconstructed attractors are not only indicative of complexity, but could also be used to detect causality, a method called Convergent Cross-Mapping (CCM) (Sugihara et al., 2012). In 2015, van Nes and colleagues used this method to find causal links in the climate system (Van Nes et al., 2015a), demonstrating that the current climate data is of sufficient quality and quantity to use CCM (something which was not the case previously (Grassberger, 1986)). This has invited for a variety of studies applying reconstruction methods to climatic time series (Hirata et al., 2016; Trauth et al., 2019; Huang et al., 2020). In **chapter 7**, I explore the use of CCM to detect causal links, with the ultimate goal of finding an explanation for the saw-tooth shape of the glacial-interglacial cycles. The currently available literature has described all links as possible causal links in the network, which makes it difficult to pinpoint the dominant drivers. Our CCM analysis suggests the presence of one dominant loop from ocean ventilation to biological productivity to climate, back to ocean ventilation. This loop provides a possible explanation for the shape of the observed oscillations. The CCM method is known to be sensitive to noise and it requires much data (Mønster et al., 2017). It is possible that as data quality increases, we will detect causal links that are currently unclear. One example in our analysis is the time series of North-East Atlantic $\delta^{13}\text{C}$ (a proxy for ocean ventilation), which was not clearly distinguishable from noise. As data quality is increasing in all domains, the method has the potential to unravel causal links

and feedbacks in systems from other scientific fields, as demonstrated by recent studies in ecology (Sugihara et al., 2012), physiology (Heskamp et al., 2014; Verma et al., 2016) and social systems (Luo et al., 2014).

Finding the best methods

For all chapters in this thesis, I made a comparison between methods. For some chapters this is very explicit, such as **chapter 2** on finding similarities and differences between DIORs and complexity quantification tools, or **chapter 5** on comparing different multivariate indicators of resilience loss. For other chapters, this choice was made behind the scenes, for example in **chapter 6 - 7** where I wanted to focus on the particular domains (language change and the carbon cycle respectively) more than on the methodological issues. I have come to realize that most people, including myself, are looking for a way to tell which method is the best to use (i.e. highest performance/most robust), but in all cases I came to the conclusion that it depends. It depends on the type of data that is available (see **chapter 5**) and on the type of system under investigation (see **chapter 2**). Perhaps, this is not a satisfying answer to the question: “What is the best method?”, but I have come to appreciate the answer, as it nicely captures the complexity of the subject and it demonstrates that no method is wrong (although some methods are useless for a certain application) (Qiao et al., 2015). This links a little bit to the ‘no free lunch’ theorem in optimization that states that averaged over all optimization problems, all optimization algorithms perform equally well (Adam et al., 2019).

In **chapter 6 - 7** I had a wide range of choices for methods to consider. The methods I used in these chapters were well-suited for the data, but other choices could have been made and comparing many different methods could have been a study by itself. For instance the language development over time studied in **chapter 6** is without a doubt an incredibly complex system (Freeman and Cameron, 2008). In this chapter, we did not aim for a true understanding of these complexities, but instead we wanted to identify patterns to form an hypothesis that could ultimately lead to fundamental understanding. Our analysis was thus an exploratory journey through the data. We found that a PCA was particularly useful to detect patterns in the data. The input for this PCA were the time series for the 5000 most used words in the English language. The principal components were relative combinations of words that together changed in time in a correlated or anticorrelated matter. We found that the PCA could disentangle the temporal behaviour of words like ‘God’, ‘mind’, and ‘love’ (words relating to intuition) from words like ‘company’, ‘committee’, and ‘duty’ (words related to institutions), without having any prior knowledge about the meaning of these words. Since these words appeal to the imagination, this inspired us to dive deeper and check the behaviour of different words related to intuitive vs rational thinking in multiple languages. We are aware of the limitations of the google books dataset (see for example Pechenick et al. (2015)), but since there is currently no better data available and the results seem robust for different language sets,

we think the patterns we detect are worth further investigation.

Considerations

Here, I like to address some philosophical considerations about science in general, that I pondered during the writing of this thesis. The three considerations are 1) the search for universal patterns, 2) the benefits of fundamental and interdisciplinary science, and 3) the remarkable human ability to find and explain non-existent patterns.

Universal patterns

Ivar Ekeland writes in ‘Mathematics and the unexpected’: “The fisherman who tries to disentangle his lines bases his hopes of succeeding on the fact that the lines were straight and clear to begin with. There is no similar a priori knowledge underpinning Ptolemy’s or Kepler’s efforts – only their faith in the hidden harmony of the cosmos” (Ekeland, 1990). Like Ptolemy and Kepler, many scientific endeavors have focused on finding universal patterns. This thesis is no exception to that rule. However, there is no law that states that such a “hidden harmony of the cosmos” (or hidden harmony to complex systems) actually exists. This raises the question why so many scientist strive for universal patterns.

Answering this question is beyond the scope of this thesis, but I would like to mention two points that helped me justify this search for universality. Firstly, one theory that explains two findings is always preferred over two separate theories for those two findings, based on the principle of parsimony, also known as occam’s razor (Blumer et al., 1987). I believe that most people would agree with me that the main impact of for example Newtons laws was that they unified observations here on earth with observations in space, and the main impact of Darwin’s work on evolution was that it provided a theory that accounted for multiple different observations in one unified theory (Darwin, 1909). Secondly, science has a component of beauty in it, not unlike the arts. Indeed, mathematical equations elicit a response in the same regions in the brain that light up when listening to music or looking at art. This even happens for people that do not have a mathematical background (Zeki et al., 2014). Furthermore, ‘elegance’ or ‘simplicity’ are often used as guiding principles for scientific studies (MacArthur, 2021). Perhaps, most beauty is experienced in unified theories.

Fundamental science and interdisciplinarity

Some scientific studies have unforeseen consequences. For example, in 1962 Osamu Shimomura and colleagues were working on the question: “why do jellyfish glow?”. During their studies, they identified a green fluorescent protein (Shimomura et al., 1962). 32 Years later, this protein allowed for the monitoring of gene expressions and protein localization

in living animals (Chalfie et al., 1994), a breakthrough for many medical applications, which resulted in a Nobel Prize in 2008. This example clearly demonstrates that curiosity driven science can have large and unpredictable impacts.

I started my introduction with mentioning climate change, biodiversity loss and science skepticism as three challenges that are currently faced by science and society. The problems have many facets and either result from or affect (or both) multiple different scientific and societal areas. It is therefore not surprising that more effort is put into creating interdisciplinary research teams that tackle these issues (Jacobs and Frickel, 2009; Okamura, 2019). The work in this thesis too, has evolved from collaborations with computer scientists, psychologists, ecologists, mathematicians, climatologists and medical specialists and I am currently unsure where I should put myself on that list (perhaps as an applied mathematician). These collaborations are not always easy, as it can take some time to find a common language, but they are certainly interesting.

Many training programs have appeared in recent years to accommodate interdisciplinary approaches to scientific problems. One example is the bachelor program *bèta-gamma* at the University of Amsterdam that I followed myself. However, in later career stages, interdisciplinary projects seem to be discouraged, as funding success decreases as a research team becomes more interdisciplinary (Bromham et al., 2016). The main benefit of fundamental science (the unforeseen applications to different topics than originally designed for) can only be given a chance if scientists are allowed the time to explore what is going on outside their own field of research.

Explaining non-existent patterns

Human brains have an extremely well-developed ability for pattern detection and processing (Mattson, 2014). This ability is so well-developed that humans are notoriously bad at detecting randomness (Williams and Griffiths, 2013). For example, when the music platform *spotify* was launched, it had an option for a random shuffle of a playlist. However, its users sent complaints that the playlists were not random at all. It turned out that listeners found patterns in the random algorithm that *spotify* used. For example, a playlist does not “feel” random if it plays two songs of the same artist in a row, something which can occur in a purely random order. Therefore, the current ‘shuffle’ option in *spotify* is not random anymore, but instead it avoids clusters, and the complaints have stopped (Poláček, 2014). Our pattern-finding ability has allowed us to even detect non-existent patterns.

Another more personal example happened to me while preparing **chapter 7** of this thesis. In the first version that I prepared, I had accidentally reversed the time axis of the dust data-set. Since the CCM method that I used for the analysis works with reconstructions of chaotic attractors, this mistake was not immediately clear to me just from the analysis. The causal network that was generated from this analysis was meaningless as it was based

on a mistake. It was an interesting experience that based on the available literature I was able to find an explanation for all the links that were mistakenly found. It seems like we are not only good at detecting non-existent patterns, but also at finding explanations for these non-existent patterns.

Statistical tests and good null-models can help to objectively identify patterns, but choosing the right null-model is not always a trivial task (Gotelli and Ulrich, 2012). Additionally, in systems where many possible relationships have been suggested to exist, such as the climate system (Daruka and Ditlevsen, 2016; Crucifix et al., 2017), any finding can be explained. Therefore, our intuition about a pattern and our ability to explain a pattern, are not always indicative of the truth.

Awareness of this problem might help to avoid the pitfall. Furthermore, it can be a good idea to have several people double-check the analysis and code (for a programming study), a practice which is perhaps more common in software engineering than in science (Bacchelli and Bird, 2013). Last, I am in favour of motivating scientists to publish their code alongside with their papers. This can feel scary, especially since most scientists are not trained as software engineers (Barnes, 2010). However, it is of paramount importance for science to be transparent and reproducible, which includes sharing code (Toelch and Ostwald, 2018). Fortunately, various initiatives have developed outlines and offer guidance (see for example the FAIR initiative, which advocates Findability, Accessibility, Interoperability and Reusability of software (Lamprecht et al., 2020)).

Prospect

The work performed in this thesis raises several questions. I here want to reflect on four of them: 1) Can knowledge about resilience also lead to knowledge on useful interventions? 2) Are the models that I used of the right level of complexity? 3) What is the effect of multiple interacting scales? 4) What can noise teach us?

Can knowledge about resilience also lead to knowledge on useful interventions?

Detecting the vulnerable variables in a multivariate system can be used as a predictive tool, as we have described in **chapter 3 - 4**. One interesting implication of having a direction where a system is most likely to leave the current equilibrium, is that this direction could also provide a direction that offers most opportunity for change in case the current equilibrium is an unwanted situation. In that sense, our studies could help to choose interventions. This idea is receiving interest in various fields. For example, in health care Timothy Buchman seems positive about this premise, as he writes: “Today’s physicians are already testing to ‘see if the network is right’; tomorrow’s physicians may well use therapies to ‘make the network right’” (Buchman, 2002). Also in studies on social-ecological systems this concept has received considerable attentions for example by influencing feedbacks in marine ecosystems (Nyström et al., 2012) or interventions on

resilience in governance (Lebel et al., 2006).

Are the models of the right level of complexity? The framework described in **chapter 2** about resilience and complexity quantification tools is a quite simplistic and general framework that might need extensions to serve the specific complexities of each particular subsystem. Evans et al. (2013b), warn that the preference for simple models has “limited progress in ecology”. This suggests that even though these simplistic theories might guide our thinking about the world from a systems perspective, it might not always lead to knowledge about the specifics of the system under study.

I encountered the mismatch between theory and practice when I learned about the work of Eus van Someren about insomnia (see for example Van Someren and Riemersma-Van Der Lek (2007)). We hypothesized that changes in sleep stages might be better understood by using the MAF analysis as described in **chapter 3**. However, the cyclic behaviour of the data, the noisiness, the fact that the surface EEG data is a proxy of the real brain activity, and the redundancy in the data all created challenges that made us decide that the current theory is not well equipped for the current data. I still think it is an interesting idea to explore sleep as a complex dynamical system, and it is exciting to see that other attempts were more successful (for example de Mooij et al. (2020)).

What is the effect of the multiple interacting scales that constitute a complex system? In the introduction, I state that complex systems have interactions on different time scales. However, in the analysis done in this thesis, I have simplified systems so far that I only investigate their behaviour on one scale. Adding multiple scales in models often results in complicated analysis tools, both for different scales in time (Koutsoyiannis, 2001) and in space (Thuiller et al., 2003). In many scientific areas the importance of different scales is acknowledged, for example in climate science (Franzke et al., 2020), human physiology (Goldberger et al., 1990; Gormally and Romero, 2020) and forest ecosystems (Blanco et al., 2020). However, in simple models, this phenomenon is underrepresented. Obviously, the effect of different scales cannot be tested in a isolated way as it relies on multiple interactions by definition. However, even in simple models these effects can be explored, as demonstrated for example by Van Der Laan and Hogeweg (1995) and Chaparro-Pedraza and de Roos (2020). Before we can successfully apply these theoretical ideas to real data, this step deserves some more attention.

What can noise teach us? The work in this thesis is largely inspired by studies that exploit the natural fluctuations, or the noise, of a system to infer some properties of the system’s behaviour (Scheffer et al., 2009). Using noise as a way to obtain information, instead of a nuisance to get rid of, is a relatively new idea which, I believe, could be expanded further. In most modelling studies, noise is implemented as Gaussian, white, and additive and it is the same for all variables (Boettiger, 2018). If this assumption is loosened, systems may show significantly different results (Arani, 2019), suggesting that the indicators seem more generic than they are, because they are often tested for one

particular noise regime. In **chapter 3** and **chapter 5** I have played around with the noise regime slightly, by changing the way it was distributed over the variables or by using multiplicative noise, which affected the resilience indicators under study, but I did not take a systematic approach to unravel the details of the possibilities of complex noise regimes. Therefore more insights in the effects of other noise distributions, other noise colors, and differently correlated noise are required, especially to strengthen the link with the real world where complex noise regimes seem to be the rule rather than the exception (Shoemaker et al., 2020).

Concluding remarks

With this thesis, I hope to have contributed to our understanding of the behaviour of complex systems. A complete understanding of complex systems is a goal that is too ambitious for one lifetime. However, the time seems ripe to embrace what Stephen Hawking has called “the century of complexity” (Hawking, 2000). The current widespread interest in complexity thinking has created a strong interdisciplinary scientific community, which has greatly inspired the work presented here, and which is ready to tackle the challenges ahead.

References

- U Rajendra Acharya, K Paul Joseph, Natarajan Kannathal, Choo Min Lim, and Jasjit S Suri. Heart rate variability: a review. *Medical and biological engineering and computing*, 44(12):1031–1051, 2006.
- Stavros P Adam, Stamatios-Aggelos N Alexandropoulos, Panos M Pardalos, and Michael N Vrahatis. No free lunch theorem: A review. *Approximation and optimization*, pages 57–82, 2019.
- Valentina Agostini, Emma Chiaramello, Lorenzo Canavese, Carla Bredariol, and Marco Knaflitz. Postural sway in volleyball players. *Human movement science*, 32(3):445–456, 2013.
- Andrew P Allen and Kevin E Thomas. A dual process account of creative thinking. *Creativity Research Journal*, 23(2):109–118, 2011.
- Richard B Alley, Jochem Marotzke, William D Nordhaus, Jonathan T Overpeck, Dorothy M Peteet, Roger A Pielke, Raymond T Pierrehumbert, Peter B Rhines, Thomas F Stocker, Lynne D Talley, et al. Abrupt climate change. *Science*, 299(5615):2005–2010, 2003.
- Zainy MH Almurad, Clément Roume, Hubert Blain, and Didier Delignières. Complexity matching: restoring the complexity of locomotion in older people through arm-in-arm walking. *Frontiers in Physiology*, 9:1766, 2018.
- RF Anderson, S Ali, LI Bradtmiller, SHH Nielsen, MQ Fleisher, BE Anderson, and LH Burckle. Wind-driven upwelling in the southern ocean and the deglacial rise in atmospheric co₂. *science*, 323(5920):1443–1448, 2009.
- David Angeli, James E Ferrell, and Eduardo D Sontag. Detection of multistability, bifurcations, and hysteresis in a large class of biological positive-feedback systems. *Proc. Natl. Acad. Sci. U.S.A.*, 101(7):1822–1827, 2004.
- Babak MS Arani. *Inferring ecosystem states and quantifying their resilience: linking theories to ecological data*. PhD thesis, Wageningen University, 2019.
- David E Archer, Gidon Eshel, Arne Winguth, Wallace Broecker, Ray Pierrehumbert, Michael Tobis, and Robert Jacob. Atmospheric pCO₂ sensitivity to the biological pump in the ocean. *Global Biogeochemical Cycles*, 14(4):1219–1230, 2000.

- Alex Arenas, Albert Diaz-Guilera, and Conrad J Pérez-Vicente. Synchronization reveals topological scales in complex networks. *Physical review letters*, 96(11):114102, 2006.
- Willem A Arrindell and Jan Van der Ende. An empirical test of the utility of the observations-to-variables ratio in factor and components analysis. *Appl. Psychol. Meas.*, 9(2):165–178, 1985.
- W Brian Arthur. Competing technologies, increasing returns, and lock-in by historical events. *Econ. J.*, 99(394):116–131, 1989.
- Matthijs Baas, Carsten KW De Dreu, and Bernard A Nijstad. A meta-analysis of 25 years of mood-creativity research: Hedonic tone, activation, or regulatory focus? *Psychological bulletin*, 134(6):779, 2008.
- Alberto Bacchelli and Christian Bird. Expectations, outcomes, and challenges of modern code review. In *2013 35th International Conference on Software Engineering (ICSE)*, pages 712–721. IEEE, 2013.
- György Baffy and Joseph Loscalzo. Complexity and network dynamics in physiological adaptation: an integrated view. *Physiology & behavior*, 131:49–56, 2014.
- David Balduzzi and Giulio Tononi. Integrated information in discrete dynamical systems: motivation and theoretical framework. *PLoS Comput Biol*, 4(6):e1000091, 2008.
- Rotem Bar-Or, Carynelisa Erlick, and Hezi Gildor. The role of dust in glacial–interglacial cycles. *Quaternary Science Reviews*, 27:201–208, 2008.
- Albert-László Barabási. Taming complexity. *Nature physics*, 1(2):68–70, 2005.
- Albert-László Barabási and Eric Bonabeau. Scale-free networks. *Scientific american*, 288(5):60–69, 2003.
- Albert-László Barabási et al. *Network science*. Cambridge university press, 2016.
- Benoît G Bardy, Olivier Oullier, Reinoud J Bootsma, and Thomas A Stoffregen. Dynamics of human postural transitions. *Journal of Experimental Psychology: Human Perception and Performance*, 28(3):499, 2002.
- Nick Barnes. Publish your computer code: it is good enough. *Nature News*, 467(7317):753–753, 2010.
- Paul T Barrett and Paul Kline. The observation to variable ratio in factor analysis. *Personality Study & Group Behaviour*, 1(1):23–33, 1981.
- Jordi Bascompte, Pedro Jordano, Carlos J Melián, and Jens M Olesen. The nested assembly of plant–animal mutualistic networks. *Proc. Natl. Acad. Sci. U.S.A.*, 100(16):9383–9387, 2003.
- Jordi Bascompte, Pedro Jordano, and Jens M Olesen. Asymmetric coevolutionary networks facilitate biodiversity maintenance. *Science*, 312(5772):431–433, 2006.
- Ugo Bastolla, Miguel A Fortuna, Alberto Pascual-García, Antonio Ferrera, Bartolo Luque,

-
- and Jordi Bascompte. The architecture of mutualistic networks minimizes competition and increases biodiversity. *Nature*, 458(7241):1018–1020, 2009.
- Sebastian Bathiany, Martin Claussen, and K Fraedrich. Detecting hotspots of atmosphere–vegetation interaction via slowing down—part 1: A stochastic approach. *Earth System Dynamics*, 4(1):63–78, 2013.
- Ivan Bautmans and Tony Mets. A fatigue resistance test for elderly persons based on grip strength: reliability and comparison with healthy young subjects. *Aging clinical and experimental research*, 17(3):217–222, 2005.
- Ivan Bautmans, Ellen Gorus, Rose Njemini, and Tony Mets. Handgrip performance in relation to self-perceived fatigue, physical functioning and circulating il-6 in elderly persons without inflammation. *BMC geriatrics*, 7(1):5, 2007.
- Brian Beckage, Louis J Gross, and Stuart Kauffman. The limits to prediction in ecological systems. *Ecosphere*, 2(11):1–12, 2011.
- Edward A Bender, Ted J Case, and Michael E Gilpin. Perturbation experiments in community ecology: theory and practice. *Ecology*, 65(1):1–13, 1984.
- Elisa Benincà, Jef Huisman, Reinhard Heerkloss, Klaus D Jöhnk, Pedro Branco, Egbert H Van Nes, Marten Scheffer, and Stephen P Ellner. Chaos in a long-term experiment with a plankton community. *Nature*, 451(7180):822–825, 2008.
- Mark D Bertness and Ragan Callaway. Positive interactions in communities. *Trends Ecol. Evol.*, 9(5):191–193, 1994.
- Reinette Biggs, Garry D Peterson, and Juan C Rocha. The regime shifts database: a framework for analyzing regime shifts in social-ecological systems. 2018.
- Juan A Blanco, Aitor Ameztegui, and Francisco Rodríguez. Modelling forest ecosystems: A crossroad between scales, techniques, and applications. *Ecol. Model.*, 425:109030, 2020.
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Occam’s razor. *Information processing letters*, 24(6):377–380, 1987.
- Maarten C Boerlijst, Thomas Oudman, and André M de Roos. Catastrophic collapse can occur without early warning: examples of silent catastrophes in structured ecological models. *PloS one*, 8(4):e62033, 2013.
- Carl Boettiger. From noise to knowledge: how randomness generates novel phenomena and reveals information. *Ecology letters*, 21(8):1255–1267, 2018.
- Carl Boettiger, Noam Ross, and Alan Hastings. Early warning signals: the charted and uncharted territories. *Theoretical ecology*, 6(3):255–264, 2013.
- Johan Bollen, Bruno Gonçalves, Ingrid van de Leemput, and Guangchen Ruan. The happiness paradox: your friends are happier than you. *EPJ Data Science*, 6(1):4, 2017.

- Frank G Borg and Gerd Laxåback. Entropy of balance-some recent results. *Journal of neuroengineering and rehabilitation*, 7(1):1–11, 2010.
- Denny Borsboom. A network theory of mental disorders. *World psychiatry*, 16(1):5–13, 2017.
- Elisabeth H Bos and Peter De Jonge. “critical slowing down in depression” is a great idea that still needs empirical proof. *Proceedings of the National Academy of Sciences*, 111(10):E878–E878, 2014.
- Nathælle Bouttes, Didier M Roche, and D Paillard. Systematic study of the impact of fresh water fluxes on the glacial carbon cycle. *Climate of the Past*, 8:589–607, 2012.
- PW Boyd and CS Law. The Southern Ocean Iron RElease Experiment (SOIREE) – Introduction and summary. *Deep-Sea Research II*, 48:2425–2438, 2001.
- Anil Bozbiyik, Marco Steinacher, Fortunat Joos, Thomas F Stocker, and Laurie Menviel. Fingerprints of changes in the terrestrial carbon cycle in response to large reorganizations in ocean circulation. *Climate of the Past*, 7:319–338, 2011.
- Anita Breuer, Todd Landman, and Dorothea Farquhar. Social media and protest mobilization: Evidence from the tunisian revolution. *Democratization*, 22(4):764–792, 2015.
- William A Brock and Stephen R Carpenter. Interacting regime shifts in ecosystems: implication for early warnings. *Ecological Monographs*, 80(3):353–367, 2010.
- Wallace Broecker and Stephen Barker. A 190 permil drop in atmosphere’s $\delta^{14}\text{C}$ during the “Mystery Interval” (17.5 to 14.5 kyr). *Earth and Planetary Science Letters*, 256(1-2):90–99, 2007.
- Wallace S Broecker. Paleocean circulation during the last deglaciation: A bipolar seesaw? *Paleoceanography*, 13:119–121, 1998.
- Wally Broecker and Elizabeth Clark. Search for a glacial-age ^{14}C -depleted ocean reservoir. *Geophysical Research Letters*, 37:L13606, 2010.
- WS Broecker and D McGee. The ^{13}C record for atmospheric CO_2 : What is it trying to tell us? *Earth and Planetary Science Letters*, 368:175–182, 2013.
- Lindell Bromham, Russell Dinnage, and Xia Hua. Interdisciplinary research has consistently lower funding success. *Nature*, 534(7609):684–687, 2016.
- Victor Brovkin, Martin Claussen, Vladimir Petoukhov, and Andrey Ganopolski. On the stability of the atmosphere-vegetation system in the sahara/sahel region. *Journal of Geophysical Research: Atmospheres*, 103(D24):31613–31624, 1998.
- Victor Brovkin, Andrey Ganopolski, David Archer, and Stefan Rahmstorf. Lowering of glacial atmospheric CO_2 in response to changes in oceanic circulation and marine biogeochemistry. *Paleoceanography and Paleoclimatology*, 22(4), 2007.
- Marc Brysbaert, Paweł Mandera, Samantha F McCormick, and Emmanuel Keuleers.

-
- Word prevalence norms for 62,000 english lemmas. *Behavior research methods*, 51(2): 467–479, 2019.
- Timothy G Buchman. The community of the self. *Nature*, 420(6912):246–251, 2002.
- Andrea Burke, Andrew L Stewart, Jess F Adkins, Raffaele Ferrari, Malte F Jansen, and Andrew F Thompson. The glacial mid-depth radiocarbon bulge and its implications for the overturning circulation. *Paleoceanography*, 20:1021–1039, 2015.
- Laura A Burkle, John C Marlin, and Tiffany M Knight. Plant-pollinator interactions over 120 years: loss of species, co-occurrence, and function. *Science*, 339(6127):1611–1615, 2013.
- Tony Cai, Jianqing Fan, and Tiefeng Jiang. Distributions of angles in random packing on spheres. *The Journal of Machine Learning Research*, 14(1):1837–1864, 2013.
- Curtis B Caldwell, Erica L Moran, and Earl R Bogoch. Fractal dimension as a measure of altered trabecular bone in experimental inflammatory arthritis. *Journal of bone and mineral research*, 13(6):978–985, 1998.
- Ragan M Callaway. Positive interactions among plants. *Bot. Rev.*, 61(4):306–349, 1995.
- Liangyue Cao. Practical method for determining the minimum embedding dimension of a scalar time series. *Physica D: Nonlinear Phenomena*, 110(1-2):43–50, 1997.
- SR Carpenter and WA Brock. Rising variance: a leading indicator of ecological transition. *Ecology letters*, 9(3):311–318, 2006.
- Stephen R Carpenter, Babak MS Arani, Paul C Hanson, Marten Scheffer, Emily H Stanley, and Egbert Van Nes. Stochastic dynamics of cyanobacteria in long-term high-frequency observations of a eutrophic lake. *Limnology and Oceanography Letters*, 5(5): 331–336, 2020.
- Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of modern physics*, 81(2):591, 2009.
- Nina R Catubig, DE Archer, Roger Francois, Peter DeMenocal, Will Howard, and Ein-Fen Yu. Global deep-sea burial rate of calcium carbonate during the last glacial maximum. *Paleoceanography*, 13:298–310, 1998.
- Carlos A Celis-Morales, Paul Welsh, Donald M Lyall, Lewis Steell, Fanny Petermann, Jana Anderson, Stamatina Iliodromiti, Anne Sillars, Nicholas Graham, Daniel F Mackay, et al. Associations of grip strength with cardiovascular, respiratory, and cancer outcomes and all cause mortality: prospective cohort study of half a million uk biobank participants. *Bmj*, 361, 2018.
- Martin Chalfie, Yuan Tu, Ghia Euskirchen, William W Ward, and Douglas C Prasher. Green fluorescent protein as a marker for gene expression. *Science*, 263(5148):802–805, 1994.

- P Catalina Chaparro-Pedraza and André M de Roos. Ecological changes with minor effect initiate evolution to delayed regime shifts. *Nature ecology & evolution*, 4(3):412–418, 2020.
- Luonan Chen, Rui Liu, Zhi-Ping Liu, Meiyi Li, and Kazuyuki Aihara. Detecting early-warning signals for sudden deterioration of complex diseases by dynamical network biomarkers. *Scientific reports*, 2, 2012.
- Shiyang Chen, Eamon B O’Dea, John M Drake, and Bogdan I Epureanu. Eigenvalues of the covariance matrix as early warning signals for critical transitions in ecological systems. *Scientific reports*, 9(1):1–14, 2019.
- Dante R Chialvo. Emergent complexity: what uphill analysis or downhill invention cannot do. *New Ideas in Psychology*, 26(2):158–173, 2008.
- Megumi O Chikamoto, Laurie Menviel, Ayako Abe-Ouchi, Rumi Ohgaito, Axel Timmermann, Yusuke Okazaki, Naomi Harada, Akira Oka, and Anne Mouchet. Variability in North Pacific intermediate and deep water ventilation during Heinrich events in two coupled climate models. *Deep-Sea Research II*, 61–64:114–126, 2012.
- James S Clark, Steven R Carpenter, Mary Barber, Scott Collins, Andy Dobson, Jonathan A Foley, David M Lodge, Mercedes Pascual, Roger Pielke, William Pizer, et al. Ecological forecasts: an emerging imperative. *Science*, 293(5530):657–660, 2001.
- Audrey Coreau, Gilles Pinay, John D Thompson, Pierre-Olivier Cheptou, and Laurent Mermet. The rise of research on futures in ecology: rebalancing scenarios and predictions. *Ecol. Lett.*, 12(12):1277–1286, 2009.
- Madalena Costa, Ary L Goldberger, and C-K Peng. Multiscale entropy analysis of complex physiologic time series. *Physical review letters*, 89(6):068102, 2002.
- Franck Courchamp, Tim Clutton-Brock, and Bryan Grenfell. Inverse density dependence and the allee effect. *Trends Ecol. Evol.*, 14(10):405–410, 1999.
- Nicholas B Cross, Jonathan C Craig, and Angela C Webster. Asking the right question and finding the right answers. *Nephrology*, 15(1):8–11, 2010.
- Thomas J Crowley and Steven K Baum. Is the greenland ice sheet bistable? *Paleoceanography*, 10(3):357–363, 1995.
- Michel Crucifix. Oscillators and relaxation phenomena in Pleistocene climate theory. *Philosophical Transactions of the Royal Society A*, 370:1140–1165, 2012.
- Michel Crucifix, Guillaume Lenoir, and Takahito Mitsui. *Challenges for Ice Age Dynamics: A Dynamical Systems Perspective*, page 1–32. Cambridge University Press, 2017. doi: 10.1017/9781316339251.002.
- John Crusius, Thomas F Pedersen, Stephanie Kienast, Lloyd Keigwin, and Laurent Labeyrie. Influence of northwest Pacific productivity on North Pacific Intermediate Water oxygen concentrations during the Bølling-Allerød interval (14.7–12.9 ka). *Geol-*

- ogy, 32:633–636, 2004.
- Judith A Curry and Peter J Webster. Climate science and the uncertainty monster. *Bulletin of the American Meteorological Society*, 92(12):1667–1682, 2011.
- William B Curry and Delia W Oppo. Glacial water mass geometry and the distribution of $\delta^{13}\text{C}$ of σCO_2 in the western Atlantic Ocean. *Paleoceanography*, 20:PA1017, 2005.
- Hélène Cyr and Michael L Face. Magnitude and patterns of herbivory in aquatic and terrestrial ecosystems. *Nature*, 361(6408):148–150, 1993.
- Lei Dai, Daan Vorselen, Kirill S Korolev, and Jeff Gore. Generic indicators for loss of resilience before a tipping point leading to population collapse. *Science*, 336(6085):1175–1177, 2012.
- Vasilis Dakos. Identifying best-indicator species for abrupt transitions in multispecies communities. *Ecological indicators*, 94:494–502, 2018.
- Vasilis Dakos and Jordi Bascompte. Critical slowing down as early warning for the onset of collapse in mutualistic communities. *Proceedings of the National Academy of Sciences*, 111(49):17546–17551, 2014.
- Vasilis Dakos and Fernando Soler-Toscano. Measuring complexity to infer changes in the dynamics of ecological systems under stress. *Ecological Complexity*, 32:144–155, 2017.
- Vasilis Dakos, Marten Scheffer, Egbert H van Nes, Victor Brovkin, Vladimir Petoukhov, and Hermann Held. Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*, 105(38):14308–14312, 2008.
- Vasilis Dakos, Stephen R Carpenter, William A Brock, Aaron M Ellison, Vishwesha Guttal, Anthony R Ives, Sonia Kefi, Valerie Livina, David A Seekell, Egbert H van Nes, et al. Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. *PloS one*, 7(7):e41010, 2012a.
- Vasilis Dakos, Egbert H Van Nes, Paolo D’Odorico, and Marten Scheffer. Robustness of variance and autocorrelation as indicators of critical slowing down. *Ecology*, 93(2):264–271, 2012b.
- Vasilis Dakos, Egbert H van Nes, and Marten Scheffer. Flickering as an early warning signal. *Theoretical ecology*, 6(3):309–317, 2013.
- Vasilis Dakos, Stephen R Carpenter, Egbert H van Nes, and Marten Scheffer. Resilience indicators: prospects and limitations for early warnings of regime shifts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1659):20130263, 2015.
- Vasilis Dakos, Sarah M Glaser, Chih-hao Hsieh, and George Sugihara. Elevated nonlinearity as an indicator of shifts in the dynamics of populations under stress. *Journal of The Royal Society Interface*, 14(128):20160845, 2017.
- Jeffrey M Dambacher, Hang-Kwang Luh, Hiram W Li, and Philippe A Rossignol. Quali-

- tative stability and ambiguity in model ecosystems. *Am. Nat.*, 161(6):876–888, 2003.
- István Daruka and Peter D Ditlevsen. A conceptual model for glacial cycles and the middle Pleistocene transition. *Climate dynamics*, 46(1-2):29–40, 2016.
- Charles Darwin. *The origin of species*. PF Collier & son New York, 1909.
- Dhurjati Prasad Datta and Santanu Raut. The arrow of time, complexity and the scale free analysis. *Chaos, Solitons & Fractals*, 28(3):581–589, 2006.
- Susanne MM de Mooij, Tessa F Blanken, Raoul PPP Grasman, Jennifer R Ramautar, Eus JW Van Someren, and Han LJ van der Maas. Dynamics of sleep: Exploring critical transitions and early warning signals. *Computer methods and programs in biomedicine*, 193:105448, 2020.
- Peter C De Ruiter, Anje-Margriet Neutel, and John C Moore. Energetics, patterns of interaction strengths, and stability in real ecosystems. *Science*, 269(5228):1257–1260, 1995.
- Julia Dehm. Highlighting inequalities in the histories of human rights: Contestations over justice, needs and rights in the 1970s. *LJIL*, 31:871, 2018.
- Didier Delignières, Marina Fortes, Grégory Ninot, et al. The fractal dynamics of self-esteem and physical self. *Nonlinear Dynamics in Psychology and Life Sciences*, 8(4):479–510, 2004.
- AJ Desbarats and R Dimitrakopoulos. Geostatistical simulation of regionalized pore-size distributions using min/max autocorrelation factors. *Mathematical Geology*, 32(8):919–942, 2000.
- Curtis Deutsch, Daniel M Sigman, Robert C Thunell, Anna Nele Meckler, and Gerald H Haug. Isotopic constraints on glacial/interglacial changes in the oceanic nitrogen budget. *Global Biogeochemical Cycles*, 18:GB4012, 2004.
- Douglas W Diamond and Philip H Dybvig. Bank runs, deposit insurance, and liquidity. *J. Political Econ.*, 91(3):401–419, 1983.
- Louis A DiBerardino III, John D Polk, Karl S Rosengren, Jesse B Spencer-Smith, and Elizabeth T Hsiao-Wecksler. Quantifying complexity and variability in phase portraits of gait. *Clinical Biomechanics*, 25(6):552–556, 2010.
- Joan E Dodgson. Quality in research: Asking the right question. *Journal of Human Lactation*, 36(1):105–108, 2020.
- Karen M Douglas, Joseph E Uscinski, Robbie M Sutton, Aleksandra Cichocka, Turkey Nefes, Chee Siang Ang, and Farzin Deravi. Understanding conspiracy theories. *Political Psychology*, 40:3–35, 2019.
- SM Downes, AS Budnick, Jorge Louis Sarmiento, and R Farneti. Impacts of wind stress on the antarctic circumpolar current fronts and associated subduction. *Geophysical*

-
- Research Letters*, 38:L11605, 2011.
- Carlos M Duarte, Lucille Chapuis, Shaun P Collin, Daniel P Costa, Reny P Devassy, Victor M Eguiluz, Christine Erbe, Timothy AC Gordon, Benjamin S Halpern, Harry R Harding, et al. The soundscape of the anthropocene ocean. *Science*, 371(6529), 2021.
- Marcos Duarte and Dagmar Sternad. Complexity of human postural control in young and older adults during prolonged standing. *Experimental brain research*, 191(3):265–276, 2008.
- Gérard Duménil and Dominique Lévy. The neoliberal (counter-) revolution. *Neoliberalism: A critical reader*, pages 9–19, 2005.
- David Easley and Jon Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press, Cambridge, 2010.
- Tarsha Eason and Heriberto Cabezas. Evaluating the sustainability of a regional system using fisher information in the san luis basin, colorado. *Journal of environmental management*, 94(1):41–49, 2012.
- Wesley Ebisuzaki. A method to estimate the statistical significance of a correlation when the data are serially correlated. *Journal of Climate*, 10(9):2147–2153, 1997.
- S Eggleston, Jochen Schmitt, Bernhard Bereiter, Robert Schneider, and Hubertus Fischer. Evolution of the stable carbon isotope composition of atmospheric CO₂ over the last glacial cycle. *Paleoceanography*, 31:434–452, 2016.
- Ivar Ekeland. *Mathematics and the Unexpected*. University of Chicago Press, 1990.
- Phoebe C Ellsworth. Appraisal theory: Old and new questions. *Emotion Review*, 5(2):125–131, 2013.
- Thomas Elmqvist, Carl Folke, Magnus Nyström, Garry Peterson, Jan Bengtsson, Brian Walker, and Jon Norberg. Response diversity, ecosystem change, and resilience. *Frontiers in Ecology and the Environment*, 1(9):488–494, 2003.
- Samuel Epstein, Ralph Buchsbaum, Heinz Lowenstam, and Harold C Urey. Carbonate-water isotopic temperature scale. *Geological Society of America Bulletin*, 62(4):417–426, 1951.
- Matthew R Evans, Ken J Norris, and Tim G Benton. Predictive ecology: systems approaches. *Philos. Trans. Royal Soc. B*, 367:163–167, 2012.
- Matthew R Evans, Mike Bithell, Stephen J Cornell, Sasha RX Dall, Sandra Díaz, Stephen Emmott, Bruno Ernande, Volker Grimm, David J Hodgson, Simon L Lewis, et al. Predictive systems ecology. *Proc. Royal Soc. B*, 280(1771):20131452, 2013a.
- Matthew R Evans, Volker Grimm, Karin Johst, Tarja Knuuttila, Rogier De Langhe, Catherine M Lessells, Martina Merz, Maureen A O’Malley, Steve H Orzack, Michael Weisberg, et al. Do simple models lead to generality in ecology? *Trends in ecology &*

- evolution*, 28(10):578–583, 2013b.
- Jasmine Fardouly and Elise Holland. Social media is not real life: The effect of attaching disclaimer-type labels to idealized social media images on women’s body image and mood. *New media & society*, 20(11):4311–4328, 2018.
- Michael JR Fasham, editor. *Ocean biogeochemistry*. Springer, Heidelberg, 2003.
- Scott L Feld. Why your friends have more friends than you do. *American Journal of Sociology*, 96(6):1464–1477, 1991.
- Raffaele Ferrari, Malte F Jansen, Jess F Adkins, Andrea Burke, Andrew L Stewart, and Andrew F Thompson. Antarctic sea ice control on ocean circulation in present and glacial climates. *Proceedings of the National Academy of Sciences*, 111:8753–8758, 2014.
- James E Ferrell Jr. Self-perpetuating states in signal transduction: positive feedback, double-negative feedback and bistability. *Curr. Opin. Cell Biol.*, 14(2):140–148, 2002.
- Hubertus Fischer, Marie-Louise Siggaard-Andersen, Urs Ruth, Regine Röthlisberger, and Eric Wolff. Glacial/interglacial changes in mineral dust and sea-salt records in polar ice cores: Sources, transport, and deposition. *Reviews of Geophysics*, 45(1), 2007.
- Aaron J Fisher, John D Medaglia, and Bertus F Jeronimus. Lack of group-to-individual generalizability is a threat to human subjects research. *Proceedings of the National Academy of Sciences*, 115(27):E6106–E6115, 2018.
- Ruben Fossion, Ana Leonor Rivera, and Bruno Estañol. A physicist’s view of homeostasis: how time series of continuous monitoring reflect the function of physiological variables in regulatory mechanisms. *Physiological measurement*, 39(8):084007, 2018.
- Christian LE Franzke, Susana Barbosa, Richard Blender, Hege-Beate Fredriksen, Thomas Laepple, Fabrice Lambert, Tine Nilsen, Kristoffer Rypdal, Martin Rypdal, Manuel G Scotto, et al. The structure of climate variability across scales. *Reviews of Geophysics*, 58(2):e2019RG000657, 2020.
- Diane Larsen Freeman and Lynne Cameron. Research methodology on language development from a complex systems perspective. *The modern language journal*, 92(2): 200–213, 2008.
- Eric D Galbraith, Samuel L Jaccard, Thomas F Pedersen, Daniel M Sigman, Gerald H Haug, Mea Cook, John R Southon, and Roger Francois. Carbon dioxide release from the North Pacific abyss during the last deglaciation. *Nature*, 449:890–893, 2007.
- Jianxi Gao, Baruch Barzel, and Albert-László Barabási. Universal resilience patterns in complex networks. *Nature*, 530(7590):307, 2016.
- Scott T Gaynor, V Robin Weersing, David J Kolko, Boris Birmaher, Jungeun Heo, and David A Brent. The prevalence and impact of large sudden improvements during adolescent therapy for depression: A comparison across cognitive-behavioral, family,

- and supportive therapy. *Journal of consulting and clinical psychology*, 71(2):386, 2003.
- IM Gelfand and AM Yaglom. Calculation of the amount of information about a random function contained in another such function. *American Mathematical Society Translation. Series*, 2(12):191–198, 1959.
- Peter R Gent. Effects of Southern Hemisphere wind changes on the Meridional Overturning Circulation in ocean models. *Annual Review of Marine Science*, 8:79–94, 2016.
- Paolo Gerbaudo and Emiliano Treré. In search of the ‘we’ of social media activism: introduction to the special issue on social media and protest identities, 2015.
- Jaboury Ghazoul. Floral diversity and the facilitation of pollination. *J. Ecol.*, 94(2):295–304, 2006.
- Sanne MW Gijzel, Ingrid A van de Leemput, Marten Scheffer, Geert EA van Bon, Vivian Weerdesteyn, Thijs MH Eijsvogels, Maria TE Hopman, Marcel GM Olde Rikkert, and René JF Melis. Dynamical indicators of resilience in postural balance time series are related to successful aging in high-functioning older adults. *The Journals of Gerontology: Series A*, 74(7):1119–1126, 2019a.
- Sanne MW Gijzel, Heather E Whitson, Ingrid A van de Leemput, Marten Scheffer, Dieneke van Asselt, Jerrald L Rector, Marcel GM Olde Rikkert, and René JF Melis. Resilience in clinical care: getting a grip on the recovery potential of older adults. *Journal of the American Geriatrics Society*, 67(12):2650–2657, 2019b.
- Isabelle M Gil, Lloyd D Keigwin, and Fatima G Abrantes. Deglacial diatom productivity and surface ocean properties over the Bermuda Rise, northeast Sargasso Sea. *Paleoceanography*, 24:PA4101, 2009.
- Malcolm Gladwell. *The tipping point: How little things can make a big difference*. Little, Brown, 2006.
- AL Golberger. Non-linear dynamics for clinicians: chaos theory, fractals, and complexity at the bedside. *The Lancet*, 347(9011):1312–1314, 1996.
- Ary L Goldberger, David R Rigney, and Bruce J West. Chaos and fractals in human physiology. *Scientific American*, 262(2):42–49, 1990.
- Albert Goldbeter. *Biochemical oscillations and cellular rhythms: the molecular bases of periodic and chaotic behaviour*. Cambridge University Press, Cambridge, 1996.
- Brenna MG Gormally and L Michael Romero. What are you actually measuring? a review of techniques that integrate the stress response on distinct time-scales. *Functional Ecology*, 34(10):2030–2044, 2020.
- Nicholas J Gotelli and Werner Ulrich. Statistical challenges in null model analysis. *Oikos*, 121(2):171–180, 2012.
- Allan Gottschalk, Mark S Bauer, and Peter C Whybrow. Evidence of chaotic mood

- variation in bipolar disorder. *Archives of general psychiatry*, 52(11):947–959, 1995.
- Jean-Luc Gouzé. Positive and negative circuits in dynamical systems. *J. Biol. Syst.*, 6(01):11–15, 1998.
- Peter Grassberger. Do climatic attractors exist? *Nature*, 323(6089):609–612, 1986.
- Peter Grassberger. Randomness, information, and complexity. *arXiv preprint arXiv:1208.3459*, 2012.
- Peter Grassberger and Itamar Procaccia. Characterization of strange attractors. *Physical review letters*, 50(5):346, 1983a.
- Peter Grassberger and Itamar Procaccia. Estimation of the kolmogorov entropy from a chaotic signal. *Physical review A*, 28(4):2591, 1983b.
- Celso Grebogi, Edward Ott, and James A Yorke. Crises, sudden changes in chaotic attractors, and transient chaos. *Physica D*, 7(1-3):181–200, 1983.
- Jessica L Green, Alan Hastings, Peter Arzberger, Francisco J Ayala, Kathryn L Cottingham, Kim Cuddington, Frank Davis, Jennifer A Dunne, Marie-Josée Fortin, Leah Gerber, et al. Complexity in ecology and conservation: mathematical, statistical, and computational challenges. *BioScience*, 55(6):501–510, 2005.
- Volker Grimm, Eloy Revilla, Uta Berger, Florian Jeltsch, Wolf M Mooij, Steven F Railsback, Hans-Hermann Thulke, Jacob Weiner, Thorsten Wiegand, and Donald L DeAngelis. Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *science*, 310(5750):987–991, 2005.
- Ross Haenfler, Brett Johnson, and Ellis Jones. Lifestyle movements: Exploring the intersection of lifestyle and social movements. *Social Movement Studies*, 11(1):1–20, 2012.
- Hermann Haken. Synergetics. In *Self-Organizing Systems*, pages 417–434. Springer, 1987.
- Shihui Han and Ernst Pöppel. *Culture and neural frames of cognition and communication*. Springer Science & Business Media, 2011.
- Sarah M Hanley, Susan E Watt, and William Coventry. Taking a break: The effect of taking a vacation from facebook and instagram on subjective well-being. *Plos one*, 14(6):e0217743, 2019.
- Ilkka Hanski. Metapopulation dynamics. *Nature*, 396(6706):41–49, 1998.
- Steven R Hare and Nathan J Mantua. Empirical evidence for north pacific regime shifts in 1977 and 1989. *Prog. Oceanogr.*, 47(2-4):103–145, 2000.
- Adam P Hasenfratz, Samuel L Jaccard, Alfredo Martínez-García, Daniel M Sigman, David A Hodell, Derek Vance, Stefano M Bernasconi, Helga Kikki F Kleiven, F Alexander Haumann, and Gerald H Haug. The residence time of Southern Ocean surface waters and the 100,000-year ice age cycle. *Science*, 363:1080–1084, 2019.
- Alan Hastings. Timescales, dynamics, and ecological understanding. *Ecology*, 91(12):

- 3471–3480, 2010.
- Alan Hastings and Thomas Powell. Chaos in a three-species food chain. *Ecology*, 72(3): 896–903, 1991.
- Alan Hastings and Derin B Wysham. Regime shifts in ecological systems can occur with no warning. *Ecol. Lett.*, 13(4):464–472, 2010.
- Jeff Hasty, David McMillen, and James J Collins. Engineered gene circuits. *Nature*, 420 (6912):224–230, 2002.
- Matz A Haugen, Bala Rajaratnam, and Paul Switzer. Extracting common time trends from concurrent time series: Maximum autocorrelation factors with application to tree ring time series data. *arXiv preprint arXiv:1502.01073*, 2015.
- Stephen Hawking. What is complexity?, 2000. Available at <http://www.complexsys.org/>.
- CT Hayes, RF Anderson, and MQ Fleisher. Opal accumulation rates in the Equatorial Pacific and mechanisms of deglaciation. *Paleoceanography*, 26:PA1207, 2011.
- Hermann Held and Thomas Kleinen. Detection of climate system bifurcations by degenerate fingerprinting. *Geophysical Research Letters*, 31(23), 2004.
- Rick Hennekam, Bregje van der Bolt, Egbert H van Nes, Gert J de Lange, Marten Scheffer, and Gert-Jan Reichart. Early-warning signals for marine anoxic events. *Geophysical Research Letters*, 47(20):e2020GL089183, 2020.
- Linda Heskamp, AS Meel-van den Abeelen, Joep Lagro, and JA Claassen. Convergent cross mapping: a promising technique for cerebral autoregulation estimation. *Int. J. Clin. Neurosci. Ment. Heal. S*, 20, 2014.
- Thomas T Hills, Eugenio Proto, Daniel Sgroi, and Chanuki Illushka Seresinha. Historical analysis of national subjective wellbeing using millions of digitized books. *Nature human behaviour*, 3(12):1271–1275, 2019.
- Yoshito Hirata, José M Amigó, Yoshiya Matsuzaka, Ryo Yokota, Hajime Mushiake, and Kazuyuki Aihara. Detecting causality by combined use of multiple methods: Climate and brain examples. *PloS one*, 11(7):e0158572, 2016.
- Marina Hirota, Milena Holmgren, Egbert H Van Nes, and Marten Scheffer. Global resilience of tropical forest and savanna to critical transitions. *Science*, 334(6053):232–235, 2011.
- Crawford Stanley Holling. Engineering resilience versus ecological resilience. *Engineering within ecological constraints*, 31(1996):32, 1996.
- Milena Holmgren, Marten Scheffer, and Michael A Huston. The interplay of facilitation and competition in plant communities. *Ecology*, 78(7):1966–1975, 1997.
- Jasper Homminga, Bert Van-Rietbergen, Eva-Maria Lochmüller, Harrie Weinans, Felix

- Eckstein, and Rik Huiskes. The osteoporotic vertebral structure is well adapted to the loads of daily life, but not to infrequent “error” loads. *Bone*, 34(3):510–516, 2004.
- Werner Horsthemke. Noise induced transitions. In *Non-Equilibrium Dynamics in Chemical Systems*, pages 150–160. Springer, 1984.
- Toke T Høye, Eric Post, Niels M Schmidt, Kristian Trøjelsgaard, and Mads C Forchhammer. Shorter flowering seasons and declining abundance of flower visitors in a warmer arctic. *Nat. Clim. Change*, 3(8):759–763, 2013.
- Yu Huang, Christian LE Franzke, Naiming Yuan, and Zuntao Fu. Systematic identification of causal relations in high-dimensional chaotic systems: application to stratosphere-troposphere coupling. *Climate Dynamics*, 55(9):2469–2481, 2020.
- Bernardo A Huberman and Tad Hogg. Complexity and adaptation. *Physica D: Nonlinear Phenomena*, 22(1-3):376–384, 1986.
- Rik Huiskes, Ronald Ruimerman, G Harry Van Lenthe, and Jan D Janssen. Effects of mechanical forces on maintenance and adaptation of form in trabecular bone. *Nature*, 405(6787):704–706, 2000.
- Jef Huisman and Franz J Weissing. Biodiversity of plankton by species oscillations and chaos. *Nature*, 402(6760):407–410, 1999.
- Harold E Hurst. The problem of long-term storage in reservoirs. *Hydrological Sciences Journal*, 1(3):13–27, 1956.
- Peter Huybers and Carl Wunsch. A depth-derived Pleistocene age model: Uncertainty estimates, sedimentation variability, and nonlinear climate change. *Paleoceanography*, 19:PA1028, 2004.
- Rumen Iliev and Robert Axelrod. Does causality matter more now? increase in the proportion of causal language in english texts. *Psychological science*, 27(5):635–643, 2016.
- Rumen Iliev, Joe Hoover, Morteza Dehghani, and Robert Axelrod. Linguistic positivity in historical texts reflects dynamic environmental and psychological factors. *Proceedings of the National Academy of Sciences*, 113(49):E7871–E7879, 2016.
- Intergovernmental Panel on Climate Change. *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, volume 5. IPCC, Geneva, 2014.
- Martin Iriondo. Patagonian dust in antarctica. *Quaternary International*, 68:83–86, 2000.
- Anthony R Ives. Measuring resilience in stochastic systems. *Ecological Monographs*, 65(2):217–233, 1995.
- Anthony R Ives, Brian Dennis, Kathryn L Cottingham, and Stephen R Carpenter. Estimating community stability and ecological interactions from time-series data. *Ecological*

- monographs*, 73(2):301–330, 2003.
- Samuel L Jaccard and Eric D Galbraith. Large climate-driven changes of oceanic oxygen concentrations during the last deglaciation. *Nature Geoscience*, 5:151–156, 2012.
- Samuel L Jaccard, Gerald H Haug, Daniel Mikhail Sigman, Thomas F Pedersen, Hans R Thierstein, and Ursula Röhl. Glacial/interglacial changes in subarctic North Pacific stratification. *Science*, 308:1003–1006, 2005.
- Samuel L Jaccard, Eric D Galbraith, Daniel Mikhail Sigman, and GH Haug. A pervasive link between Antarctic ice core and subarctic Pacific sediment records over the past 800 kyrs. *Quaternary Science Reviews*, 29:206–212, 2010.
- Samuel L Jaccard, Christopher T Hayes, Alfredo Martinez-Garcia, David A Hodell, Robert F Anderson, Daniel M Sigman, and GH Haug. Two modes of changes in Southern Ocean productivity over the past million years. *Science*, 339:1419–1423, 2013.
- Jerry A Jacobs and Scott Frickel. Interdisciplinarity: A critical assessment. *Annual review of Sociology*, 35, 2009.
- Malte F Jansen. Glacial ocean circulation and stratification explained by reduced atmospheric temperature. *Proceedings of the National Academy of Sciences*, 114:45–50, 2017.
- Richard Jenkins. Disenchantment, enchantment and re-enchantment: Max weber at the millennium. *Max Weber Studies*, pages 11–32, 2000.
- Junjie Jiang, Zi-Gang Huang, Thomas P Seager, Wei Lin, Celso Grebogi, Alan Hastings, and Ying-Cheng Lai. Predicting tipping points in mutualistic networks through dimension reduction. *Proc. Natl. Acad. Sci. U.S.A.*, 115(4):E639–E647, 2018.
- Pedro Jordano. Patterns of mutualistic interactions in pollination and seed dispersal: connectance, dependence asymmetries, and coevolution. *Am. Nat.*, 129(5):657–677, 1987.
- Daniel Kahneman. Thinking, fast and slow, 2017.
- Hyun Gu Kang, Madalena D Costa, Attila A Priplata, Olga V Starobinets, Ary L Goldberger, Chung-Kang Peng, Dan K Kiely, L Adrienne Cupples, and Lewis A Lipsitz. Frailty and the degradation of complex balance dynamics during a dual-task protocol. *Journals of Gerontology Series A: Biomedical Sciences and Medical Sciences*, 64(12):1304–1311, 2009.
- Peter M Kareiva, Joel G Kingsolver, Raymond B Huey, et al. *Biotic interactions and global change*. Sinauer Associates Inc., Sunderland, MA, 1993.
- Sabine Kasten, Ralf R Haese, Matthias Zabel, Carsten Rühlemann, and Horst D Schulz. Barium peaks at glacial terminations in sediments of the equatorial Atlantic Ocean – Relicts of deglacial productivity pulses? *Chemical Geology*, 175:635–651, 2001.

- Steve Keen. Commentary: measuring complexity–puzzles and tentative solutions. *Frontiers of evolutionary economics: competition, self-organization and innovation policy*, 2001.
- Sonia Kéfi, Milena Holmgren, and Marten Scheffer. When can positive interactions cause alternative stable states in ecosystems? *Functional Ecology*, 30(1):88–97, 2016.
- Pelin Kesebir and Selin Kesebir. The cultural salience of moral character and virtue declined in twentieth century america. *The Journal of Positive Psychology*, 7(6):471–480, 2012.
- E Toby Kiers, Marie Duhamel, Yugandhar Beesetty, Jerry A Mensah, Oscar Franken, Erik Verbruggen, Carl R Fellbaum, George A Kowalchuk, Miranda M Hart, Alberto Bago, et al. Reciprocal rewards stabilize cooperation in the mycorrhizal symbiosis. *Science*, 333(6044):880–882, 2011.
- Ava Kofman. Bruno latour, the post-truth philosopher, mounts a defense of science. *New York Times Magazine*, 25:2018, 2018.
- Karen E Kohfeld and Zanna Chase. Controls on deglacial changes in biogenic fluxes in the North Pacific Ocean. *Quaternary Science Reviews*, 30:3350–3363, 2011.
- Karen E Kohfeld, Corinne Le Quéré, Sandy P Harrison, and Robert F Anderson. Role of marine biology in glacial-interglacial CO₂ cycles. *Science*, 308(5718):74–78, 2005.
- Karen E Kohfeld, Robert M Graham, Agatha M De Boer, Louise C Sime, Eric W Wolff, Corinne Le Quéré, and Laurent Bopp. Southern Hemisphere westerly wind changes during the Last Glacial Maximum: Paleo-data synthesis. *Quaternary Science Reviews*, 68:76–95, 2013.
- Demetris Koutsoyiannis. Coupling stochastic models of different timescales. *Water Resources Research*, 37(2):379–391, 2001.
- Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.
- Douglas E Lake, Joshua S Richman, M Pamela Griffin, and J Randall Moorman. Sample entropy analysis of neonatal heart rate variability. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 283(3):R789–R797, 2002.
- Fabrice Lambert, Barbara Delmonte, Jean-Robert Petit, Matthias Bigler, Patrick R Kaufmann, Manuel A Hutterli, Thomas F Stocker, Urs Ruth, Jørgen Peder Steffensen, and Valter Maggi. Dust-climate couplings over the past 800,000 years from the EPICA Dome C ice core. *Nature*, 452:616–619, 2008.
- Munkhnasan Lamchin, Woo-Kyun Lee, Seong Woo Jeon, Sonam Wangyel Wang, Chul Hee Lim, Cholho Song, and Minjun Sung. Long-term trend and correlation between vegetation greenness and climate variables in asia based on satellite data. *Science of the Total Environment*, 618:1089–1095, 2018.

- Anna-Lena Lamprecht, Leyla Garcia, Mateusz Kuzak, Carlos Martinez, Ricardo Arcila, Eva Martin Del Pico, Victoria Dominguez Del Angel, Stephanie van de Sandt, Jon Ison, Paula Andrea Martinez, et al. Towards fair principles for research software. *Data Science*, 3(1):37–59, 2020.
- Boon Leong Lan, Yew Wai Liew, Mikito Toda, and Suraya Hani Kamsani. Flickering of cardiac state before the onset and termination of atrial fibrillation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(5):053137, 2020.
- Louis Lebel, John M Anderies, Bruce Campbell, Carl Folke, Steve Hatfield-Dodds, Terry P Hughes, and James Wilson. Governance and the capacity to manage resilience in regional social-ecological systems. *Ecology and society*, 11(1), 2006.
- Choong Ho Lee and Hyung-Jin Yoon. Medical big data: promise and challenges. *Kidney research and clinical practice*, 36(1):3, 2017.
- Tong Ihn Lee, Nicola J Rinaldi, François Robert, Duncan T Odom, Ziv Bar-Joseph, Georg K Gerber, Nancy M Hannett, Christopher T Harbison, Craig M Thompson, Itamar Simon, et al. Transcriptional regulatory networks in *saccharomyces cerevisiae*. *Science*, 298(5594):799–804, 2002.
- Ingrid A van de Leemput, Marieke Wichers, Angélique OJ Cramer, Denny Borsboom, Francis Tuerlinckx, Peter Kuppens, Egbert H van Nes, Wolfgang Viechtbauer, Erik J Giltay, Steven H Aggen, et al. Critical slowing down as early warning for the onset and termination of depression. *Proceedings of the National Academy of Sciences*, 111(1):87–92, 2014.
- Timothy M Lenton, Hermann Held, Elmar Kriegler, Jim W Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhuber. Tipping elements in the earth’s climate system. *Proceedings of the National Academy of Sciences*, 105(6):1786–1793, 2008.
- E Lespessailles, S Poupon, R Niamane, S Loiseau-Peres, G Derommelaere, R Harba, D Courteix, and CL Benhamou. Fractal analysis of trabecular bone texture on calcaneus radiographs: effects of age, time since menopause and hormone replacement therapy. *Osteoporosis international*, 13(5):366–372, 2002.
- J Jelle Lever, Egbert H van Nes, Marten Scheffer, and Jordi Bascompte. The sudden collapse of pollinator communities. *Ecology letters*, 17(3):350–359, 2014.
- J Jelle Lever, Ingrid A van de Leemput, Els Weinans, Rick Quax, Vasilis Dakos, Egbert H van Nes, Jordi Bascompte, and Marten Scheffer. Foreseeing the future of mutualistic communities beyond collapse. *Ecology letters*, 23(1):2–15, 2020.
- Simon A Levin. The problem of pattern and scale in ecology: the robert h. macarthur award lecture. *Ecology*, 73(6):1943–1967, 1992.
- Richard Levins. Discussion paper: the qualitative analysis of partially specified systems. *Annals of the New York Academy of Sciences*, 231(1):123–138, 1974.

- Anna Lichtwarck-Aschoff and Marieke MJW van Rooij. Are changes in children's communication patterns predictive of treatment outcomes for children with anxiety? *Clinical psychology & psychotherapy*, 26(5):572–585, 2019.
- Yuri Lin, Jean-Baptiste Michel, Erez Lieberman Aiden, Jon Orwant, William Brockman, and Slav Petrov. Syntactic annotations for the google books ngram corpus. 2012.
- Lorraine E Lisiecki. A benthic $\delta^{13}\text{C}$ -based proxy for atmospheric pCO_2 over the last 1.5 Myr. *Geophysical Research Letters*, 37:L21708, 2010a.
- Lorraine E Lisiecki. A simple mixing explanation for late Pleistocene changes in the Pacific-South Atlantic benthic $\delta^{13}\text{C}$ gradient. *Climate of the Past*, 6(3):305–314, 2010b.
- Lorraine E Lisiecki and Maureen E Raymo. A Pliocene-Pleistocene stack of 57 globally distributed benthic $\delta^{18}\text{O}$ records. *Paleoceanography*, 20:PA1003, 2005.
- Rui Liu, Meiyi Li, Zhi-Ping Liu, Jiarui Wu, Luonan Chen, and Kazuyuki Aihara. Identifying critical transitions and their leading biomolecular networks in complex diseases. *Scientific reports*, 2:813, 2012.
- ZhengYu Liu, Sang-Ik Shin, Robert S Webb, William Lewis, and Bette L Otto-Bliesner. Atmospheric CO_2 forcing on glacial thermohaline circulation and climate. *Geophysical Research Letters*, 32:8753–8758, 2005.
- Edward N Lorenz. Deterministic nonperiodic flow. *Journal of the atmospheric sciences*, 20(2):130–141, 1963.
- Chuan Luo, Xiaolong Zheng, and Daniel Zeng. Causal inference in social media using convergent cross mapping. In *Intelligence and Security Informatics Conference (JISIC), 2014 IEEE Joint*, pages 260–263. IEEE, 2014.
- Dieter Lüthi, Martine Le Floch, Bernhard Bereiter, Thomas Blunier, Jean-Marc Barnola, Urs Siegenthaler, Dominique Raynaud, Jean Jouzel, Hubertus Fischer, Kenji Kawamura, et al. High-resolution carbon dioxide concentration record 650,000–800,000 years before present. *Nature*, 453(7193):379, 2008.
- Jean Lynch-Stieglitz, Jess F Adkins, William B Curry, Trond Dokken, Ian R Hall, Juan Carlos Herguera, Joël J-M Hirschi, Elena V Ivanova, Catherine Kissel, Olivier Marchal, et al. Atlantic Meridional Overturning Circulation during the Last Glacial Maximum. *Science*, 316:66–69, 2007.
- Ben D MacArthur. Truth and beauty in physics and biology. *Nature Physics*, pages 1–3, 2021.
- Robert C MacCallum, Keith F Widaman, Shaobo Zhang, and Sehee Hong. Sample size in factor analysis. *Psychol. Methods*, 4(1):84–99, 1999.
- Fernando T Maestre, Ragan M Callaway, Fernando Valladares, and Christopher J Lortie. Refining the stress-gradient hypothesis for competition and facilitation in plant communities. *J. Ecol.*, 97(2):199–205, 2009.

- Benoit B Mandelbrot. *The fractal geometry of nature*, volume 2. WH freeman New York, 1982.
- Marc Mangel and Phillip S Levin. Regime, phase and paradigm shifts: making community ecology the basic science for fisheries. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1453):95–105, 2005.
- Brad Manor, Madalena D Costa, Kun Hu, Elizabeth Newton, Olga Starobinets, Hyun Gu Kang, CK Peng, Vera Novak, and Lewis A Lipsitz. Physiological complexity and system adaptability: evidence from postural control dynamics of older adults. *Journal of Applied Physiology*, 109(6):1786–1791, 2010.
- Thomas M Marchitto, Scott J Lehman, Joseph D Ortiz, Jacqueline Flückiger, and Alexander van Geen. Marine radiocarbon evidence for the mechanism of deglacial atmospheric CO₂ rise. *Science*, 316(5830):1456–1459, 2007.
- John H Martin. Glacial-interglacial CO₂ change: The iron hypothesis. *Paleoceanography*, 5:1–13, 1990.
- Katsumi Matsumoto and Yusuke Yokoyama. Atmospheric $\Delta^{14}\text{C}$ reduction in simulations of Atlantic overturning circulation shutdown. *Global Biogeochemical Cycles*, 27:296–304, 2013.
- Mark P Mattson. Superior pattern processing is the essence of the evolved human brain. *Frontiers in neuroscience*, 8:265, 2014.
- Robert M May. Biological populations with nonoverlapping generations: stable points, stable cycles, and chaos. *Science*, 186(4164):645–647, 1974.
- Robert M May. Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature*, 269(5628):471–477, 1977.
- Kevin McCann. Protecting biostructure. *Nature*, 446(7131):29, 2007.
- Kevin McCann, Alan Hastings, and Gary R Huxel. Weak trophic interactions and the balance of nature. *Nature*, 395(6704):794–798, 1998.
- Kevin Shear McCann. The diversity–stability debate. *Nature*, 405(6783):228–233, 2000.
- David McGee, Wallace S Broecker, and Gisela Winckler. Gustiness: The driver of glacial dustiness? *Quaternary Science Reviews*, 29:687–705, 2010.
- Larry McNicol, David A Story, Kate Leslie, Paul S Myles, Michael Fink, Andrew C Shelton, Ornella Clavisi, and Stephanie J Poustie. Postoperative complications and mortality in older patients having non-cardiac surgery at three melbourne teaching hospitals. *Medical journal of Australia*, 186(9):447–452, 2007.
- Anna Nele Meckler, Daniel M Sigman, Kelly A Gibson, Roger François, Alfredo Martinez-Garcia, Samuel L Jaccard, Ursula Röhl, Larry C Peterson, Ralf Tiedemann, and Gerald H Haug. Deglacial pulses of deep-ocean silicate into the subtropical North Atlantic

- Ocean. *Nature*, 495:495–498, 2013.
- Jane Memmott, Nickolas M Waser, and Mary V Price. Tolerance of pollination networks to species extinctions. *Proc. R. Soc. London, Ser. B*, 271(1557):2605–2611, 2004.
- Jane Memmott, Paul G Craze, Nickolas M Waser, and Mary V Price. Global warming and the disruption of plant–pollinator interactions. *Ecol. Lett.*, 10(8):710–717, 2007.
- L Menviel, Matthew Heathcote England, KJ Meissner, Anne Mouchet, and Jimin Yu. Atlantic-Pacific seesaw and its role in outgassing CO₂ during Heinrich events. *Paleoceanography*, 29:58–70, 2014.
- L. Menviel, P. Spence, J. Yu, M. A. Chamberlain, R. J. Matear, K. J. Meissner, and M. H. England. Southern Hemisphere westerlies as a driver of the early deglacial atmospheric CO₂ rise. *Nature Communications*, 9(1):2503, 2018.
- Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, et al. Quantitative analysis of culture using millions of digitized books. *science*, 331(6014):176–182, 2011.
- Takahiro Mieda, Kanako Taku, and Atsushi Oshio. Dichotomous thinking and cognitive ability. *Personality and Individual Differences*, 169:110008, 2021.
- Millenium Ecosystem Assessment. *Ecosystems and human well-being*, volume 5. Island Press, Washington, DC, 2005.
- Fredrik Moberg and Carl Folke. Ecological goods and services of coral reef ecosystems. *Ecol. Econ.*, 29(2):215–233, 1999.
- David A Moeller. Facilitative interactions among plants via shared pollinators. *Ecology*, 85(12):3289–3301, 2004.
- Dan Mønster, Riccardo Fusaroli, Kristian Tylén, Andreas Roepstorff, and Jacob F Sherson. Causal inference from noisy time-series data—testing the convergent cross-mapping algorithm in the presence of noise and external influence. *Future Generation Computer Systems*, 73:52–62, 2017.
- José M Montoya, Stuart L Pimm, and Ricard V Solé. Ecological networks and their fragility. *Nature*, 442(7100):259–264, 2006.
- Carolina L Morales and Anna Traveset. Interspecific pollen transfer: magnitude, prevalence and consequences for plant fitness. *Crit. Rev. Plant Sci.*, 27(4):221–238, 2008.
- Irving O Morales, Emmanuel Landa, Carlos Calderon Angeles, Juan C Toledo, Ana Leonor Rivera, Joel Mendoza Temis, and Alejandro Frank. Behavior of early warnings near the critical temperature in the two-dimensional ising model. *PloS one*, 10(6):e0130751, 2015.
- Carey K Morewedge and Daniel Kahneman. Associative processes in intuitive judgment.

- Trends in cognitive sciences*, 14(10):435–440, 2010.
- IAN Morris and K Farrell. Photosynthetic rates, gross patterns of carbon dioxide assimilation and activities of ribulose diphosphate carboxylase in marine algae grown at different temperatures. *Physiologia Plantarum*, 25(3):372–377, 1971.
- Adele K Morrison and Andrew McC. Hogg. On the relationship between Southern Ocean Overturning and ACC transport. *Journal of Physical Oceanography*, 43:140–148, 2013.
- Paul E Nachtigall and Alexander Ya Supin. A false killer whale reduces its hearing sensitivity when a loud sound is preceded by a warning. *Journal of Experimental Biology*, 216(16):3062–3070, 2013.
- Berend Nannes, Rick Quax, Hiroshi Ashikaga, Méléze Hocini, Remi Dubois, Olivier Bernus, and Michel Haïssaguerre. Early signs of critical slowing down in heart surface electrograms of ventricular fibrillation victims. In *International Conference on Computational Science*, pages 334–347. Springer, 2020.
- Egbert H van Nes. Grind for matlab, 2017. URL <http://www.sparcs-center.org/grind>.
- Egbert H van Nes and Marten Scheffer. Large species shifts triggered by small forces. *The American Naturalist*, 164(2):255–266, 2004.
- Egbert H van Nes and Marten Scheffer. Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *Am. Nat.*, 169(6):738–747, 2007.
- Michael G Neubert and Hal Caswell. Alternatives to resilience for measuring the responses of ecological systems to perturbations. *Ecology*, 78(3):653–665, 1997.
- Anje-Margriet Neutel and Michael AS Thorne. Interaction strengths in balanced carbon cycles and the absence of a relation between ecosystem complexity and stability. *Ecol. Lett.*, 17(6):651–661, 2014.
- Anje-Margriet Neutel, Johan AP Heesterbeek, and Peter C de Ruiter. Stability in real food webs: weak links in long loops. *Science*, 296(5570):1120–1123, 2002.
- Barry Newell. Simple models, powerful ideas: Towards effective integrative practice. *Global Environmental Change*, 22(3):776–783, 2012.
- Finn Årup Nielsen. A new anew: Evaluation of a word list for sentiment analysis in microblogs. *arXiv preprint arXiv:1103.2903*, 2011.
- Mark Novak, J Timothy Wootton, Daniel F Doak, Mark Emmerson, James A Estes, and M Timothy Tinker. Predicting community responses to perturbations in the face of imperfect knowledge and network complexity. *Ecology*, 92(4):836–846, 2011.
- Magnus Nyström, Albert V Norström, Thorsten Blenckner, Maricela de la Torre-Castro, Johan S Eklöf, Carl Folke, Henrik Österblom, Robert S Steneck, Matilda Thyresson, and Max Troell. Confronting feedbacks of degraded marine ecosystems. *Ecosystems*,

- 15(5):695–710, 2012.
- Keith Oatley and Jennifer M Jenkins. Human emotions: Function and dysfunction. *Annual review of psychology*, 43(1):55–85, 1992.
- Atsushi Obata. Climate-carbon cycle model response to freshwater discharge into the North Atlantic. *Journal of Climate*, 19:5479–5499, 2007.
- Keisuke Okamura. Interdisciplinarity revisited: evidence for research impact and dynamism. *Palgrave Communications*, 5(1):1–9, 2019.
- Toshinori Okuyama and J Nathaniel Holland. Network structural properties mediate the stability of mutualistic communities. *Ecol. Lett.*, 11(3):208–216, 2008.
- Marcel GM Olde Rikkert, Vasilis Dakos, Timothy G Buchman, Rob de Boer, Leon Glass, Angélique OJ Cramer, Simon Levin, Egbert Van Nes, George Sugihara, Michel D Ferrari, et al. Slowing down of recovery as generic risk marker for acute severity transitions in chronic diseases. *Critical care medicine*, 44(3):601–606, 2016.
- Merlijn Olthof, Fred Hasselman, and Anna Lichtwarck-Aschoff. Complexity in psychological self-ratings: Implications for research and practice. 2020a.
- Merlijn Olthof, Fred Hasselman, Guido Strunk, Marieke van Rooij, Benjamin Aas, Marieke A Helmich, Günter Schiepek, and Anna Lichtwarck-Aschoff. Critical fluctuations as an early-warning signal for sudden gains and losses in patients receiving psychotherapy for mood disorders. *Clinical Psychological Science*, 8(1):25–35, 2020b.
- Anne Willem Omta, George AK van Voorn, Rosalind EM Rickaby, and Michael J Follows. On the potential role of marine calcifiers in glacial-interglacial dynamics. *Global Biogeochemical Cycles*, 27:692–704, 2013.
- Anne Willem Omta, Bob Kooi, George AK van Voorn, Rosalind EM Rickaby, and Michael J Follows. Inherent characteristics of sawtooth cycles can explain different glacial periodicities. *Climate Dynamics*, 46:557–569, 2016.
- Anne Willem Omta, Raffaele Ferrari, and David McGee. An analytical framework for the steady state impact of carbonate compensation on atmospheric CO_2 . *Global Biogeochemical Cycles*, 32(4):720–735, 2018.
- Naomi Oreskes, Kristin Shrader-Frechette, and Kenneth Belitz. Verification, validation, and confirmation of numerical models in the earth sciences. *Science*, 263(5147):641–646, 1994.
- Atsushi Oshio, Takahiro Mieda, and Kanako Taku. Younger people, and stronger effects of all-or-nothing thoughts on aggression: Moderating effects of age on the relationships between dichotomous thinking and aggression. *Cogent Psychology*, 3(1):1244874, 2016.
- Julio M Ottino. Complex systems. *AIChE Journal*, 49(2):292–299, 2003.
- Robert T Paine, Mia J Tegner, and Edward A Johnson. Compounded perturbations yield

- ecological surprises. *Ecosystems*, 1(6):535–545, 1998.
- Gianfranco Parati, Juan E Ochoa, Carolina Lombardi, and Grzegorz Bilo. Assessment and management of blood-pressure variability. *Nature Reviews Cardiology*, 10(3):143, 2013.
- Adina Paytan and Elizabeth M Griffith. Marine barite: Recorder of variations in ocean export productivity. *Deep-Sea Research II*, 54:687–705, 2007.
- Eitan Adam Pechenick, Christopher M Danforth, and Peter Sheridan Dodds. Characterizing the google books corpus: Strong limits to inferences of socio-cultural and linguistic evolution. *PloS one*, 10(10):e0137041, 2015.
- James W Pennebaker, Martha E Francis, and Roger J Booth. Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates*, 71(2001):2001, 2001.
- Jacob Perkins. *Python 3 text processing with NLTK 3 cookbook*. Packt Publishing Ltd, 2014.
- Owen L Petchey, Mikael Pontarp, Thomas M Massie, Sonia Kéfi, Arpat Ozgul, Maja Weilenmann, Gian Marco Palamara, Florian Altermatt, Blake Matthews, Jonathan M Levine, et al. The ecological forecast horizon, and examples of its uses and determinants. *Ecol. Lett.*, 18(7):597–611, 2015.
- Arthur C Petersen. Philosophy of climate science. *Bulletin of the American Meteorological Society*, 81(2):265–272, 2000.
- Carlye D Peterson and Loraine E Lisiecki. Deglacial carbon cycle changes observed in a compilation of 127 benthic $\delta^{13}\text{C}$ time series (20–6 ka). *Climate of the Past*, 14:1229–1252, 2018.
- Michael Pettit. Historical time in the age of big data: Cultural psychology, historical change, and the google books ngram viewer. *History of psychology*, 19(2):141, 2016.
- Jonathan D Phillips. Sources of nonlinearity and complexity in geomorphic systems. *Progress in physical geography*, 27(1):1–23, 2003.
- Steven M Pincus. Approximate entropy as a measure of system complexity. *Proceedings of the National Academy of Sciences*, 88(6):2297–2301, 1991.
- Steven M Pincus and Richard R Viscarello. Approximate entropy: a regularity measure for fetal heart rate analysis. *Obstet Gynecol*, 79(2):249–255, 1992.
- Lukáš Poláček. How to shuffle songs?, 2014. Available at <https://engineering.atspotify.com/2014/02/28/how-to-shuffle-songs/>.
- John Porter, Peter Arzberger, Hans-Werner Braun, Pablo Bryant, Stuart Gage, Todd Hansen, Paul Hanson, Chau-Chin Lin, Fang-Pang Lin, Timothy Kratz, et al. Wireless sensor networks for ecology. *AIBS Bulletin*, 55(7):561–572, 2005.
- Carol Pride, Robert Thunell, Daniel Sigman, Lloyd Keigwin, Mark Altabet, and Eric

- Tappa. Nitrogen isotopic variations in the Gulf of California since the last glaciation: Response to global climate change. *Paleoceanography*, 14:397–409, 1999.
- Ilya Romanovič Prigogine, Isabelle Stengers, and Martinus Petrus Maria Franssen. *Orde uit chaos: de nieuwe dialoog tussen de mens en de natuur*. Bakker, 1993.
- Charles J Puccia and Richard Levins. *Qualitative modeling of complex systemsan introduction to loop analysis and time averaging*. Harvard University Press, Cambridge, MA, 1985.
- Drew Purves, Jörn PW Scharlemann, Mike Harfoot, Tim Newbold, Derek P Tittensor, Jon Hutton, and Stephen Emmott. Ecosystems: time to model all life on earth. *Nature*, 493(7432):295, 2013.
- Huijie Qiao, Jorge Soberon, and Andrew Townsend Peterson. No silver bullets in correlative ecological niche modelling: insights from testing among many potential algorithms for niche estimation. *Methods in Ecology and Evolution*, 6(10):1126–1136, 2015.
- Rick Quax, Andrea Apolloni, and Peter MA Sloot. The diminishing role of hubs in dynamical processes on complex networks. *Journal of The Royal Society Interface*, 10(88):20130568, 2013a.
- Rick Quax, Drona Kandhai, and Peter MA Sloot. Information dissipation as an early-warning signal for the lehman brothers collapse in financial time series. *Scientific reports*, 3, 2013b.
- Markus Raitzsch, Ed C Hathorne, Henning Kuhnert, Jeroen Groeneveld, and Torsten Bickert. Modern and late Pleistocene B/Ca ratios of the benthic foraminifer *Planulina wuellerstorfi* determined with laser ablation ICP-MS. *Geology*, 39:1039–1042, 2011.
- Douglas S Ramsay and Stephen C Woods. Clarifying the roles of homeostasis and allostasis in physiological regulation. *Psychological review*, 121(2):225, 2014.
- Suresh IS Rattan. Homeostasis, homeodynamics, and aging. *Encyclopedia of Gerontology. UK, Elsevier Inc*, pages 696–699, 2007.
- Suresh IS Rattan. Biological health and homeodynamic space. In *Explaining Health Across the Sciences*, pages 43–51. Springer, 2020.
- Shahra Razavi, Christina Behrendt, Mira Bierbaum, Ian Orton, and Lou Tessier. Rein-vigorating the social contract and strengthening social cohesion: Social protection responses to covid-19. *International Social Security Review*, 73(3):55–80, 2020.
- Jerrald L Rector, Sanne MW Gijzel, Ingrid A van de Leemput, Fokke B van Meulen, Marcel GM Olde Rikkert, and René JF Melis. Dynamical indicators of resilience from physiological time series in geriatric inpatients: Lessons learned. *Experimental Gerontology*, 149:111341, 2021.
- Jaime Redondo, Isabel Fraga, Isabel Padrón, and Montserrat Comesaña. The spanish adaptation of anew (affective norms for english words). *Behavior research methods*, 39

- (3):600–605, 2007.
- Enrico L Rezende, Jessica E Lavabre, Paulo R Guimarães, Pedro Jordano, and Jordi Bascompte. Non-random coextinctions in phylogenetically structured mutualistic networks. *Nature*, 448(7156):925–928, 2007.
- Joshua S Richman and J Randall Moorman. Physiological time-series analysis using approximate entropy and sample entropy. *American Journal of Physiology-Heart and Circulatory Physiology*, 278(6):H2039–H2049, 2000.
- Dean Rickles, Penelope Hawe, and Alan Shiell. A simple guide to chaos and complexity. *Journal of Epidemiology & Community Health*, 61(11):933–937, 2007.
- Andy J Ridgwell. Dust in the earth system: the biogeochemical linking of land, air and sea. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 360(1801):2905–2924, 2002.
- Max Rietkerk and Johan Van de Koppel. Alternate stable states and threshold effects in semi-arid grazing systems. *Oikos*, 79(1):69–76, 1997.
- Max Rietkerk, Stefan C Dekker, Peter C de Ruiter, and Johan van de Koppel. Self-organized patchiness and catastrophic shifts in ecosystems. *Science*, 305(5692):1926–1929, 2004.
- Johan Rockström, Will L Steffen, Kevin Noone, Åsa Persson, F Stuart Chapin III, Eric Lambin, Timothy M Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber, et al. Planetary boundaries: exploring the safe operating space for humanity. *Ecol. Soc.*, 14(2):32, 2009.
- Pablo Rodríguez-Sánchez, Egbert H Van Nes, and Marten Scheffer. Climbing escher’s stairs: A way to approximate stability landscapes in multidimensional systems. *PLOS Computational Biology*, 16(4):e1007788, 2020.
- Rudolf P Rohr, Serguei Saavedra, and Jordi Bascompte. On the structural stability of mutualistic systems. *Science*, 345(6195):1253497, 2014.
- Oscar E Romero, Jung-Hyun Kim, and Barbara Donner. Submillennial-to-millennial variability of diatom production off Mauretania, NW Africa, during the last glacial cycle. *Paleoceanography*, 23:PA3218, 2008.
- Daniel Rosenfeld, Yoram J Kaufman, and Ilan Koren. Switching cloud cover and dynamical regimes from open to closed Benard cells in response to aerosols suppressing precipitation. *Atmospheric Chemistry and Physics*, 6:2503–2511, 2006.
- Adam Rowe. Traditional publishers are selling way more non-fiction than fiction, 2018. Available at <https://www.forbes.com/sites/adamrowe1/2018/08/30/traditional-publishers-are-selling-way-more-non-fiction-than-fiction/?sh=49900f5856d0>.
- Colin F Royse, Stanton Newman, Frances Chung, Jan Stygall, Rachel E McKay, Joachim

- Boldt, Frederique S Servin, Ignacio Hurtado, Raafat Hannallah, Buwei Yu, et al. Development and feasibility of a scale to assess postoperative recovery: the post-operative quality recovery scale. *The Journal of the American Society of Anesthesiologists*, 113 (4):892–905, 2010.
- Denis Rubin, Tomer Fekete, and Lilianne R Mujica-Parodi. Optimizing complexity measures for fmri data: algorithm, artifact, and sensitivity. *PLoS One*, 8(5):e63448, 2013.
- William F Ruddiman. *Earth’s Climate: Past and future*. Macmillan, 2001.
- Fayez F Safadi, Mary F Barbe, Samir M Abdelmagid, Mario C Rico, Rulla A Aswad, Judith Litvin, and Steven N Popoff. Bone structure, development and bone biology. In *Bone pathology*, pages 1–50. Springer, 2009.
- Hassan Saif, Miriam Fernandez, Yulan He, and Harith Alani. On stopwords, filtering and data sparsity for sentiment analysis of twitter. 2014.
- Paul A Samuelson. Some implications of” linearity.”. *The Review of Economic Studies*, 15(2):88–90, 1947.
- Jorge L Sarmiento and Nicolas Gruber. *Ocean biogeochemical dynamics*. Princeton University Press, Princeton, 2006.
- Jorge L Sarmiento and JR Toggweiler. A new model for the role of the oceans in determining atmospheric pCO₂. *Nature*, 308:621–624, 1984.
- Marten Scheffer. *Ecology of shallow lakes*, volume 22. Springer Science & Business Media, 1997.
- Marten Scheffer. *Critical transitions in nature and society*, volume 16. Princeton University Press, 2009.
- Marten Scheffer and Erik Jeppesen. Alternative stable states. In *The structuring role of submerged macrophytes in lakes*, pages 397–406. Springer, 1998.
- Marten Scheffer, SH Hosper, ML Meijer, Brian Moss, and Erik Jeppesen. Alternative equilibria in shallow lakes. *Trends in ecology & evolution*, 8(8):275–279, 1993.
- Marten Scheffer, William Brock, and Frances Westley. Socioeconomic mechanisms preventing optimum use of ecosystem services: an interdisciplinary theoretical analysis. *Ecosystems*, 3(5):451–471, 2000.
- Marten Scheffer, Steve Carpenter, Jonathan A Foley, Carl Folke, and Brian Walker. Catastrophic shifts in ecosystems. *Nature*, 413(6856):591–596, 2001.
- Marten Scheffer, Jordi Bascompte, William A Brock, Victor Brovkin, Stephen R Carpenter, Vasilis Dakos, Hermann Held, Egbert H Van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.
- Marten Scheffer, Stephen R Carpenter, Timothy M Lenton, Jordi Bascompte, William Brock, Vasilis Dakos, Johan Van de Koppel, Ingrid A Van de Leemput, Simon A Levin,

- Egbert H Van Nes, et al. Anticipating critical transitions. *science*, 338(6105):344–348, 2012.
- Marten Scheffer, J Elizabeth Bolhuis, Denny Borsboom, Timothy G Buchman, Sanne MW Gijzel, Dave Goulson, Jan E Kammenga, Bas Kemp, Ingrid A van de Leemput, Simon Levin, et al. Quantifying resilience of humans and other animals. *Proceedings of the National Academy of Sciences*, 115(47):11883–11890, 2018.
- Michael F Scheier, Karen A Matthews, Jane F Owens, George J Magovern, R Craig Lefebvre, R Anne Abbott, and Charles S Carver. Dispositional optimism and recovery from coronary artery bypass surgery: the beneficial effects on physical and psychological well-being. *Journal of personality and social psychology*, 57(6):1024, 1989.
- Bernhard Schink. Synergistic interactions in the microbial world. *Antonie van Leeuwenhoek*, 81(1-4):257–261, 2002.
- David S Schmidtke, Tobias Schröder, Arthur M Jacobs, and Markus Conrad. Angst: Affective norms for german sentiment terms, derived from the affective norms for english words. *Behavior research methods*, 46(4):1108–1118, 2014.
- Jochen Schmitt, Robert Schneider, Joachim Elsig, Daiana Leuenberger, Anna Laurantou, Jérôme Chappellaz, Peter Köhler, Fortunat Joos, Thomas F. Stocker, Markus Leuenberger, and Hubertus Fischer. Carbon isotope constraints of the deglacial CO₂ rise from ice cores. *Science*, 336:711–714, 2012.
- A Schmittner and DC Lund. Early deglacial atlantic overturning decline and its role in atmospheric CO₂ rise inferred from carbon isotopes ($\delta^{13}\text{C}$). *Climate of the Past*, 11(2):135–152, 2015.
- Andreas Schmittner, Edward J Brook, and Jinho Ahn. Impact of the ocean’s overturning circulation on atmospheric CO₂. In A. Schmittner, J. Chiang, and S. Hemming, editors, *Ocean circulation: Mechanisms and impacts*, pages 315–334, Washington, 2007. AGU.
- Norbert Schwarz. Self-reports: how the questions shape the answers. *American psychologist*, 54(2):93, 1999.
- S Shahid Shaukat, Toqeer Ahmed Rao, and Moazzam A Khan. Impact of sample size on principal component analysis ordination of an environmental data set: effects on eigenstructure. *Ekológia (Bratislava)*, 35(2):173–190, 2016.
- RJ Shephard. The prediction of maximum oxygen intake from post-exercise pulse readings. *Internationale Zeitschrift für angewandte Physiologie einschließlich Arbeitsphysiologie*, 24(1):31–38, 1967.
- Osamu Shimomura, Frank H Johnson, and Yo Saiga. Extraction, purification and properties of aequorin, a bioluminescent protein from the luminous hydromedusan, aequorea. *Journal of cellular and comparative physiology*, 59(3):223–239, 1962.
- Sang-Ik Shin, Zhensheng Liu, B. Otto-Bliesner, Esther Brady, John Kutzbach, and Sandy

- Harrison. A simulation of the last glacial maximum using the NCAR-CCSM. *Climate Dynamics*, 20:127–151, 2003.
- Lauren G Shoemaker, Lauren L Sullivan, Ian Donohue, Juliano S Cabral, Ryan J Williams, Margaret M Mayfield, Jonathan M Chase, Chengjin Chu, W Stanley Harpole, Andreas Huth, et al. Integrating the underlying structure of stochasticity into community ecology. *Ecology*, 101(2):e02922, 2020.
- U Siegenthaler and T Wenk. Rapid atmospheric CO₂ variations and ocean circulation. *Nature*, 308:624–626, 1984.
- Daniel M Sigman and Ed A Boyle. Glacial-interglacial variations in atmospheric carbon dioxide. *Nature*, 407:859–869, 2000.
- Daniel M Sigman, Mathis P Hain, and Gerald H Haug. The polar ocean and glacial cycles in atmospheric co₂ concentration. *Nature*, 466(7302):47, 2010.
- Herbert A Simon. The architecture of complexity. In *Facets of systems science*, pages 457–476. Springer, 1991.
- El Houssine Snoussi. Necessary conditions for multistationarity and stable periodicity. *J. Biol. Syst.*, 6(01):3–9, 1998.
- Ricard V Solé and M^a Montoya. Complexity and fragility in ecological networks. *Proc. Royal Soc. B*, 268(1480):2039–2045, 2001.
- Ricard V Solé, Susanna Manrubia Cuevas, Bartolo Luque, Jordi Delgado, and Jordi Bascompte. Phase transitions and complex systems: Simple, nonlinear models capture complex systems at the edge of chaos. 1996.
- Brian G Southwell, Emily A Thorson, and Laura Sheble. *Misinformation and mass audiences*. University of Texas Press, 2018.
- Arie Staal, Stefan C Dekker, Marina Hirota, and Egbert H van Nes. Synergistic effects of drought and deforestation on the resilience of the south-eastern amazon rainforest. *Ecological Complexity*, 22:65–75, 2015.
- Will Steffen, Katherine Richardson, Johan Rockström, Sarah E Cornell, Ingo Fetzer, Elena M Bennett, Reinette Biggs, Stephen R Carpenter, Wim De Vries, Cynthia A De Wit, et al. Planetary boundaries: Guiding human development on a changing planet. *Science*, 347(6223):1259855, 2015.
- Jessica S Stella, Morgan S Pratchett, Pat A Hutchings, and Geoffrey P Jones. Coral-associated invertebrates: diversity, ecological importance and vulnerability to disturbance. *Oceanogr. Mar. Biol.*, 49:43–104, 2011.
- Philip A Stephens, William J Sutherland, and Robert P Freckleton. What is the allee effect? *Oikos*, 87(1):185–190, 1999.
- Peter Sterling. Allostasis: a model of predictive regulation. *Physiology & behavior*, 106

- (1):5–15, 2012.
- Sergey Stolyar, Steve Van Dien, Kristina Linnea Hillesland, Nicolas Pinel, Thomas J Lie, John A Leigh, and David A Stahl. Metabolic modeling of a mutualistic microbial community. *Mol. Syst. Biol.*, 3(1):92, 2007.
- Jason Ananda Josephson Storm. *The myth of disenchantment: Magic, modernity, and the birth of the human sciences*. University of Chicago Press, 2017.
- Steven H Strogatz. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Westview press, 2014.
- Anja S Studer, Daniel M Sigman, Alfredo Martínez-García, Verena Benz, Gisela Winckler, Gerhard Kuhn, Oliver Esper, Frank Lamy, Samuel L Jaccard, Lukas Wacker, Sergey Oleynik, Rainer Gersonde, and Gerald H Haug. Antarctic Zone nutrient conditions during the last two glacial cycles. *Paleoceanography*, 30:845–862, 2015.
- George Sugihara. Nonlinear forecasting for the classification of natural time series. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 348(1688):477–495, 1994.
- George Sugihara, Robert May, Hao Ye, Chih-hao Hsieh, Ethan Deyle, Michael Fogarty, and Stephan Munch. Detecting causality in complex ecosystems. *science*, 338(6106):496–500, 2012.
- William J Sutherland. Predicting the ecological consequences of environmental change: a review of the methods. *J. Appl. Ecol.*, 43(4):599–616, 2006.
- KB Suttle, Meredith A Thomsen, and Mary E Power. Species interactions reverse grassland responses to changing climate. *Science*, 315(5812):640–642, 2007.
- Samir Suweis and Paolo D’Odorico. Early warning signs in social-ecological networks. *PloS one*, 9(7):e101851, 2014.
- Paul Switzer and Andrew A Green. Min/max autocorrelation factors for multivariate spatial imagery. *Computer science and statistics*, pages 13–16, 1984.
- Floris Takens. Detecting strange attractors in turbulence. In *Dynamical systems and turbulence, Warwick 1980*, pages 366–381. Springer, 1981.
- Esther Tanck, Jasper Homminga, G Harry van Lenthe, and Rik Huiskes. Increase in bone volume fraction precedes architectural adaptation in growing bone. *Bone*, 28(6):650–654, 2001.
- Si Tang and Stefano Allesina. Reactivity and stability of large ecosystems. *Frontiers in Ecology and Evolution*, 2:21, 2014.
- Ina Tegen. Modeling the mineral dust aerosol cycle in the climate system. *Quaternary Science Reviews*, 22(18-19):1821–1834, 2003.
- René Thom. *Structural stability and morphogenesis*. CRC press, 2018.

- René Thomas. On the relation between the logical structure of systems and their ability to generate multiple steady states or sustained oscillations. In *Numerical methods in the study of critical phenomena*, pages 180–193. Springer-Verlag, New York, NY, 1981.
- Wilfried Thuiller, Miguel B Araújo, and Sandra Lavorel. Generalized models vs. classification tree analysis: predicting spatial distributions of plant species at different scales. *Journal of Vegetation Science*, 14(5):669–680, 2003.
- Robert C Thunell and Alisa B Kepple. Glacial-Holocene $\delta^{15}\text{N}$ record from the Gulf of Tehuantepec, Mexico: Implications for denitrification in the Eastern Equatorial Pacific and changes in atmospheric N_2O . *Global Biogeochemical Cycles*, 18:GB1001, 2004.
- Ulf Toelch and Dirk Ostwald. Digital open science—teaching digital tools for reproducible and transparent research. *PLoS biology*, 16(7):e2006022, 2018.
- J. R. Toggweiler, Joellen L Russell, and SR Carson. Midlatitude westerlies, atmospheric CO_2 , and climate change during the ice ages. *Paleoceanography*, 21:PA2005, 2006.
- JR Toggweiler. Variation of atmospheric CO_2 by ventilation of the ocean’s deepest water. *Paleoceanography*, 14:571–588, 1999.
- Martin H Trauth, Asfawossen Asrat, Walter Duesing, Verena Foerster, K Hauke Kraemer, Norbert Marwan, Mark A Maslin, and Frank Schaebitz. Classifying past climate change in the chew bahir basin, southern ethiopia, using recurrence quantification analysis. *Climate Dynamics*, 53(5):2557–2572, 2019.
- Tobias Tschumi, Fortunat Joos, M Gehlen, and C Heinze. Deep ocean ventilation, carbon isotopes, marine sedimentation and the deglacial CO_2 rise. *Climate of the Past*, 6(5):1895–1958, 2010.
- Cristina Tur, Agustín Sáez, Anna Traveset, and Marcelo Adrian Aizen. Evaluating the effects of pollinator-mediated interactions using pollen transfer networks: evidence of widespread facilitation in south andean plant communities. *Ecol. Lett.*, 19(5):576–586, 2016.
- Jason M Tylianakis, Raphael K Didham, Jordi Bascompte, and David A Wardle. Global change and species interactions in terrestrial ecosystems. *Ecol. Lett.*, 11(12):1351–1363, 2008.
- John J Tyson, Katherine C Chen, and Bela Novak. Sniffers, buzzers, toggles and blinkers: dynamics of regulatory and signaling pathways in the cell. *Curr. Opin. Cell Biol.*, 15(2):221–231, 2003.
- Jacob Usinowicz and Jonathan M Levine. Species persistence under climate change: a geographical scale coexistence problem. *Ecol. Lett.*, 21(11):1589–1603, 2018.
- Jim Van Belzen, Johan Van De Koppel, Matthew L Kirwan, Daphne Van Der Wal, Peter MJ Herman, Vasilis Dakos, Sonia Kéfi, Marten Scheffer, Glenn R Guntenspergen, and Tjeerd J Bouma. Vegetation recovery in tidal marshes reveals critical slowing down

- under increased inundation. *Nature Communications*, 8:15811, 2017.
- Jan D Van Der Laan and Pauline Hogeweg. Predator—prey coevolution: interactions across different timescales. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 259(1354):35–42, 1995.
- Egbert H Van Nes, Marten Scheffer, Victor Brovkin, Timothy M Lenton, Hao Ye, Ethan Deyle, and George Sugihara. Causal feedbacks in climate change. *Nature Climate Change*, 5(5):445–448, 2015a.
- Egbert H Van Nes, Marten Scheffer, Victor Brovkin, Timothy M Lenton, Hao Ye, Ethan Deyle, and George Sugihara. Causal feedbacks in climate change. *Nature Climate Change*, 5(5):445, 2015b.
- Guy C Van Orden, Heidi Kloos, and Sebastian Wallot. Living in the pink: Intentionality, wellbeing, and complexity. In *Philosophy of complex systems*, pages 629–672. Elsevier, 2011.
- Jan-Willem Van Prooijen and Karen M Douglas. Conspiracy theories as part of history: The role of societal crisis situations. *Memory studies*, 10(3):323–333, 2017.
- Jan-Willem Van Prooijen, André PM Krouwel, and Thomas V Pollet. Political extremism predicts belief in conspiracy theories. *Social Psychological and Personality Science*, 6(5):570–578, 2015.
- Eus JW Van Someren and Rixt F Riemersma-Van Der Lek. Live to the rhythm, slave to the rhythm. *Sleep medicine reviews*, 11(6):465–484, 2007.
- R Varadhan, CL Seplaki, QL Xue, K Bandeen-Roche, and LP Fried. Stimulus-response paradigm for characterizing the loss of resilience in homeostatic regulation associated with frailty. *Mechanisms of ageing and development*, 129(11):666–670, 2008.
- Wayne F Velicer and Joseph L Fava. Affects of variable and subject sampling on factor pattern recovery. *Psychol. Methods*, 3(2):231–251, 1998.
- Annelies J Veraart, Elisabeth J Faassen, Vasilis Dakos, Egbert H van Nes, Miquel Lüring, and Marten Scheffer. Recovery rates reflect distance to a tipping point in a living system. *Nature*, 481(7381):357–359, 2012.
- Ajay K Verma, Amanmeet Garg, Andrew Blaber, Reza Fazel-Rezai, and Kouhyar Tavakolian. Analysis of causal cardio-postural interaction under orthostatic stress using convergent cross mapping. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2319–2322. IEEE, 2016.
- Greg Versteeg. Non-parametric entropy estimation toolbox. <https://github.com/gregversteeg/NPEET>, 2014.
- Peter M Vitousek, Harold A Mooney, Jane Lubchenco, and Jerry M Melillo. Human domination of earth’s ecosystems. *Science*, 277(5325):494–499, 1997.

- Soroush Vosoughi, Deb Roy, and Sinan Aral. The spread of true and false news online. *Science*, 359(6380):1146–1151, 2018.
- Liew Yew Wai (2019). *Dynamics preceding onset and termination of atrial fibrillation and ventricular tachycardia*. PhD thesis, Monash University, 2019.
- Brian Walker, Crawford S Holling, Stephen R Carpenter, and Ann Kinzig. Resilience, adaptability and transformability in social–ecological systems. *Ecology and society*, 9(2), 2004.
- Raymond Walters. Present trends in nonfiction. *College English*, 15(3):141–147, 1953.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods*, 45(4):1191–1207, 2013.
- Andrew J Watson, Geoffrey K Vallis, and Maxim Nikurashin. Southern ocean buoyancy forcing of ocean ventilation and glacial atmospheric CO₂. *Nature Geoscience*, 8(11):861, 2015.
- Gabriel Weimann et al. *New terrorism and new media*, volume 2. Commons Lab of the Woodrow Wilson International Center for Scholars . . . , 2014.
- Els Weinans, J Jelle Lever, Sebastian Bathiany, Rick Quax, Jordi Bascompte, Egbert H Van Nes, Marten Scheffer, and Ingrid A Van De Leemput. Finding the direction of lowest resilience in multivariate complex systems. *Journal of the Royal Society Interface*, 16(159):20190629, 2019.
- Els Weinans, Anne Willem Omta, George AK van Voorn, and Egbert H van Nes. A potential feedback loop underlying glacial-interglacial cycles. *Climate Dynamics*, pages 1–13, 2021.
- Harrie Weinans, Rik Huiskes, and HJ Grootenboer. The behavior of adaptive bone-remodeling simulation models. *Journal of biomechanics*, 25(12):1425–1441, 1992.
- M Wichers. The dynamic nature of depression: a new micro-level perspective of mental disorder that meets current challenges. *Psychological medicine*, 44(7):1349–1360, 2014.
- Marieke Wichers, Peter C Groot, ESM Psychosystems, EWS Group, et al. Critical slowing down as a personalized early warning signal for depression. *Psychotherapy and psychosomatics*, 85(2):114–116, 2016.
- Kurt Wiesenfeld. Noisy precursors of nonlinear instabilities. *Journal of Statistical Physics*, 38(5-6):1071–1097, 1985.
- Grace M Wilkinson, Stephen R Carpenter, Jonathan J Cole, Michael L Pace, Ryan D Batt, Cal D Buelo, and Jason T Kurtzweil. Early warning signals precede cyanobacterial blooms in multiple whole-lake experiments. *Ecological Monographs*, 88(2):188–203, 2018.

- Joseph J Williams and Thomas L Griffiths. Why are people bad at detecting randomness? a statistical argument. *Journal of experimental psychology: learning, memory, and cognition*, 39(5):1473, 2013.
- Richard G Williams and Michael J Follows. *Ocean Dynamics and the Carbon Cycle*. Cambridge University Press, Cambridge, 2011.
- J Bastow Wilson and Andrew DQ Agnew. Positive-feedback switches in plant communities. In *Advances in Ecological Research*, volume 23, pages 263–336. Elsevier, Amsterdam, 1992.
- Shaun K Wilson, Nicholas AJ Graham, Morgan S Pratchett, Geoffrey P Jones, and Nicholas VC Polunin. Multiple disturbances and the global degradation of coral reefs: are reef fishes at risk or resilient? *Global Change Biol.*, 12(11):2220–2234, 2006.
- C Wissel. A universal law of the characteristic return time near thresholds. *Oecologia*, 65(1):101–107, 1984.
- H Arthur Woods and J Keaton Wilson. An information hypothesis for the evolution of homeostasis. *Trends in ecology & evolution*, 28(5):283–289, 2013.
- Chen Ning Yang. The spontaneous magnetization of a two-dimensional ising model. *Physical Review*, 85(5):808, 1952.
- Ibrar Yaqoob, Ibrahim Abaker Targio Hashem, Abdullah Gani, Salimah Mokhtar, Ejaz Ahmed, Nor Badrul Anuar, and Athanasios V Vasilakos. Big data: From beginning to future. *International Journal of Information Management*, 36(6):1231–1247, 2016.
- F Eugene Yates. Order and complexity in dynamical systems: homeodynamics as a generalized mechanics for biology. *Mathematical and computer modelling*, 19(6-8):49–74, 1994.
- Hao Ye, Richard J Beamish, Sarah M Glaser, Sue CH Grant, Chih-hao Hsieh, Laura J Richards, Jon T Schnute, and George Sugihara. Equation-free mechanistic ecosystem forecasting using empirical dynamic modeling. *Proceedings of the National Academy of Sciences*, 112(13):E1569–E1576, 2015a.
- Hao Ye, Ethan R Deyle, Luis J Gilarranz, and George Sugihara. Distinguishing time-delayed causal interactions using convergent cross mapping. *Scientific reports*, 5:14750, 2015b.
- Soon-Chang Yoon, Jae-Gwang Won, Ali H Omar, Sang-Woo Kim, and Byung-Ju Sohn. Estimation of the radiative forcing by key aerosol types in worldwide locations using a column model and AERONET data. *Atmospheric Environment*, 39:6620–6630, 2005.
- Jimin Yu, Wally S Broecker, Harry Elderfield, Zhangdong Jin, Jerry McManus, and Fei Zhang. Loss of carbon from the deep sea since the Last Glacial Maximum. *Science*, 330:1084–1087, 2010.
- Jimin Yu, Robert F Anderson, Zhangdong Jin, Laurie Menviel, Fei Zhang, Fredrick J

-
- Ryerson, and Eelco J Rohling. Deep South Atlantic carbonate chemistry and increased interocean deep water exchange during last deglaciation. *Quaternary Science Reviews*, 90:80–89, 2014.
- Richard E Zeebe and Dieter Wolf-Gladrow. *CO₂ in seawater: Equilibrium, kinetics, isotopes*. Elsevier, Amsterdam, 2001.
- Semir Zeki, John Paul Romaya, Dionigi MT Benincasa, and Michael F Atiyah. The experience of mathematical beauty and its neural correlates. *Frontiers in human neuroscience*, 8:68, 2014.
- Sarah Zhang. The pitfalls of using google ngram to study language, 2015. Available at <https://www.wired.com/2015/10/pitfalls-of-studying-language-with-google-ngram/>.
- Yi-Cheng Zhang. Complexity and 1/f noise. a phase space approach. *Journal de Physique I*, 1(7):971–977, 1991.
- Zeyu Zheng, Boris Podobnik, Ling Feng, and Baowen Li. Changes in cross-correlations as an indicator for systemic risk. *Scientific reports*, 2:888, 2012.
- Junhong Zhou, Daniel Habtemariam, Ikechukwu Iloputaife, Lewis A Lipsitz, and Brad Manor. The complexity of standing postural sway associates with future falls in community-dwelling older adults: the mobilize boston study. *Scientific reports*, 7(1): 1–8, 2017.

Summary

Complex systems are systems whose behaviour arises from the interaction between different elements. The dynamics of complex systems are highly unpredictable and are characterized by interactions on different scales and nonlinear responses. In **chapter 1** I point to how these properties complicate the analysis of complex systems. Additionally, I explain that the analysis of complex systems requires a range of approaches, such as simplistic models, realistic models, experiments, and time series analysis. The latter is the focus of this thesis.

Two popular tools to analyse complex systems are resilience indicators and complexity indicators. In **chapter 2** I address the use of these indicators, with applications to systems related to the human system, such as the heart, mood, bones and balance. I explore how resilience indicators can be used to infer a systems capacity to recover from small perturbations and can be used to signal upcoming ‘tipping points’ related to some diseases, such as atrial fibrillation or depression. Complexity indicators can be used to infer when a system loses responsiveness and can be used to infer a loss of complexity which is related to old age and some diseases such as congestive heart failure. I propose that the variables that constitute ‘the human system’ have different functions. Some variables aim for homeostasis, i.e. they seek an equilibrium. For these variables, their capacity to recover (their resilience) seems like a good statistic to quantify their functioning. Other variables aim for high responsiveness, i.e. they need to adapt quickly in order to keep the first variables within their healthy range. For these variables, their complexity seems like a good statistic to quantify their functioning. Lastly, I observe that some variables are balancing between these different functions, and therefore both indicators play an important role to assess a person’s health.

Currently, the most popular methods to infer a system’s resilience are lag-1 autocorrelation and the variance of a time series. As these quantification tools are based on single time series, it is not always clear how to adapt these tools to multivariate complex systems, since multivariate systems can show a different recovery pattern to different perturbations. In **Chapter 3** I propose a novel usage of a known statistical tool called Min/Max Autocorrelation Factors (MAF). This tool was developed as an alternative to PCA, but instead of finding the direction of highest variance (as with PCA), it finds the

direction of the highest autocorrelation. I propose that this direction of highest autocorrelation can be used to tell which perturbation in the system is most dangerous, in the sense that perturbations in this direction will lead to the slowest recovery of the system. Furthermore, if the system is subject to tipping points, this ‘dangerous direction’ will likely also be the direction where the system can most easily shift to another state.

An obvious next question is what this future state might look like. This question is addressed in **chapter 4**. We use the fact that most complex behaviour, such as oscillations or reactivity, arise from delayed negative feedbacks. Therefore, positive feedbacks, which are at the core of mutualistic networks, are expected to give rise to relatively simple dynamics. We find that this relative simplicity allows for extrapolation of the direction of lowest resilience, found by PCA in this chapter, to predict the future state after a tipping point has passed. Therefore, this tool is a valuable addition to our toolbox to analyse complex systems as it provides a way to predict not just when something is about to happen, but also what might happen.

A clear comparison between different multivariate indicators of resilience (of which the indicators in **chapters 3-4** are examples) is lacking. Therefore in **chapter 5** I investigate how different methods relate to one another, if there are methods that are preferred over others and under which conditions the different methods are expected to correctly predict an upcoming tipping point. I demonstrate that there is not one best indicator to warn for an upcoming transition, but that instead it depends on the scenario that the system is subject to. For instance, when data resolution is low, variance-based indicators outperform autocorrelation-based indicators, but for systems with a complex noise regime autocorrelation-based indicators outperform the variance-based indicators. Lastly, this chapter demonstrates that signals can become unreliable when not all variables can be observed. As this scenario is extremely relevant for empirical studies, where one can never be sure that all variables are included, this suggests a cautious interpretation of all work on multivariate resilience indicators.

In **chapters 6-7** I explore the applications of time series analysis tools for complex systems to two real world datasets. In **chapter 6**, I systematically analyze word-use in millions of books from 1850-2019. I demonstrate with PCA that there are two dominant modes of change. The first captures the general trends of world popularity over time, with on one side words like ‘crown’, ‘iron’ and ‘wit’ and on the other side words like ‘computer’, ‘privacy’, and ‘taxi’. The second mode seems to disentangle human nature related words, such as pronouns and emotions, from words related to rational decision, procedures and systems. The behaviour on this second mode is remarkably similar to the sentiment of language over time. I demonstrate that the rational/procedures/systems words show a steady increase from 1850 onwards. This increase stagnates around 1980 and afterwards these types of words show a clear drop. The words related to humans/emotions/sapiens show the exact opposite behaviour. We propose that the increase of sentiment laden words

we observe in the past 20 years could be a reaction to several decades of rational thinking. The fact that the strong increase in sentiment laden words is accompanied by an increase in the use of facebook, suggests that this shift from rational thinking to intuitive thinking is strengthened by the increasing popularity of social media.

In **chapter 7** I use climate data of the past 800.000 years to infer causal links in the carbon cycle. I use sediment cores to determine Ba/Fe (a proxy for biological productivity), $\delta^{18}\text{O}$ (a proxy for climate and ice cover), and $\delta^{13}\text{C}$ (a proxy for ocean ventilation) and ice cores to determine dust and CO_2 (a proxy for climate and alkalinity). One mystery of the glacial-interglacial cycles is their saw-tooth shape of slow cooling and rapid warming. This behaviour hints at the existence of nonlinear processes in the system, for which a possible mechanism could be a feedback loop. As all possible links have been described, it is hard to pinpoint the dominant drivers. Here, I demonstrate that a causal detection method based on Taken's theorem called convergent cross-mapping (CCM) can elucidate causal links in the system and results in one dominant causal loop from ocean ventilation to biological productivity to climate back to ocean ventilation. This loop forms a potential explanation for the shape of the glacial-interglacial cycles.

In **chapter 8** I reflect on the findings of the previous chapters. Furthermore, I share my thoughts on three philosophical considerations. Firstly, I notice how the work in this thesis, as many scientific endeavors, looks for universal patterns even though there is no rule that states that those universal patterns exist. I speculate that perhaps scientists are looking not only for 'truth' but also for 'beauty', and experience most beauty in universal patterns. Secondly, I describe how many scientists are driven by curiosity and that curiosity-based research can lead to unexpectedly far-reaching findings. However, I observe that even though training programs seem to emphasize fundamental research and interdisciplinarity, funding success decreases for interdisciplinary teams. I hope this changes in the future, as all the major issues that are currently faced by science and society require an interdisciplinary team of creative minds to be solved. Lastly, I remark how humans are skilled in finding and explaining patterns, such that they even find and explain patterns that do not exist. I propose that science can learn from software developers in the sense that we should thoroughly check and test every step of the process.

Acknowledgements

The finishing of this thesis would not have been possible without the help and assistance that I received.

I would first like to thank my co-promotor and daily supervisor Ingrid van de Leemput. Ingrid, I first met you in 2015 when I was looking for a place and a project for my master thesis and from that time on you have become my mentor and friend. Ingrid, you are an inspiration to me. You are not only a great scientist who comes up with great questions and builds bridges between different subjects, but you are also a pleasure to work with. I think you can work with anyone on anything, you are a master of creating a nice atmosphere and I am grateful that I could come to you with any problem that I encountered or any crazy idea that passed my mind.

Secondly I want to thank my promotor Marten Scheffer and co-promotor Egbert van Nes for their support, insightful feedback and great ideas during my PhD. Marten, your creativity always motivated me to think outside the box and your positive attitude helped me to stay enthusiastic. Egbert, I have learned that when it comes to the behaviour of complex systems I cannot trust my own intuition but I can trust yours. Thank you for sharing your wisdom on so many occasions.

Three other people played a major role in this thesis. Firstly, Sebastian Bathiany, my co-supervisor during the first two years. Sebastian, you seemed to trust in me when I was doubting myself and you inspired me to keep trying, exploring, learning and failing and you even taught me to enjoy this process. Secondly, Jelle Lever, my supervisor during my master thesis and collaborator on several projects during my PhD. Jelle, your drive for excellence and your perseverance are inspiring. Where others might lose interest in projects during the final stage, you are able to keep improving on them. Lastly, my supervisor during my master thesis and collaborator during various projects of my PhD Rick Quax. Rick, I quickly noticed that my work could benefit from your critical insights and analytic skills and I feel like your contribution to this thesis has sharpened the work that I get to present.

Obviously, all co-authors of this thesis have made valuable contributions and I would like to express my gratitude to all of you: Jerrald Rector, Merlijn Olthof, Jordi Bascompte,

Vasilis Dakos, Johan Bollen, Anne Willem Omta & George van Voorn. George, I am looking forward to our continuing collaboration in the next 1.5 years.

One activity that I enjoyed tremendously during the first years of my PhD were the ‘Thursday theoreticians meetings’ (or GOSSIP meetings). Arie, Babak, Bregje, Pablo, Sanne G. & Usman, thank you so much for the fun discussions that we had. All of you have graduated now, which leaves me some big shoes to fill.

Thanks to Shruti for the delicious Indian dinners, to Merel for pragmatic life advice, to Nancy for helping me with stuff that makes me nervous (like filling in forms), and all other AEW-ers for great Thursday lunch discussions and for just being there.

Apart from training my brain these past years, I also trained my muscles every now and then. Thanks to all my climbing buddies for keeping me safe (and sane), especially to Hazimah and Noel. Noel, you taught me how to climb which improved 50% of my free time and 95% of my holidays ever since. Hazi, you were my main climbing buddy, together we tried, pushed ourselves and each other, fell and overcame our fears. Furthermore, I greatly value the chats that we had driving to and from the gym.


My other sane-keeping activity was playing the guitar. Thank you Ger for teaching and motivating me and thank you Jetse, Jouke and Thomas for playing together. Those guitar ensemble sessions were often the highlight of my week and I have missed it so much during these last months.

I want to thank my friends in the *I, computationalist* group. Thanks for making me laugh, for the beers and for the fun we had. I hope that after I move to Utrecht I can see you more often again.

My friends from *B15*, I still feel grateful for ending up in one room together with you 10 years ago and I am glad we kept in touch throughout the years.

Dear *softijsjes*, I know you guys about half of my life now and I hope you stay in my life much longer and that this year we can make up for the lost weekeindjes of last year. Special thanks to Adriana, Benthe, Iris, Larissa & Margriet. You supported me with hugs, board games, hikes and wine on more than one occasion and I am looking forward to the next.

Linda and Christina, we go all the way back to my time in Oud-Beijerland and I am very happy that I can still call you my friends.

Harrie, Inge, Anne & Valentijn, you are simply the best .

Lastly, I want to thank Victor. Victor, you are a complex system that I will probably never fully understand, but having you by my side has been the best part of these past 4.5 years.

Peer-reviewed journal publications

Els Weinans, J. Jelle Lever, Sebastian Bathiany, Rick Quax, Jordi Bascompte, Egbert H. Van Nes, Marten Scheffer, and Ingrid A. Van De Leemput (2019). Finding the direction of lowest resilience in multivariate complex systems. *Journal of the Royal Society Interface*, 16(159), 20190629.

J. Jelle Lever, Ingrid A. van de Leemput, **Els Weinans**, Rick Quax, Vasilis Dakos, Egbert H. van Nes, Jordi Bascompte, and Marten Scheffer (2020). Foreseeing the future of mutualistic communities beyond collapse. *Ecology letters*, 23(1), 2-15.

Els Weinans, Anne Willem Omta, George A.K. van Voorn, and Egbert H. van Nes. A potential feedback loop underlying glacial-interglacial cycles. *Climate Dynamics*, (2021): 1-13.

Els Weinans, Rick Quax, Egbert H. van Nes, and Ingrid A. van de Leemput. Evaluating the performance of multivariate indicators of resilience loss. *Scientific reports*, 11, no. 1 (2021): 1-11.



*Netherlands Research School for the
Socio-Economic and Natural Sciences of the Environment*

D I P L O M A

for specialised PhD training

The Netherlands research school for the
Socio-Economic and Natural Sciences of the Environment
(SENSE) declares that

Els Weinans

born on 28 June 1992, Nijmegen, The Netherlands

has successfully fulfilled all requirements of the
educational PhD programme of SENSE.

Wageningen, 18 October 2021

Chair of the SENSE board



Prof. dr. Martin Wassen

The SENSE Director



Prof. Philipp Pattberg

The SENSE Research School has been accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW)



K O N I N K L I J K E N E D E R L A N D S E
A K A D E M I E V A N W E T E N S C H A P P E N



The SENSE Research School declares that **Els Weinans** has successfully fulfilled all requirements of the educational PhD programme of SENSE with a work load of 41.3 EC, including the following activities:

SENSE PhD Courses

- o Environmental research in context (2016)
- o Research in context activity: 'Teaching about my PhD project at a primary school' (2019)

Other PhD and Advanced MSc Courses

- o The essentials of scientific writing and presenting, Wageningen Graduate Schools (2016)
- o The Urbino summer school in paleoclimatology, Urbino, Italy (2017)
- o Climate time series analysis, Climate Risk analysis, Heckenbeck, Germany (2017)
- o Scientific writing, Wageningen Graduate Schools (2017)
- o Lake Como summer school: Complex Networks: Theory, Methods and Applications, Lake Como School of advanced studies, Como, Italy (2018)
- o Scientific Artwork - Vector graphics & images, Wageningen Graduate Schools (2018)
- o Teaching and supervising thesis students, Wageningen Graduate Schools (2018)
- o Scientific publishing, Wageningen Graduate Schools (2019)
- o Workshop 'discussing new developments in dynamical systems science related to the climate', Wageningen, The Netherlands

Management and Didactic Skills Training

- o Teaching assistant in the BSc course 'Modeling Biological Systems' (2017-2019)
- o Teaching in the MSc course 'Complexity in Ecological Systems' (2017-2019)

Oral Presentations

- o *Detecting instabilities in complex systems.* NESSC days 2018, 14 March 2018, Veldhoven, The Netherlands
- o *Finding the direction of lowest resilience in complex networks.* Gerontology (GERO) & Frailty in Ageing research (FRIA) Department, VUB, 10 October 2019, Brussels, Belgium
- o *A potential feedback loop underlying glacial-interglacial cycles.* Department of Earth, Atmospheric and Planetary Sciences, 14 August 2020, Cambridge, United States of America (online presentation)
- o *On order, disorder and the role for complexity and resilience thinking.* Radboud Interfaculty Complexity Hub, 31 August 2020, Nijmegen, The Netherlands (online presentation)

SENSE coordinator PhD education

Dr. ir. Peter Vermeulen

This work was carried out under the program of the Netherlands Earth System Science Centre (NESSC), financially supported by the Ministry of Education, Culture and Science (OCW). Printing of this thesis was financially supported by the Aquatic Ecology and Water Quality Management Group of Wageningen University.

Cover design: Larissa van der Laan - ANDA design

Printing: Ridderprint | www.ridderprint.nl

