Detecting and distinguishing transitions in ecological systems: model and data-driven approaches

by

Thomas Bury

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Mathematics

Waterloo, Ontario, Canada, 2020

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

Supervisors:	Chris Bauch Professor, Dept. of Applied Mathematics, University of Waterloo.
	Madhur Anand Professor, School of Environmental Sciences, University of Guelph.
Internal Members:	Matthew Scott Associate professor, Dept. of Applied Mathematics, University of Waterloo.
	Sri Namachchivaya Professor, Dept. of Applied Mathematics, University of Waterloo.
Internal-External Member:	Kim Cuddington Associate professor, Dept. of Biology, University of Waterloo.
External Examiner:	Louis Gross Professor, Ecology & Evolutionary Biology, University of Tennessee.

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

- Chapter 2: The work in this chapter was performed by Thomas Bury (TB) under the supervision of Dr. Madhur Anand (MA) and Dr. Chris Bauch (CB), resulting in a co-authored paper published in PLOS Computational Biology [30]. Contributions were as follows: MA conceived of the study. TB curated the data, analysed the model and wrote the first draft. All authors developed the model and revised the manuscript.

- Chapter 4: This chapter is based on a manuscript under review with PNAS [29]. The work was conceived of and designed by TB, CB and MA. TB developed the EWS methodology and software, performed the analysis and wrote the first draft. All authors revised and commented on the manuscript.

- Chapter 5: This chapter is based on a manuscript in preparation [31]. MA, CB and Dr. Ryan Norris conceived of the study. TB developed the model, performed the analysis and wrote the draft.

Abstract

There exists a plethora of systems that have the capacity to undergo sudden transitions that result in a significantly different state or dynamic. Consider the collapse of fisheries, outbreak of disease or transition to a 'Hothouse Earth' to name a few. The common factor among these transitions is mathematical - they are the result of crossing a bifurcation point. This thesis is concerned with the detection and description of these bifurcations from time series data, and the mechanisms that lead to these transitions. We begin in the domain of climate change, where models of the climate system are extremely sophisticated, but those that incorporate social dynamics and its two-way coupling with climate dynamics are lacking. In developing a simple socio-climate model, we show how mechanisms such as social learning, social norms, and perceived mitigation costs play a major role in climate change trajectories. These social effects can strongly determine the predicted peak global temperature anomaly, how quickly human populations respond to a changing climate, and how we can chart optimal pathways to climate change mitigation. However, we also show that if the climate model is subject to a tipping point, the climate can transition to a new state before mitigating behaviour becomes sufficiently widespread to prevent the transition. This motivates a need for early warning signals (EWS) of tipping points. Hence, in the next chapter we focus on the development of EWS in time series data that can be used to detect an upcoming bifurcation. This thesis develops two 'spectral EWS', which are derived from the power spectrum. We show that the peak in the power spectrum provides a more sensitive and conservative EWS when compared to conventional metrics, and the shape of the power spectrum, quantified using AIC weights, provides clues as to the type of approaching bifurcation. We validate these spectral EWS with empirical data from a predator-prey system. Finally we focus on EWS for population extinction, where we study the efficacy of EWS in seasonal environments. We find that conventional EWS prevail under seasonal environments, however asymmetries exist in higher-order metrics such as skewness and kurtosis that could be used to distinguish the driver of extinction. To conclude, nonlinear behaviour arising from social learning and social norms yield bifurcations that have profound impacts on future trajectories of climate change, and bifurcations can be anticipated across a wide range of systems using spectral EWS, that also provide information on the type of bifurcation. The further development of generic and system-specific EWS will play an important role in preserving healthy ecosystem functioning in the Anthropocene.

Acknowledgements

I would like to thank first and foremost my supervisors Dr Chris Bauch and Dr Madhur Anand for their guidance, support and insightful discussions over the past years. The annual lab dinners were also a lot of fun.

Thank you to Dr Zoran Miskovic, Dr Matt Scott, Dr Sri Namachchivaya, Dr Kim Cuddington and Dr Louis Gross for sitting on my PhD advisory and examination committee.

Thank you to my collaborators Dr Gregor Fussmann, Joey Burant and Dr Ryan Norris for sharing their biological expertise and experimental data.

Thank you to my friends inside and outside of the department for times well spent.

Finally, I would like to thank my close family and Steph for their continued love and support.

Table of Contents

List of Tables

List of Figures

Glossary

Alternative stable states The presence of multiple attractors in system under a fixed parameter configuration. A necessary ingredient for a critical transition. 30

AR(1) process Autoregressive model of order one. A stochastic difference equation describing a mean-reverting random walk. Discrete-time analogue of the continuous-time Ornstein-Uhlenbeck process. xvii, 142

Basin of attraction The set of initial conditions that evolve to a specified equilibrium. 40

Bifurcation A topological change in the behaviour a dynamical system as a parameter varies smoothly past a point. 28

Bistability A property exhibited by a system with two alternative stable states. 29

Catastrophic bifurcation A bifurcation at which point the change in system behaviour results in an irreversible transition to an alternative state. 28

Codimension The number of independent conditions (parameters) determining a bifurcation. 47

Critical slowing down The phenomenon that as a bifurcation is approached, the return time to equilibrium following disturbance increases. 28, 30

Critical transition / regime shift The abrupt change in state induced as conditions cross a catastrophic bifurcation, or following a perturbation that pushes the system across the border of the basin of attraction. 5, 27, 29

Early warning signal A statistical metric that undergoes recognisable change prior to a bifurcation. Also termed 'resilience indicator' [57]. 27

Eigenvalues Quantify the local stability of a system in equilibrium. The system is stable so long as all eigenvalues are negative (magnitude less than one) in continuous (discrete)-time systems. 31

Flip/Period-doubling bifurcation A local bifurcation involving a transition to behaviour twice the period of previous behaviour. Occurs in discrete-time systems. Characterised by a real eigenvalue passing through negative one. 31, 143

Fold/Saddle-node bifurcation A local bifurcation involving collision of a stable and unstable fixed point, resulting in a discontinuous transition to an alternative stable state. Characterised by a real eigenvalue passing through zero (one) in continuous (discrete)-time systems. Often associated with a 'tipping point'. 31, 137

Gaussian white noise A random signal consisting of Gaussian distributed points with zero temporal correlation. 5, 47, 138

Global bifurcation Induced by changes in global system properties such as the collision of basin boundaries. Not detectable from local stability analysis. 40, 41

Hopf bifurcation A local bifurcation involving a transition from a stable equilibrium to stable oscillations. Occurs in continuous-time systems. Characterised by a complex-conjugate pair of eigenvalues attaining positive real part. A supercritical Hopf bifurcation involves a stable limit cycle emerging from the fixed point, resulting in oscillations with gradually increasing amplitude. A subcritical Hopf bifurcation involves an unstable limit cycle engulfing the fixed point, resulting in a sudden transition to large amplitude oscillations. 31, 138, 144

Hysteresis A phenomena that refers to the irreversibility of transitions in a bistable system. 5, 28

Jacobian matrix Describes the local (linear) dynamics of a system about an equilibrium state. Also known as the 'Community matrix' in ecology. 31

Local bifurcation Induced by changes in local stability (eigenvalues of the Jacobian matrix). Detectable from local stability analysis. 30, 47, 130

Neimark-Sacker bifurcation A local bifurcation involving a transition from a stable equilibrium to stable oscillations. Occurs in discrete-time systems (analogous to the Hopf). Characterised by a complex-conjugate pair of eigenvalues attaining magnitude greater than one. 31, 143, 144

Noise-induced transition A shift in state as the system is pushed out of the basin of attraction due to high noise intensity. 40

Normal form The simplest form of a bifurcation to which all others of the same type are locally topologically equivalent. 47, 137, 138, 142

Ornstein-Uhlenbeck process A stochastic differential equation describing a mean-reverting random walk. Continuous-time analogue of the discrete-time AR(1) process. xv, xviii

Pitchfork bifurcation A local bifurcation involving a transition between one and three equilibria. Characterised by a real eigenvalue passing through zero (one) in continuous (discrete)-time systems. Occurs in systems with symmetry. 31, 137

Positive feedback loop Refers to system processes that stimulate further their own activity, therefore causing instability. 30

Random matrix theory Random matrix theory studies the properties of random matrices, which are matrices whose elements are random variables. 130

Stationary process A stochastic process whose statistical properties do not change over time. In mathematical terms, the statistical moments are constant and the autocorrelation between two times is only a function of the time difference. 34

Stochastic process A process that generates a random variable through time, that can be characterised by the mean, variance, and higher statistical moments. They are widely used as mathematical models for systems that involve 'randomness'. 37

Tipping point A threshold where a system undergoes an abrupt transition to an alternative dynamic regime, that is not immediately reversible. 1

Transcritical bifurcation A local bifurcation involving a smooth transition between colliding stable states. Characterised by a real eigenvalue passing through zero (one) in continuous (discrete)-time systems. 31, 137

VAR(1) process Vector autoregressive model of order one. A stochastic difference equation describing a mean-reverting random walk in multiple dimensions. Discrete-time analogue of the multi-variate Ornstein-Uhlenbeck process. 144

Abbreviations

EWS Early warning signal(s) 3

OUp Ornstein-Uhlenbeck process 32

PSD Power spectral density 50

Chapter 1

Introduction

Nonlinear relationships exist in almost every aspect of life. Take coffee as a simple example. Doubling my regular caffeine intake to write this thesis may have made me twice as productive, but if I doubled it again, there would probably be negative consequences; certainly not another two-fold increase in productivity. Doubling the temperature of lukewarm coffee might make it twice as pleasurable to drink, but doubling it again would make it impossible to drink. The relationships among these properties of coffee are therefore nonlinear. The consequence of nonlinearity in this example is rather trivial, however in more complex systems, nonlinearities can lead to surprising results. Consider for example, an ecosystem. Here, we have many interacting components, such as different species, habitat quality and harvesting rates, which are variables that depend on each other via nonlinear relationships. In fact, often it is the rate of change of a variable that has a nonlinear relationship between the other variables, in which case we call this a nonlinear process. The nonlinear processes that make up this ecosystem can give rise to remarkable behaviour such as oscillations, tipping points and even chaotic dynamics. As we will see throughout this thesis, these nonlinear systems can undergo a wide variety of transitions, which are elucidated by mathematical models. Let us first discuss the most widely talked about transition - the tipping point [77].

A tipping point refers to a threshold in a system, beyond which runaway change propels the system to a new state [191]. Though the concept of a tipping point has been around for decades in scientific literature (e.g. [185]), the term 'tipping point' recently gained widespread acceptance, following Malcolm Gladwell's international best-seller [77]. Indeed, the number of research articles studying tipping points appears to itself have crossed a

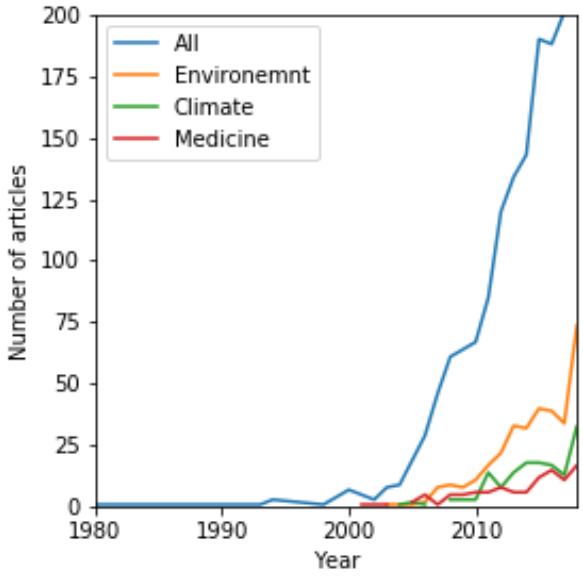

Figure 1.1. The rise in use of the term 'tipping point' among scientific journals. Data obtained from the *Web of Science*. The topic of each article was classified using *Web of Science* categories.

tipping point (Figure 1.1). Tipping points have been observed in empirical data and conjectured by mathematical models in a wide variety of complex systems including social media [146], ecosystems [167] and the climate [53], to name a few. You may even have experienced one personally, be it capsizing a canoe, breaking the ice over a puddle or investing in Bitcoin. The three features common to tipping points are 1) they occur abruptly 2) they result in a transition to a far-away state, and 3) they are not immediately reversible. Given the large consequences of such transitions in systems that we rely upon, there is great need to better understand and predict tipping points. However, it should be noted that tipping points are just one of the many types of transition that can abruptly occur in a dynamical system.

The wealth of possible transitions that a system can undergo can be appreciated from bifurcation theory [181, 114]. A bifurcation is a point in parameter space where the stability properties of the system undergo a qualitative change. For example there can be an exchange of stability between two colliding states (Transcritical), a loss of stability to a stable limit cycle (Hopf), or annihilation between a stable and unstable state resulting to a transition to another attractor (Fold). A tipping point is often manifested as Fold

bifurcation in mathematical models, however this is not always the case as we will see below in the model of social tipping. Given the variety of different bifurcations (and therefore transitions) that can occur, we do not refine ourselves to transitions that correspond to tipping points, but speak of transitions more generally that correspond to the crossing of a bifurcation point. Examples of transitions that do not correspond to tipping points include the transition to oscillations in predator-prey communities via a Hopf bifurcation [70], the transition to cardiac arrhythmias via a period-doubling bifurcation [153] and the transition to species extinction via a Transcritical bifurcation [61].

This thesis addresses transitions from two different paradigms of applied mathematics. First, there is the realm of mathematical modelling, where one makes mechanistic assumptions about a system, written in the form of equations, and tries to explain observed behaviour or predict future behaviour based on the model, although the uses of mathematical modelling go far beyond this [66]. In this thesis, we will be using simple models, with the primary goal of explaining the effect of different mechanisms on transitions, and assessing the behaviour of the model in parameter regimes outside of what is currently observed. These models give insight into the type of dynamics that a system can exhibit for a range of different parameter values - they are not making predictions per se. This leads us to the second paradigm of data-driven approaches, that are largely independent of model assumptions.

Models can only take us so far in detecting upcoming transitions in dynamical systems. They are well-suited to describing the type of dynamics that can exist and under which conditions, though their ability to locate a transition is hampered by uncertainty in model parameter values and assumed mechanisms. Recent work has shown that data-driven approaches that monitor changes in statistical properties over time can help in detecting transitions [168, 53, 190, 61, 50]. These statistical changes have been referred to as 'early warning signals' (EWS), a term which we will use throughout the thesis. Research on EWS for tipping points exploded after a popular review in 2009 [168] (also seen in Figure Figure 1.1), now with a multitude of different metrics in both spatial and temporal data. We provide a review of this relatively new approach in Chapter 3, with a focus on the current limitations, that this thesis addresses.

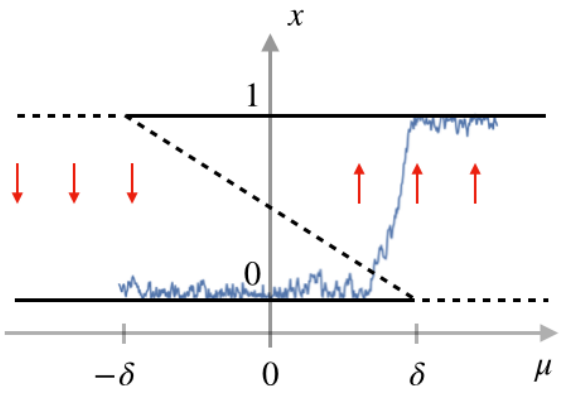

Figure 1.2. Bifurcation diagram and single realisation for the simple social model. Solid (dashed) black lines show stable (unstable) states. Blue line shows a stochastic simulation of the model in Equation 1.5. Red arrows indicate the direction of flow that a trajectory would undergo with fixed parameter values. The model exhibits two Transcritical bifurcations at $\mu = \delta, -\delta$ which both result in critical transitions to the alternative state. The simulation has μ increasing linearly from $[-1, 1.5]$ over 50 time units, and parameter values $\delta = 1$, $\kappa = 1$ and $\sigma = 0.1$.

1.1 Transitions in a simple social system

It is interesting to see how transitions can occur in social systems from the perspective of a simple model. Consider a population, where individuals may choose between two strategies A and B. As an example, let's take the choice of being a smoker (A) or a non-smoker (B). Let x be the proportion of individuals who adopt strategy B (so $1 - x$ adopt strategy A). Each strategy has a payoff, which we assign as

$$e_A = -\mu(t) + \delta(1 - x) \tag{1.1}$$

$$e_B = \delta x \tag{1.2}$$

where $\mu(t)$ is a general time-dependent cost that an individual incurs to adopt strategy A. In our example, this could be some function of the perceived health risk of smoking, which started increasing during the 50s as studies began reporting links between smoking and lung cancer. The δ terms represent payoffs based on the current proportion of individuals with each strategy, capturing the tendency for individuals to conform to social norms [142, 116]. For instance, the payoff to adopt strategy B, δx, increases with x, the propor-

tion of individuals who adopt this strategy. Now incorporating a social learning dynamic from evolutionary game theory [92], which we provide more detail on in Chapter 2, the population evolves according to

$$\dot{x} = \kappa x(1-x)(e_B - e_A) \tag{1.3}$$

$$= \kappa x(1-x)(\mu(t) + \delta(2x-1)) \tag{1.4}$$

where κ is a social learning rate. The bifurcation diagram for this simple model (Figure 1.2) shows how the popular strategy among the population can abruptly shift to the alternate strategy as $\mu(t)$ crosses a threshold. Moreover, once the system has transitioned, it is 'locked in' by social norms, in the sense that reversing the change in $\mu(t)$ does not recover the original state - an effect known as hysteresis. A simple way to model random effects in the system is to include a white noise term, yielding the stochastic model

$$\dot{x} = \kappa x(1-x)(\mu(t) + \delta(2x-1)) + \sigma\xi(t) \tag{1.5}$$

where σ is a noise amplitude and $\xi(t)$ is a Gaussian white noise process. A stochastic simulation of this system with gradually increasing μ is shown in Figure 1.2, demonstrating a critical transition.

In the context of smoking where $\mu(t)$ is related to factors such as perceived health risk, scientific evidence of the dangers of smoking was clear as early as 1964, but declines in smoking abundance did not take place until over a decade later [74], in part due to rigid social norms. This phenomena is captured by the simple model wherein increases in μ have little effect on the state of the system until it exceeds a threshold (the strength of social norms, δ), at which point the system transitions relatively quickly to a majority adopting the alternate strategy. Although a large simplification of human behaviour, this model demonstrates a very general phenomenon of how social norms can hold populations in a certain way of thinking, until a new idea comes along, significant enough to exceed the threshold set by social norms.

In Chapter 2, this thesis will explore the consequences of these dynamic social processes in the context of climate change and mitigative behaviour among individuals. The setup for social dynamics is similar, however factors within the payoffs will now be dependent on the climate, which is itself an evolving system. A fundamental question is how does the interaction between social processes and the climate system effect global climate predictions? We address this question from a simple modelling framework. The combination of models

for social behaviour and geophysical models of the Earth's climate system is something that only recently became apparent in the scientific literature [11], and early results have revealed how nontrivial, nonlinear behaviour can arise when the two are coupled. We contribute to this emerging literature with a simple model that captures the salient features of social and climate dynamics.

1.2 Thesis outline

This thesis is organised as follows. We begin in the modelling paradigm in Chapter 2, where we use a simple socio-climate model to better understand the impact of various social mechanisms (e.g social norms, social learning) on future trajectories of climate change. The modelling of social behaviour adopts the framework outlined earlier, except now payoffs are coupled to the climate system, which is itself evolving over time. Modelling the climate system with a two-way coupling to a social component is a relatively new idea [11] and the work in this thesis is the first to model the effects of endogenous social processes on climate change. This chapter has been published in PLOS Computational Biology [30].

In Chapter 3, we move to the data-driven paradigm, and introduce the key concepts and theory behind critical transitions and their corresponding model-free EWS. This chapter serves as an introduction to the area, but also a review of the key literature that concerns EWS in time series data. We provide an example of EWS in action, and also highlight the current limitations of EWS, that the contributions in this thesis address.

Chapter 4 develops a genre of EWS that we have termed 'spectral EWS' since they use information in the power spectrum to infer an upcoming bifurcation. The paper [29] shows how spectral EWS can provide not only advanced warning of a transition compared to traditional EWS, but also information on the type of upcoming bifurcation. It is currently in submission with PNAS.

Chapter 5 is concerned with the behaviour of EWS prior to extinction in seasonal populations. The work shows how EWS vary depending on the process that is driving extinction, and also how time series data of populations from different seasons can be used together to provide an enhanced composite EWS of extinction. This paper is being prepared for submission.

The thesis concludes in Chapter 6, where we discuss the broader implications of our results

and draw on further research directions that merit pursuit. Note that a glossary of technical terms is provided at the start of the document.

Chapter 2

Charting pathways to climate change mitigation in a coupled socio-climate model

This chapter is based on the paper: Thomas M Bury, Chris T Bauch, and Madhur Anand. Charting pathways to climate change mitigation in a coupled socio-climate model. *PLOS computational biology*, 15 (6):e1007000, 2019.

2.1 Abstract

Geophysical models of climate change are becoming increasingly sophisticated, yet less effort is devoted to modelling the human systems causing climate change and how the two systems are coupled. Here, we develop a simple socio-climate model by coupling an Earth system model to a social dynamics model. We treat social processes endogenously - emerging from rules governing how individuals learn socially and how social norms develop - as well as being influenced by climate change and mitigation costs. Our goal is to gain qualitative insights into scenarios of potential socio-climate dynamics and to illustrate how such models can generate new research questions. We find that the social learning rate is strongly influential, to the point that variation of its value within empirically plausible ranges changes the peak global temperature anomaly by more than 1°C. Conversely, social norms reinforce majority behaviour and therefore may not provide help when we most need it because they suppress the early spread of mitigative behaviour. Finally, exploring the model's parameter space for mitigation cost and social learning suggests optimal intervention pathways for climate change mitigation. We find that prioritising an increase in social learning as a first step, followed by a reduction in mitigation costs provides the most efficient route to a reduced peak temperature anomaly. We conclude that socio-climate models should be included in the ensemble of models used to project climate change.

2.2 Introduction

According to many ancient myths, humans did not invent fire-making de novo but rather learned it from personalities like Prometheus and subsequently spread the practice amongst themselves. These stories reveal how ancient myth-makers already grasped the fundamental importance of social learning - the process whereby individuals learn new behaviours, values and opinions from others [155]. Social learning is no less relevant in the era of human-environment challenges [91, 10, 71]. The importance of social learning and social processes more generally in climate change mitigation and adaptation is well recognised [36, 41, 24, 7]. Increasingly sophisticated geophysical climate models are helping us understand the impacts of anthropogenic greenhouse gas (GHG) emissions [25, 46, 177], and the importance of these models is hard to understate. However, climate projections depend strongly on the assumed trajectory of GHG emissions [192]. This trajectory is determined by human behaviour and yet climate models generally do not incorporate dynamic so-

cial processes relevant to GHG emissions. Rather, GHG emissions are assumed to follow some specified trajectory. These trajectories are constructed with socio-economic factors in mind, (see Representative Concentration Pathways [192] and Shared Socioeconomic Pathways [141] for instance), but are not coupled to climate dynamics and do not capture human responses to climate change in a mechanistic way.

Just as human behaviour influences climate trends, climate change in turn influences human behaviour concerning GHG emissions, including both climate change mitigation and adaptation [36, 41, 7, 198, 95]. Individuals in places with rising average temperatures are more likely to perceive climate change [95], and social effects are apparent when individuals take steps in response to such shifting perceptions [41, 7]. There is also an important distinction between social learning and social norms - socially accepted and widely practised modes of conduct [3]. Social norms are known to have a strong influence on human behaviour [40] including aspects relating to climate change [24, 3, 139] and therefore play an important role in determining emission trajectories [41]. Multiple studies show a tendency for individuals to conform to emerging norms in support of climate change mitigation [24, 3]. Moreover, it appears that individuals are often not consciously aware of the importance of social norms in their decision-making and instead falsely ascribe their decisions to other factors [139]. However, it is important to note that social norms do not automatically promote socially beneficial outcomes. They can equally well force conformity to a destructive norm such as political extremism [174]. This also happens in the context of climate change behaviour, where it has been found that individuals may also conform to a norm of non-mitigation, by adjusting their habits to match those of less environmentally friendly neighbours [139, 180].

Hence, Earth's climate and human subsystems are part of a single coupled system where social dynamics play a vital role. Yet, models of Earth's coupled climate-behaviour system remain essentially undeveloped. One such approach [11] couples a sophisticated climate model [177] to a model for individual behavioural change based on the theory of planned behaviour - a dominant paradigm in psychology [1]. The authors find that the sensitivity of global temperature change to human factors such as response to extreme events, social norms and perceived ability to adopt mitigative strategies is of a similar magnitude to its sensitivity to geophysical factors. They deduce that quantifying behavioural uncertainty and physical uncertainty in climate projections deserve equal attention. The model focuses on how individual psychology and behaviour are influenced by extreme weather events. The model for individual psychology is scaled up to the population level in a way that does not

account for endogenous social processes such as social learning. As such, individuals do not learn behaviour or opinions from one another, and social norms are treated as a fixed effect that does not depend on the population's current composition of attitudes.

Here, we treat social learning and social norms endogenously, by modelling their dynamics as they emerge from rules governing how individuals interact, learn and behave. Our first objective is to develop qualitative insights into how different aspects of the system - endogenous social processes, temperature trends, and mitigation costs - separately and together determine possible dynamics of the larger socio-climate system. Our second objective is to illustrate potential uses of coupled socio-climate models to chart social and economic policy pathways that mitigate climate change as quickly as possible. To meet these objectives, we sought to develop a model that (1) could capture a range of IPCC climate change scenarios, ranging from 4 degrees of warming by 2100 (RCP 8.5 scenario) to sub 2 degrees of warming (RCP 2.6 scenario), (2) was simple enough to analyse so that we could learn which mechanisms drive the predicted socio-climate dynamics, (3) was based on existing approaches for modelling social dynamics and climate dynamics, and (4) captured the salient features of social and climate systems. Given the model's simplicity, it is primed for insights as to how social and climate processes interact, though limited in its predictive capacity due to the complexity of the socio-climate system. The development of more complex socio-climate models will be an important research avenue, once the mechanisms of socio-climate dynamics are better understood.

2.3 Methods

Model

Geophysical models in the climate science literature span a wide range of different complexities depending on the associated research objective. Highly complex models are the state-of-the-art for weather and climate prediction [177, 107, 113], whereas simple models allow us to assess processes and feedbacks, thereby improving our intuition of climate system dynamics [117, 47, 48, 134, 125]. Likewise, the behavioural sciences have benefited from a variety of modelling approaches, that address the diverse set of social processes that take place on the individual and societal level [88]. Here, we use minimal models for both social and climate dynamics. Starting simple allows us to build intuition on the effect

of socio-climate feedbacks that have yet been considered in the climate change literature. The social model is widespread and, despite its simplicity, captures the salient aspects of social dynamics [88, 96, 91]. Moreover, the simple Earth system model that we use [117] accurately follows the projections of the state-of-the-art CMIP5 models when forced with the IPCC emission scenarios (Figure A.1).

Behaviour dynamics

Individuals in our model are either 'mitigators' or 'non-mitigators' and they learn these behaviours from others at a specified social learning rate. They switch to the more attractive behaviour according to a utility function governed by the costs of climate change mitigation (such as the cost of installing solar panels or buying gas-electric hybrid vehicles), the costs imposed on non-mitigative behaviour (such as a carbon tax), the costs associated with the average global temperature anomaly, and the utility associated with social norms that reinforce the majority behaviour, whether it be mitigation or non-mitigation. Social norms strengthen as the majority behaviour becomes more prevalent, consistent with empirical studies [24, 139]. Formally, the utility of being a mitigator is taken as

$$e_M = -\alpha + c\tilde{f}(T_f) + \delta x \tag{2.1}$$

where α is the cost of adopting mitigative strategies, c is a proportionality constant that regulates the extent to which climate change costs influence incentive to mitigate, $\tilde{f}(T)$ is the cost associated with with a temperature anomaly of T degrees Celsius, T_f is a projected temperature anomaly (see below), δ is the strength of social norms and x is the proportion of mitigators in the population. The utility of being a non-mitigator is taken as

$$e_N = -\gamma - \tilde{f}(T_f) + \delta(1 - x), \tag{2.2}$$

where γ is the cost of non-mitigative behaviour, representing for example a carbon tax. Note that the utility due to social norms here is instead weighted by the proportion of non-mitigators, $1 - x$.

Social learning is implemented as follows. Each individual samples other members of the population at a fixed rate κ. If 'individual A' samples 'individual B' and B's strategy has a higher payoff that of A, then A will switch to B's strategy with a probability proportional to the difference in payoffs. Thus if there are x mitigators in the population, the rate at which

a non-mitigator encounters a mitigator is κx. Should there be a utility gain in switching ($e_M > e_N$), the non-mitigator will switch their strategy with a probability proportional to the payoff difference, $e_M - e_N$. Since there are a total of $1 - x$ non-mitigators in the population, the total rate at which non-mitigators switch to being mitigators is

$$r_{N\to M} = \kappa x(1-x)\max\{e_M - e_N, 0\}. \tag{2.3}$$

Using similar arguments, one can show

$$r_{M\to N} = \kappa x(1-x)\max\{e_N - e_M, 0\} \tag{2.4}$$

and so the net rate of change in mitigators is

$$\frac{dx}{dt} = r_{N\to M} - r_{M\to N} = \kappa x(1-x)(e_M - e_N), \tag{2.5}$$

which has an equivalent form to the replicator equations of evolutionary game theory [92]. Writing the payoff functions explicitly gives

$$\frac{dx}{dt} = \kappa x(1-x)\big[\gamma - \alpha + (c+1)\tilde{f}(T_f) + \delta(2x-1)\big], \tag{2.6}$$

which can be reduced to

$$\frac{dx}{dt} = \kappa x(1-x)\big[-\beta + f(T_f) + \delta(2x-1)\big], \tag{2.7}$$

by introducing the new parameter $\beta = \alpha - \gamma$ (the net cost to mitigate) scaled function $f(T) = (c+1)\tilde{f}(T)$ (the net temperature associated gain to mitigate).

Temperature projection

Because long-term climate forecasts are known to influence individual decisions on whether to support mitigation [198], the utility function assumes that individuals base decisions on long-term extrapolations of recently experienced climate trends. The projected temperature takes the form

$$T_f(t) = T(t) + \left(\frac{t_f}{t_p}\right)\left(T(t) - T(t-t_p)\right) \tag{2.8}$$

where t_p is the number of years back to extrapolate from and t_f is the number of years forward to extrapolate to.

Perceived costs associated with climate change

We assume that the costs associated with climate change have a sigmoidal relationship with the global temperature anomaly. This is motivated by the slow mitigative response to global warming over the past decade (Figure A.8 - Figure A.10) and the anticipated non-linear alterations in the Earth system with increasing temperature [118]. Specifically

$$\tilde{f}(T) = \frac{\tilde{f}_{\max}}{1 + e^{-\omega(T - T_c)}}, \tag{2.9}$$

where $\tilde{f}_{\max}$ corresponds to a maximum cost, ω is the degree of nonlinearity of the sigmoid, and T_c is the critical temperature about which costs are most sensitive to change (Figure A.5). We combine c and $\tilde{f}_{\max}$ into the single parameter

$$f_{\max} = (c + 1)\tilde{f}_{\max}, \tag{2.10}$$

which yields the functional form

$$f(T) = \frac{f_{\max}}{1 + e^{-\omega(T - T_c)}}, \tag{2.11}$$

as used in (2.7).

Earth system model

We couple the social model to an Earth system model [117] with reduced ocean dynamics [133]. Dynamics for atmospheric CO_2 is modified to include an anthropogenic emission term, dependent on the proportion of non-mitigators. Specifically

$$\frac{dC_{\text{at}}}{dt} = \epsilon(t)\frac{1 - x}{1 - x_0} - P + R_{\text{veg}} + R_{\text{so}} - F_{\text{oc}} \tag{2.12}$$

where C_{at} is the deviation in atmospheric CO_2 from pre-industrial values, $\epsilon(t)$ is the baseline rate of CO_2 emissions in the absence of mitigation, x_0 is the initial proportion of mitigators,

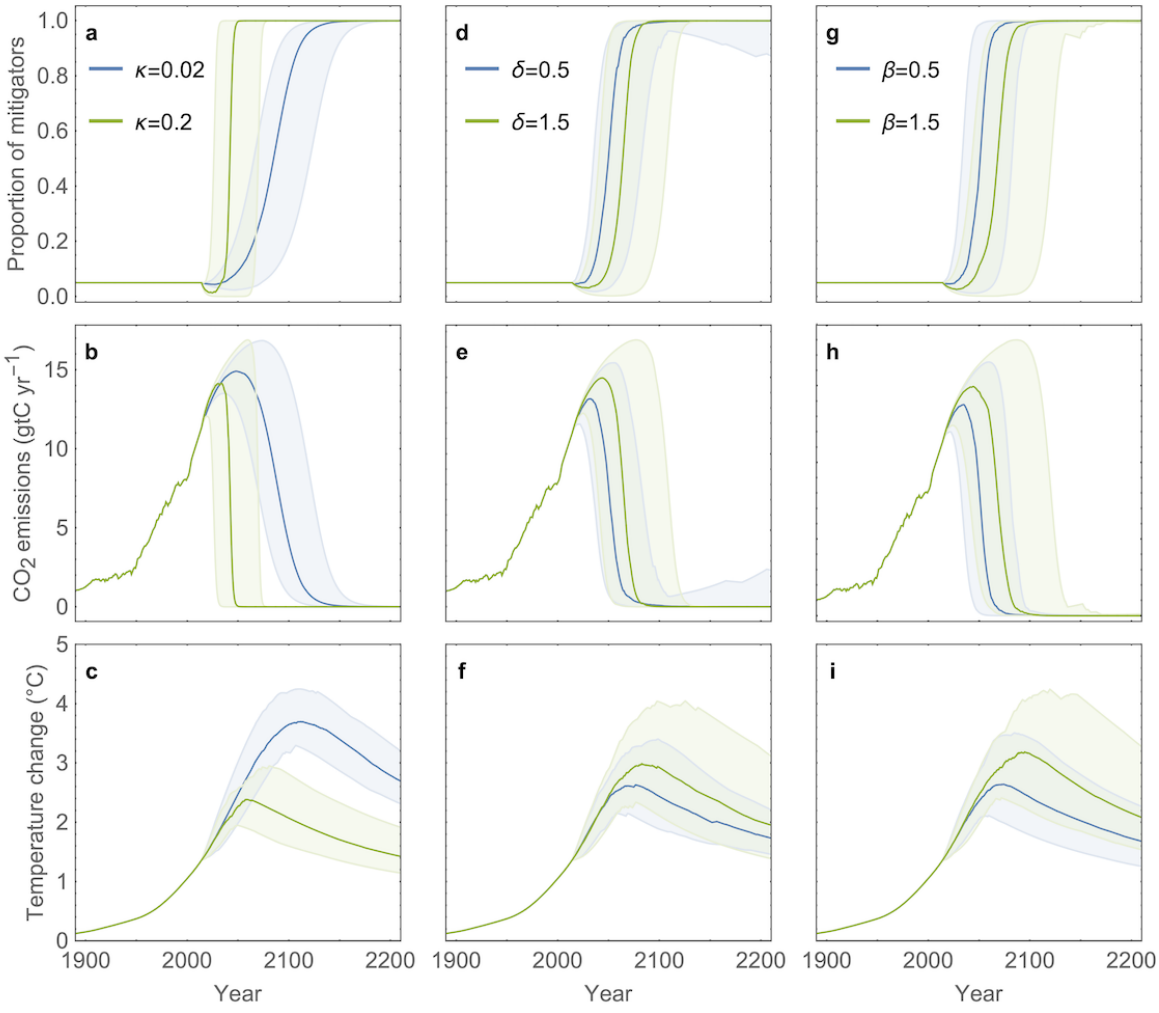

Figure 2.1. Endogenous social dynamics influence climate trajectories. Shown are ensembles of model simulations comparing two parameter values for the social learning rate (κ) (a-c), strength of social norms (δ) (d-f), and costs associated with mitigation (β) (g-i). All other parameters are drawn from triangular probability distributions with defined upper, lower and baseline values (Table A.2). Displayed are median trajectories with 95% confidence intervals from an ensemble of 100 realisations. Additional details in Methods.

P is the rate of carbon uptake via photosynthesis, R_{veg} is the outward carbon flux via plant respiration, R_{so} is the outward carbon flux via soil respiration and F_{oc} is the net uptake of carbon by the oceans. Functional forms for the transfer of carbon via these processes are provided in Supplementary Information. Global surface temperature is assumed to evolve with the carbon cycle according to [117]

$$c\frac{dT}{dt} = (F_d - \sigma T^4)a_E \tag{2.13}$$

where T is the deviation in global surface temperature from pre-industrial values, c is the thermal capacity specific heat capacity of the Earth's surface, F_d is the net downward flux of radiation absorbed at the planet's surface (which depends on the opacity of CO_2), σ is the Stefan-Boltzmann constant, and a_E is the Earth's surface area. Full details are provided in Appendix A.

Simulation

Over the period from 1800 to 2014, the socio-climate model is simulated with a fixed social component, forced with historical anthropogenic carbon emissions. Initial conditions for all climate variables are zero since they represent deviations from pre-industrial values. Social dynamics are initiated in 2014 with an initial proportion of mitigators $x_0 = 0.05$. The ensuing dynamics of $\epsilon(t)$ follow an increasing but saturating trend corresponding to the world's increasing but saturating population size and energy demands. Specifically

$$\epsilon(t) = \begin{cases} \text{linear interpolation of historical emissions} & t \leq 2014 \\ \epsilon_{2014} + \frac{(t-2014)\epsilon_{\max}}{t-2014+s} & t \geq 2014 \end{cases} \tag{2.14}$$

where $\epsilon_{\max}$ is the saturating value, and s the half-saturation constant, of $\epsilon(t)$. This expression is shown graphically in Figure A.7. The system of (delay) ordinary differential equations is simulated using the NDSolve package in Wolfram Mathematica. Historical CO_2 emissions were obtained from the CDIAC data repository [17].

Parameters and sensitivity analysis.

Baseline climate parameters are obtained from the original Earth system model [117] where they were fitted to obtain historical trends of temperature and carbon dynamics. Social parameters are more speculative and so are given wide upper and lower bounds. The relative cost of warming ($f(T)$) with respect to the net cost of mitigation (β) is chosen in accordance with the argument that the costs of preventative action will be far less than the cost implied otherwise by global warming [178]. For sensitivity analyses we draw parameters from triangular distributions that peak at baseline values and extend to upper and lower bounds (Table A.2). Parameters are kept fixed preceding 2014 to retain historical trends in the simulations.

2.4 Results

The model demonstrates how the social learning rate can strongly determine temperature trends. We first consider a null hypothesis where adaptive behaviour is removed from the model by forcing the proportion of mitigators in the population to remain constant. In this case of fixed behaviour, emissions saturate and the temperature anomaly increases indefinitely (Figure A.2). However, once social learning is added and the proportion of mitigators is allowed to evolve dynamically as in our baseline model, the predicted average global temperature anomaly can peak anywhere from 2.2°C, near the Intergovernmental Panel on Climate Change (IPCC) limit [170] (in the case of very rapid social learning) to 3.5°C (in the more realistic case where social learning unfolds on a generational timescale) (Figure 2.1a-c). Whether people discuss climate change more or less often can therefore strongly influence temperature trends. Because we model social norms as something that tends to reinforce majority behaviour and attitudes - whatever they might be - one might think that social norms act as a double-edged sword. In fact, they operate more like an unhelpful scimitar, as illustrated by comparing cases of low and high strength of social norms. Because the population starts off from a state of largely non-mitigating behaviour, increasing the strength of social norms suppresses the spread of mitigating behaviour for decades by entrenching non-mitigation as a norm, even when rising temperatures strongly justify an immediate shift (Figure 2.1d-f). (This model dynamic echoes not only current climate norms reinforcing non-mitigation [180] but also past social shifts occurring on decadal timescales, such as evolving social norms about when and where smoking is

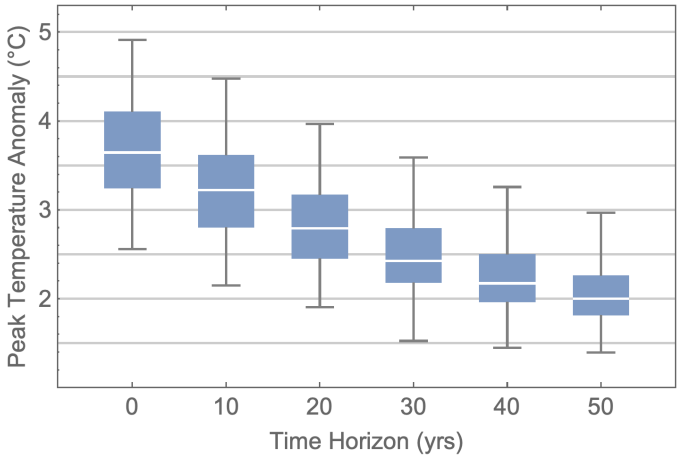

Figure 2.2. Peak temperature anomaly strongly depends on individual time horizon. Box-and-whisker plot shows the peak temperature anomaly measured over 100 realisations for fixed time horizons (t_f) and all other parameters drawn from triangular probability distributions with defined upper, lower and baseline values (Table A.2). Boxes span the interquartile range, whiskers span the entire set of realisations, and box dividers mark the median. Additional details in Methods.

acceptable.) However, when mitigating behaviour eventually does become widespread, a higher strength of social norms does not significantly accelerate its spread. Rather, the two curves for cases of high and low social norm strength simply move in parallel to one another because by this time, the utility function that determines behaviour change is dominated by the large temperature anomaly (Figure 2.1d). In this parameter regime, social norms generate a perverse asymmetry, in contrast to findings from other socio-climate models that assume social norms can only support climate change mitigation [180].

The model also shows how a reduction in net mitigation cost can significantly accelerate the onset of social change. For instance, a 67% reduction in the mitigation cost increases the percentage of mitigators by 2060 from 10% to 90% (Figure 2.1g-i). Therefore, policies that reduce the cost of mitigation (through e.g. subsidies, tax cuts) will benefit from the accelerating effects of social learning and must be timed correctly.

Our baseline model assumes that individuals' perceived cost of climate change impacts depends on a linear extrapolation of the recent temperature anomaly over the previous ten years (Methods). If individuals instead base their decisions only on the current temperature

anomaly, the simulated global temperature anomaly lies well above the 2°C target set by the IPCC, and exhibits wide variation in sensitivity analysis (Figure 2.2). This contrasts with our baseline model where the population movement towards mitigative strategies ignites earlier, significantly reducing the global temperature anomaly. This predicted dynamic stems from the multi-decadal lag between GHG emissions and the consequent global temperature rise [86].

Our model predicts medium-term GHG emission trajectories (Figure 2.1b,e,h) that are qualitatively similar to those often assumed under various future emissions scenarios. This raises the question of how such models can be useful. The socio-climate model enables us to explore how socio-climate dynamics might respond to changes that are under the control of policymakers. For instance, it is possible to compare social and economic policy interventions by considering the effects of simultaneous parameter changes, instead of one at a time. This enables us to chart out the quickest pathways from highest to lowest temperature anomalies. The relative merits of increasing the social learning rate vs. reducing the net cost of mitigation are illustrated with a contour plot where the contours represent peak temperature anomaly as a function of the two parameters (Figure 2.3). Increasing the social learning rate (e.g. through media coverage and public fora devoted to climate change) is particularly effective when social learning is slow, but has saturating benefits, as indicated by the increasing vertical spacing of contour lines for at higher learning rates. In contrast, reducing the net mitigation cost (e.g. through tax breaks) drives a more linear response in peak temperature anomaly. Crucially, it should be noted that both a reduction of net mitigation cost and an increase in the social learning rate are required to achieve the IPCC target. The arrows in Figure 2.3 show the 'path of steepest descent' - the most efficient combination of the two measures. Starting from a situation of high projected temperature anomalies, the model predicts that increasing the social learning rate should first be prioritised, followed by a reduction in net mitigation cost once the benefits of social learning begin to saturate. This approach gets us to the region of parameter space corresponding to the IPCC target faster than alternative trajectories.

A sensitivity analysis reveals the relative influence of each parameter on the peak temperature anomaly (Figure 2.4). The time horizon of individuals' temperature projection, social learning rate and costs of mitigation are major factors, all of which may be influenced by appropriate intervention. The importance of social parameter uncertainties in determining climate predictions indicated by our model has also been predicted by other socio-climate models[11]. Interestingly, the system is relatively insensitive to the initial proportion of

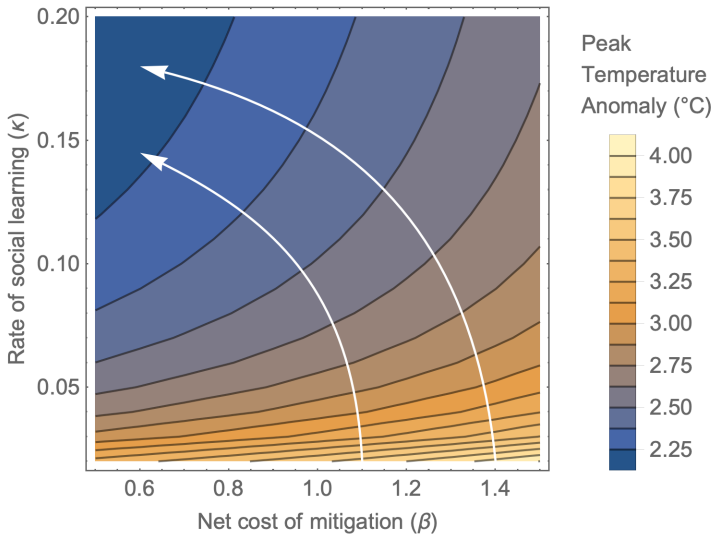

Figure 2.3. Optimal pathways to mitigation via increasing social learning and reducing mitigation costs. Contour plot showing peak temperature anomaly at specified values of the net cost of mitigation (β) and rate of social learning (κ). Arrows indicate the direction of steepest descent - the most efficient combination of the two measures to reduce temperature anomalies. All other parameters are fixed at baseline values (Table A.2). Additional details in Methods.

mitigators, suggesting that the mediation of social processes, as opposed to the current social state, is key to guiding the socio-climate system to a trajectory of reduced emissions. Sensitivity analyses such as these can help investigators determine priorities for data collection: the parameters exhibiting the greatest influence on predictions should be targeted for data collection so we can best reduce model uncertainty.

A striking feature revealed by the sensitivity analysis is the asymmetry in many of the parameter dependencies. Consider the three parameters with highest impact on the peak temperature anomaly (concerning forecast horizon, learning rate and global warming costs). A decrease in these parameters is more detrimental than an increase is beneficial. For example, a forecast horizon 10 years over baseline value results in a 0.6 degree decrease in peak temperature anomaly, whereas a forecast horizon 10 years below baseline value results in a 1 degree increase. This imbalance is a manifestation of the nonlinear interactions between and within each of the social and climate system.

The sensitivity analysis also reveals non-monotonic relationships between the peak tem-

perature anomaly and the parameters. For example, both an increase and a decrease in solar flux results in a higher peak temperature anomaly. Interestingly, this is not the case if the climate subsystem is considered in isolation. For a fixed emissions scenario, a higher (lower) solar flux will always result in a higher (lower) peak temperature anomaly, since the solar flux is proportional to the net downward radiation absorbed by the planet's surface. The coupling to social dynamics fundamentally alters this relationship. In the socio-climate system, a reduced solar flux results in a slower increase in surface temperature. As a consequence, individuals are less incentivised to mitigate, causing the social system to maintain a regime of non-mitigative behaviour. The accompanying high rate of CO_2 emissions quickly overcompensates for the reduced solar flux, yielding a higher peak temperature anomaly. Thus seemingly useful interventions to the physical system can actually end up doing more harm than good when there is strong coupling to a social system, as is the case for global warming.

2.5 Discussion

This study has shown how social processes can influence climate dynamics, according to one possible way of modelling social dynamics and norms. However, other frameworks for modelling human behaviour could yield different predictions. For instance, the socio-climate model of Ref. [11] does not include social learning. Individuals respond directly to changes in the climate, and not through interactions with one another. As a consequence, the rate at which individuals adopt mitigative strategies only varies with the current climate situation, and not with current population consensus. Mitigation efforts can therefore be expected to closely follow the severity of climate change in the model. In our model, social learning manifests as a feedback within the social system, resulting in qualitatively different socio-climate trajectories. Mitigative behaviour is initially suppressed–even as temperatures rise to levels that should incentivise mitigation–due to low numbers of mitigating individuals and therefore little turnover of behaviour in the population. However, social learning creates a positive feedback loop once there is a net positive utility to mitigate, and so as the numbers of mitigators increases, so too does the rate at which non-mitigators switch to being mitigators. This results in a sharp non-linear increase in mitigators, as a combined outcome of both the social and the climate system dynamics. We note that, all else being equal, adding social learning to a model has the effect of slowing down behaviour change in the human population (since a process takes

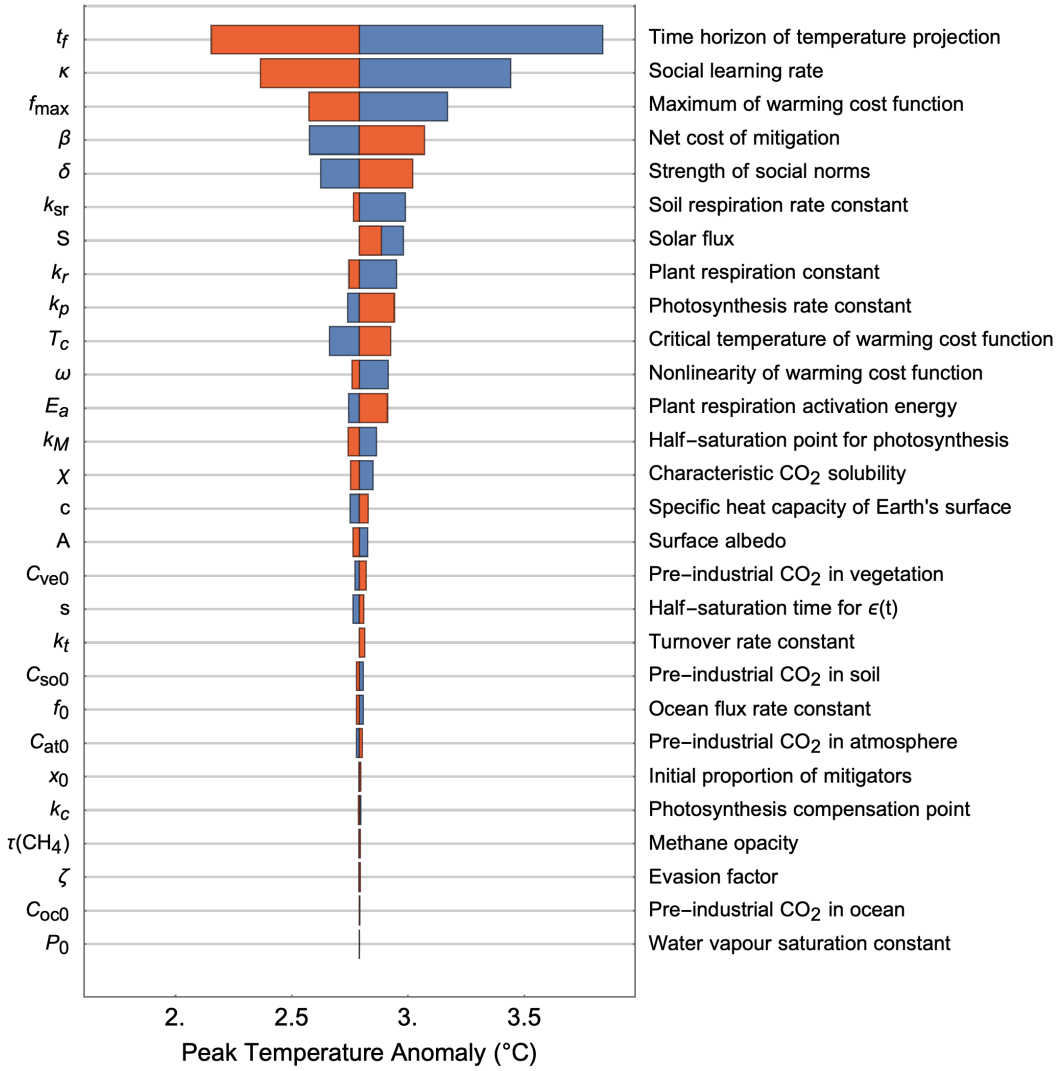

Figure 2.4. Peak temperature anomaly is sensitive to both social and climate parameters. Tornado plot showing the deviation in peak temperature anomaly when parameters are varied individually. A red (blue) bar indicates the deviation when the specified parameter adopts its upper (lower) bound, while all other parameters are fixed at baseline values (Table A.2). Parameters ordered by relative impact from top to bottom. Additional details in Methods.

time, by definition), and therefore the mitigation response of human populations.

Conversely, the case of very rapid social learning recovers a 'best response' model similar to those assumed in classical economics, where individuals immediately adopt the highest payoff strategy without learning the behaviour from others. Whether or not this assumption can approximate behaviour in real human populations hinges upon how fast social learning occurs - individuals would need to sample others rapidly enough to enable complete population behaviour change within 5 years for this approximation to work in our model, which seems implausible (Figure 2.1a-c).

In a different vein, we assumed a homogeneous population with respect to mixing and individual utilities. The model's social dynamics capture interactions at the individual level, though there are many different scales of social organisation that the model does not consider, from families/neighbourhoods to cities/states and up to interacting countries. Future models could include this more hierarchical social structure. Similarly, these models could include different types of individual with correspondingly different utilities. For instance, the model could include industrial corporations with utilities biased toward shareholder profit, and social institutions (such as laws, taxes, the education system) that reflect the current governmental stance. Social learning may also take on different forms due to diverse individual psychologies and values [16, 45, 97]. Such heterogeneities are known to affect the dynamics of a wide variety of systems [62] and can prevent population consensus by permitting development of echo chambers [59]. Our model also makes the simplifying assumption that individuals base their temperature projection on linear extrapolation of past temperatures. This could be generalised to a non-linear extrapolation to reflect an individual's perception of 'accelerating' change. Extending socio-climate models to include these finer details should prove valuable in further investigations.

Climate change is a manifestation of coupled human-environment dynamics and therefore we should start coupling climate models to social models [36, 187]. Our simple coupled socio-climate model shows that the rate at which individuals learn socially strongly influences the peak global temperature anomaly, to the point that variation of this parameter within plausible ranges changes the peak temperature anomaly by more than 1°C. Therefore, it matters whether social processes cause slow or fast uptake of climate change mitigation measures. We found that social norms may not provide help when we most need it, although this finding could be nuanced by adding social heterogeneity. Finally, we illustrated how exploring the parameter space of socio-climate models suggests optimal paths for mitigating climate change. A more sophisticated policy impact assessment model

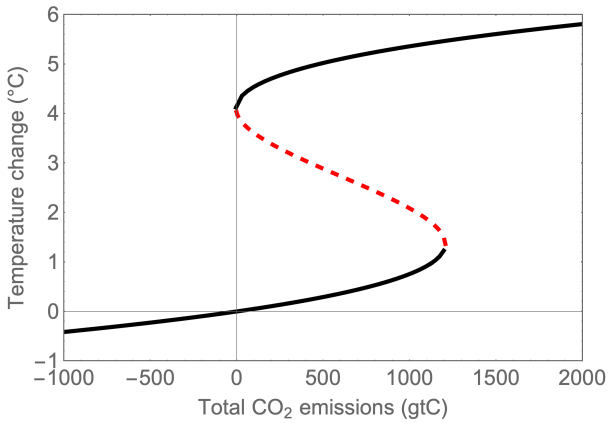

Figure 2.5. Modelling a tipping point in the climate system. Bifurcation diagram for model of global temperature given in (2.16), treating total anthropogenic CO_2 emissions as a parameter. Solid (dashed) lines correspond to stable (unstable) states. Parameters are $a = 0.015$, $b = 0.018$, $T_1 = 4$, $T_2 = 4.2$.

based on a coupled socio-climate approach could therefore be useful to decision-makers facing a mandate to reduce GHG emissions with a fixed budget. In summary, it is essential for climate change research to account for dynamic social processes in order to generate accurate predictions of future climate trends, and the paradigm of coupled socio-climate modelling could help us address this challenge.

2.5.1 Climate tipping

The majority of Earth system models, including the one used in this study, assume a quasilinear relationship between global temperature and cumulative carbon dioxide emissions [179]. However, there exist certain biogeophysical feedbacks that constitute positive feedback loops, which could yield tipping points within the Earth system at certain concentrations of CO_2. It is suggested that some of the elements giving rise to these feedbacks are vulnerable to tipping within just a 1°C to 3°C increase in global temperature [176]. An example of a positive feedback loop is the thawing of permafrost [165], which releases carbon dioxide and methane into the atmosphere, resulting in increased temperatures and therefore more thawing of permafrost. The combination of these positive feedback loops or 'tipping elements' [118] could lead to a climate trajectory that does not stabilise at an intermediate temperature, but continue to increase even as anthropogenic emissions are

reduced, resulting in a transition to a so-called 'Hothouse Earth' [176].

To investigate the scenario of a climate tipping point within a socio-climate framework, we construct the modified model

$$\dot{x} = \kappa x(1-x)(-\beta + f(T_f) + \delta(2x-1)), \tag{2.15}$$

$$\dot{T} = a\left[-T(T-T_1)(T-T_2) + bC\right], \tag{2.16}$$

$$\dot{C} = \epsilon(t)\frac{1-x}{1-x_0} \tag{2.17}$$

where social dynamics are as before, but now global temperature is modelled by a generic form (2.16) that exhibits bistability, and a Fold bifurcation (Figure 2.5) that is driven by the total anthropogenic CO_2 emissions, C. In this setup, the Earth system components are modelled implicitly within the differential equation for temperature. A simulation of this modified model is shown in Figure 2.6 and compared to a simulation of the original model at baseline parameter values. The simulation of the modified model shows the scenario of the climate system crossing the Fold bifurcation before mitigative behaviour has become sufficiently widespread to prevent the transition. The global temperature shows slow increase up to the tipping point, beyond which the climate system accelerates rapidly towards a new undesirable state of high global temperature. Due to the minor increases in temperature prior to the tipping point, the social system does not respond quickly enough to prevent the transition. Therefore techniques to predict upcoming tipping points are urgently required. This motivates the following chapters, which study early warning signals - statistical changes in time series data that could be used to infer upcoming bifurcations, and therefore tipping points.

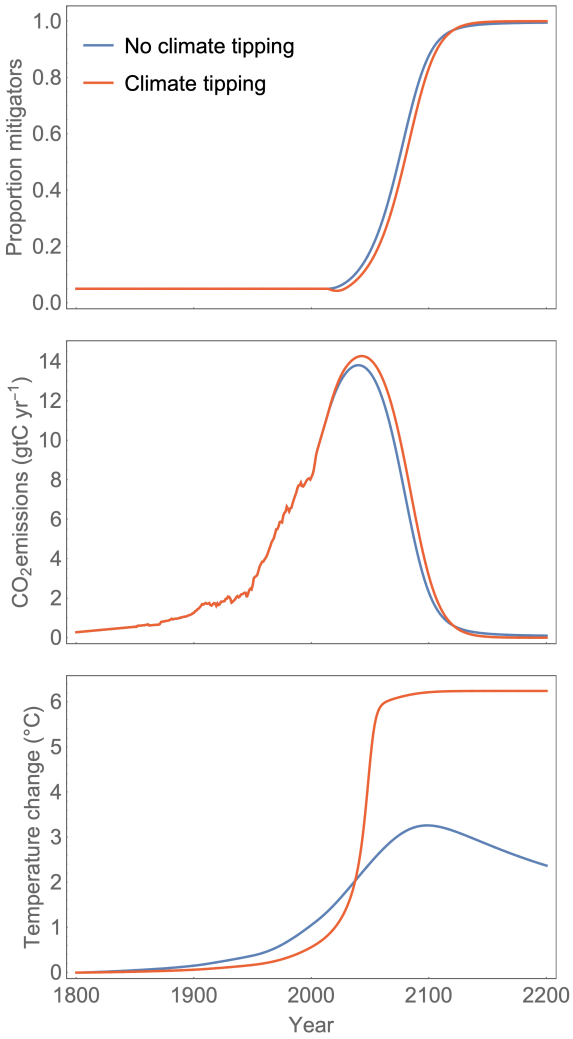

Figure 2.6. Comparing socio-climate trajectories with and without a tipping point. Blue line shows a simulation of the original socio-climate model at baseline parameter values. Red line shows a simulation of the modified model that possesses a tipping point in the dynamics of global temperature as CO_2 levels increase. In this configuration, the climate flips to a new state before mitigative behaviour has become widespread enough to avert the tipping point. Parameters of the original socio-climate model are taken at baseline values. Parameters for the climate tipping model are $a = 0.015$, $b = 0.018$, $T_1 = 4$, $T_2 = 4.2$, with baseline social parameters.

Chapter 3

Early warning signals for critical transitions: a review

3.1 Preface

Mechanistic models have long been used as tools to investigate the stability of systems, be that an ecosystem, a financial market, or a suspension bridge. Many systems such as these possess thresholds at which the system suddenly loses stability and transitions to a distinctly different state, an event known as a critical transition. In the realm of mathematical models, these thresholds correspond to bifurcations, and the surrounding theory [181, 114] can be used to get insight into the type of behaviour the system exhibits after crossing the threshold. However, for this information to be relevant, a mathematical model must accurately represent the system in question i.e. it must capture the appropriate system mechanisms. Moreover, to get a precise estimate on the location of bifurcations, the model should have accurate parameter values. Fortunately for suspension bridges, the governing mechanisms are well established from physics and engineering (e.g. Newton's laws of motion), and parameter values can be measured directly. As such, engineers can compute the load a suspension bridge can take before it becomes unstable. In contrast, the mechanisms and parameter values of an ecosystem or financial market are less certain, and models of these systems, while useful (see e.g. [66]), do not generally provide precise values for where stability will be lost. To help with this, a recent growth of literature has been devoted to the development of so-called early warning signals (EWS), that are independent

of any model assumptions or parametrisation [168]. Instead, they capture generic changes in statistical properties that occur prior to an instability.

This chapter provides the reader with a basic understanding of critical transitions and how EWS can be used to aid in their detection. We begin by describing the variety of transitions that can occur in a system due to different types of bifurcation, and what makes a transition 'critical'. We then introduce the phenomena of critical slowing down, which the majority of EWS are based upon. We show how critical slowing down and the accompanying EWS can be understood intuitively using stability landscapes, and provide a more mathematical justification by deriving analytical approximations for them in a simple stochastic model. There exist a plethora of EWS that have been developed, and so we restrict our review to EWS that are designed for time series data (as opposed to spatial data [54, 103, 159]), with a particular focus on their ability to predict certain types of bifurcation and make use of multiple sets of time series data, since it is these aspects that this thesis contributes to in Chapter 4 and Chapter 5. Finally we discuss the limitations of EWS and the circumstances under which they should be applied.

3.2 Bifurcations and critical transitions

In systems with nonlinear feedback mechanisms, bifurcations are likely to occur. These are points in parameter space where the stability properties of the system undergo a qualitative change. Bifurcations come in a variety of types, with each one resulting in a specific change in dynamics, such as the onset of oscillations or a transition to a far-away attractor. For a full description of the different types of bifurcation, there exist many good books on the subject [181, 114].

An important property of bifurcations, is whether or not the transition they yield is reversible. Some bifurcations yield a 'smooth' transition i.e. the system remains within a local neighbourhood as the bifurcation is crossed (e.g Pitchfork, supercritical Hopf). As such, the original state can be restored upon returning the bifurcation parameter to its original value. In contrast, other bifurcations result in a 'critical' transition whereby the system moves to an attractor outside of the local neighbourhood (e.g. Fold bifurcation, subcritical Hopf). We refer to these bifurcations as catastrophic bifurcations. In these cases, the original state is not recovered upon returning the bifurcation parameter to its original value - an effect known as hysteresis. Thus, a critical transition not only results

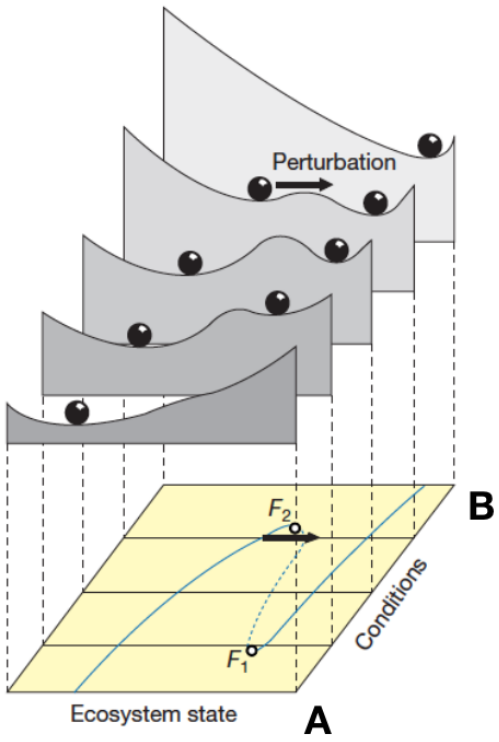

Figure 3.1. Schematic of alternative stable states and a critical transition upon changing external conditions. Lower panel shows the bifurcation diagram representing a Fold bifurcation. Upper slices show the stability landscape fixed values of external conditions. In conditions at point A, the system possesses a single stable state, and so will return to this state upon a perturbation of any size. As conditions vary towards point B the system develops bistability, allowing the possibility of noise to knock the system into the alternative stable state. As the conditions get closer to B the edge of the basin of attraction gets closer, making it more likely for noise to push the system out. For conditions at B, the original stable state no longer exists due to the Fold bifurcation, and so the system will have necessarily transitioned at this point. The effect of hysteresis can be seen in that returning the conditions to the bistable configuration does not recover the original state. Figure adapted from [167].

in a large shift to an alternative dynamical regime, but is also not immediately reversible. In the context of real-world systems that humans rely upon such as ecosystems, predicting these transitions is clearly of significant importance.

It should be noted that there exist many terms in the literature referring to 'abrupt transitions from small changes'. Examples include tipping points, break points, regime shifts and critical transitions. In this thesis, we reserve the term critical transition for transitions that occur due to the approach of a catastrophic bifurcation. In reality, systems will often transition prior to the bifurcation being crossed, due to the noise magnitude becoming on the order of the basin width. We use the term transition more generally, to describe a qualitative change in system dynamics, that may or may not be the result of a catastrophic bifurcation. For example, many predator prey models undergo the onset of oscillations via a supercritical Hopf bifurcation upon an increased level of resources [70] (the so-called paradox of enrichment [163]). In contrast, significant reduction of resources can result in

extinction via a Transcritical bifurcation. Neither of these bifurcations are catastrophic, although they do result in a qualitative change of dynamics, so we refer to them simply as 'transitions'.

A necessary ingredient for a catastrophic bifurcation, and thus a critical transition, is that of a positive feedback loop which refers to a change that it self-reinforcing [104]. An alarming example of this is the Earth's albedo effect, whereby the melting of ice promotes greater warming, which in turn melts more ice. A system with positive feedback loops then has the potential to exhibit alternative stable states, meaning that under the same parameter configurations, the system could sit in more than one stable state. This can be visualised with stability landscapes as in Figure 3.1, where valleys correspond to stable states and the slope represents the rate of change of the system. A prototypical example of a systems with alternative stable states is that of freshwater lakes, which may reside in either an oligotrophic or a eutrophic state under the same environmental conditions [167]. Such contrasting states under the same environmental conditions are remarkably common in nature [166], and suggest the presence of underlying bifurcations in these systems.

3.3 Critical slowing down

At the heart of EWS theory is the phenomenon of critical slowing down. When a system resides in a stable state, it will return to equilibrium following some small perturbation. However, as a system gets close to a local bifurcation, the time it takes for the system to recover will increase due to a weakening of restoring forces. Critical slowing down refers to the process of these restoring forces tending to zero as a local bifurcation is approached. In continuous-time mathematical models, this corresponds to an eigenvalue of the Jacobian matrix obtaining positive real part, or in discrete-time systems, an eigenvalue obtaining magnitude greater than one. The type of bifurcation that occurs depends on the way the eigenvalues achieve this (Figure 3.2).

The appealing feature of critical slowing down is that it is generic to local bifurcations [181] i.e. in the approach to every local bifurcation, the local recovery rate tends to zero. The first suggestion of this concept as an EWS was by Wissel [204] who proposed an increase in return time as a signal for thresholds in ecosystems. This was later demonstrated in a living system [194], by perturbing a cyanobacterial population at different light levels, and showing that recovery time increased at levels closer to the tipping point. However,

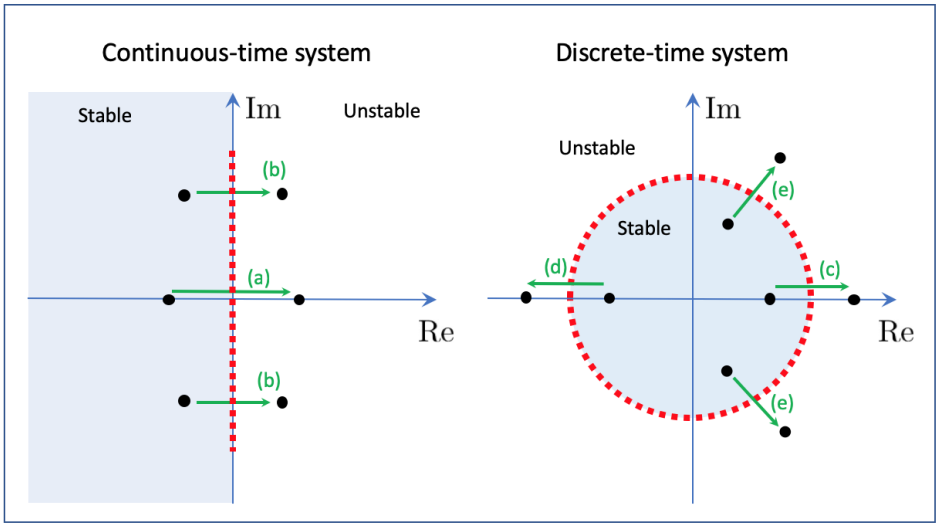

Figure 3.2. Eigenvalue evolution that accompanies each type of local bifurcation. The schematics show eigenvalues of the Jacobian matrix plotted on the complex plane, and the generic ways in which they can evolve to result in a local bifurcation. When all eigenvalues lie within the shaded region, the system is stable (it returns to equilibrium following a small displacement). In a continuous-time system, **a.** a real eigenvalue becoming positive accompanies either a Fold, Transcritical or Pitchfork bifurcation, **b.** a complex-conjugate pair of eigenvalues attaining positive real part accompanies a Hopf bifurcation. In a discrete-time system, **c.** a real eigenvalue becoming greater than one accompanies either a Fold, Transcritical or Pitchfork bifurcation, **d.** a real eigenvalue becoming less than negative one accompanies a Flip bifurcation, **e.** a complex conjugate pair of eigenvalues attaining a magnitude greater than one accompanies a Neimark-Sacker bifurcation [114].

perturbing a system and monitoring its recovery is often not possible. Fortunately, critical slowing down is manifested indirectly via certain statistical metrics in systems subject to noise [98]. As the restoring forces diminish, the system becomes increasingly more like its previous state, i.e. the lag-1 autocorrelation increases [106]. Moreover, the noise is able to drive the system further from the equilibrium state, resulting in an increase in variance [32]. Rising variance and lag-1 autocorrelation are the two most commonly used EWS have been detected in many empirical systems prior to tipping points [50, 38, 196, 206, 53].

The effects of critical slowing down can be appreciated mathematically by deriving the statistical indicators analytically. We demonstrate this here with the simple one-dimensional

system with additive white noise,

$$\dot{x} = f(x) + \sigma\xi(t), \tag{3.1}$$

where $\xi(t)$ is a Gaussian white noise process. We assume the system sits in a stable equilibrium $x = x^*$, so $f(x^*) = 0$ and $f'(x^*) < 0$. Expanding about the equilibrium point with the substitution $x(t) = x^* + \eta(t)$, we have

$$\begin{aligned}
\dot{\eta} &= f(x^* + \eta) + \sigma\xi(t) && (3.2)\\
&= f(x^*) + \eta f'(x^*) + O(\eta^2) + \sigma\xi(t) && (3.3)\\
&= \lambda\eta + O(\eta^2) + \sigma\xi(t), && (3.4)
\end{aligned}$$

where $\lambda = f'(x^*)$ is the local recovery rate. Assuming noise is sufficiently small, and so the system stays close to equilibrium, we may drop the $O(\eta^2)$ term. At a local bifurcation is approached, the recovery rate λ tends to zero. Note that the stability landscape of this system is

$$V(x) = \frac{1}{2}|\lambda|x^2 \tag{3.5}$$

demonstrating how it flatter as the bifurcation is approached ($\lambda \to 0$). The residual dynamics $\eta(t)$ satisfy the stochastic differential equation

$$\dot{\eta} = \lambda\eta + \sigma\eta(t) \tag{3.6}$$

which is the famous Ornstein-Uhlenbeck process (OUp) [188]. We can derive the stationary variance and autocorrelation using standard methods [73] as follows. Integrating the expression in (3.6) gives

$$\eta(t) = \eta(0)e^{\lambda t} + \sigma e^{\lambda t}\int_0^t \xi(t')e^{-\lambda t'}dt', \tag{3.7}$$

where $\eta(0)$ is the initial position deviation from equilibrium. We are interested in the statistics of the process at stationarity, and so may safely drop the transient term, to arrive at

$$\eta(t) = \sigma e^{\lambda t}\int_0^t \xi(t')e^{-\lambda t'}dt'. \tag{3.8}$$

Note that in the limit as $\lambda \to 0$, the restoring forces vanish resulting in a random walk

$$\eta(t) = \sigma \int_0^t \xi(t')dt'. \tag{3.9}$$

Do demonstrate the statistical changes that result form critical slowing down, we compute the autocovariance function, which is defined as

$$\phi(\tau) = \langle \eta(t)\eta(t+\tau) \rangle \tag{3.10}$$

where τ is the lag time, and $\langle ... \rangle$ denotes the expectation of a random variable. Inserting Eqn (B.4) into this expression yields

$$\phi(\tau) = \sigma^2 e^{2\lambda t + \lambda\tau} \int_0^t dt' \int_0^{t+\tau} dt'' \langle \xi(t')\xi(t'') \rangle e^{-\lambda(t'+t'')}. \tag{3.11}$$

The white noise process satisfies

$$\langle \xi(t')\xi(t'') \rangle = \delta(t' - t'') \tag{3.12}$$

where δ is the Dirac Delta function. The autocovariance may then be simplified and integrated to get

$$\phi(\tau) = \frac{-\sigma^2}{2\lambda} e^{\lambda\tau}(1 - e^{2\lambda t}) \tag{3.13}$$

which, in the stationary regime ($t \to \infty$) is

$$\phi(\tau) = \frac{-\sigma^2}{2\lambda} e^{\lambda\tau}. \tag{3.14}$$

This yields the following expression for the variance

$$\mathrm{Var}(\eta) = \phi(0) = \frac{-\sigma^2}{2\lambda} \tag{3.15}$$

and the lag-τ autocorrelation

$$\rho(\tau) = \frac{\phi(\tau)}{\phi(0)} = e^{\lambda\tau}. \tag{3.16}$$

It is clear that as a bifurcation is approached in this general system ($\lambda \to 0^-$), the variance diverges and the lag-1 autocorrelation approaches unity. This is a consequence of critical slowing down.

3.4 Computing EWS from time series data

We have seen theoretically the statistical changes that are expected to take place as a local bifurcation is approached. But how do we go about detecting these changes in empirical data, which is often non-stationary, noisy, and of low resolution? A thorough methodological review is provided elsewhere [55], but here we describe two important procedures that are required prior to computing EWS from empirical data, and the system-specific parameters that need to be chosen.

EWS are concerned with changes in fluctuations about equilibrium, not how the equilibrium itself changes over time. The first step to computing EWS is therefore to detrend the data, which is typically done using Gaussian smoothing or a Lowess filter. Both detrending techniques take a parameter (bandwidth and span, respectively) that specifies the size of the interval used to fit the smoothing function. This interval needs to be small enough to capture changes in the trend, but large enough such that it does not filter out the stochastic effects that we are hoping to capture (Figure 3.3). Choosing an appropriate span is therefore dependent on time scales of the system and the amount of available data, and the sensitivity of EWS to this parameter should be assessed.

Once the data has been detrended, the residual dynamics are obtained by taking the difference between the raw time series and the trend. It is the residual time series that is used for computing EWS. In order to detect changes in the statistical metrics, a rolling window is used, whereby the indicator is computed from the time series within the rolling window at each location. The size of the rolling window is another parameter to be specified and should be chosen carefully. The rolling window size should be small enough such that the time series can be considered approximately stationary within it, but large enough such that there is enough data within the window to provide a reasonable estimate of the statistic.

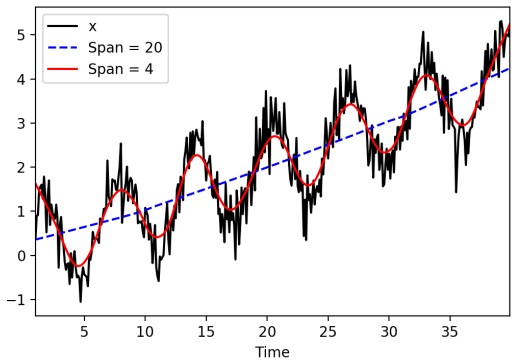

Figure 3.3. Comparing the Lowess span used for detrending a time series. The appropriate span for detrending a time series is system dependent. The detrending should capture any directional / periodic trends without filtering out the stochastic component that is relevant to the EWS. Here, although a span of 20 captures the linear trend in this time series, it does not capture the periodicity, which would significantly impact the EWS. A span of 4 is more appropriate, as it captures both the linear and the periodic trends. This time series is generated from $x_t = 0.1\,t + \sin(t) + \epsilon_t$, where ϵ_t is a normal random variable of mean zero and standard deviation 0.4.

3.5 Example: May's harvesting model

We demonstrate the application of EWS using output from May's harvesting model [124], which is given by

$$\frac{dN}{dt} = rN\left(1 - \frac{N}{K}\right) - \frac{\beta N^2}{N_0^2 + N^2} \tag{3.17}$$

where N represents the abundance of some population. The population undergoes logistic growth, with growth rate r and carrying capacity K, but also undergoes harvesting, which occurs according to the familiar 'Type III' consumption function whereby consumption increases faster that linearly for small population values, and eventually saturates for to a constant for N above some characteristic value N_0.

We use the rescaled form with the dimensionless variables $x = N/K$, $\tilde{t} = rt$,

$$\frac{dx}{d\tilde{t}} = x(1-x) - \frac{hx^2}{s^2 + x^2}, \tag{3.18}$$

where $s = N_0/K$, $h = \beta/rK$. It can be shown that this system exhibits alternative stable

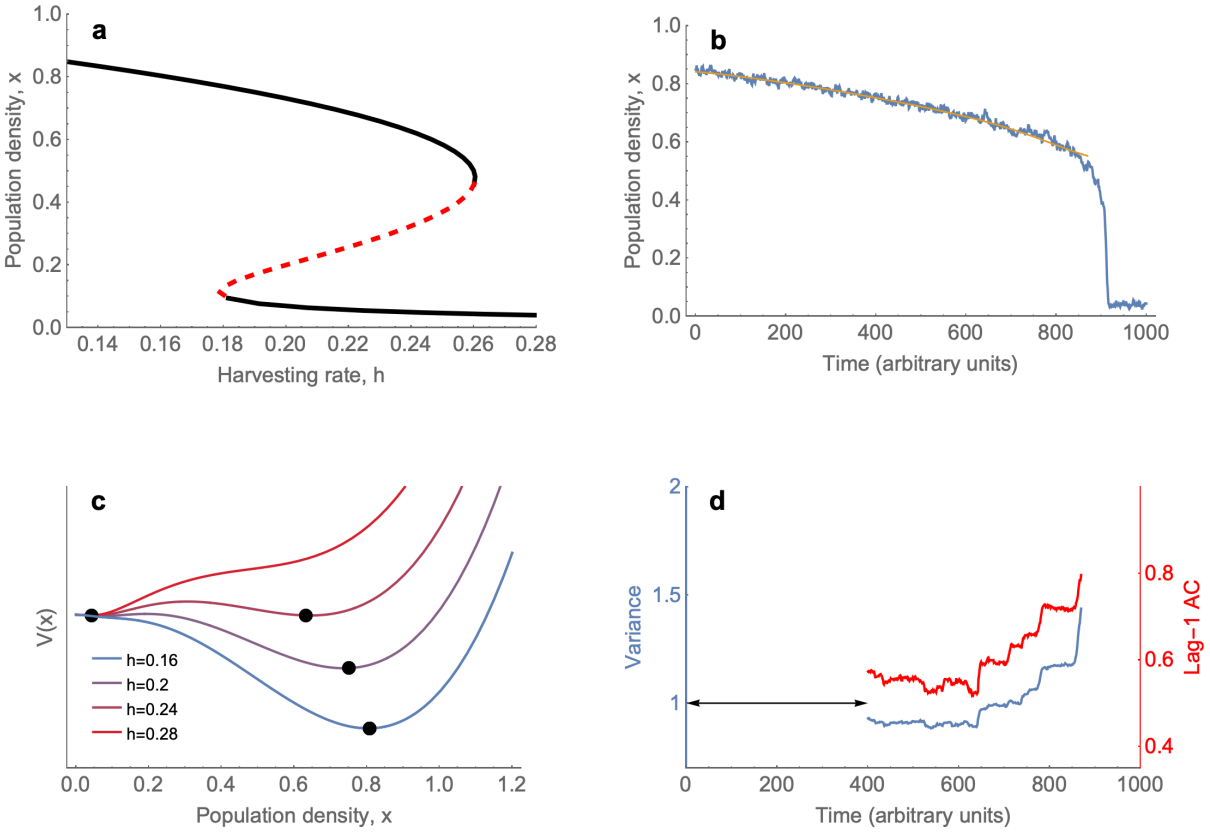

Figure 3.4. EWS prior to the critical transition in May's harvesting model. This simple ecological model [124] (3.18) illustrates how under simple mechanistic assumptions, a population can transition from high abundance to low abundance via a Fold bifurcation. For real populations, it is hard to pinpoint the location of a bifurcation due to uncertainty of parameters and model assumptions. Model-free EWS such as variance and lag-1 autocorrelation can help predict the approach of bifurcations, as demonstrated here. **a.** Bifurcation diagram where black lines indicate stable states, red dashed indicates unstable state. **b.** A realisation of the stochastic model (blue) in (3.20), with h increasing linearly over the interval $[0.13, 0.28]$. The trend (orange) is captured using a Gaussian filter of bandwidth 0.2 times the length of the time series. **c.** Stability landscape for the deterministic model at different harvesting rates h. **d.** Variance and lag-1 autocorrelation computed from the stochastic realisation, using a rolling window as illustrated by the arrow. Parameters values are $s = 0.1$, $\sigma = 0.01$.

states if and only if $s < 1/3\sqrt{3}$ [124]. The bifurcation diagram for varying h and fixed $s = 0.1$ is shown in Figure 3.4a. Note the critical transition that occurs from high to low population abundance via a Fold bifurcation, as the scaled harvesting rate h increases. The stability landscape can be computed from the dynamical equation $\dot{x} = f(x)$ as

$$V(x) = -\int^{x} f(t)dt, \tag{3.19}$$

which is shown in Figure 3.4c. Note how as h increases, a stable state at low population densities arises, and then at a critical value of h the equilibrium state at high population densities vanishes.

To frame this model in a stochastic setting, we include additive white noise. Dropping the tilde for notational convenience, the system may be written as the following stochastic process,

$$\frac{dx}{dt} = x(1-x) - \frac{hx^2}{s^2+x^2} + \sigma\xi(t) \tag{3.20}$$

where σ is a noise amplitude and $\xi(t)$ is a Gaussian white noise process. A simulation for linearly increasing h over the interval $[0.13, 0.28]$ is shown in Figure 3.4b. The variance and lag-1 autocorrelation for this trajectory are computed as outlined in the previous section and plotted in Figure 3.4d. Note how both of these metrics increase prior to the critical transition, providing and advanced warning, as is to be expected for the approach of a local bifurcation.

3.6 A plethora of metrics

We have shown theoretically and with a simple model how variance and lag-1 autocorrelation are expected to increase prior to a critical transition, potentially serving as an EWS. These are just two metrics amongst an abundance of others that have been proposed over the past two decades [169]. The full suite of EWS can be partitioned into those that apply to time series data and those that apply to spatial data. We restrict our focus to EWS relevant to time series data. Within this subset of EWS, there exists another divide between model- and metric-based indicators. Model-based indicators assume a statistical model, which is fit to data, before inferring changes in stability and providing an EWS. Metric-based indicators on the other hand, make no initial assumptions about the structure of

the data, and are based solely on properties of the time series. We restrict our review to metric-based indicators that apply to time series data and their current limitations, since this is where contributions of this thesis have been made.

Table 3.1 lists the metric-based EWS for time series data developed to date. The sheer number of different metrics highlights the interest that this field has generated. The table also provides information on two important properties: 1) whether the EWS can provide information on the type of bifurcation that the approaching instability corresponds to, and 2) whether the EWS provides its signal based on multi-variate information i.e. multiple sets of time series data. The EWS originally proposed, namely autocorrelation and (univariate) statistical moments [204, 106, 32] showed promise in detecting upcoming instabilities, however it soon became evident that they would not determine the type of bifurcation [102], and therefore the nature of the transition to be expected. Moreover, the use of only a single time series to detect an upcoming transition has been shown to be inadequate in certain cases [18, 137], since critical slowing down occurs in a direction of phase space that may be orthogonal to the state variable being monitored. In such cases one cannot expect an EWS. As such, it has been recognised that the development of EWS that are bifurcation specific and/or draw insights from multi-variate data is an important endeavour [57].

We can see from Table 3.1 that some of the more recently proposed EWS provide bifurcation-specific information and make use of multivariate data. However, the development and empirical testing of these EWS is still in its naissance. To the best of our knowledge, there exists only the following bifurcation-specific EWS, and each one has its own specific limitations and room for improvement. Eigenvalue reconstruction [203] takes in noisy, multi-variate time series data and provides an approximation to the Jacobian matrix that generated the data. Considering the nature of the eigenvalues and how they change over time, one can make some deductions about the type of bifurcation that is approaching. However, in order to reconstruct the Jacobian using the proposed methods, time series data for each variable in the system is required, which is often not the case. Secondly, the scaling of variance has been shown to vary depending on the type of bifurcation [129]. However detecting this scaling in an applied setting is only possible with controlled experimental setups, where the bifurcation parameter can be accurately measured. Thirdly, large perturbation monitoring [119] has shown promise in determining the type of bifurcation. However, as the name suggests, this approach requires the system to undergo a large perturbation, which is often infeasible. Finally, critical attractor growth [100] is an EWS that

Early warning signal	*Bifurcation-specific*	*Multi-variate data*	*Ref.*
Recovery time (from perturbation)			[190]
Variance / Standard deviation / C.V.			[32]
Lag-1 autocorrelation			[106]
Skewness			[82]
Kurtosis			[15]
Conditional heteroskedasticity			[171]
Spectral reddening			[106]
Spectral exponent			[55]
Spectral ratio			[15]
BDS test			[34]
Detrended fluctuation analysis			[120]
Quickest detection method			[35]
Eigenvalue reconstruction	*	*	[203]
Scaling of variance	*		[129]
Large perturbation monitoring	*		[119]
Covariance matrix (dominant eigenvalue)		*	[33]
Cross correlation		*	[39]
Fisher information		*	[64]
Information dissipation length		*	[154]
Kolmogorov complexity			[27, 52]
Shannon entropy			[27]
Nonlinearity		*	[58]
Critical attractor growth	*		[100]
Dynamical network marker		*	[39]

Table 3.1. Review of metric-based EWS for detecting of bifurcations in time series data. Asterisks indicate whether the EWS provides bifurcation-specific information and/or makes use of multi-variate time series data. Shorthand notation: C.V., coefficient of variation.

has been suggested as specific to detecting global bifurcations, not for distinguishing local bifurcations. There is therefore need for complementary methods that provide information on incipient bifurcations from noisy time series data, something this thesis addresses with 'spectral EWS' in Chapter 4.

There are also recent developments in EWS that use multiple sets of time series data corresponding to different variables of the same system [33, 39, 99, 183, 65, 51]. For example, properties of the covariance matrix, such as the dominant eigenvalue [33] and the cross correlation [39] can provide a more reliable EWS than using just the variance of the time series of a single variable of the system. An aspect of this will be addressed in Chapter 5, where we demonstrate how using EWS from time series data of populations in their breeding and non-breeding habitats can warn of a transition to extinction and the process that is driving the extinction in the first place. Knowing the cause of population decline is required for developing effective intervention strategies. I believe there are many opportunities ahead with regards to the development of EWS in multi-variate time series data.

3.7 Limitations and prospects for EWS

Before EWS are used in an empirical setting, it is important to understand their limitations and recognise the circumstances under which they are appropriate. Many of these limitations may be overcome in time as the field progresses. The 'ideal' conditions for the use of EWS are as follows:

1. **Slow approach to the bifurcation.** The system needs to be in a so-called quasi-stationary state whereby the bifurcation parameter varies slowly with respect to the characteristic timescale of the system. Formally $\dot{\alpha} \ll 1/|\lambda|$ where λ is the local recovery rate and $\alpha(t)$ is the slowly varying parameter [203]. Under conditions where a parameter varies faster than the recovery time, the system may undergo rate-induced tipping [5], which need not be associated with a bifurcation. Early warning signals for rate-induced tipping are a current avenue of research [161].

2. **Small noise amplitude.** Noise must be small relative to the width of the basin of attraction, otherwise the system risks undergoing a purely noise-induced transition, which is not preceded by EWS [22]. Moreover, conventional EWS are based on

the behaviour of linearised dynamics, which are only a good approximation for small noise. However recent research is developing EWS based on large perturbation theory [119, 75], that applicable to systems with large noise and have the advantage of seeing a larger portion of the stability landscape, therefore obtaining more information about possible bifurcations.

3. **A local bifurcation.** Local bifurcations are preceded by critical slowing down since they are always accompanied by a loss in local stability [181]. However, critical transitions may occur due to global bifurcations whereby a larger invariant set such as a periodic orbit or a chaotic attractor collides with the current equilibrium. Such transitions are not preceded by conventional EWS [87], although there may be EWS not based on critical slowing down that could warn of such transitions [100, 203]. Another limitation with conventional EWS is the inability to determine the type of local bifurcation that is approaching, which is something this thesis addresses.

4. **Measurement of an appropriate variable.** Critical slowing down occurs along some direction of phase space (given by the dominant eigenvector(s) of the Jacobian matrix), and so does not necessarily get manifested in all system variables [18]. Determining which variables are appropriate to provide EWS in multi-dimensional systems is an active area of research [137, 51, 172, 199].

5. **White, additive noise.** Most modelling studies of EWS assume additive white noise for simplicity. However, environmental noise often contains some temporal correlation and scales with the size of the system, making it multiplicative. Multiplicative noise has been shown to distort EWS in some systems [56], as can correlated noise [152]. More work is needed know which indicators perform best under these scenarios, and which ones fail.

6. **High resolution data.** For EWS to be detected in time series data, the spacing between data points should be small compared to that of the recovery time and period of intrinsic oscillations if they exist [55]. Mathematically we should have $\Delta t < 1/|\lambda|$, $1/\mathrm{Im}(\lambda)$ where Δt is the spacing between data points and λ is the local recovery rate. Fortunately methods for data collection are improving across fields such as ecology, medicine and climate science, making the prospect of EWS more favourable.

Unfortunately, field data in ecology rarely adheres to the above limitations for the ideal detection of EWS. Therefore, the majority of empirical studies that have successfully de-

tected EWS prior to transitions have either been in disciplines with an abundance of data, such as engineering [78] and medicine [129], or in ecological studies that involve a controlled experimental setup [61, 50, 38]. A challenge with ecological data is that it is often subject to noise that is both multiplicative and correlated in time [147], which falls outside the ideal conditions for EWS, at least for now. Studies with ecological field data have found this noise, along with the lack of high frequency data, to hamper traditional EWS [147, 81]. However, we argue that there is plenty of room for optimism.

Theoretical and *in silico* studies on the effect of correlated [152, 207, 63] and multiplicative [143, 56, 173] noise are emerging. While these studies caution that these properties of noise can distort EWS, they also find robustness in certain indicators, and point to the need for EWS that are tailored to the system at hand, given knowledge of the type of noise occurring in the system. Moreover, methods are being developed, including a contribution from this thesis (Chapter 4), that better cope with multiplicative noise when predicting bifurcations. Perhaps the greatest restriction on the use of of EWS is the requirement for high resolution data, however massive datasets are becoming more available with advances in techniques such as remote sensing and automatic measuring schedules. As such, the possibility of using EWS to predict critical transitions in real complex systems is becoming ever more concrete. It will be important that theoretical advances accompany this influx of data, in order to maximise its potential for describing the fate of a system. Our first theoretical advance in the field of EWS is presented in the following chapter.

Chapter 4

Detecting and distinguishing tipping points using spectral early warning signals

This chapter is based on the paper: Thomas M Bury, Chris T Bauch, and Madhur Anand. Detecting and distinguishing tipping points using spectral early warning signals. *In review*, 2019.

4.1 Abstract

Theory and observation tell us that many complex systems exhibit tipping points–thresholds involving an abrupt and irreversible transition to a contrasting dynamical regime. Such events are commonly referred to as critical transitions. Current research is focusing on the development of early warning signals (EWS) of critical transitions that could help prevent undesirable events such as ecosystem collapse. As of yet, EWS are generic, and therefore do not indicate the type of transition. For instance, conventional EWS may fail to distinguish non-critical transitions, such as the onset of oscillations (e.g., Hopf bifurcation) from critical ones (e.g., Fold bifurcation). Moreover, EWS are less reliable in systems with density-dependent noise. Here, we use Ornstein-Uhlenbeck theory to knit together different signals present in the power spectrum to create 'spectral EWS', which provide greater sensitivity to transition proximity and therefore more advanced warning, higher robustness to density-dependent noise, and clues to the type of approaching transition. We demonstrate the advantage of using spectral EWS in concert with conventional indicators using a population model and data from a predator-prey chemostat experiment. The ability to better infer and differentiate the nature of upcoming transitions in complex systems will help humanity manage critical transitions in the Anthropocene.

4.2 Introduction

The understanding that complex systems can possess thresholds marking a sudden shift to an alternative dynamical regime has been around for a long time (e.g in ecology [124, 93]). Such a threshold may be referred to directly as a tipping point/catastrophic bifurcation, or by its inferred dynamics, as a critical transition/regime shift. Predicting tipping points and their ensuing dynamics remains a significant challenge, since the observable state of a system may show little change right up until it is too late. Even where data is abundant, parameterised models based on biological principles are rarely able to pinpoint tipping points due to uncertainty in system parameters and mechanisms. However, a new wave of research is targeting stochasticity as a possible treasure trove of information on the otherwise hidden, and often surprising, dynamics of complex systems [20, 43].

A significant development in this area is that of early warning signals (EWS), which are a suite of statistical metrics that are expected to undergo observable change prior to a tipping

point [168, 55]. Most EWS are grounded on the phenomenon of 'critical slowing down', which is a generic feature of local bifurcations [181]. It involves the degradation of restoring forces along some dimension of the system's state space, resulting in a longer return time to equilibrium following a perturbation. In stochastic systems, this manifests as an increase in variance [32], higher correlations in time [89] and space [54], lower frequencies in the power spectrum [106], and notable changes in several other statistical metrics [43].

The generality of critical slowing down is both a blessing and a curse. On one hand, it allows EWS to be applied to a wide range of systems including socio-ecological [10, 146], neurological [129], financial [60] and climate [19] systems. On the other hand, bifurcations come in assorted forms which each possess their own unique dynamics [114], such as a smooth transition to an intersecting state, the onset of oscillations, or an abrupt departure to a far away attractor (e.g. Transcritical, Hopf, Fold, respectively). These local bifurcations are all accompanied by critical slowing down [181] and so cannot be distinguished by the conventional critical slowing down methodology [102]. As such, EWS that are specific to each type of bifurcation are required to predict ensuing dynamics.

Moreover, the majority of EWS research has revolved around the Fold bifurcation (associated with a critical tipping point), due to its presence and adverse consequences in many simple models. However, we show from random matrix theory (Appendix A) that high-dimensional systems are likely candidates for bifurcations with an oscillatory component (e.g., Hopf), due to their approximately 'circular' distribution of eigenvalues in their Jacobian.

The theory of stochastic processes has much to offer the comparatively new field of EWS [110, 73, 20, 143], though most studies do not go beyond the rules of thumb (provided by critical slowing down) to seek more specific stochastic signatures. This is well-illustrated by the general framework of a system of variables $\vec{s}$ that evolves in time according to

$$d\vec{s} = f(\vec{s})dt + \sigma d\vec{W}(t), \tag{4.1}$$

where f captures the within-system dynamics and $d\vec{W}(t)$ is a vector of Wiener processes representing environmental noise with amplitudes given in the matrix σ. For relatively small noise, the dynamics about an equilibrium state are well approximated by

$$d\vec{x} = A\vec{x}dt + \sigma d\vec{W}(t) \tag{4.2}$$

where $\vec{x}$ is the deviation of the state from equilibrium and A is the Jacobian matrix of

f at equilibrium, which contains the local interaction terms between the variables. This process, originally studied in physics to model Brownian motion [188], is an Ornstein-Uhlenbeck process (OUp), for which general statistical properties can be derived [73, 56]. Therefore, given a system that fits into this framework, analytical expressions for EWS can be derived in terms of system parameters (within A) and relative noise strengths and correlations (within σ). This way, one can move beyond generic indicators of an upcoming transition (which do not always behave as expected), towards more reliable indicators that are specific to the system being modelled. This analytical approach has been adopted in previous studies to investigate EWS behaviour in particular models [136, 158], and under different regimes of noise [143, 152]. It has also been used to study different bifurcations [143, 201], however a complete set of EWS approximations prior to each type of local bifurcation is lacking.

Here, we use the theory behind the OUp (Eqn. [1]) and the Vector autoregression (discrete-time analogue of the OUp) [26] to derive EWS approximations specific to each type of local bifurcation in discrete and continuous-time systems (Supplementary Table 1). This provides a more complete framework for which EWS to expect preceding each type of bifurcation. The autocorrelation function and power spectrum are of particular interest, as they possess certain distinguishing features of bifurcations that are not exploited in the conventional application of EWS. In deriving metrics to exploit these features, we focus on the power spectrum, due to its intuitive properties and rich history in time series analysis [151].

The power spectrum describes the relative importance of the various frequency components within a time series [151]. It is a useful tool to quantify underlying frequencies of a system that may arise from both external and internal processes [127, 156]. It is also a complete representation of a (stationary) time series, and so contains a great deal more information than conventional EWS, that generally constitute a single metric. An important question then is: how can this information be harnessed to inform a time series observer about upcoming bifurcations? Previous work on this matter includes an early study from the physics literature [201], which shows how the power spectrum varies prior to instabilities of periodic orbits. Later studies have harnessed the phenomena of 'spectral reddening' [106, 15], a trend towards lower frequencies in the spectrum, to anticipate a Fold bifurcation. This trend can be captured by the 'spectral ratio' [15] a metric which considers the ratio of power at a low frequency to power at a high frequency, and therefore increases prior to a Fold bifurcation. However, this metric is subjective as the relative frequencies must

be chosen in advance. Another metric derived from the power spectrum is the power law decay rate, which again, captures the effect of spectral reddening [150]. We use the term spectral EWS to refer to metrics that have been derived from the power spectrum to detect upcoming bifurcations.

This chapter is arranged as follows. We begin by outlining important observations from the analytical derivations of standard EWS prior to each bifurcation. We then introduce two 'spectral EWS', that have been inspired from the analytical derivations, and explain how they extend conventional EWS in detecting bifurcations. We then demonstrate their application to a harvesting model and an empirical dataset, that contain various bifurcations. Finally, we discuss the broader implications of these techniques, and in which circumstances they are best applied, relative to alternate EWS techniques.

4.3 Theory

4.3.1 Deriving bifurcation-specific EWS

We derive bifurcation-specific EWS using the normal form of each bifurcation [114] with additive white noise. This provides a reasonable approximation to how EWS should behave prior to bifurcations in general since, near any bifurcation, there exists a mapping to its corresponding normal form [114]. Here, we show the derivation of EWS prior to the Fold bifurcation in discrete and continuous time for illustration of the techniques used. Derivations for the other local codimension-1 bifurcations are provided in Appendix B. The list of analytical expressions for EWS in each case are provided in Table 4.1, and visual display of their behaviour is provided in Figure 4.1.

Continuous-time framework. In a continuous-time framework, the normal form of the Fold bifurcation with additive white noise is given by

$$\frac{dx}{dt} = \alpha - x^2 + \sigma\xi(t), \tag{4.3}$$

where α is the bifurcation parameter, $\xi(t)$ is a white noise process and σ is the noise amplitude. The Fold bifurcation occurs at $\alpha = 0$. For sufficiently small σ, dynamics about

the stable equilibrium branch ($x = \sqrt{\alpha}$) may be approximated by the linearised form

$$\frac{d\eta}{dt} = \lambda\eta + \sigma\xi(t) \tag{4.4}$$

where $\eta = x - \sqrt{\alpha}$, $\lambda = -\sqrt{\alpha}$, and we have dropped terms of order $O(\eta^2)$. Note that the recovery rate $\lambda \to 0^-$ as the bifurcation is approached, which corresponds to critical slowing down. Equation (4.4) is an Ornstein-Uhlenbeck process, for which statistical properties may be derived, and serve as EWS for the bifurcation. Integrating the expression gives

$$\eta(t) = \eta(0)e^{\lambda t} + \sigma e^{\lambda t}\int_0^t \xi(t')e^{-\lambda t'}dt', \tag{4.5}$$

where $\eta(0)$ is the initial position of the process. Since we are interested in stationary dynamics about equilibrium, we may safely drop the transient term here (note that it decays since $\lambda < 0$). This yields

$$\eta(t) = \sigma e^{\lambda t}\int_0^t \xi(t')e^{-\lambda t'}dt'. \tag{4.6}$$

Note that in the limit as the bifurcation is approached ($\lambda \to 0^-$), the restoring forces vanish resulting in a random walk

$$x(t) = \sigma\int_0^t \xi(t')dt', \tag{4.7}$$

at which point the stationarity assumption breaks down. Therefore, these expressions lose their validity at the bifurcation point. We now compute the autocovariance function of $\eta(t)$, which is defined as

$$\phi(\tau) = \langle\eta(t)\eta(t+\tau)\rangle \tag{4.8}$$

where τ is a lag time. Using the expression (4.5) we obtain

$$\phi(\tau) = \sigma^2 e^{2\lambda t+\lambda\tau}\int_0^t dt'\int_0^{t+\tau} dt''\langle\xi(t')\xi(t'')\rangle e^{-\lambda(t'+t'')}. \tag{4.9}$$

Since white noise is uncorrelated, it satisfies [73]

$$\langle\xi(t')\xi(t'')\rangle = \delta(t'-t'') \tag{4.10}$$

where δ is the Dirac Delta function. The autocovariance function may then be simplified and integrated to give

$$\phi(\tau) = \frac{-\sigma^2}{2\lambda} e^{\lambda\tau}(1 - e^{2\lambda t}) \tag{4.11}$$

which, in the stationary regime ($t \to \infty$) is

$$\phi(\tau) = \frac{-\sigma^2}{2\lambda} e^{\lambda\tau}. \tag{4.12}$$

This yields the familiar expressions for the variance and autocorrelation of the OUp, given by

$$\mathrm{Var}(\eta) = \phi(0) = \frac{-\sigma^2}{2\lambda} \tag{4.13}$$

and

$$\rho(\tau) = \frac{\phi(\tau)}{\phi(0)} = e^{\lambda\tau}. \tag{4.14}$$

These expressions show that as the bifurcation is approached, the variance increases unboundedly (until the approximations are no longer valid) and the autocorrelation increases to unity.

The power spectrum of the process can be computed using the Wiener-Khinchin Theorem, which states that for a stationary stochastic process, the power spectrum is the Fourier Transform of the autocovariance function [26]. We use the convention

$$S(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{i\omega\tau} \phi(\tau) d\tau, \tag{4.15}$$

which for a real stochastic process (and given $\phi(\tau) = \phi(-\tau)$) may be reduced to

$$S(\omega) = \frac{1}{\pi} \int_{0}^{\infty} \cos(\omega\tau) \phi(\tau) d\tau. \tag{4.16}$$

Inserting the autocorrelation from Eqn. (4.12) we get

$$S(\omega) = -\frac{\sigma^2}{2\lambda} \frac{1}{\pi} \int_{0}^{\infty} \cos(\omega\tau) e^{\lambda\tau} d\tau \tag{4.17}$$

which integrates to give

$$S(\omega) = \frac{\sigma^2}{2\pi} \frac{1}{\omega^2 + \lambda^2}. \tag{4.18}$$

In certain instances it is useful to consider the power spectral density (PSD), which is based on the spectrum of autocorrelations $\rho(\tau)$ rather than the spectrum of autocovariances $\phi(\tau)$ [26]. This is equivalent to normalising the power spectrum by the variance of the process, meaning that any changes in the power spectral density are a result of changes in the frequency *distribution*, and not simply higher variance in the system. Thus the power spectral density is

$$\tilde{S}(\omega) = \frac{1}{\pi} \frac{-\lambda}{\omega^2 + \lambda^2} \tag{4.19}$$

which is in fact a Cauchy distribution with scale $|\lambda|$.

Discrete-time framework. It is important to also assess the behaviour of EWS in a discrete-time framework, since discrete models can display a range of behaviour not exhibited by the corresponding continuous-time models. A striking example is the onset of period-doubling (Flip) bifurcations to chaos in the discrete logistic model [126].

The normal form for the Fold bifurcation in a discrete-time framework is

$$x_{t+1} = x_t + \alpha - x_t^2 + \sigma\epsilon_t \tag{4.20}$$

where now ϵ_t is a normal random variable of mean zero and unit variance. The bifurcation occurs $\alpha = 0$. Linearising about the stable equilibrium branch ($x = \sqrt{\alpha}$) yields

$$\eta_{t+1} = \lambda\eta_t + \sigma\epsilon_t \tag{4.21}$$

where $\eta_t = x_t - \sqrt{\alpha}$, and $\lambda = 1 - 2\sqrt{\alpha}$. Recall that in discrete-time system, local bifurcations occur when the eigenvalue of the equation dynamics reaches a magnitude of 1 (crosses the unit circle in the complex plane). The process in equation (4.21) is an AR(1) process, of which statistical properties can be derived [26]. In higher dimensions, this would be a VAR process, which is analysed in Appendix B.

The variance of this process can be computed by squaring both sides of (4.21) and taking

	Bifurcation	Dominant eigenvalue(s)	Variance	Lag-τ AC, $\rho(\tau)$	Power Spectrum, $S(\omega)$
Continuous-time $t, \tau \in \mathbb{R}$	Fold, TC, PF	$\lambda \in \mathbb{R}, \quad \lambda \to 0^-$	$-\frac{\sigma^2}{2\lambda}$	$e^{\lambda\lvert\tau\rvert}$	$\frac{\sigma^2}{2\pi}\left(\frac{1}{\omega^2+\lambda^2}\right)$
	Hopf	$\lambda_{1,2} = \mu \pm i\omega_0, \quad \mu \to 0^-$	$-\frac{\sigma_1^2}{2\mu}$	$e^{\mu\lvert\tau\rvert}\cos\omega_0\tau$	$\frac{\sigma_1^2}{4\pi}\left(\frac{1}{(\omega-\omega_0)^2+\mu^2}+\frac{1}{(\omega+\omega_0)^2+\mu^2}\right)$
Discrete-time $t, \tau \in \mathbb{Z}$	Fold, TC, PF	$\lambda \in \mathbb{R}, \quad \lambda \to 1^-$	$\frac{\sigma^2}{1-\lambda^2}$	$\lambda^{\lvert\tau\rvert}$	$\frac{\sigma^2}{2\pi}\left(\frac{1}{1+\lambda^2-2\lambda\cos(\omega)}\right)$
	Flip	$\lambda \in \mathbb{R}, \quad \lambda \to -1^+$	$\frac{\sigma^2}{1-\lambda^2}$	$\lambda^{\lvert\tau\rvert}$	$\frac{\sigma^2}{2\pi}\left(\frac{1}{1+\lambda^2-2\lambda\cos(\omega)}\right)$
	Neimark-Sacker	$\lambda_{1,2} = re^{\pm i\theta}, \quad r \to 1^-$	$\frac{\sigma_1^2}{1-r^2}$	$r^{\lvert\tau\rvert}\cos(\theta\tau)$	$\frac{\sigma_1^2}{4\pi}\left(\frac{1}{1+r^2-2r\cos(\omega-\theta)}+\frac{1}{1+r^2-2r\cos(\omega+\theta)}\right)$

Table 4.1. Analytical approximations for EWS preceding each local, codimension-1 bifurcation. Derivations come from the normal form of each bifurcation [114] subject to additive Gaussian white noise. Expressions for each statistic are derived using techniques from stochastic process theory [73], and provided in Appendix B. The derivations rely on the assumption of small noise (only the linear dynamics preceding the bifurcation are considered), and stationarity (the bifurcation must be approached slowly enough such that time-series within a sufficiently large rolling window may be considered stationary). Shorthand notation includes TC: Transcritical, PF: Pitchfork. The Hopf and Neimark-Sacker bifurcations may be either super or subcritical.

the expectation, giving

$$\langle \eta_{t+1}^2 \rangle = \lambda^2 \langle \eta_t^2 \rangle + 2\lambda\sigma\langle \epsilon_t \eta_t \rangle + \sigma^2 \langle \epsilon_t^2 \rangle. \tag{4.22}$$

Assuming the process is stationary, the second moment does not change over time, and so $\langle x_{t+1}^2 \rangle = \langle x_t^2 \rangle$. The white noise term is a normal random variable, independent of x_t and with unit variance, hence $\langle \epsilon_t x_t \rangle = 0$, and $\langle \epsilon_t^2 \rangle = 1$. The expression can then be rearranged to give the variance of the process

$$\mathrm{Var}(\eta) = \langle \eta_t^2 \rangle = \frac{\sigma^2}{1-\lambda^2}, \tag{4.23}$$

which diverges as the bifurcation is approached ($\lambda \to 1^-$), as with the continuous-time system

The lag-1 autocorrelation can be computed by multiplying (4.21) by η_t and taking expec-

tations. This yields

$$\langle \eta_{t+1}\eta_t \rangle = \lambda \langle \eta_t^2 \rangle + \sigma \langle \epsilon_t \eta_t \rangle \tag{4.24}$$

which rearranges to

$$\rho_1 = \frac{\langle \eta_{t+1}\eta_t \rangle}{\langle \eta_t^2 \rangle} = \lambda. \tag{4.25}$$

One can obtain the autocorrelation function by applying this technique recursively, to obtain

$$\rho_\tau = \lambda^{|\tau|}, \tag{4.26}$$

which increases monotonically to unity, similarly to the continuous-time case.

The power spectrum can be computed by taking the Discrete-Fourier-Transform of the autocovariance function

$$\phi_\tau = \langle x_t x_{t+\tau} \rangle = \frac{\sigma^2 \lambda^{|\tau|}}{1-\lambda^2}, \tag{4.27}$$

which gives

$$S(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} e^{-i\omega\tau} \phi_\tau \tag{4.28}$$

$$= \frac{1}{2\pi} \frac{\sigma^2}{1-\lambda^2} \sum_{\tau=-\infty}^{\infty} e^{-i\omega\tau} \lambda^{|\tau|}. \tag{4.29}$$

Noting that the infinite sum can be split into two geometric series, one obtains

$$S(\omega) = \frac{\sigma^2}{2\pi} \frac{1}{1+\lambda^2-2\lambda\cos(\omega)}. \tag{4.30}$$

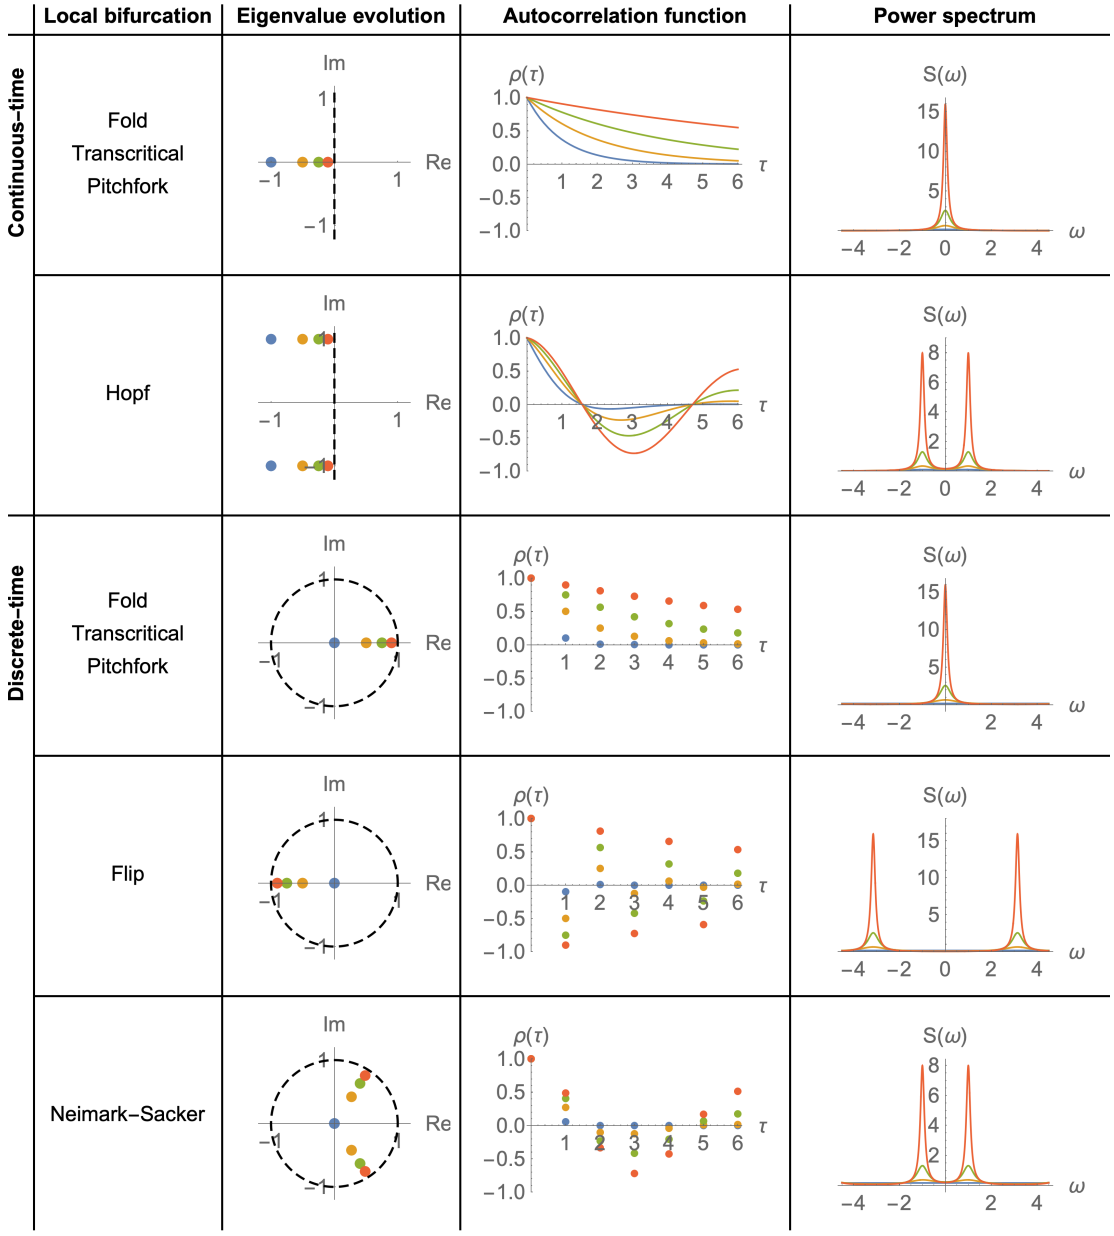

Figure 4.1. Graphical representation of EWS prior to each local codimension-1 bifurcation. We show the analytical form of the autocorrelation function $\rho(\tau)$ and power spectrum $S(\omega)$ at different proximities to each bifurcation given by eigenvalues of the Jacobian matrix. For the continuous-time bifurcations, we show $\mathrm{Re}(\lambda) = \{-1, -0.5, -0.25, -0.1\}$. For the discrete-time bifurcations we display $|\lambda| = \{0, 0.5, 0.75, 0.9\}$. Other parameters include the frequency at the Hopf/Neimark-Sacker bifurcation ($\omega_0 = 1$), and the noise amplitudes ($\sigma_1 = \sigma_2 = 1$).

4.4 Results

4.4.1 Insights from analytical expressions

The analytical expressions for EWS reveal characteristic features that can be used to distinguish certain bifurcations, such as the Fold and Hopf bifurcation (Figure 1). Specifically, local bifurcations without an oscillatory component (Fold/Transcritical/Pitchfork) yield increasing autocorrelation, conforming to the expected behaviour of critical slowing down (Figure 1b). In contrast, local bifurcations with an oscillatory component (Hopf/Flip/Neimark-Sacker), yield a trend in autocorrelation that depends on the relationship between the lag time τ, and the underlying period of oscillations T (Figure 1e). At sufficiently low lags ($\tau < T/4$), autocorrelation increases as the bifurcation is approached. However, for lag times closer to half the period of oscillations, the autocorrelation decreases. This can be understood intuitively by noting that at a lag time of $T/2$, one is computing the correlation between peaks and troughs of the underlying frequency, which become more pronounced as the bifurcation is approached. Since these points occur on opposite sides of the trajectory mean, they possess negative correlation. This finding corroborates previous studies that have found decreasing autocorrelation in both empirical [78] and model [58] studies preceding Hopf bifurcations. Studies that have found increasing autocorrelation prior to Hopf bifurcations [102, 9] have used lag times much smaller than the period of oscillations. Generally lag-1 autocorrelation is used, in which case the lag time simply the time between successive data points.

The power spectrum perhaps provides the most intuitive picture of the approaching bifurcation. Prior to non-oscillatory bifurcations it trends to lower frequencies, which is a manifestation of critical slowing down. Prior to oscillatory bifurcations it trends to the frequency of oscillations that occur at the bifurcation. The fact that this frequency is observed in the power spectrum prior to the bifurcation is due to resonant amplification of stochasticity [127, 72, 21], which occurs in systems possessing damped oscillations subject to environmental or demographic noise. The advantage to having analytical expressions for the behaviour of these functions is that we can now construct metrics that exploit their distinguishing features for bifurcation detection.

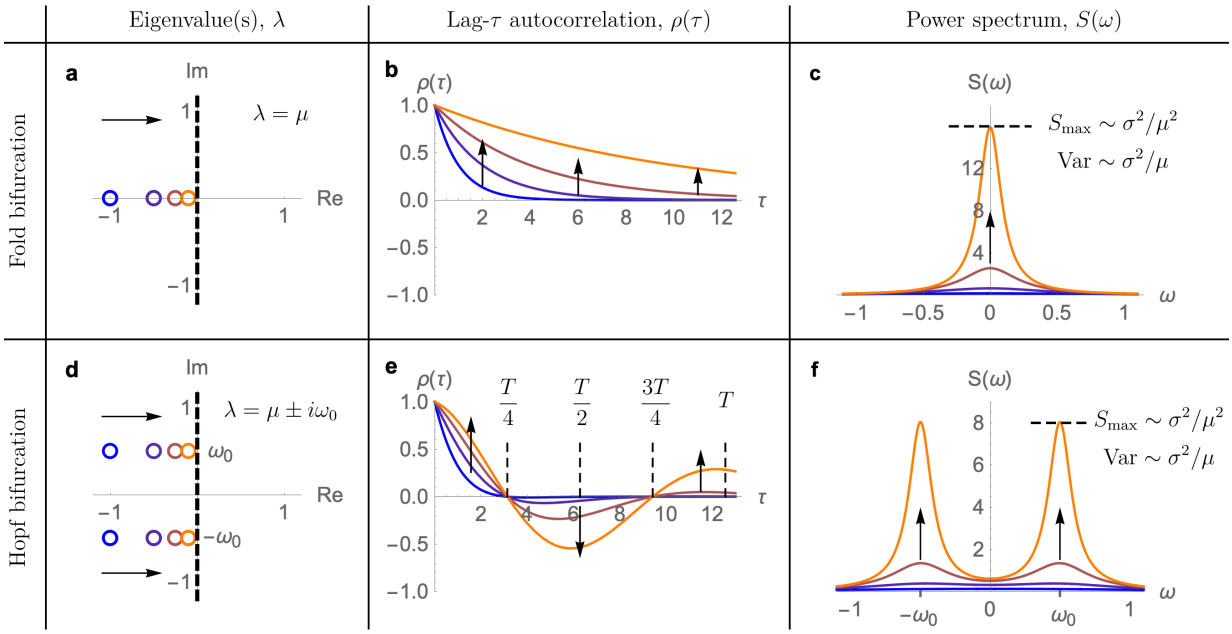

Figure 4.2. Contrasting EWS preceding the Fold and the Hopf bifurcation. Using theory from stochastic processes [73], we obtain analytical approximations EWS preceding each type of local codimension-1 bifurcation (Supplementary Table 1). Displayed are the analytical forms for the autocorrelation function and power spectrum that precede the Fold and Hopf bifurcation, respectively. They are plotted at instances where the real part of the dominant eigenvalue takes the values $\mu = \{-1, -0.5, -0.25, -0.1\}$. **a,** The Fold bifurcation involves a single real eigenvalue becoming positive. **b,** Autocorrelation at all lag times increases, although higher lags yield a delayed warning. **c,** The power spectrum trends to lower frequencies with a peak amplitude that increases asymptotically faster than the variance (as $\mu \to 0$). **d,** A Hopf bifurcation involves a complex-conjugate pair of eigenvalues obtaining positive real part. The imaginary part of the eigenvalues (ω_0) corresponds to the frequency of oscillations that occur at the bifurcation. **e,** The trend of autocorrelation depends on how the lag time compares with the underlying period of oscillations ($T = 2\pi/\omega_0$). **f,** The power spectrum trends to the underlying frequency of oscillations ω_0, with the peak amplitude increasing asymptotically as before. Fixed parameters are $\omega_0 = 0.5$, $\sigma = 1$. Derivation of analytical approximations are provided in Supplementary Methods.

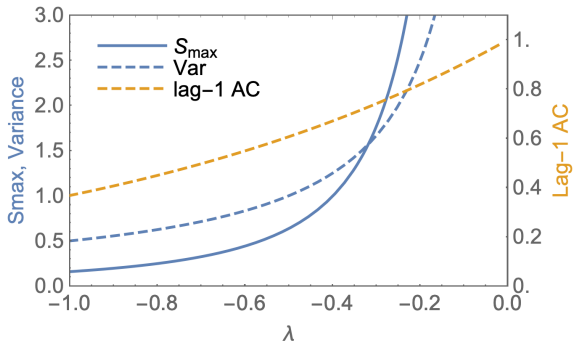

Figure 4.3. S_{max} is more sensitive to changes in bifurcation-proximity than variance and lag-1 autocorrelation. Shown are the analytical approximations for variance, $S_{\max}$ and lag-1 autocorrelation prior to the normal form Fold bifurcation, where λ is the eigenvalue of the corresponding Jacobian matrix. $S_{\max}$ has twice the relative rate of increase compared to variance, and exceeds the relative rate of increase of lag-1 autocorrelation as early as $\lambda = -2$.

4.4.2 Advantages of spectral EWS

We construct two spectral EWS that capture the important features of the power spectrum that relate to the bifurcations. First, we use the peak in the power spectrum ($S_{\max}$) as an indicator of bifurcation proximity. Like other EWS, it increases prior to local bifurcations due to critical slowing down. However its asymptotic behaviour in the approach to the bifurcation makes it advantageous, (Figure 4.3). Consider the analytical form for the power spectrum prior to the normal form Fold bifurcation. It goes like

$$S(\omega) = \frac{\sigma^2}{2\pi} \frac{1}{\omega^2 + \lambda^2}, \tag{4.31}$$

which has a peak at $\omega = 0$, taking the value

$$S_{\max} = \frac{\sigma^2}{2\pi\lambda^2} \sim \frac{1}{\lambda^2} \quad \text{as} \quad \lambda \to 0^-. \tag{4.32}$$

The asymptotic behaviour of variance however is

$$\text{Var} = -\frac{\sigma^2}{2\lambda} \sim \frac{1}{|\lambda|} \quad \text{as} \quad \lambda \to 0^-, \tag{4.33}$$

which has half the relative rate of increase of $S_{\max}$, i.e.

$$\frac{(\text{Var})'}{\text{Var}} = \frac{1}{2} \frac{(S_{\max})'}{S_{\max}} \tag{4.34}$$

where $'$ denotes differentiation with respect to λ. Thus, halving the distance to the bifurcation results in a two-fold increase in variance, but a four-fold increase in $S_{\max}$. Given that it is the *increase* in these metrics, that provide the EWS, we expect that $S_{\max}$ will serve as a more sensitive indicator to changes in bifurcation proximity. One can also compute the relative rate of increase in autocorrelation

$$\frac{(\text{lag-}\tau\text{ AC})'}{\text{lag-}\tau\text{ AC}} = \frac{\tau e^{\lambda\tau}}{e^{\lambda\tau}} = \tau \tag{4.35}$$

which is simply 1 for lag-1 AC. $S_{\max}$ exceeds a relative rate of increase of 1 as early as $\lambda = -2$. We show that this asymptotic behaviour of $S_{\max}$ holds prior to all local codimension-1 bifurcations in Appendix B.

Another advantage to this metric is its ability to still provide an early warning signal in systems where multiplicative muffles the early warning signal in variance [56]. Generalising the EWS expressions to systems with multiplicative noise (Appendix B) yields

$$\text{Var} \sim \frac{\sigma^2 \left[x^*(\lambda)\right]^2}{|\lambda|}, \quad S_{\max} \sim \frac{\sigma^2 \left[x^*(\lambda)\right]^2}{\lambda^2}, \quad \text{as } \lambda \to 0^-, \tag{4.36}$$

where $x^*(\lambda)$ is the equilibrium state of the system as a function of λ. Therefore, whether or not these metrics increase prior to the bifurcation, depends on how the equilibrium state of the system varies with λ. If the equilibrium state is invariant up to the bifurcation ($x^*(\lambda) = c$), such as the Flip bifurcation in the Ricker model, then clearly both variance and $S_{\max}$ will increase up to the transition. However, if the equilibrium state decreases up to the bifurcation (e.g. $x^*(\lambda) = c\lambda^{1/2}$, as for the Fold bifurcation), then Var $\sim \sigma^2$ and $S_{\max} \sim \sigma^2/\lambda$, i.e. $S_{\max}$ will provide an early warning signal and variance will not.

The second spectral EWS we propose comes from AIC weights [195]–a metric that determines the relative parsimony of a set of models fitted to a dataset–to determine which bifurcation the measured power spectrum corresponds to. The 'models' here are the analytical forms for the power spectra preceding each bifurcation, and a flat power spectrum to serve as a null model (white noise). These spectral EWS should be used together as they provide complementary information about an upcoming transition. $S_{\max}$ warns of an upcoming bifurcation, and the AIC weights provide information on the type of transition to expect by indicating which analytical form most parsimoniously fits the power spectra.

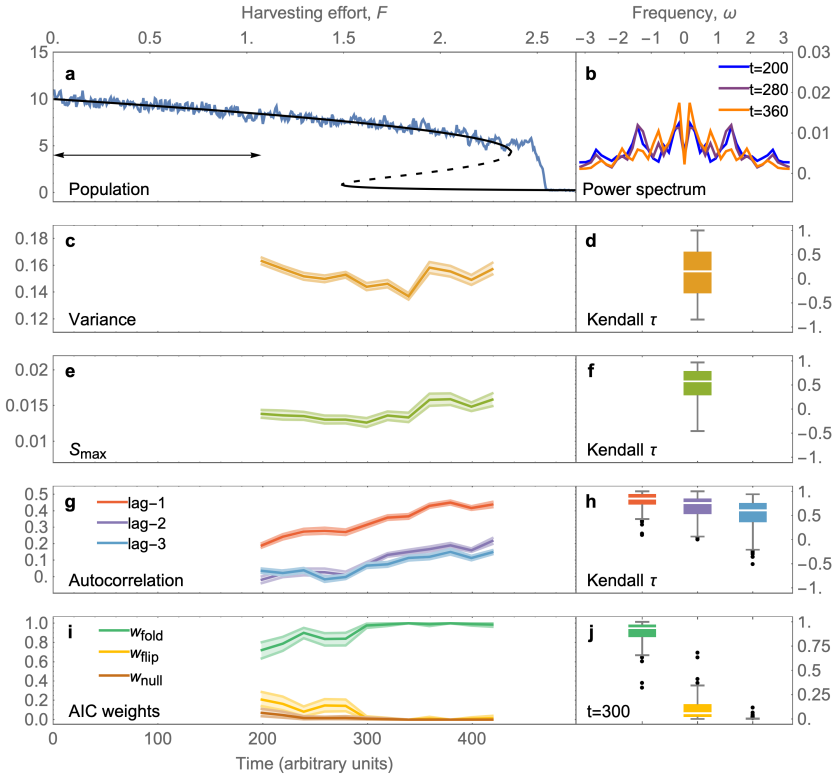

Figure 4.4. Characteristic early warning signals preceding the Fold bifurcation in the Ricker model. a, Simulated response of population abundance (blue) to increasing harvesting effort. Bifurcation diagram (black) of the deterministic model (solid - stable state; dashed - unstable state). The arrow illustrates the rolling window (40% of the time series) used for computing EWS. **b,** Power spectrum becomes dominated by lower frequencies over time signalling a potential Fold bifurcation. **c,** Variance provides no early warning due to density-dependent noise and a decreasing population abundance [56]. **d,** Kendall tau values from 100 realisations confirm no consistent trend in variance. **e,** $S_{\max}$ provides an early warning signal as early as $t = 300$ and **f,** displays an increasing trend for the majority of realisations. **g,** Autocorrelation provides early warning, **h,** with stronger signals at lower lag times. **i,** The AIC weight favours the Fold bifurcation, with increasing confidence once $S_{\max}$ begins to increase. The spectral metrics combined therefore provide a characteristic early warning signal. **j,** At time $t = 300$, $\mathrm{w}_{\mathrm{fold}}$ predicts the correct bifurcation for 98 of 100 realisations. Early warning signals are displayed as means with 95% confidence intervals over 100 bootstrapped samples from the time series in (**a**), after removing the trend with a Lowess filter. Kendall tau values are shown using box-whisker diagrams, where boxes mark the median and span the interquartile range and whiskers show the full range excluding outliers (black dots) from the ensemble of realisations.

4.4.3 Characteristic EWS in a model system

We test the spectral EWS on the classical Ricker model of a logistically growing population that is subject to exploitation (Methods), as it exhibits different bifurcations depending on the parameter settings [58]. Increasing harvesting effort yields a Fold bifurcation to a diminished population state, whereas increasing the intrinsic growth rate of the population (which can occur through size-selective harvesting [140]) yields a Flip (period-doubling) bifurcation to an oscillatory regime. Note that whereas the Fold bifurcation yields a critical transition, the Flip bifurcation yields a smooth transition. As such, distinguishing between the bifurcations in advance is important. The EWS are computed on bootstrapped samples from segments of the time series within a rolling window (Methods), which provides the error bars in the EWS metrics. We emphasise that these error bars are not from multiple simulations of the system, but from the single trajectory as shown. We do this since, in reality, one usually only has a single realisation to work with. To verify their trends over multiple simulations, we use Kendall tau values, which serve as a measure of increasing or decreasing trend. EWS without bootstrapping are shown in Appendix E, Figure B.7 and Figure B.8.

In the scenario undergoing the Fold bifurcation (Figure 4.4) the trajectory shows sign of slowing down (increased autocorrelation), but an overall decrease in variance, failing to warn of the critical transition. This is a known problem in systems that are subject to density-dependent noise and a decrease in population size prior to the bifurcation [56]. However, since $S_{\max}$ is a more sensitive metric to changes in bifurcation proximity than variance, it still provides a signal prior to the bifurcation. The AIC weights correctly identify the approaching Fold bifurcation and do so by time $t = 300$ for 98 of the 100 realisations.

The scenario undergoing the Flip bifurcation (Figure 4.5) looks somewhat similar when one only considers the time series trend prior to the transition. The spectral EWS, however, tell a very different story, in agreement with the theory. Note that a Flip bifurcation from equilibrium yields oscillations of period $T = 2$ and so one expects an underlying frequency of $\omega = 2\pi/T = \pi$ prior to the transition, as observed in the power spectrum. Correspondingly, autocorrelation decreases at lag-1 and increases at lag-2 (as expected from Figure 1e). $S_{\max}$ shows a marked increase (stronger than that of variance), and the AIC weights correctly identify the Flip bifurcation, well before it arrives. The spectral EWS together therefore provide a characteristic EWS on the Flip bifurcation that can be

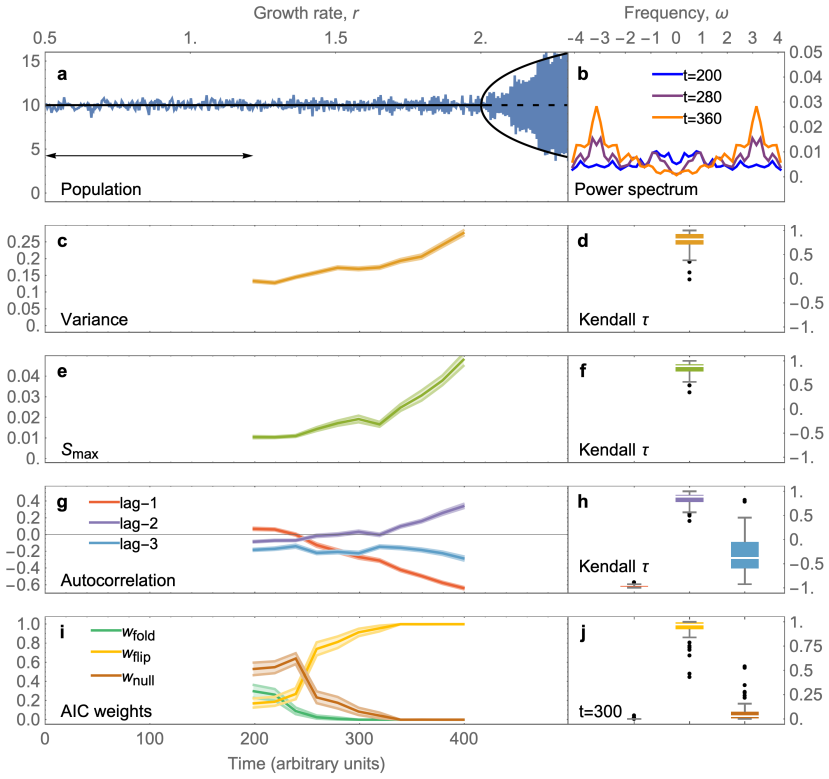

Figure 4.5. Characteristic early warning signals preceding the Flip bifurcation in the Ricker model. a, Simulated response of population abundance (blue) to increasing growth rate. Bifurcation diagram (black) of the deterministic model (solid - stable state/limit cycle; dashed - unstable state). **b,** Power spectrum becomes dominated by a non-zero frequency over time signalling a potential Flip bifurcation. **c** Variance provides an early warning signal, and **d,** is confirmed to have an increasing trend over an ensemble of 100 realisations. **e,** $S_{\max}$ yields an early warning signal with **f,** a stronger increasing trend than variance. Note that while variance increases two-fold, $S_{\max}$ increases four-fold, consistent with the asymptotic relationships of these metrics with the bifurcation parameter (Figure 1). **g,** Autocorrelation trend depends on the relationship between lag time and underlying period of oscillations ($T = 2$) (Figure 1). Lag-1 autocorrelation decreases, since in this case the lag is half of the period. **h,** Higher lags yield less pronounced signals on average. **i,** The AIC weight favours the Flip bifurcation once $S_{\max}$ begins to increase, which is well before the transition. The spectral metrics combined therefore provide a characteristic early warning signal. **j,** At time $t = 300$, $\mathrm{w}_{\mathrm{flip}}$ predicts the correct bifurcation for 98 of 100 realisations. Computation and display details are as in Figure 2.

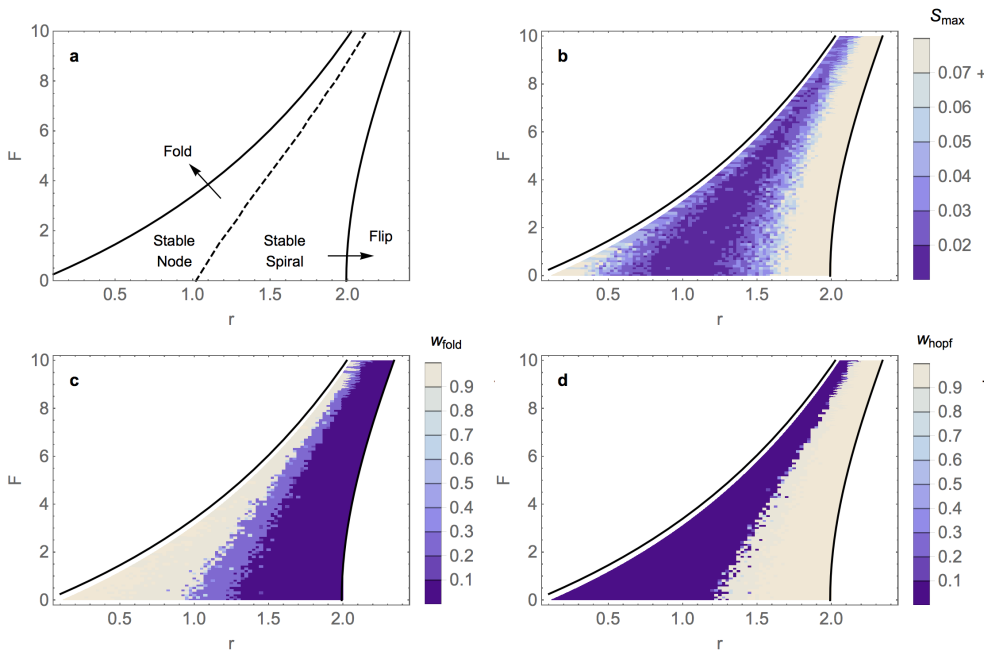

Figure 4.6. Spectral EWS across the parameter plane of the Ricker model. For fixed values of (r, F) within the stable region of the bifurcation plane, we simulate the Ricker model and compute the EWS metrics. **a.** Regions of stability and bifurcation curves. Crossing the Flip/Fold bifurcation curve results in a transition to oscillations/an alternative stable state. Perturbations in the stable node/stable spiral regime result in a monotonic/oscillatory recovery trajectory. **b.** Peak in the power spectrum $S_{\max}$. **c.** Fold AIC weight. **d.** Hopf AIC weight.

distinguished from the approaching Fold bifurcation.

To assess the efficacy of the EWS under drivers varying simultaneously, we assess the behaviour of the EWS over all combinations of the bifurcation parameters. We find that the spectral EWS distinguish the two bifurcations along any path in the plane of parameter values (Figure 4.6), with a sharp change in the AIC weights when the stable state transitions from a stable node to a stable spiral. The temporal EWS behave as expected for different combinations of parameter values (Figure B.9), with the coefficient of variation increasing prior to the both bifurcation thresholds and lag-1 autocorrelation increasing prior to the Fold threshold, but decreasing prior to the Flip bifurcation.

4.4.4 Characteristic EWS in empirical data

We use chemostat data from a predator-prey experiment conducted in a previous study [70] (Figure 4.11) to provide empirical validation of the spectral EWS. We also verify that the same EWS occur in the corresponding model (Appendix C). The authors showed from both experiments and a parametrised model that the system exhibits two Hopf bifurcations (H1, H2) as the dilution rate (controlling nutrient uptake) is varied (Figure 4.7a). The observed Hopf bifurcations occur at slightly higher dilution rates than the model predicts [70], highlighting the difficulty of locating bifurcations in real systems, and the importance of having additional predictive tools such as EWS.

We compute EWS for each experimental time series and find good agreement with theoretical expectations (Figure 4.7b-d). Preceding both Hopf bifurcations, the spectral EWS provide a characteristic warning of a Hopf bifurcation – $S_{\max}$ increases, and w_{Hopf} is the significant AIC weight. Moreover, the power spectrum in the pre-bifurcation regime of H1 ($\delta = 0.04$) shows a dominant frequency, $\omega_0 \approx 1/3$ (Figure 4.8) corresponding to an underlying period of oscillations, $T = 2\pi(1/3)^{-1} \approx 19$ days, which approximates the period observed in the oscillatory regime. This suggests that in addition to determining the type of bifurcation, the power spectrum can provide an early estimate of the period of oscillations at the bifurcation. Power spectra for every empirical time series are shown for *Chlorella* and *Brachionus* in Figure B.11 and Figure B.12, respectively.

4.5 Discussion

The ability to not only detect, but characterise bifurcations is essential for obtaining knowledge of upcoming qualitative changes in a system's dynamics. In this paper, we have derived analytical expressions for EWS prior to each type of local bifurcation, enabling us to construct spectral EWS that distinguish oscillatory from non-oscillatory bifurcations and provide a more sensitive warning of changes in bifurcation proximity. We demonstrated these tools with a well-known population model, showing how the onset of collapse can be distinguished from the onset of oscillations, and applied them to empirical data of a population exhibiting a Hopf bifurcation. This paper shows that spectral EWS offer complementary information to conventional EWS and should added to the repertoire of tools for predicting tipping points in real systems.

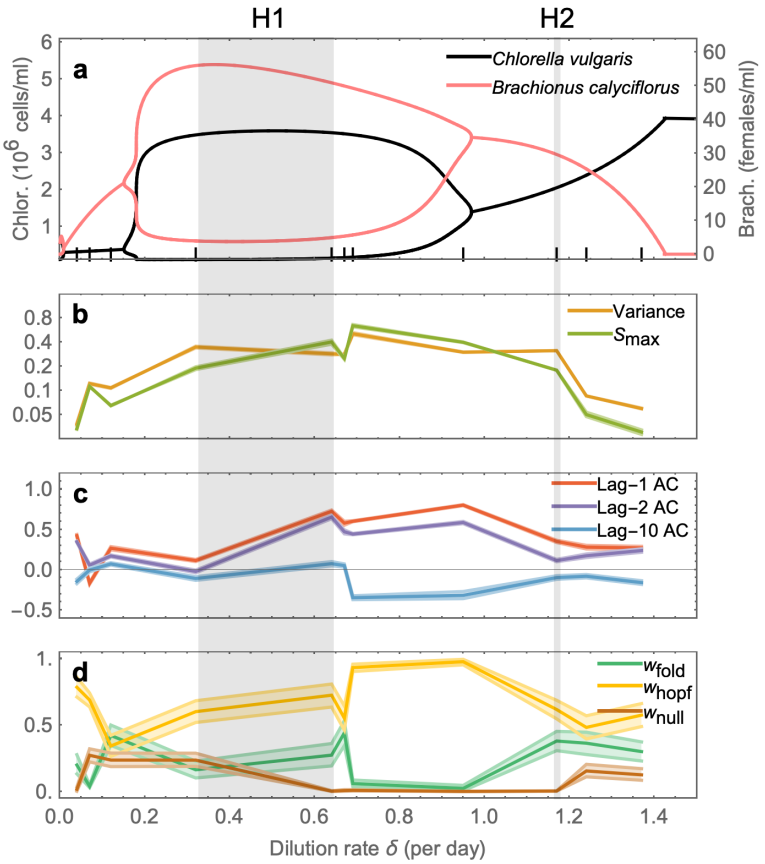

Figure 4.7. Characteristic early warning signals in an empirical predator-prey system. a, Bifurcation diagram of a parametrised model [70] for the predator-prey dynamics between *Brachionus calyciflorus* and *Chlorella vulgaris*. Lines show stable states/limit cycles, indicating two Hopf bifurcations. The Hopf bifurcations in the experiments (H1, H2) were observed within the grey regions, slightly higher dilution rates than predictions from the model [70]. Vertical black lines show the dilution rate of each chemostat experiment. **b,** Variance and $S_{\max}$ show increasing trends prior to the Hopf bifurcations. **c,** Autocorrelation at low (1,2) and high (10) lag times provides no clear early warning signal. **d,** The Hopf AIC weight dominates in the vicinity of the bifurcation, thus paired with $S_{\max}$, provides a characteristic early warning signal. Early warning signals are computed for 100 bootstrapped samples of each chemostat time series (fixed dilution rate), and displayed as means with 95% confidence intervals, across the samples and across the two species. Spectral EWS are also present in the the time series of each species considered separately (Figure B.10).

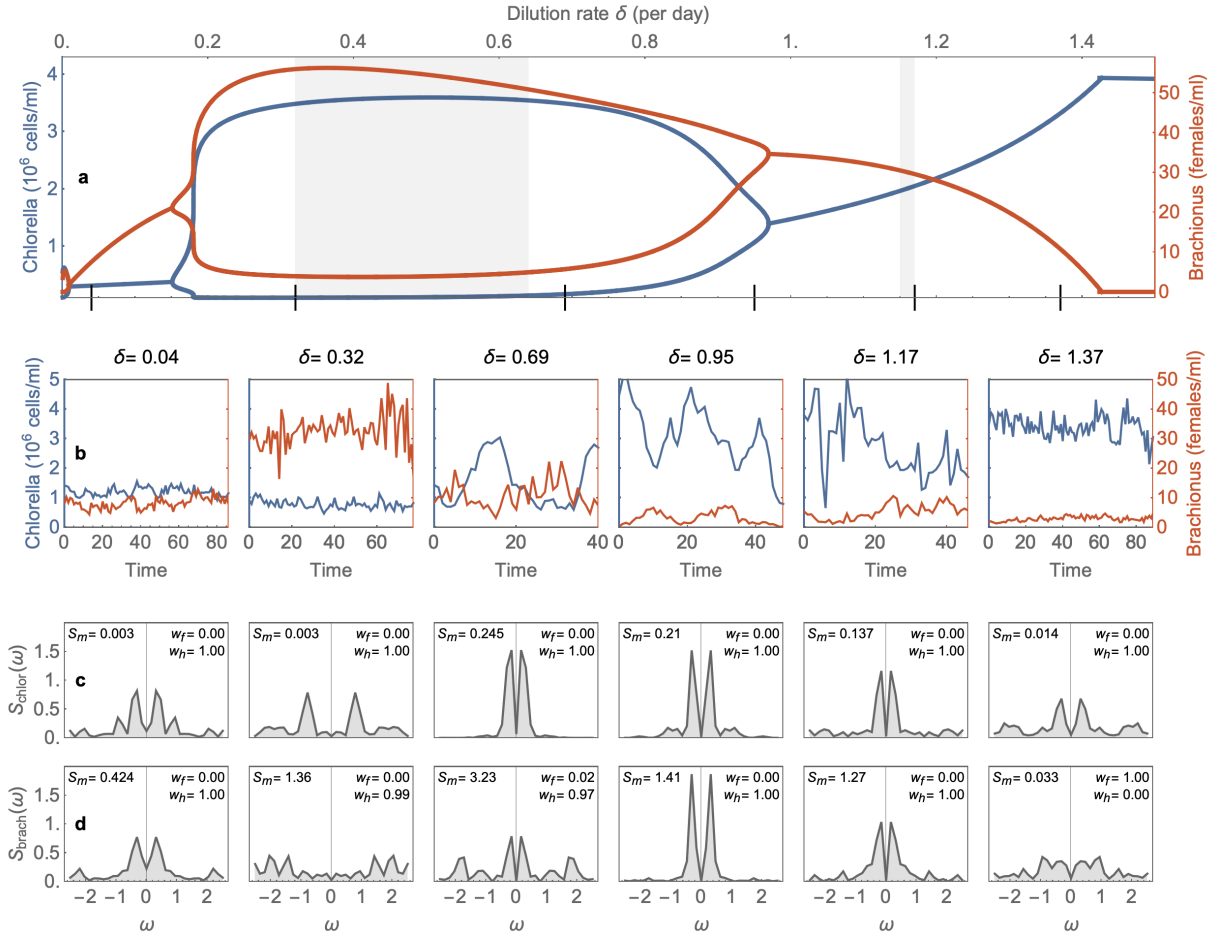

Figure 4.8. Spectral EWS for the empirical predator-prey system. a. Bifurcation diagram of the model system showing equilibria / extreme points of limit cycles for the prey (*Chlorella*, blue) and predator (*Brachionus*, red). The two grey regions mark where the two Hopf bifurcations occur in the empirical system. Black dashes mark the dilution rates shown in b. **b.** Empirical time-series data from chemostats with the given dilution rate (δ). **c-d,** Normalised power spectra of the time-series data in b for *Chlorella* and *Brachionus* respectively. Data insets give the peak power ($S_{\max}$), the Fold AIC weight (w_{f}) and the Hopf AIC weight (w_{h}).

To some extent, models provide information on the type of bifurcations that arise, however they are conditioned on a set of mechanistic assumptions, which are often debated, or at best, simplifications of reality. As such, bifurcations that occur in model systems may not be the same as the bifurcations that occur in reality. Model-free approaches such as the statistical EWS developed here, provide information on the bifurcations without assuming any underlying model, making them well suited to systems with uncertain mechanisms. For example, the transition from quiescence to the spiking of neurons in the mammalian cortex has been modelled as both a Fold and Hopf bifurcation, depending on the underlying model assumptions [129]. Computing spectral EWS of empirical data, one could discern which bifurcation the transition corresponds to, providing stronger validation for one model over the other and therefore learning something about the underlying mechanisms of the system.

There exist other recent approaches to tipping point detection that, under appropriate circumstances, show potential to predict specific bifurcations. Eigenvalue reconstruction [203] allows one to obtain an approximation of the Jacobian matrix from time series data. This method permits the monitoring of stability loss via a diminishing eigenvalue, and the underlying frequency via the eigenvalue's imaginary part (Figure 4.2 a,d). This approach, however requires time series data from multiple variables in order to make assertions on the type of bifurcation, whereas the power spectrum can be computed from an individual time series. Another approach considers the scaling of critical slowing down as a bifurcation is approached [129], however this requires a controlled setting whereby changes in the bifurcation parameter are able to be measured. Finally, promising studies in large perturbation theory [119, 75] demonstrate how information on the bifurcation type and distance can be obtained by carefully monitoring a system's recovery trajectory following a perturbation. This again requires a controlled environment, and a system whereby large perturbations are feasible.

The findings of this study have to be seen in light of some limitations, which could be addressed in future studies. First, the spectral EWS proposed only distinguish between oscillatory and non-oscillatory bifurcations, and not between bifurcations within these groups. For example, a supercritical and subcritical Hopf bifurcation, whilst giving rise to the same spectral EWS, result in structurally different transitions. Whereas the supercritical Hopf yields a smooth, reversible transition to small oscillations, the subcritical Hopf yields a critical, irreversible transition to a distant limit cycle. Distinguishing between such bifurcations requires consideration of nonlinear terms, which this study has assumed to be

negligible - a valid assumption for systems subject to small noise relative to the width of the basin of attraction. Future studies should investigate how nonlinear terms alter the power spectrum prior to a bifurcation, though this will require mathematical methods beyond the scope of this paper.

Second, this study only considers transitions arising due to local codimension-1 bifurcations. However, complex systems may lose stability via global bifurcations, which are not accompanied by critical slowing down, resulting in the failure of corresponding EWS [87, 23]. In such situations, EWS based on other dynamical features are required, such as critical attractor growth, which has been shown to precede interior crises of excitable systems [100]. Given that the spectral EWS here are derived from changes in local dynamics, they cannot be expected to provide warning of global bifurcations. In addition, bifurcations of a higher codimension (number of parameters that must be varied for the bifurcation to occur), such as the Bautin bifurcation, remain little explored in the EWS literature yet arise in many models (e.g [164]).

Third, spectral EWS, as with conventional EWS, may be distorted in cases of multiplicative noise, multiple scales of correlation and sparsely sampled data [147]. Further research should investigate the robustness of spectral EWS to these factors, and how the power spectrum varies in each case. For example, one would expect correlated environmental noise to create an additional peak in the power spectrum at the characteristic frequency of the noise. Moreover, statistical issues may arise in computing the power spectrum if the data is sparse relative to the timescale of dynamics. Finally, the work here does not consider spatial systems, where a power spectrum can be computed in both space and time. Investigating the behaviour of the power spectrum prior to different bifurcations in these systems may inspire further spectral EWS to predict and characterise transitions in spatial systems.

The spectral EWS developed here have several implications for science and policy. Competing hypotheses can lead to models with different types of bifurcations [197, 129]. Monitoring time-series data with spectral EWS can be used to infer particular models based on their bifurcation structure, and therefore deepen our understanding of the system itself. To avoid a regime shift, action must be taken well in advance of the bifurcation [15], requiring indicators that respond quickly to changes in bifurcation proximity. Incorporating $S_{\max}$ increases the likelihood of a sufficiently early warning due to its high sensitivity to bifurcation proximity. Moreover, the ability to distinguish oscillatory bifurcations (Hopf/Flip/Neimark-Sacker) from non-oscillatory ones is crucial, since their inferred dy-

namics produce contrasting outcomes. Our work here furthers methodology to learn useful information from stochasticity [20] and offers ready-to-go tools for further application.

4.6 Methods

4.6.1 Computing EWS from time series data

All procedures used to compute and analyse EWS in this paper may be implemented using the Python package *ewstools*, which is available and maintained at https://github.com/ThomasMBury/ewstools.

Temporal EWS. We compute temporal EWS including variance, coefficient of variation, autocorrelation, skewness and kurtosis according to common practices [55]. We initially detrend the data using a Lowess function with a specified span. The EWS metrics are then computed over a rolling window of size chosen such that the system can be considered approximately stationary within that window. Given the time series $\{x\}_{j=1}^{n}$, the variance is computed as

$$\text{Var} = \frac{1}{n}\sum_{j=1}^{n}(x_j - \mu)^2 \tag{4.37}$$

where $\mu = \sum_{j=1}^{n} x_j$, the mean value of the data. The coefficient of variation is the ratio of the standard deviation with the mean, thus

$$\text{CV} = \frac{\sqrt{\text{Var}}}{\mu}. \tag{4.38}$$

The autocovariance at lag τ is computed as

$$\phi(\tau) = \frac{1}{n}\sum_{j=1}^{n-\tau}(x_j - \mu)(x_{j+\tau} - \mu) \tag{4.39}$$

from which the autocorrelation at lag τ can be computed by normalising by the variance:

$$\rho(\tau) = \frac{\phi(\tau)}{\text{Var}}. \tag{4.40}$$

The sample skew of the time-series is computed as

$$\text{Skew} = \frac{\frac{1}{n}\sum_{j=1}^{n}(x_j - \mu)^3}{\left(\frac{1}{n}\sum_{j=1}^{n}(x_j - \mu)^2\right)^{3/2}} \tag{4.41}$$

and the kurtosis as

$$\text{Kurt} = \frac{\frac{1}{n}\sum_{j=1}^{n}(x_j - \mu)^4}{\left(\frac{1}{n}\sum_{j=1}^{n}(x_j - \mu)^2\right)^{2}} - 3. \tag{4.42}$$

Spectral EWS. To obtain spectral EWS, we first detrend the original time series and work with the residuals over a rolling window as before. We then approximate the power spectrum the data within the rolling window using Welch's method [200, 12], as illustrated in Figure 4.9. Welch's method involves computing the periodogram over many overlapping Hamming windows, and then taking the average. This gives a more consistent, unbiased estimate of the power spectrum than a single periodogram over the full length of the rolling window [200]. The periodogram is computed as $P(k) = |\tilde{x}(k)|^2$, where

$$\tilde{x}(k) = \frac{1}{\sqrt{n}}\sum_{j=1}^{n} x_j e^{\frac{2\pi i(j-1)(k-1)}{n}}, \tag{4.43}$$

is the Discrete Fourier Transform of $\{x_j\}_{j=1}^{n}$, the data points weighted by the Hamming window according to

$$\text{Hamming}(t) = \frac{21}{46}\cos\left(\frac{2\pi t}{n}\right) + \frac{25}{46}, \qquad |t| \leq \frac{n}{2}. \tag{4.44}$$

The peak in the power spectrum $S_{\max}$ is computed from the power spectrum $\{S(k)\}_{k=1}^{n}$ as

$$S_{\max} = \max_k S(k). \tag{4.45}$$

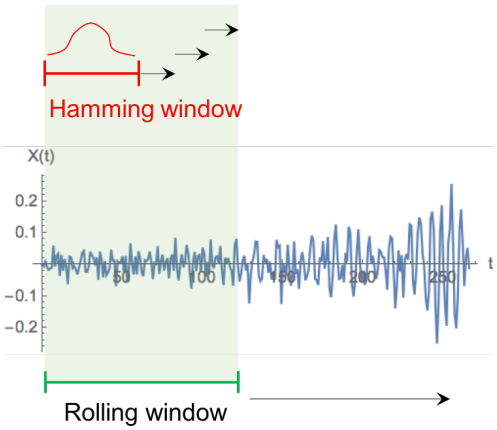

Figure 4.9. Computing the power spectrum over a rolling window. A rolling window moves incrementally across the residual time-series. At each increment, the power spectrum of the data within the window is estimated using Welch's method [200]. This involves computing the periodogram across several Hamming windows with a specified length and offset, and computing their average.

The AIC weights [2] are computed from the fitting of the canonical power spectrum forms

$$S_{\text{null}}(\omega;,\sigma) = \frac{\sigma^2}{2\pi}, \tag{4.46}$$

$$S_{\text{fold}}(\omega;\sigma,\lambda) = \frac{\sigma^2}{2\pi}\frac{1}{\omega^2+\lambda^2}, \tag{4.47}$$

$$S_{\text{hopf}}(\omega;\sigma,\mu,\omega_0) = \frac{\sigma^2}{4\pi}\left(\frac{1}{(\omega-\omega_0)^2+\mu^2}+\frac{1}{(\omega+\omega_0)^2+\mu^2}\right). \tag{4.48}$$

to the measured power spectrum, as illustrated in Figure 4.10. S_{null} is the power spectrum for white noise. Systems far from a bifurcation under a strong white noise regime have a relatively flat spectrum, best captured by this fit.

The models are fit to the measured power spectrum using a nonlinear optimisation algorithm (Levenberg–Marquardt) from the open-source Python package LMFIT [138]. Further details such as setting initialisation parameters details are provided in Appendix E. The AIC scores are given by

$$\phi_i = -2\text{Log}L_i + 2V_i, \qquad i = 1\ldots n \tag{4.49}$$

where L_i is the maximum likelihood for the candidate model i, and V_i is the number of free parameters that model i possesses. The model with the lowest AIC score is the one which statistically best explains the data. Note that models are penalised for having extra parameters, so although S_{hopf} can fit a unimodal power spectrum just as well as S_{fold} (by setting $\omega_0 = 0$) it obtains a higher AIC score due to its extra parameter. This procedure for computing AIC scores is based on the principle of entropy maximisation from information

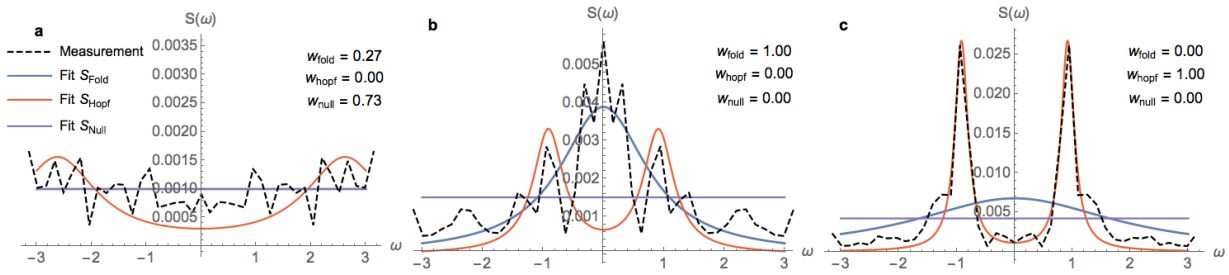

Figure 4.10. Illustration of computing AIC weights in different dynamic regimes. For each power spectrum measurement, the analytical forms S_{fold}, S_{hopf} and S_{null} are fitted using nonlinear optimisation procedures. The AIC weights for each fit are then computed (goodness of fit, relative to the other models). **a.** Equilibrium system far from a bifurcation with strong white noise. **b.** Power spectrum in the vicinity of a Fold bifurcation. **c.** Power spectrum in the vicinity of a Hopf bifurcation.

theory [2].

The AIC weights provide a normalised metric that discriminates between the models [195]. They are computed as follows. First, we obtain the difference between each score ϕ_i, and the score corresponding the model with the best fit [195]

$$\Delta_i = \phi_i - \min_i \phi_i. \tag{4.50}$$

The AIC weights are then determined by

$$w_i = \frac{e^{-\frac{1}{2}\Delta_i}}{\sum_{i=1}^{n} e^{-\frac{1}{2}\Delta_i}}. \tag{4.51}$$

Note that the AIC weights are normalised and so sum up to 1. Values of a weight closer to 1 favour the corresponding model.

Bootstrapping time series data. Prior to computing EWS, we use block-bootstrapping to generate an ensemble of samples. This involves detrending the data using a Lowess filter, followed by using a rolling window to obtain overlapping segments that can be considered approximately stationary. For each segment, we generate 100 bootstrapped samples. Each sample is built by selecting blocks of the time series randomly with displacement. The

block length is drawn from a geometric distribution [149], with an average large enough such that significant temporal correlations in the time series are retained.

4.6.2 Model system and empirical data

The Ricker model with harvesting. We use a Ricker-type model that describes the logistic growth of a population subject to harvesting [58]. The model reads

$$N_{t+1} = N_t e^{(r(1-N_t/K)+\sigma\epsilon_t)} - F\frac{N_t^2}{N_t^2+h^2} \tag{4.52}$$

where N_t is the population size at time t, r is the intrinsic growth rate, K is the carrying capacity, F is the maximum rate of harvesting, h is a half-saturation constant, σ is the noise amplitude, and ϵ_t is a normal random variable with zero mean and unit variance. Baseline parameters are $r = 0.75$, $K = 10$, $F = 0$, $h = 0.75$, $\sigma = 0.04$. The model exhibits a Fold bifurcation at $F = 2.36$, and a Flip bifurcation at $r = 2.00$ followed by a sequence of further Flip bifurcations to chaos.

We simulate two model scenarios, similar to a previous study [58]. In one, the harvesting rate F increases linearly over $[0, 2.7]$, resulting in a Fold bifurcation. In the other, the growth rate r increases linearly over the interval $[0.5, 2.3]$ resulting in a Flip bifurcation. All other parameters remain fixed. Both scenarios are simulated for 500 time-steps. Negative population values arising from noise are reset to zero.

Predator prey data

Chemostat data was available for 14 different dilution rates [70], shown in Figure 4.11. For computing EWS, we considered only time series with greater than 25 data points, of which there were 11. The data was initially detrended with a Lowess filter with an 80 day span to account for any unintentional drift in the dilution rate.

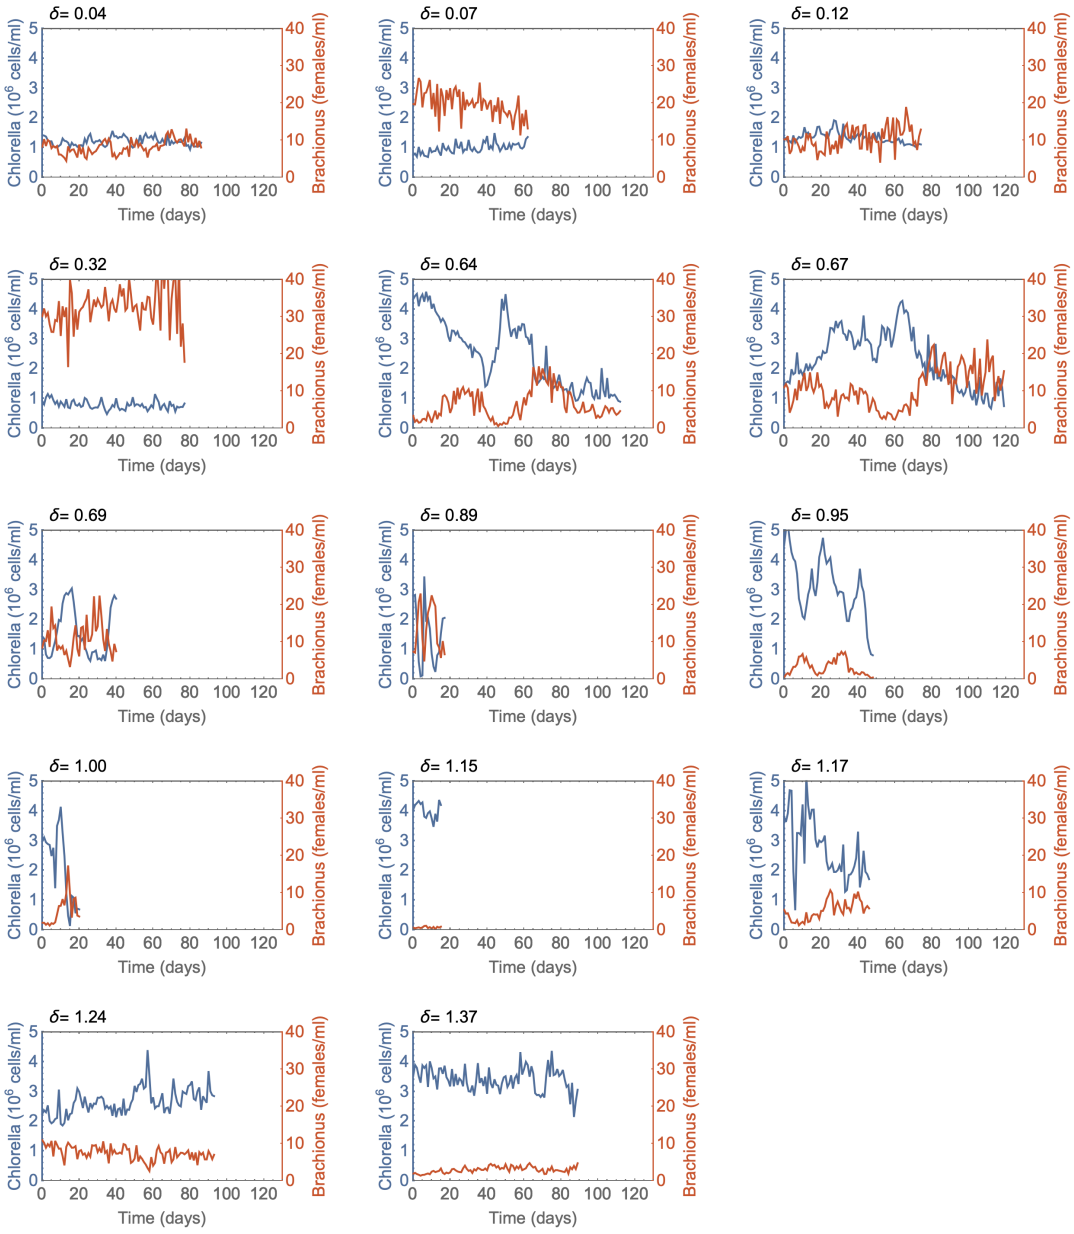

Figure 4.11. Predator-prey trajectories from chemostat experiments (adapted from Fussmann et al [70]). Time series data of the prey species (*Chlorella*, blue) and the predator species (*Brachionus*, red) from chemostat experiments run at different dilution rates (δ). The sampling frequency is one measurement per day. A Hopf bifurcation was conjectured to occur in the range $0.32 < \delta < 0.64$ and at $\delta \approx 1.16$, based on the system's transition from equilibrium to oscillatory behaviour either side of these thresholds.

Chapter 5

Early warning signals for population extinction in seasonal environments

This chapter is based on the paper in progress: Thomas M Bury, Joseph Burant, Ryan Norris, Chris T Bauch, and Madhur Anand. Early warning signals for population extinction in seasonal environments. *In progress*, 2019.

5.1 Abstract

Population decline and extinctions are occurring at an alarming rate across a number of different taxa. Predicting population extinction, and working out the drivers causing it, are therefore a research priority. A growing body of work is focused on the development of early warning signals for events such as population extinction, based on generic properties that a dynamical system exhibits prior to a bifurcation. These EWS have shown success at predicting population extinction in controlled experimental environments. However, little is known about how EWS behave in populations subject to seasonal variation, which is inherent to the majority of populations. Here, using a simple population model, we show that conventional EWS are still present under a seasonal environment, although higher order moments like skewness and kurtosis display interesting asymmetries depending on whether degradation occurs in the breeding habitat, or the non-breeding habitat. For example, time series from a population in its breeding period showed increasing skewness as habitat was degraded in the non-breeding period, but decreasing skewness upon habitat degradation in the breeding period. Such asymmetries could be used to infer the driver of the population decline, as well as being used alongside increasing variance and lag-1 autocorrelation as signals of approaching extinction.

5.2 Introduction

Environmental deterioration is a major threat to a large number of endangered bird, mammal and amphibian species [6]. Leading causes of deterioration include habitat loss and degradation [83], invasion of non-indigenous species [193], climate change [145] and pollution [105]. Estimating the risk of population extinction is therefore a central problem in theoretical ecology. However, most theoretical work has focused on extinction due to demographic and environmental stochasticity in stationary environments [144, 115, 122]. Given the current rates of environmental degradation in populations, improving methods to detect population extinction due to environmental degradation is a priority for research.

For the majority of biologically sensible models, extinction due to environmental degradation corresponds to the crossing of a bifurcation (e.g. [109]). In populations with an Allee effect, this bifurcation is often a Fold bifurcation, resulting a sharp transition from a significant population abundance to extinction. Without Allee effects the transition to extinction

is often a Transcritical bifurcation - a smooth decline in population abundance until there are no individuals left. As we have seen from Chapter 3, when a system approaches a bifurcation, it undergoes generic changes in its stochastic properties, referred to as early warning signals (EWS), that could in theory be used to warn of the transition [168]. Recent studies have developed and tested EWS on experimental populations [50, 61, 194], showing their success at providing an advanced warning of population extinction in controlled environments.

An often-overlooked feature in models for population extinction is that of seasonality, where populations alternate between a breeding and a non-breeding season. Typically, the breeding season will coincide with a season that provides an abundance of resources, whereas the non-breeding season will coincide with a season of harsher environmental conditions. The effect of seasonality on population dynamics has been studied theoretically with population models of low [108, 13] and high [121] complexity, in population experiments [13, 14, 28], and from observations of populations in the wild [182, 205]. These studies have generally found that seasonality does have profound effects of the vital rates of the population, and consequently the overall dynamics. However, the effect of seasonality on EWS of population extinction is not yet understood.

Here, we extend a commonly used population model to include a breeding and non-breeding season, to investigate the effects that seasonality has on EWS of extinction, and the dynamical properties of the system more generally. We derive analytical approximations for the EWS based on small-noise and stationarity assumptions, and use stochastic simulations to verify their behaviour and investigate the behaviour of higher-order EWS which are difficult to compute analytically. Finally we suggest, specific metrics that may be able to distinguish between breeding and non-breeding habitat degradation based on the statistical properties population data.

5.3 Methods

5.3.1 The non-seasonal Ricker model

The Ricker model is a suitable framework for modelling discrete-time population dynamics with non-overlapping generations [157]. We choose the Ricker model over the discrete logistic model, since the latter can result in a jump to negative population densities [126].

The Ricker model may be written

$$N_{t+1} = N_t e^{r-\alpha N_t} \tag{5.1}$$

where r is the intrinsic growth rate, and α controls the strength of density dependence. This system undergoes the famous period-doubling route to chaos as the growth rate is increased (Figure 5.1). At low growth rates, further reduction results in extinction via a Transcritical bifurcation.

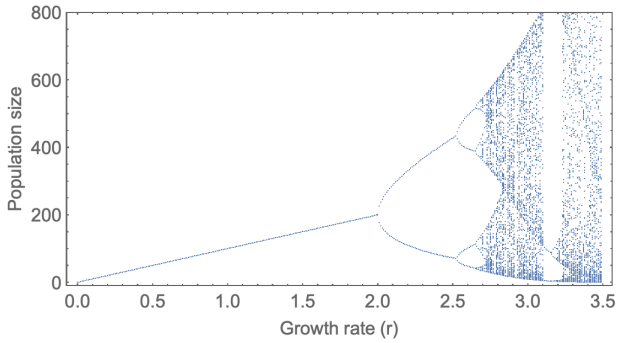

Figure 5.1. Bifurcation diagram for the Ricker model for variation in growth rate r. The system undergoes a period-doubling bifurcation at $r = 2$, followed by further period-doubling bifurcations go chaos, as r increases. The system undergoes a Transcritical bifurcation to extinction at $r = 0$. Parameter values are $\alpha = 0.01$.

There exists an alternate form for the Ricker model, widely used in the literature, which incorporates a 'carrying capacity' explicitly as a parameter. This model takes the form $N_{t+1} = N_t e^{r(1-N_t/K)}$, where K is the carrying capacity. It serves as useful conceptual model, however can become biologically unrealistic when considering variation in the growth and density-dependent parameters [112, 85]. For example, the model implies that populations with high r converge to their carrying capacity faster than populations with low r, which is counterintuitive for populations temporarily above their carrying capacity. Moreover, it doesn't make much sense to define a carrying capacity in environments where populations experience negative growth rates (i.e. more death than reproduction), and the model is not defined for a carrying capacity of zero. When we extend the model to seasonal environments, we will need to consider negative growth rates. Therefore, we use the more classical form of the Ricker model given in (5.1), where a parameter that governs the 'strength of density dependent effects' is used instead of the carrying capacity. This form has been shown to have more biologically relevant properties when considering variation

in growth rates and environment [112, 85].

We incorporate demographic and environmental stochasticity into the model as

$$N_{t+1} = N_t e^{r-\alpha N_t} + \sigma_d \sqrt{N_t}\epsilon_t^{(d)} + \sigma_e N_t \epsilon_t^{(e)}, \tag{5.2}$$

where σ_d and σ_e are noise amplitudes for the demographic and environmental noise processes respectively, and $\epsilon_t^{(d)}$ and $\epsilon_t^{(e)}$ correspond to white noise processes. The magnitude of demographic noise is taken as proportional to the square root of the population size, consistent with the Van Kampen system size expansion [189]. The magnitude of the environmental noise is taken as proportional to the population size, since it is assumed that fluctuations in environmental quality is common to all the individuals [84]. The continuous-time version of this equation has been referred to as the canonical model for stochastic population dynamics [84] and has been studied extensively [144].

Previous experiments with *Daphnia Magna* [80] showed that a reduction in habitat quality corresponded to a reduction in both r and K using an r-K formulation of the Ricker model. An equivalent effect is captured with the formulation in (5.1) upon decreasing r, since the effective growth rate and carrying capacity in this model is r and r/α, respectively. We therefore simulate habitat degradation in this model by decreasing r over time. Further justification for this approach, is that reducing r results in critical slowing down features such as increasing C.V., lag-1 autocorrelation and skewness, which was observed in population experiments, also of *Daphnia Magna*, that were subject to habitat degradation. Changes in α do not result in critical slowing down - only a change in equilibrium (the local recovery rate of the model can be computed to be $\lambda = 1 - r$, which is independent of α).

5.3.2 The biseasonal Ricker model

We construct a model that incorporates two different seasons. One season, where environmental conditions are suitable for breeding, such that the net growth rate is positive, and another where conditions do not permit successful breeding, giving a negative net growth rate. We follow methods by Betini et al. [13], only difference being the form of the Ricker model, where we use the form outlined in the previous section. The structure of the model

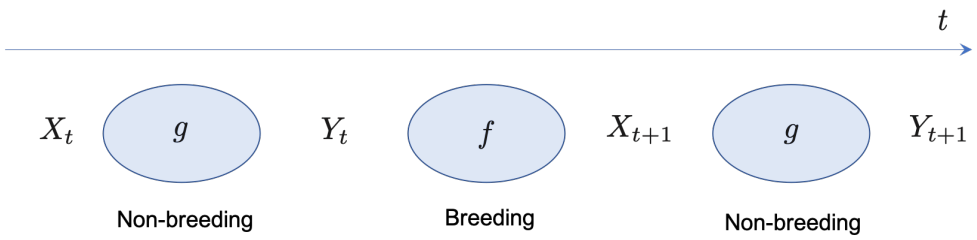

Figure 5.2. Schematic of structure of biseasonal Ricker model. The non-breeding population size X_t enters the non-breeding period which has dynamics governed by the function g. The survivors make up the breeding population size Y_t. They enter the breeding period with dynamics governed by f. The offspring of this population make up X_{t+1}, the non-breeding population size of the next generation.

is illustrated in Figure 5.2. The governing difference equations are

$$X_{t+1} = f(X_t, Y_t) = Y_t e^{r_b - cX_t - \alpha_b Y_t}, \tag{5.3a}$$

$$Y_{t+1} = g(X_{t+1}) = X_{t+1} e^{r_{nb} - \alpha_{nb} X_{t+1}} \tag{5.3b}$$

where X_t is the non-breeding population size, Y_t is the breeding population size, c is the strength of carry-over effects, r_b and r_{nb} correspond to the growth rates in the breeding and non-breeding seasons respectively, and similarly, α_b and α_{nb} control the density-dependent effects. The model used by Betini et al. [13] is recovered upon setting $\alpha_i = r_i/K_i$, where K_i is the carrying capacity of the breeding/non-breeding period, and so it is straightforward to map the parameters used in their study (which were fit to data) to the parameters of this model. Baseline parameters [13] are given by

$$r_b = 2.24, \quad r_{nb} = -0.0568, \quad \alpha_b = 0.01, \quad \alpha_{nb} = 6.72 \times 10^{-4}, \quad c = 0.001, \tag{5.4}$$

which are used in model simulations unless otherwise stated.

We incorporate demographic and environmental stochasticity as before, except now it is included in two distinct periods. The stochastic model is given by

$$X_{t+1} = f(X_t, Y_t) + \sigma_d \sqrt{Y_t} \epsilon_t^{(d)} + \sigma_e Y_t \epsilon_t^{(e)} \tag{5.5a}$$

$$Y_{t+1} = g(X_{t+1}) + \sigma_d \sqrt{X_{t+1}} \delta_t^{(d)} + \sigma_e X_{t+1} \delta_{t+1}^{(e)}, \tag{5.5b}$$

where $\epsilon_t^{(d)}$, $\delta_t^{(d)}$, and $\epsilon_t^{(e)}$, $\delta_t^{(e)}$ are Gaussian white noise processes corresponding to demographic and environmental noise respectively. We investigate extinction events due to decreasing growth rate in the breeding (r_b) and non-breeding (r_{nb}) habitats.

5.3.3 Deriving EWS approximations

Analytical derivations follow a similar procedure to the discrete-time derivations in Chapter 4 however the seasonal model here uses two compartments and multiplicative noise so we explain the extended approach here. Full derivations for the aseasonal and seasonal model are provided in Appendix A and B respectively. Consider the general framework as in (5.5a), given by

$$X_{t+1} = f(X_t, Y_t) + \sigma_d \sqrt{Y_t}\epsilon_t^{(d)} + \sigma_e Y_t \epsilon_t^{(e)} \tag{5.6}$$

$$Y_{t+1} = g(X_{t+1}) + \sigma_d \sqrt{X_{t+1}}\delta_t^{(d)} + \sigma_e X_{t+1}\delta_{t+1}^{(e)}. \tag{5.7}$$

Note that the noise term for each compartment can be combined, since the weighted sum of two normal distributions is itself a normal distribution. Therefore we can write

$$\sigma_d \sqrt{Y_t}\epsilon_t^{(d)} + \sigma_e Y_t \epsilon_t^{(e)} = \sigma(Y_t)\epsilon_t \tag{5.8}$$

where ϵ_t is a normal random variable with zero mean and unit variance and

$$\sigma(Y_t) = \sqrt{\sigma_d^2 Y_t + \sigma_e Y_t^2} \tag{5.9}$$

is the effective noise amplitude which is dependent on Y_t. This reduces the difference equations to

$$X_{t+1} = f(X_t, Y_t) + \sigma(Y_t)\epsilon_t^X, \tag{5.10a}$$

$$Y_{t+1} = g(X_{t+1}) + \sigma(X_{t+1})\epsilon_t^Y, \tag{5.10b}$$

where ϵ_t^X and ϵ_t^Y are normal random variables with mean zero and unit variance. To explore the stochastic behaviour of this system, we consider perturbations about some equilibrium

(X^*, Y^*) given by

$$\xi_t = X_t - X^* \tag{5.11}$$
$$\eta_t = Y_t - Y^*. \tag{5.12}$$

Substituting these into the dynamical equations of (5.10) and keeping just the linear terms provides a good approximation to the residual dynamics. The resulting equations are simple enough to compute statistical metrics directly (Appendix C.2).

5.4 Results

5.4.1 EWS of population extinction in the Ricker model

To evaluate EWS in the Ricker model, we adopt two approaches. One, where we derive approximations directly from the model equations given in (5.2), and another where we simulate the system and compute EWS from the time series data. The advantage of the former, is that it allows us to obtain a relationship between the EWS and the parameters of the model, including the noise amplitude, providing more general information. However, this approach relies on the assumptions made during their derivation, namely a small noise limit and stationarity. The simulations provide a way to compute the EWS in a nonlinear, nonstationary setting, and allow us to compute 'higher-order' EWS such as skewness and kurtosis, which we do not have analytical approximations for.

Approximations for the variance, autocorrelation function and power spectrum of the Ricker model are derived, and are given by

$$\text{Var} = \frac{\sigma^2}{r(2-r)}, \tag{5.13}$$
$$\rho(\tau) = (1-r)^\tau, \tag{5.14}$$
$$S(\omega) = \frac{\sigma^2}{2\pi}\left(\frac{1}{1+(1-r)^2-2(1-r)\cos(\omega)}\right), \tag{5.15}$$

where σ is the combined amplitude of demographic and environmental noise, given by

$$\sigma = \sqrt{\sigma_d^2 N^* + \sigma_e^2 (N^*)^2}, \quad \text{where} \quad N^* = \frac{r}{\alpha}, \tag{5.16}$$

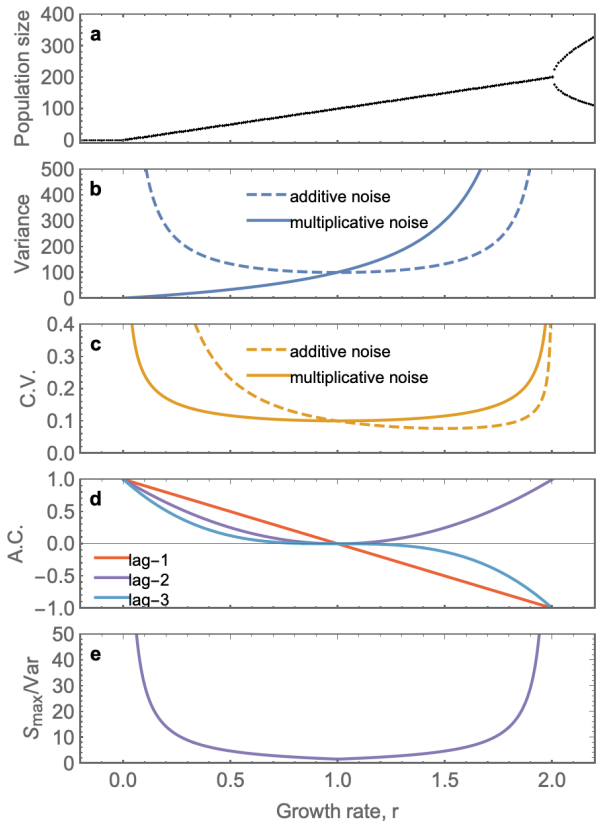

Figure 5.3. Analytical approximations for EWS in the Ricker model. a. Bifurcation diagram showing the Transcritical ($r = 0$) and Flip ($r = 2$) bifurcation. **b.** Variance as a function of r in the additive and multiplicative noise regime. Multiplicative noise suppresses the EWS in variance prior to population extinction. **c.** Coefficient of variation provides a warning prior to extinction for both types of noise. **d.** Autocorrelation at different lag times. All provide a warning prior to extinction, though show different behaviour prior to the Flip bifurcation. Autocorrelation is independent of whether additive or multiplicative noise is considered. **e.** The ratio of the peak in the power spectrum and the variance provides a warning prior to both bifurcations, and is independent of external noise. Derivation of the analytical approximations are provided in Appendix C.1. Parameters here are $\alpha = 0.01$, $\sigma = 0.1$.

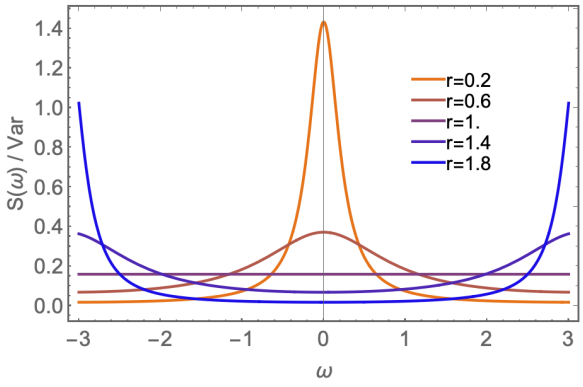

Figure 5.4. Analytical approximation for the power spectral density of the Ricker model at different growth rates. The power spectrum is derived analytically from the stochastic difference equation (5.2) using standard methods (Appendix C.1). Shown is the power spectral density, which is the power spectrum normalised by the variance. Near to the Flip bifurcation the power spectrum peaks at $\omega = \pm\pi$, whereas near to the Transcritical bifurcation (extinction) the power spectrum peaks at $\omega = 0$. Prior to both bifurcations, the peak of the power spectral density ($S_{\max}/\text{Var}$) increases, serving as a potential EWS. Parameters are $\alpha = 0.01$, $\sigma = 0.1$.

the equilibrium state. Derivations are provided in Appendix C.1. From these metrics one can readily derive several others, such as the coefficient of variation (C.V.) and the peak in the power spectrum ($S_{\max}$). We plot the EWS approximations and power spectrum in Figure 5.3 and Figure 5.4 respectively. EWS computed directly from simulations of the Ricker model approaching extinction are shown in Figure 5.5, and show good agreement with the analytical approximations. Several observations are in order.

Using the variance of a time series is a popular EWS in the literature [169]. However it can fail under regimes of multiplicative noise, since in such cases the variance is affected by changes in the system size [55]. In the case of population extinction, the declining population density will naturally lead to smaller noise amplitudes, simply because less individuals are experiencing the external perturbations. Therefore variance can decrease as a bifurcation is approached in cases where the system size also decreases. This is seen in both the analytical and simulated EWS. Note however that if additive noise is assumed (σ constant), as is the case in many stochastic models, then it is clear from (5.13) that the variance diverges as $r \to 0$, the familiar consequence of critical slowing down.

The coefficient of variation (C.V.), defined as the ratio of the standard deviation of the time series and its mean value, is a useful alternative to the variance as an EWS. This is because it compensates for the diminishing impact of environmental noise as the mean of the state variable approaches zero. Whereas environmental noise is proportional to the

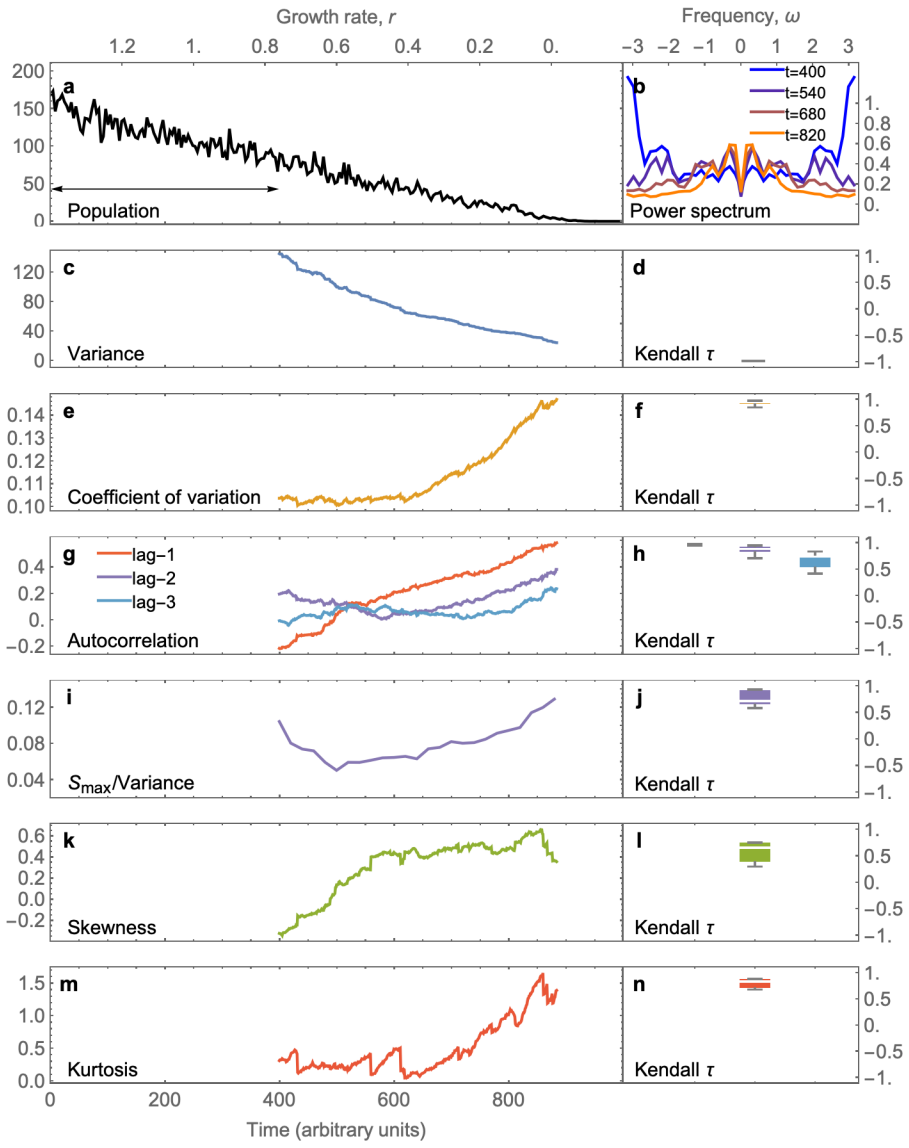

Figure 5.5. EWS in simulations of the Ricker model with decreasing growth rate. Growth rate (r) is decreased linearly over the interval $[-0.2, 1.4]$ during 1000 time increments. We show EWS from a single realisation, and Kendall tau values computed beyond $t = 600$ from an ensemble of 100 realisations. This behaviour has been observed in experiments with *Daphnia magna* where increasing skewness, C.V and lag-1 autocorrelation were observed as nutrients were reduced [61]. Fixed model parameters are $\alpha = 0.01$, $\sigma = 0.01$.

size of the system ($\propto N^*$), the coefficient of variation is normalised by the size of the system ($\propto 1/N^*$) and so these effects cancel out. As a consequence, the C.V. is able to show the effects of critical slowing down, without being distorted. Moreover, the ratio of demographic noise strength to system size scales like $1/\sqrt{N}$, and so this also increases as extinction is approached, contributing to the increase in C.V. For the Ricker model, it can be seen that the C.V increases in both additive and multiplicative noise regimes prior to the Flip and the Transcritical bifurcation (Figure 5.3). It should be noted however that an initial decrease is expected if the system is moving towards extinction, but began near to the Flip bifurcation.

The autocorrelation increases prior to extinction, for each lag time considered, although longer lags yield a less reliable signal in simulations. Near to the Flip bifurcation the autocorrelation behaves very differently depending on the lag, as expected from theoretical results in 4. At lag-1, autocorrelation decreases in the approach to a Flip bifurcation, but increases at lag-2, for example. The analytical and simulated power spectrum also behave as expected, with a peak at $\omega = \pi$ near the Flip bifurcation and a peak at $\omega = 0$ closer to the Transcritical bifurcation. The peak in the power spectral density goes like

$$\frac{S_{\max}}{\text{Var}} = \frac{2-r}{2\pi r} \sim \frac{1}{r} \quad \text{as} \quad r \to 0, \tag{5.17}$$

which is independent of noise amplitude so serves as a good indicator.

Finally, we compute two higher-order metrics, namely the skewness and kurtosis, which have been suggested as EWSs [82]. We refer to these metrics as higher-order, since they consider the shape of the stability landscape, beyond the region where a linear approximation is valid. We therefore do not have analytical approximations for them, since our methods involve taking a linear approximation. We do however compute them from simulated data Figure 5.5, and find that they both increase prior to extinction.

5.4.2 Impacts of seasonality on dynamics and EWS

Adding a second compartment to the model that corresponds to a non-breeding period has a remarkable effect on the stability properties of the system, as shown by the bifurcation diagrams (Figure 5.6). First, the Flip bifurcation that was present in the standard Ricker model has been pushed to a higher growth rate ($r_b \approx 2.4$), presumably due to compensation from the negative growth rate in the non-breeding season. Perhaps more surprising, is

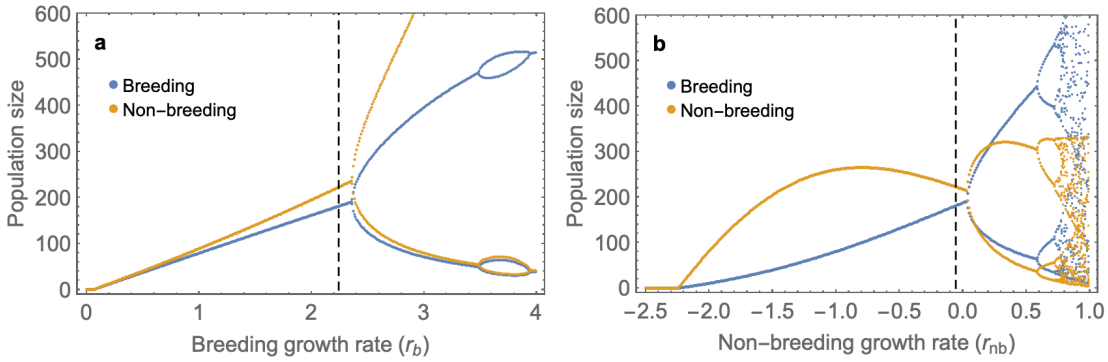

Figure 5.6. Bifurcation diagrams of biseasonal Ricker model for varying growth rates. Solid lines show stable state/limit cycle. Dashed vertical line shows empirically measured value for fruit flies [13]. Increasing r_b yields instability via Flip bifurcation (paradox of enrichment [163]), whereas decreasing r_b yields a declining population size to extinction via a Transcritical bifurcation. Increasing r_{nb} to positive values yields Flip bifurcations to chaos, however since r_{nb} represents growth rate in the non-breeding period it should be negative. Decreasing r_{nb} initially causes an increase in the non-breeding population size (due to less competition during the breeding period), followed by a decrease to extinction via a Transcritical bifurcation.

the absence of a chaotic regime as the breeding growth rate increases. Whereas the non-seasonal Ricker model breaks into chaotic dynamics for $r \gtrapprox 2.7$, the corresponding seasonal model remains in periodic motion. As the breeding growth rate increases, the system undergoes three Flip bifurcations which result in period doubling, followed by three more Flip bifurcations which result in period-halving, resulting in a return to equilibrium at higher values for r_b (Figure C.1). The inclusion of a non-breeding season, where the net growth rate is negative, appears to suppress the onset of chaos. Note that for positive growth rates in the non-breeding period ($r_{nb} > 0$) the system does go into chaos, however this region is not biologically relevant.

With seasonality in the model, a Transcritical bifurcation to extinction may now occur in two different ways. The growth rate in the breeding period may decrease (corresponding to habitat degradation in the breeding period, and therefore less offspring per capita), or the growth rate in the non-breeding period may decrease (corresponding to habitat degradation in the non-breeding period, and therefore a lower chance that individuals survive). It can be seen from the bifurcation diagrams that the trend to extinction is slightly different depending on which habitat undergoes deterioration. For decreasing r_b, both the breeding

and non-breeding population densities approach zero in a monotonic fashion. However, a decrease in r_{nb} causes an initial divergence in the population trends, with the non-breeding population size initially increasing. This is due to less competition in the breeding season, hence a larger offspring and a larger non-breeding population size. This was observed recently in seasonal population experiments [28].

Having seen how the stability properties change with addition of a non-breeding period to the Ricker model, it is worth exploring how the EWS change. The behaviour of EWS in this system at differing values of the growth rates is shown for the breeding and non-breeding populations in Figure 5.7 and Figure 5.8, respectively. Note that as was the case of the Ricker model, the breeding population exhibits decreasing variance, increasing C.V. and increasing autocorrelation at low lags as the breeding growth rate approaches zero. One would expect this due to the generic properties of critical slowing down that occur at a Transcritical bifurcation. However the skewness, which increased prior to extinction in the Ricker, now decreases prior to extinction as the breeding growth rate is decreased. This goes to show that even though the trajectory to extinction of a non-seasonal vs. seasonal population may look similar, their higher-order statistical properties such as the skewness may be very different. These differences could perhaps be used to obtain information about the system from data, which we investigate in the following section.

5.4.3 Using EWS to infer the driver of extinction

Extinction in populations may occur through both degradation in the breeding habitat and degradation in the non-breeding habitat, though it is often difficult to determine the root cause of the population decline. From the biseasonal Ricker model, we see that there exist asymmetries in some of the EWS, specifically the skewness, kurtosis and variance, that could be used to infer the type of habitat degradation that the population is undergoing. In the breeding population (Figure 5.7), the skewness decreases prior to extinction from breeding habitat degradation, but increases prior to extinction from non-breeding habitat degradation. The inverse behaviour occurs for the kurtosis.

In the non-breeding population (Figure 5.8), we observe decreasing variance prior to extinction from breeding habitat degradation, but a sharp increase in variance prior to extinction from non-breeding habitat degradation (after an initial decrease as the system moves away from the Flip bifurcation). The skewness and kurtosis show roughly opposite trends as was for the breeding population. The kurtosis increases as the system moves away from

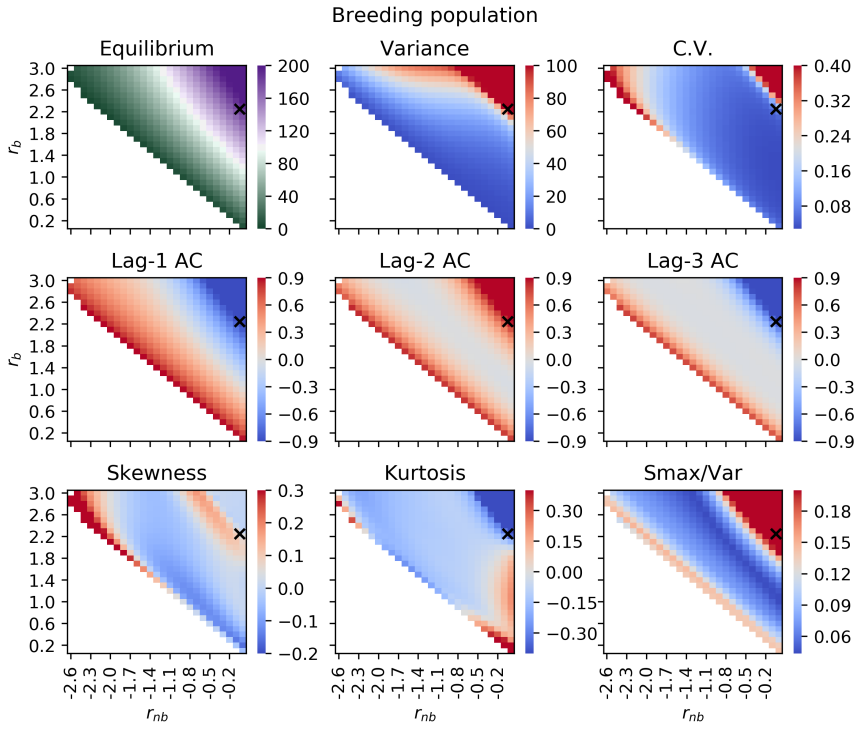

Figure 5.7. EWS in the breeding population density (Y_t) for the biseasonal Ricker model. The biseasonal Ricker model is simulated at different combinations of growth parameters r_b and r_{nb}, and EWS are computed from the resulting time series. Other model parameters take baseline values given in (5.4). Noise amplitude is taken as $\sigma_d = \sigma_e = 0.02$. The cross marks the growth rates observed in fruit fly experiments with fixed habitats [13].

the Flip bifurcation, though to a much larger extent upon degradation in the non-breeding habitat. Prior to non-breeding habitat induced extinction, the kurtosis then drops sharply to negative values, contrary to extinction from breeding habitat reduction, where it remains at high values. So in theory, the differences in these statistical metrics could be used to inform the region of habitat degradation causing population declines.

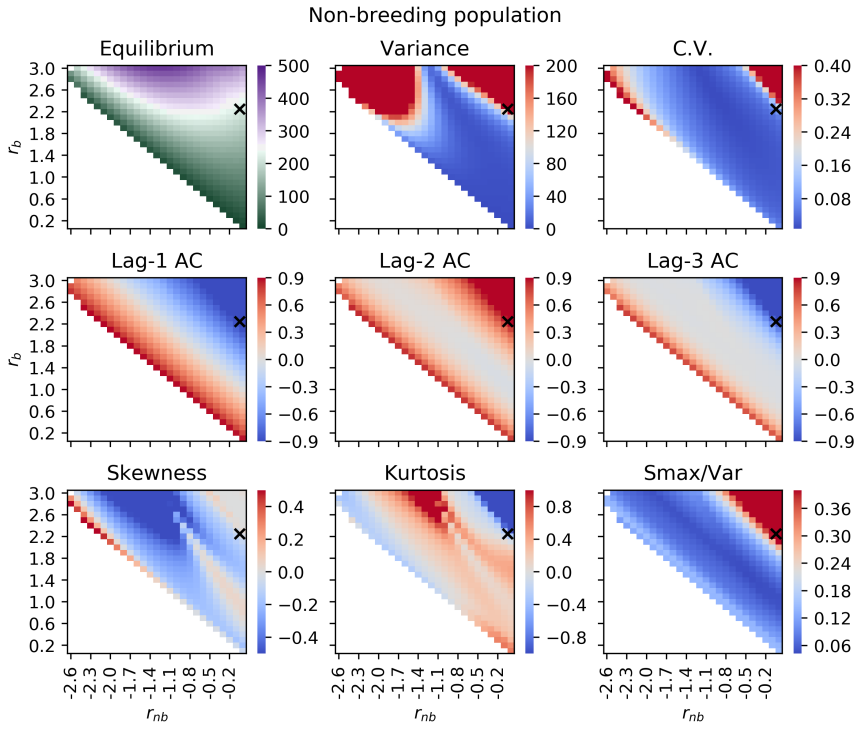

Figure 5.8. EWS in the breeding population density (X_t) for the biseasonal Ricker model. The biseasonal Ricker model is simulated at different combinations of growth parameters r_b and r_{nb}, and EWS are computed from the resulting time series. Other model parameters take baseline values given in (5.4). Noise amplitude is taken as $\sigma_d = \sigma_e = 0.02$. The cross marks the growth rates observed in fruit fly experiments with fixed habitats [13].

5.5 Discussion

With a commonly used non-seasonal and corresponding biseasonal population model, we have shown how seasonality can effect EWS for population extinction. In the non-seasonal model, we derived analytical approximations for EWS, which demonstrated how the EWS vary along the stable branch between the Transcritical bifurcation (to extinction) at low growth rate, and the Flip bifurcation (to oscillations) at a high growth rate. We also showed how the power spectrum displays qualitatively different features prior to each bifurcation, in line with theoretical expectations [29]. In the biseasonal model, we found qualitative differences in the behaviour of the higher-order EWS such as skewness and kurtosis, depending on whether the breeding or non-breeding habitat was degraded. Fi-

nally, we showed how, in theory, one could distinguish between breeding and non-breeding habitat destruction from the EWS in the time series.

The higher-order statistical moments like skewness and kurtosis display rather intricate behaviour for varying growth rates in the biseasonal model, and so inferring a populations distance to the different bifurcations using solely these metrics would be difficult. The advantages lie in using several EWS metrics combined. For example, the lag-1 AC displays a very consistent trend of increasing from negative values near the Flip bifurcation to positive values as either of the habitats is degraded. In pairing this metric with skewness for example, one can then start to make judgements as to which habitat is being degraded.

Further work should test the EWS considered here on empirical population datasets that contain breeding and non-breeding population densities, such as migratory bird data [148], and see if similar asymmetries arise based on the type of habitat degradation. There is also scope for the development of composite EWS [61], which combine several EWS to provide an enhanced signal. It would be interesting to see, in the context of seasonal environments, which combinations of EWS in the breeding and non-breeding period provide the most robust signal of upcoming extinction. EWS that include trait variables such as body size and age of maturation have also shown to be effective [44, 42], and there exist long term empirical studies that have observed trait changes prior to population decline [67, 8]. Finally, future work should check the robustness of these results to model complexity by computing EWS in more detailed seasonal population models.

Chapter 6

Conclusion

Using both modelling and data-driven approaches, this thesis set out to investigate the impact of dynamic social processes on climate change, and develop tools to better predict bifurcations and their implied dynamics from ecological data. In line with a modelling study that used an alternate framework for social dynamics [11], we found that uncertainty in social processes has a major impact on climate change projections. Particular to our study, was the incorporation of endogenous social processes, namely social learning and dynamic social norms, which we found to have their own unique impact on the socio-climate trajectories. The rate of social learning, when varied between plausible ranges, varied the peak temperature anomaly by over 1°C, with higher values resulting in a faster transition to mitigative behaviour among the population. Social norms on the other hand, suppressed the early spread of mitigative behaviour resulting in higher temperature anomalies. This study highlighted the importance of coupling social processes to climate dynamics when it comes to developing policies to steer the socio-climate system in the right direction.

Changes in the climate system and increasing anthropogenic pressures are putting many ecosystems at risk of transitions to an alternative undesirable state [162]. A significant part of this thesis was devoted to constructing novel early warning signals (EWS) for these transitions, that are both more sensitive than traditional metrics, but also capable of extracting information on the type of transition to expect. This goal was achieved in part with the development of spectral EWS - metrics derived from the power spectrum of time series data. These metrics allowed us to distinguish between two contrasting transitions in a simple population model, and were applied to an empirical dataset of predator-prey dynamics to provide characteristic EWS of a Hopf bifurcation. These metrics capture

generic dynamical processes that occur prior to specific bifurcations and so can be applied to noisy time-series across the wide range of disciplines that are subject to bifurcations.

The behaviour of EWS in seasonal environments is not well studied, and so this thesis investigated EWS in the context of population dynamics with the goal of identifying EWS that were suitable for detecting population extinction in seasonal environments. Moreover, whether extinction is being driven by breeding or non-breeding habitat degradation is typically unknown, and important information for guiding conservation efforts. A question this project asked was whether the time series of the population in the breeding and non-breeding period could show signatures of decline that were specific to the type of environmental degradation. We found that whilst conventional EWS like lag-1 autocorrelation and coefficient of variation behaved as expected prior to both types of extinction (an increase prior to the bifurcation), higher-order metrics like skewness and kurtosis displayed interesting asymmetries that could be exploited to inform of the type of habitat degradation. For example, time series from a population in its breeding period showed increasing skewness as habitat was degraded in the non-breeding period, but decreasing skewness upon habitat degradation in the breeding period. Therefore, combining higher-order metrics with conventional EWS could be a promising way forward for not only detecting the approach of a transition, but describing the key driver of the approaching transition. Moreover we found that spectral EWS can be used to distinguish between the Flip bifurcation to oscillations and the Transcritical bifurcation to extinction. Given the varying advantages of different EWS, it is clear that in order to obtain maximal information about imminent transitions in these systems, a multitude of different EWS should be used.

There is often debate as to how complex a model should be to best serve its purpose. Decisions have to be made on the number of variables and parameters to include, and whether factors such as stochasticity, spatial dynamics, or heterogeneity are required. The research in this thesis worked with both deterministic and stochastic models, though all of them were relatively simple, in the sense that a minimal number of variables were considered. In the climate study, we developed a simple deterministic model, as our goals were to investigate the effect of a few salient social mechanisms with clarity. Although there is undoubtably a stochastic component to human behaviour and climate dynamics, incorporating stochasticity would be unlikely to distinctly modify the trends in the trajectories, which are the primary focus of the study. Moreover, despite the simplicity of the climate component to the model, its predictions were similar enough to state-of-the-art complex models to serve our purposes.

Future studies should build on the socio-climate framework developed here to include more intricate aspects of the social system, such as heterogeneity among agents and a more diverse set of strategies. The complexity of the climate model should match that of the social model. In our study, we used an Earth system model [117] which models carbon transfer between the atmosphere, ocean and land at a global scale with low complexity. Our social model was also of low complexity, so the Earth system model was sufficient for our purposes. A climate model of intermediate complexity such as C-ROADS [177] would allow for a more intricate coupling between the social and climate system, and the investigation of more specific policies. In comparison to the Earth system model, C-ROADS includes finer details such as a wider array of GHGs and a multilayer ocean model that incorporates pH and sea level rise. Outputs such as sea level rise as well as temperature could be coupled to the human system, and the model includes a term for afforestation to investigate carbon sequestration policies. At the highest level of complexity are the general circulation models (GCMs). Current GCMs use a spatial resolution of between 100 and 200 cubic kilometres and take into account factors such as ice sheet dynamics. Using these spatially explicit models would be more appropriate for modelling human behaviour across nations, where climate change impacts will differ significantly.

Conclusions drawn from our work in socio-climate modelling should be tested against different model frameworks for social dynamics, as is done for models of the climate system where ensembles are taken to provide the most accurate predictions [107]. Popular alternative frameworks for modelling social dynamics include 'best response' models, and threshold models. Best response models, first developed in classical economic theory, have played a role modelling human-environment systems such as fisheries [69] and forests [90]. These models assume that individuals adopt the strategy that maximises their utility, and immediately switch when an alternate strategy obtains a higher utility. This approach therefore does not capture social learning, where individuals take time to learn strategies from one another, a process that our study was focused around. Threshold models for collective behaviour, originally proposed by Granovetter [79], provide a framework where each individual possesses a threshold indicating the popularity of a strategy required for them to switch. An important feature of this model is the heterogeneity of agents - a distribution of thresholds is assumed. Therefore one can model both 'rebels' and 'conservatives' as individuals possessing thresholds at opposite ends of the spectrum. Such a framework could provide a good base for exploring the effect of agent heterogeneity on socio-climate trajectories.

In the development and testing of EWS, it was essential that we use stochastic models, as it is stochasticity that gives rise to EWS in the first place (along with the dynamic phenomena of critical slowing down). In deriving theoretical approximations for EWS, we used the normal form of each bifurcation accompanied by additive white noise. Each of these forms can be thought of as a simple model for a more complex system that is in the neighbourhood of the corresponding bifurcation. The justification for the simplicity here is that near to local bifurcations, the behaviour of dynamical systems is well-approximated by the corresponding normal form of the bifurcation in a lower-dimensional phase space [114, 132]. Therefore, results derived for the normal forms are generalisable to higher-dimensional systems, provided one can work out the subspace of the higher-dimensional system that is undergoing critical slowing down, which may be possible if data is available for each variable [199]. However, the use of additive white noise as opposed to multiplicative or coloured noise is a simplification that should be relaxed in future theoretical studies, as it has been suggested that most environmental noise in ecology is slightly 'red' (positive temporal correlation).

Predicting the type of approaching bifurcation from time series data alone remains a challenge. The spectral EWS developed in this thesis are able to distinguish between oscillatory and non-oscillatory bifurcations, and approximate the period of oscillations that occurs at the bifurcation. However there still exist bifurcations that cannot be distinguished using this method, such as the Fold and Transcritical bifurcations, which do not have an oscillatory component. The main issue here is that these bifurcations exhibit the same linearised behaviour about the branch that leads up to the bifurcation. Since most EWS are derived in a limit of small noise, such that only linear terms are relevant, these EWS are expected to exhibit the same behaviour prior to each of these bifurcations. Only when one considers dynamics outside of the local linear regime can information be obtained on the type of these non-oscillatory bifurcations, and whether they yield critical or smooth transitions. Developing EWS that can distinguish between these bifurcations remains an important endeavour given the drastically varying consequences between critical and smooth transitions.

In this thesis, analytical derivations of EWS use a 'small noise limit', which yields expressions for EWS that are approximate in a regime of small noise. This allows one to restrict attention to linear dynamics, making closed, analytical expressions feasible. However, this approximation breaks down when deviations from equilibrium become large, as is the case very near to a bifurcation. A field adept to deal with noisy systems outside of

the small noise limit is that of random dynamical systems (RDS). For a recent overview of RDS and review of major results, see [135]. Here, instead of describing dynamics by an equilibrium state perturbed by small noise, one derives dynamics in terms of invariant measures, which are probability distributions describing the possible states of the system. A 'stochastic bifurcation', is then a point in parameter space, where this invariant measure undergoes a qualitative change. Treating noisy dynamical systems in this way, although more demanding mathematically, allows one to obtain rigorous results on the behaviour of a system in large noise situations and obtain approximations for stochastic dynamics in high-dimensional systems. The integration of RDS theory with EWS may inspire new metrics, and provide insight to EWS behaviour under circumstances of large noise.

The research in this thesis has several implications for science and policy. The simple socio-climate model shows how variation in strength of each social process has distinct effects on socio-climate trajectories. Moreover the order and timing of how they are varied matters. The model suggests that social learning should be prioritised for most immediate gains, followed by a reduction in mitigation costs. Policies can address these aspects of the social system. For example, social learning rate could be increased through an increase in media coverage of climate-induced extreme events, the organisation of climate marches, and widespread coverage of the international climate reports. Individual incentive to adopt mitigative behaviour could be increased through a carbon tax or discounted environmentally friendly products. Therefore socio-climate models should be used for policy impact assessments, and should be continued to be developed at different levels of complexity. Early warning signals, though still at the research stage, will be of benefit to policy makers that have access to high resolution data. The spectral EWS developed here provide a key advantage in being more sensitive to approaching transitions and providing information on the type of transition. Combined with model-based approaches and conventional EWS, detecting transitions in real systems with high-frequency output is becoming more feasible.

Nonlinear processes are the rule rather than the exception in complex systems. Mathematical models that capture these processes can elucidate the effect they have on the system as a whole. The interplay between nonlinearities in social dynamics and the climate system are only beginning to be included in models. Given the urgency of the climate crises, it is crucial that collaboration among the social and climate sciences takes place to further develop socio-climate models in view of implementing the most effective policies. Alongside this, the continuous development and testing of EWS in empirical systems is required, to bring this set of tools closer to practical application. The ability to detect and

describe transitions in the Anthropocene, be they in the socio-climate system or the ecosystems upon which we rely, is an urgent need. Continued work in EWS and socio-climate modelling will contribute to this endeavour.

References

[1] Icek Ajzen. The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2):179–211, 1991. 10

[2] Hirotugu Akaike. Factor analysis and aic. In *Selected Papers of Hirotugu Akaike*, pages 371–386. Springer, 1987. 69, 70

[3] Hunt Allcott. Social norms and energy conservation. *Journal of public Economics*, 95(9-10):1082–1095, 2011. 10

[4] Stefano Allesina and Si Tang. Stability criteria for complex ecosystems. *Nature*, 483 (7388):205–208, 2012. 133, 136, 137, 138

[5] Peter Ashwin, Sebastian Wieczorek, Renato Vitolo, and Peter Cox. Tipping points in open systems: bifurcation, noise-induced and rate-dependent examples in the climate system. *Phil. Trans. R. Soc. A*, 370(1962):1166–1184, 2012. 40

[6] Jonathan EM Baillie, Craig Hilton-Taylor, Simon N Stuart, et al. 2004 iucn red list of threatened species. 2004. 74

[7] Julia Baird, Ryan Plummer, Constanze Haug, and Dave Huitema. Learning effects of interactive decision-making processes for climate change adaptation. *Global Environmental Change*, 27:51–63, 2014. 9, 10

[8] Javier Balbontın, Vincenzo Penteriani, and Miguel Ferrer. Variations in the age of mates as an early warning signal of changes in population trends? the case of bonelli's eagle in andalusia. *Biological Conservation*, 109(3):417–423, 2003. 89

[9] Ryan D Batt, William A Brock, Stephen R Carpenter, Jonathan J Cole, Michael L Pace, and David A Seekell. Asymmetric response of early warning indicators of

phytoplankton transition to and from cycles. *Theoretical ecology*, 6(3):285–293, 2013. 54

[10] Chris T Bauch, Ram Sigdel, Joe Pharaon, and Madhur Anand. Early warning signals of regime shifts in coupled human–environment systems. *Proceedings of the National Academy of Sciences*, page 201604978, 2016. 9, 45

[11] Brian Beckage, Louis J Gross, Katherine Lacasse, Eric Carr, Sara S Metcalf, Jonathan M Winter, Peter D Howe, Nina Fefferman, Travis Franck, Asim Zia, et al. Linking models of human behaviour and climate alters projected climate change. *Nature Climate Change*, page 1, 2018. 6, 10, 19, 21, 90

[12] Elisa Beninca, Vasilis Dakos, Egbert H Van Nes, Jef Huisman, and Marten Scheffer. Resonance of plankton communities with temperature fluctuations. *The American Naturalist*, 178(4):E85–E95, 2011. 68

[13] Gustavo S Betini, Cortland K Griswold, and D Ryan Norris. Carry-over effects, sequential density dependence and the dynamics of populations in a seasonal environment. *Proceedings of the Royal Society of London B: Biological Sciences*, 280 (1759):20130110, 2013. 75, 77, 78, 85, 87, 88, 169

[14] Gustavo S Betini, Cortland K Griswold, Livia Prodan, and D Ryan Norris. Body size, carry-over effects and survival in a seasonal environment: consequences for population dynamics. *Journal of Animal Ecology*, 83(6):1313–1321, 2014. 75

[15] Reinette Biggs, Stephen R Carpenter, and William A Brock. Turning back from the brink: detecting an impending regime shift in time to avert it. *Proceedings of the National academy of Sciences*, 106(3):826–831, 2009. 39, 46, 66

[16] Chris Blackmore, Raymond Ison, and Janice Jiggins. Social learning: an alternative policy instrument for managing in the context of europe's water. *Environmental Science & Policy*, 10(6):493–498, 2007. 23

[17] TA Boden, RJ Andres, and Gregg Marland. Global, regional, and national fossil-fuel co2 emissions (1751-2014)(v. 2017). Technical report, Carbon Dioxide Information Analysis Center (CDIAC), Oak Ridge National . . . , 2017. 16, 129, 130

[18] 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. 38, 41

[19] Niklas Boers. Early-warning signals for dansgaard-oeschger events in a high-resolution ice core record. *Nature communications*, 9(1):2556, 2018. 45

[20] Carl Boettiger. From noise to knowledge: how randomness generates novel phenomena and reveals information. *Ecology letters*, 2018. 44, 45, 67, 152, 167

[21] Carl Boettiger and Ryan Batt. Bifurcation or state tipping: assessing transition type in a model trophic cascade. *Journal of Mathematical Biology*, pages 1–13, 2019. 54

[22] Carl Boettiger and Alan Hastings. No early warning signals for stochastic transitions: insights from large deviation theory. *Proceedings of the Royal Society of London B: Biological Sciences*, 280(1766):20131372, 2013. 40

[23] Carl Boettiger, Noam Ross, and Alan Hastings. Early warning signals: the charted and uncharted territories. *Theoretical ecology*, 6(3):255–264, 2013. 66

[24] Bryan Bollinger and Kenneth Gillingham. Peer effects in the diffusion of solar photovoltaic panels. *Marketing Science*, 31(6):900–912, 2012. 9, 10, 12

[25] Laurent Bopp, L Resplandy, James C Orr, Scott C Doney, John P Dunne, M Gehlen, P Halloran, Christoph Heinze, Tatiana Ilyina, Roland Seferian, et al. Multiple stressors of ocean ecosystems in the 21st century: projections with cmip5 models. *Biogeosciences*, 10:6225–6245, 2013. 9

[26] George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015. 46, 49, 50, 145, 154

[27] Tobias S Brett, John M Drake, and Pejman Rohani. Anticipating the emergence of infectious diseases. *Journal of The Royal Society Interface*, 14(132):20170115, 2017. 39

[28] Joseph B Burant, Gustavo S Betini, and D Ryan Norris. Simple signals indicate which period of the annual cycle drives declines in seasonal populations. *Ecology Letters*, 2019. 75, 86

[29] Thomas M Bury, Chris T Bauch, and Madhur Anand. Detecting and distinguishing tipping points using spectral early warning signals. *In review*, 2019. iv, 6, 88

[30] Thomas M Bury, Chris T Bauch, and Madhur Anand. Charting pathways to climate change mitigation in a coupled socio-climate model. *PLOS computational biology*, 15 (6):e1007000, 2019. iv, 6, 168

[31] Thomas M Bury, Joseph Burant, Ryan Norris, Chris T Bauch, and Madhur Anand. Early warning signals for population extinction in seasonal environments. *In progress*, 2019. iv

[32] SR Carpenter and WA Brock. Rising variance: a leading indicator of ecological transition. *Ecology letters*, 9(3):311–318, 2006. 31, 38, 39, 45

[33] SR Carpenter, WA Brock, JJ Cole, JF Kitchell, and ML Pace. Leading indicators of trophic cascades. *Ecology letters*, 11(2):128–138, 2008. 39, 40

[34] Stephen R Carpenter, Jonathan J Cole, Michael L Pace, Ryan Batt, WA Brock, Timmothy Cline, Jim Coloso, James R Hodgson, Jim F Kitchell, David A Seekell, et al. Early warnings of regime shifts: a whole-ecosystem experiment. *Science*, 332 (6033):1079–1082, 2011. 39

[35] Stephen R Carpenter, William A Brock, Jonathan J Cole, and Michael L Pace. A new approach for rapid detection of nearby thresholds in ecosystem time series. *Oikos*, 123(3):290–297, 2014. 39

[36] Noel Castree, William M Adams, John Barry, Daniel Brockington, Bram Büscher, Esteve Corbera, David Demeritt, Rosaleen Duffy, Ulrike Felt, Katja Neves, et al. Changing the intellectual climate. *Nature climate change*, 4(9):763, 2014. 9, 10, 23

[37] Jérôme Chappellaz, JM Barnola, D Raynaud, Ye S Korotkevich, and C Lorius. Ice-core record of atmospheric methane over the past 160,000 years. *Nature*, 345(6271): 127, 1990. 132

[38] Andrew Chen, Alvaro Sanchez, Lei Dai, and Jeff Gore. Dynamics of a producer-freeloader ecosystem on the brink of collapse. *Nature communications*, 5, 2014. 31, 42

[39] 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. 39, 40

[40] Robert B Cialdini, Carl A Kallgren, and Raymond R Reno. A focus theory of normative conduct: A theoretical refinement and reevaluation of the role of norms in human behavior. In *Advances in experimental social psychology*, volume 24, pages 201–234. Elsevier, 1991. 10

[41] Susan Clayton, Patrick Devine-Wright, Paul C Stern, Lorraine Whitmarsh, Amanda Carrico, Linda Steg, Janet Swim, and Mirilia Bonnes. Psychological research and global climate change. *Nature Climate Change*, 5(7):640, 2015. 9, 10

[42] Christopher F Clements and Arpat Ozgul. Including trait-based early warning signals helps predict population collapse. *Nature communications*, 7, 2016. 89

[43] Christopher F Clements and Arpat Ozgul. Indicators of transitions in biological systems. *Ecology letters*, 2018. 44, 45

[44] Christopher F Clements, Julia L Blanchard, Kirsty L Nash, Mark A Hindell, and Arpat Ozgul. Body size shifts and early warning signals precede the historic collapse of whale stocks. *Nature ecology & evolution*, 1(7):0188, 2017. 89

[45] Kevin Collins and Ray Ison. Living with environmental change: adaptation as social learning. *Environmental Policy and Governance*, 19(6):351–357, 2009. 23

[46] WJ Collins, Nicolas Bellouin, M Doutriaux-Boucher, N Gedney, P Halloran, T Hinton, J Hughes, CD Jones, M Joshi, S Liddicoat, et al. Development and evaluation of an earth-system model-hadgem2. *Geoscientific Model Development*, 4(4):1051, 2011. 9

[47] Peter M Cox, Richard A Betts, Chris D Jones, Steven A Spall, and Ian J Totterdell. Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model. *Nature*, 408(6809):184, 2000. 11

[48] Peter M Cox, Chris Huntingford, and Chris D Jones. *Conditions for sink-to-source transitions and runaway feedbacks from the land carbon cycle*. Cambridge University Press, 2006. 11

[49] Katharine Z Coyte, Jonas Schluter, and Kevin R Foster. The ecology of the microbiome: networks, competition, and stability. *Science*, 350(6261):663–666, 2015. 136

[50] 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. 3, 31, 42, 75

[51] Vasilis Dakos. Identifying best-indicator species for abrupt transitions in multispecies communities. *Ecological Indicators*, 2017. 40, 41

[52] Vasilis Dakos and Fernando Soler-Toscano. Measuring complexity to infer changes in the dynamics of ecological systems under stress. *Ecological Complexity*, 2016. 39

[53] 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. 2, 3, 31

[54] Vasilis Dakos, Egbert H van Nes, Raúl Donangelo, Hugo Fort, and Marten Scheffer. Spatial correlation as leading indicator of catastrophic shifts. *Theoretical Ecology*, 3 (3):163–174, 2010. 28, 45

[55] Vasilis Dakos, Stephen R Carpenter, William A Brock, Aaron M Ellison, Vishwesha Guttal, Anthony R Ives, Sonia Kéfi, 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, 2012. 34, 39, 41, 45, 67, 82, 174

[56] 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, 2012. 41, 42, 46, 57, 58, 59

[57] 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. xvi, 38

[58] 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. 39, 54, 59, 71

[59] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3): 554–559, 2016. 23

[60] Cees Diks, Cars Hommes, and Juanxi Wang. Critical slowing down as an early warning signal for financial crises? *Empirical Economics*, pages 1–28, 2015. 45

[61] John M Drake and Blaine D Griffen. Early warning signals of extinction in deteriorating environments. *Nature*, 467(7314):456–459, 2010. 3, 42, 75, 83, 89

[62] Richard Durrett and Simon Levin. The importance of being discrete (and spatial). *Theoretical population biology*, 46(3):363–394, 1994. 23

[63] Partha Sharathi Dutta, Yogita Sharma, and Karen C Abbott. Robustness of early warning signals for catastrophic and non-catastrophic transitions. *Oikos*, 2018. 42

[64] Tarsha Eason, Ahjond S Garmestani, and Heriberto Cabezas. Managing for resilience: early detection of regime shifts in complex systems. *Clean Technologies and Environmental Policy*, 16(4):773–783, 2014. 39

[65] Tarsha Eason, Ahjond S Garmestani, Craig A Stow, Carmen Rojo, Miguel Alvarez-Cobelas, and Heriberto Cabezas. Managing for resilience: an information theory-based approach to assessing ecosystems. *Journal of Applied Ecology*, 53(3):656–665, 2016. 40

[66] Joshua M. Epstein. Why model? *Journal of Artificial Societies and Social Simulation*, 11(4):12, 2008. 3, 27

[67] Miguel Ferrer, Vincenzo Penteriani, Javier Balbontın, and Massimo Pandolfi. The proportion of immature breeders as a reliable early warning signal of population decline: evidence from the spanish imperial eagle in doñana. *Biological Conservation*, 114(3):463–466, 2003. 89

[68] AD Friend, AK Stevens, RG Knox, and MGR Cannell. A process-based, terrestrial biosphere model of ecosystem dynamics (hybrid v3. 0). *Ecological Modelling*, 95(2-3): 249–287, 1997. 132

[69] John M Fryxell, Ray Hilborn, Carling Bieg, Katrine Turgeon, Amanda Caskenette, and Kevin S McCann. Supply and demand drive a critical transition to dysfunctional fisheries. *Proceedings of the National Academy of Sciences*, 114(46):12333–12337, 2017. 92

[70] Gregor F Fussmann, Stephen P Ellner, Kyle W Shertzer, and Nelson G Hairston Jr. Crossing the hopf bifurcation in a live predator-prey system. *Science*, 290(5495): 1358–1360, 2000. xiii, 3, 29, 62, 63, 71, 72, 153, 154, 164

[71] Alison P Galvani, Chris T Bauch, Madhur Anand, Burton H Singer, and Simon A Levin. Human–environment interactions in population and ecosystem health. *Proceedings of the National Academy of Sciences*, 2016. 9

[72] Hu Gang, T Ditzinger, CZ Ning, and H Haken. Stochastic resonance without external periodic force. *Physical Review Letters*, 71(6):807, 1993. 54

[73] Crispin W Gardiner et al. *Handbook of stochastic methods*, volume 3. Springer Berlin, 1985. 32, 45, 46, 48, 51, 55, 142

[74] Surgeon General. The health consequences of smoking—50 years of progress: a report of the surgeon general. In *US Department of Health and Human Services*. Citeseer, 2014. 5

[75] Amin Ghadami, Eleni Gourgou, and Bogdan I Epureanu. Rate of recovery from perturbations as a means to forecast future stability of living systems. *Scientific Reports*, 8(1):9271, 2018. 41, 65

[76] Roger M Gifford. Implications of co 2 effects on vegetation for the global carbon budget. In *The global carbon cycle*, pages 159–199. Springer, 1993. 132

[77] Malcolm Gladwell. *The tipping point: How little things can make a big difference*. Little, Brown, 2006. 1

[78] EA Gopalakrishnan, Yogita Sharma, Tony John, Partha Sharathi Dutta, and RI Sujith. Early warning signals for critical transitions in a thermoacoustic system. *Scientific reports*, 6, 2016. 42, 54

[79] Mark Granovetter. Threshold models of collective behavior. *American journal of sociology*, 83(6):1420–1443, 1978. 92

[80] Blaine D Griffen and John M Drake. Effects of habitat quality and size on extinction in experimental populations. *Proceedings of the Royal Society B: Biological Sciences*, 275(1648):2251–2256, 2008. 77

[81] Alena Sonia Gsell, Ulrike Scharfenberger, Deniz Özkundakci, Annika Walters, Lars-Anders Hansson, Annette BG Janssen, Peeter Nõges, Philip C Reid, Daniel E Schindler, Ellen Van Donk, et al. Evaluating early-warning indicators of critical transitions in natural aquatic ecosystems. *Proceedings of the National Academy of Sciences*, page 201608242, 2016. 42

[82] Vishwesha Guttal and Ciriyam Jayaprakash. Changing skewness: an early warning signal of regime shifts in ecosystems. *Ecology letters*, 11(5):450–460, 2008. 39, 84

[83] Nick M Haddad, Lars A Brudvig, Jean Clobert, Kendi F Davies, Andrew Gonzalez, Robert D Holt, Thomas E Lovejoy, Joseph O Sexton, Mike P Austin, Cathy D Collins, et al. Habitat fragmentation and its lasting impact on earth's ecosystems. *Science advances*, 1(2):e1500052, 2015. 74

[84] Hiroshi Hakoyama and Yoh Iwasa. Extinction risk of a density-dependent population estimated from a time series of population size. *Journal of Theoretical Biology*, 204 (3):337–359, 2000. 77

[85] TG Hallam and CE Clark. Non-autonomous logistic equations as models of populations in a deteriorating environment. *Journal of Theoretical Biology*, 93(2):303–311, 1981. 76, 77

[86] James Hansen, Larissa Nazarenko, Reto Ruedy, Makiko Sato, Josh Willis, Anthony Del Genio, Dorothy Koch, Andrew Lacis, Ken Lo, Surabi Menon, et al. Earth's energy imbalance: Confirmation and implications. *science*, 308(5727):1431–1435, 2005. 19

[87] Alan Hastings and Derin B Wysham. Regime shifts in ecological systems can occur with no warning. *Ecology letters*, 13(4):464–472, 2010. 41, 66

[88] Dirk Helbing. *Quantitative sociodynamics: stochastic methods and models of social interaction processes*. Springer Science & Business Media, 2010. 11, 12, 119

[89] Hermann Held and Thomas Kleinen. Detection of climate system bifurcations by degenerate fingerprinting. *Geophysical Research Letters*, 31(23), 2004. 45

[90] Kirsten A Henderson, Madhur Anand, and Chris T Bauch. Carrot or stick? modelling how landowner behavioural responses can cause incentive-based forest governance to backfire. *PloS one*, 8(10), 2013. 92

[91] Kirsten A Henderson, Chris T Bauch, and Madhur Anand. Alternative stable states and the sustainability of forests, grasslands, and agriculture. *Proceedings of the National Academy of Sciences*, 113(51):14552–14559, 2016. 9, 12, 119, 122

[92] Josef Hofbauer and Karl Sigmund. *Evolutionary games and population dynamics*. Cambridge university press, 1998. 5, 13

[93] Crawford S Holling. Resilience and stability of ecological systems. *Annual review of ecology and systematics*, pages 1–23, 1973. 44

[94] John T Houghton. *Climate change 1995: The science of climate change: contribution of working group I to the second assessment report of the Intergovernmental Panel on Climate Change*, volume 2. Cambridge University Press, 1996. 132

[95] Peter D Howe, Ezra M Markowitz, Tien Ming Lee, Chia-Ying Ko, and Anthony Leiserowitz. Global perceptions of local temperature change. *Nature Climate Change*, 3(4):352, 2013. 10

[96] Clinton Innes, Madhur Anand, and Chris T Bauch. The impact of human-environment interactions on the stability of forest-grassland mosaic ecosystems. *Scientific reports*, 3:2689, 2013. 12, 119, 122

[97] Raymond L Ison, Kevin B Collins, and Philip J Wallis. Institutionalising social learning: Towards systemic and adaptive governance. *Environmental Science & Policy*, 53:105–117, 2015. 23

[98] Anthony R Ives. Measuring resilience in stochastic systems. *Ecological Monographs*, 65(2):217–233, 1995. 31

[99] Anthony R Ives and Vasilis Dakos. Detecting dynamical changes in nonlinear time series using locally linear state-space models. *Ecosphere*, 3(6):1–15, 2012. 40

[100] Rajat Karnatak, Holger Kantz, and Stephan Bialonski. Early warning signal for interior crises in excitable systems. *Physical Review E*, 96(4):042211, 2017. 38, 39, 41, 66

[101] James F Kasting, Daniel P Whitmire, and Ray T Reynolds. Habitable zones around main sequence stars. *Icarus*, 101(1):108–128, 1993. 132

[102] Sonia Kéfi, Vasilis Dakos, Marten Scheffer, Egbert H Van Nes, and Max Rietkerk. Early warning signals also precede non-catastrophic transitions. *Oikos*, 122(5):641–648, 2013. 38, 45, 54

[103] Sonia Kefi, Vishwesha Guttal, William A Brock, Stephen R Carpenter, Aaron M Ellison, Valerie N Livina, David A Seekell, Marten Scheffer, Egbert H van Nes, and Vasilis Dakos. Early warning signals of ecological transitions: methods for spatial patterns. *PloS one*, 9(3):e92097, 2014. 28

[104] Sonia Kéfi, Milena Holmgren, and Marten Scheffer. When can positive interactions cause alternative stable states in ecosystems? *Functional Ecology*, 30(1):88–97, 2016. 30

[105] Jeremy T Kerr and David J Currie. Effects of human activity on global extinction risk. *Conservation Biology*, 9(6):1528–1538, 1995. 74

[106] Thomas Kleinen, Hermann Held, and Gerhard Petschel-Held. The potential role of spectral properties in detecting thresholds in the earth system: application to the thermohaline circulation. *Ocean Dynamics*, 53(2):53–63, 2003. 31, 38, 39, 45, 46

[107] Reto Knutti and Jan Sedláček. Robustness and uncertainties in the new cmip5 climate model projections. *Nature Climate Change*, 3(4):369, 2013. 11, 92

[108] M Kot and WM Schaffer. The effects of seasonality on discrete models of population growth. *Theoretical Population Biology*, 26(3):340–360, 1984. 75

[109] Mark Kot. *Elements of mathematical ecology*. Cambridge University Press, 2001. 74

[110] Christian Kuehn. A mathematical framework for critical transitions: normal forms, variance and applications. *Journal of nonlinear science*, 23(3):457–510, 2013. 45

[111] Christian Kuehn, Gerd Zschaler, and Thilo Gross. Early warning signs for saddle-escape transitions in complex networks. *Scientific reports*, 5, 2015. 136

[112] Eizi Kuno. Some strange properties of the logistic equation defined with r and k: Inherent defects or artifacts? *Population Ecology*, 33(1):33–39, 1991. 76, 77

[113] Shoji Kusunoki. Is the global atmospheric model mri-agcm3. 2 better than the cmip5 atmospheric models in simulating precipitation over east asia? *Climate Dynamics*, pages 1–22, 2016. 11

[114] Yuri A Kuznetsov. *Elements of applied bifurcation theory*, volume 112. Springer Science & Business Media, 2013. 2, 27, 28, 31, 45, 47, 51, 93, 140, 146

[115] Russell Lande, Steinar Engen, and Bernt-Erik Saether. *Stochastic population dynamics in ecology and conservation*. Oxford University Press on Demand, 2003. 74

[116] Maria Knight Lapinski and Rajiv N Rimal. An explication of social norms. *Communication theory*, 15(2):127–147, 2005. 4

[117] Timothy M Lenton. Land and ocean carbon cycle feedback effects on global warming in a simple earth system model. *Tellus B: Chemical and Physical Meteorology*, 52 (5):1159–1188, 2000. 11, 12, 14, 16, 17, 92, 119, 121, 132

[118] 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. 14, 24, 122

[119] Joosup Lim and Bogdan I Epureanu. Forecasting a class of bifurcations: Theory and experiment. *Physical Review E*, 83(1):016203, 2011. 38, 39, 41, 65

[120] Valerie N Livina and Timothy M Lenton. A modified method for detecting incipient bifurcations in a dynamical system. *Geophysical Research Letters*, 34(3), 2007. 39

[121] Eduardo Liz. Effects of strength and timing of harvest on seasonal population models: stability switches and catastrophic shifts. *Theoretical Ecology*, 10(2):235–244, 2017. 75

[122] Donald Ludwig. The distribution of population survival times. *The American Naturalist*, 147(4):506–526, 1996. 74

[123] Robert M May. Will a large complex system be stable? *Nature*, 238(5364):413–414, 1972. 133

[124] Robert M May. Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature*, 269(5628):471–477, 1977. 35, 36, 37, 44

[125] Robert M May. Uses and abuses of mathematics in biology. *Science*, 303(5659): 790–793, 2004. 11

[126] Robert M May et al. Biological populations with nonoverlapping generations: stable points, stable cycles, and chaos. *Science*, 186(4164):645–647, 1974. 50, 75

[127] Alan J McKane and Timothy J Newman. Predator-prey cycles from resonant amplification of demographic stochasticity. *Physical review letters*, 94(21):218102, 2005. 46, 54

[128] Madan Lal Mehta. *Random matrices*, volume 142. Elsevier, 2004. 133

[129] Christian Meisel, Andreas Klaus, Christian Kuehn, and Dietmar Plenz. Critical slowing down governs the transition to neuron spiking. *PLoS Comput Biol*, 11(2): e1004097, 2015. 38, 39, 42, 45, 65, 66

[130] Jose M Montoya and Ricard V Solé. Topological properties of food webs: from real data to community assembly models. *Oikos*, 102(3):614–622, 2003. 134

[131] Akihiko Mougi and Michio Kondoh. Diversity of interaction types and ecological community stability. *Science*, 337(6092):349–351, 2012. 133, 136

[132] James Murdock. *Normal forms and unfoldings for local dynamical systems*. Springer Science & Business Media, 2006. 93

[133] KE Muryshev, AV Eliseev, II Mokhov, and AV Timazhev. A lag between temperature and atmospheric co2 concentration based on a simple coupled model of climate and the carbon cycle. In *Doklady Earth Sciences*, volume 463, pages 863–867. Springer, 2015. 14, 119, 120, 121, 132

[134] Shinichi Nakajima, Yoshi-Yuki Hayashi, and Yutaka Abe. A study on the "runaway greenhouse effect" with a one-dimensional radiative–convective equilibrium model. *Journal of the Atmospheric Sciences*, 49(23):2256–2266, 1992. 11, 132

[135] Navaratnam Sri Namachchivaya. Random dynamical systems: addressing uncertainty, nonlinearity and predictability. *Meccanica*, 51(12):2975–2995, 2016. 94

[136] Ehsan Negahbani, D Alistair Steyn-Ross, Moira L Steyn-Ross, Marcus T Wilson, and Jamie W Sleigh. Noise-induced precursors of state transitions in the stochastic wilson–cowan model. *The Journal of Mathematical Neuroscience (JMN)*, 5(1):9, 2015. 46

[137] Ehsan Negahbani, D Alistair Steyn-Ross, Moira L Steyn-Ross, and Luis A Aguirre. Which system variables carry robust early signs of upcoming phase transition? an ecological example. *PloS one*, 11(9):e0163003, 2016. 38, 41

[138] Matthew Newville, Till Stensitzki, Daniel B Allen, Michal Rawlik, Antonino Ingargiola, and Andrew Nelson. Lmfit: non-linear least-square minimization and curve-fitting for python. *Astrophysics Source Code Library*, 2016. 69, 158

[139] Jessica M Nolan, P Wesley Schultz, Robert B Cialdini, Noah J Goldstein, and Vladas Griskevicius. Normative social influence is underdetected. *Personality and social psychology bulletin*, 34(7):913–923, 2008. 10, 12

[140] Esben M Olsen, Mikko Heino, George R Lilly, M Joanne Morgan, John Brattey, Bruno Ernande, and Ulf Dieckmann. Maturation trends indicative of rapid evolution preceded the collapse of northern cod. *Nature*, 428(6986):932, 2004. 59

[141] Brian C O'Neill, Elmar Kriegler, Keywan Riahi, Kristie L Ebi, Stephane Hallegatte, Timothy R Carter, Ritu Mathur, and Detlef P van Vuuren. A new scenario framework for climate change research: the concept of shared socioeconomic pathways. *Climatic Change*, 122(3):387–400, 2014. 10

[142] Tamer Oraby, Vivek Thampi, and Chris T Bauch. The influence of social norms on the dynamics of vaccinating behaviour for paediatric infectious diseases. *Proceedings of the Royal Society B: Biological Sciences*, 281(1780):20133172, 2014. 4

[143] Suzanne M O'Regan and Danielle L Burton. How stochasticity influences leading indicators of critical transitions. *Bulletin of Mathematical Biology*, pages 1–25, 2018. 42, 45, 46

[144] Otso Ovaskainen and Baruch Meerson. Stochastic models of population extinction. *Trends in ecology & evolution*, 25(11):643–652, 2010. 74, 77

[145] Michela Pacifici, Wendy B Foden, Piero Visconti, James EM Watson, Stuart HM Butchart, Kit M Kovacs, Brett R Scheffers, David G Hole, Tara G Martin, H Resit Akcakaya, et al. Assessing species vulnerability to climate change. *Nature climate change*, 5(3):215–224, 2015. 74

[146] A Demetri Pananos, Thomas M Bury, Clara Wang, Justin Schonfeld, Sharada P Mohanty, Brendan Nyhan, Marcel Salathé, and Chris T Bauch. Critical dynamics in population vaccinating behavior. *Proceedings of the National Academy of Sciences*, page 201704093, 2017. 2, 45

[147] Charles T Perretti and Stephan B Munch. Regime shift indicators fail under noise levels commonly observed in ecological systems. *Ecological Applications*, 22(6):1772–1779, 2012. 42, 66

[148] Bruce G Peterjohn, John R Sauer, and Chandler S Robbins. Population trends from the north american breeding bird survey. 1995. 89

[149] Dimitris N Politis and Joseph P Romano. The stationary bootstrap. *Journal of the American Statistical association*, 89(428):1303–1313, 1994. 71

[150] Joshua Prettyman, Tobias Kuna, and Valerie Livina. A novel scaling indicator of early warning signals helps anticipate tropical cyclones. *EPL (Europhysics Letters)*, 121(1):10002, 2018. 47

[151] Maurice Bertram Priestley. *Spectral analysis and time series*, volume 1. Academic press London, 1981. 46

[152] Shanshan Qin and Chao Tang. Early-warning signals of critical transition: Effect of extrinsic noise. *Physical Review E*, 97(3):032406, 2018. 41, 42, 46

[153] Thomas Quail, Alvin Shrier, and Leon Glass. Predicting the onset of period-doubling bifurcations in noisy cardiac systems. *Proceedings of the National Academy of Sciences*, 112(30):9358–9363, 2015. 3

[154] 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, 2013. 39

[155] Mark Reed, Anna Evely, Georgina Cundill, Ioan Fazey, Jayne Glass, Adele Laing, Jens Newig, Brad Parrish, Christina Prell, Chris Raymond, et al. What is social learning? *Ecology and Society*, 15(4), 2010. 9

[156] Daniel C Reuman, Robert A Desharnais, Robert F Costantino, Omar S Ahmad, and Joel E Cohen. Power spectra reveal the influence of stochasticity on nonlinear population dynamics. *Proceedings of the National Academy of Sciences*, 103(49): 18860–18865, 2006. 46

[157] William Edwin Ricker. Stock and recruitment. *Journal of the Fisheries Board of Canada*, 11(5):559–623, 1954. 75

[158] Luca Ridolfi, Paolo D'Odorico, and Francesco Laio. Indicators of collapse in systems undergoing unsustainable growth. *Bulletin of mathematical biology*, 77(2):339–347, 2015. 46

[159] Luca Rindi, Martina Dal Bello, Lei Dai, Jeff Gore, and Lisandro Benedetti-Cecchi. Direct observation of increasing recovery length before collapse of a marine benthic ecosystem. *Nature Ecology & Evolution*, 1(6):s41559–017, 2017. 28

[160] Jörgen Ripa and Per Lundberg. Noise colour and the risk of population extinctions. *Proc. R. Soc. Lond. B*, 263(1377):1751–1753, 1996. 167

[161] Paul Ritchie and Jan Sieber. Early-warning indicators for rate-induced tipping. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 26(9):093116, 2016. 40

[162] Johan Rockström, Will Steffen, Kevin Noone, Åsa Persson, F Stuart Chapin, Eric F Lambin, Timothy M Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber, et al. A safe operating space for humanity. *nature*, 461(7263):472–475, 2009. 90

[163] Michael L Rosenzweig. Paradox of enrichment: destabilization of exploitation ecosystems in ecological time. *Science*, 171(3969):385–387, 1971. 29, 85

[164] Lucia Russo, Konstantinos Spiliotis, Francesco Giannino, Stefano Mazzoleni, and Constantinos Siettos. Bautin bifurcations in a forest-grassland ecosystem with human-environment interactions. *Scientific Reports*, 9(1):2665, 2019. 66

[165] Kevin Schaefer, Hugues Lantuit, Vladimir E Romanovsky, Edward AG Schuur, and Ronald Witt. The impact of the permafrost carbon feedback on global climate. *Environmental Research Letters*, 9(8):085003, 2014. 24

[166] Marten Scheffer. *Critical transitions in nature and society*. Princeton University Press, 2009. 30

[167] Marten Scheffer, Steve Carpenter, Jonathan A Foley, Carl Folke, and Brian Walker. Catastrophic shifts in ecosystems. *Nature*, 413(6856):591–596, 2001. 2, 29, 30

[168] 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. 3, 28, 45, 75

[169] Marten Scheffer, Stephen R Carpenter, Vasilis Dakos, and Egbert H van Nes. Generic indicators of ecological resilience: inferring the chance of a critical transition. *Annual Review of Ecology, Evolution, and Systematics*, 46:145–167, 2015. 37, 82

[170] Jerald L Schnoor. The ipcc fourth assessment, 2007. 17

[171] David A Seekell, Stephen R Carpenter, and Michael L Pace. Conditional heteroscedasticity as a leading indicator of ecological regime shifts. *The American Naturalist*, 178(4):442–451, 2011. 39

[172] Nahal Sharafi, Marc Timme, and Sarah Hallerberg. Critical transitions and perturbation growth directions. *arXiv preprint arXiv:1707.08082*, 2017. 41

[173] Yogita Sharma, Partha Sharathi Dutta, and AK Gupta. Anticipating regime shifts in gene expression: The case of an autoactivating positive feedback loop. *Physical Review E*, 93(3):032404, 2016. 42

[174] Eliot R Smith, Diane M Mackie, and Heather M Claypool. *Social psychology*. Psychology Press, 2014. 10

[175] HJ Sommers, A Crisanti, Haim Sompolinsky, and Y Stein. Spectrum of large random asymmetric matrices. *Physical review letters*, 60(19):1895, 1988. 134, 137, 138

[176] Will Steffen, Johan Rockström, Katherine Richardson, Timothy M Lenton, Carl Folke, Diana Liverman, Colin P Summerhayes, Anthony D Barnosky, Sarah E Cornell, Michel Crucifix, et al. Trajectories of the earth system in the anthropocene. *Proceedings of the National Academy of Sciences*, 115(33):8252–8259, 2018. 24, 25

[177] John Sterman, Thomas Fiddaman, Travis Franck, Andrew Jones, Stephanie McCauley, Philip Rice, Elizabeth Sawin, and Lori Siegel. Climate interactive: the c-roads climate policy model. *System Dynamics Review*, 28(3):295–305, 2012. 9, 10, 11, 92

[178] Nicholas Stern, Siobhan Peters, Vicki Bakhshi, Alex Bowen, Catherine Cameron, Sebastian Catovsky, Diane Crane, Sophie Cruickshank, Simon Dietz, Nicola Edmonson, et al. *Stern Review: The economics of climate change*, volume 30. HM treasury London, 2006. 17

[179] Thomas F Stocker, Dahe Qin, G-K Plattner, Melinda MB Tignor, Simon K Allen, Judith Boschung, Alexander Nauels, Yu Xia, Vincent Bex, and Pauline M Midgley. Climate change 2013: The physical science basis. contribution of working group i to the fifth assessment report of ipcc the intergovernmental panel on climate change, 2014. 24, 124

[180] Susanne Stoll-Kleemann, Tim O'Riordan, and Carlo C Jaeger. The psychology of denial concerning climate mitigation measures: evidence from swiss focus groups. *Global environmental change*, 11(2):107–117, 2001. 10, 17, 18

[181] Steven H Strogatz. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Westview press, 2014. 2, 27, 28, 30, 41, 45

[182] Alex O Sutton, Dan Strickland, Nikole E Freeman, Amy EM Newman, and D Ryan Norris. Autumn freeze-thaw events carry over to depress late-winter reproductive performance in canada jays. *Royal Society Open Science*, 6(4):181754, 2019. 75

[183] Samir Suweis and Paolo D'Odorico. Early warning signs in social-ecological networks. *PloS one*, 9(7):e101851, 2014. 40

[184] Terence Tao and Van Vu. Random matrices: the circular law. *Communications in Contemporary Mathematics*, 10(02):261–307, 2008. 134, 135

[185] Rene Thom. Structural stability, catastrophe theory, and applied mathematics. *SIAM review*, 19(2):189–201, 1977. 1

[186] JR Toggweiler and Jorge L Sarmiento. Glacial to interglacial changes in atmospheric carbon dioxide: The critical role of ocean surface water in high latitudes. *The carbon cycle and atmospheric CO2: natural variations archean to present*, 32:163–184, 1985. 132

[187] Billie Lee Turner, Pamela A Matson, James J McCarthy, Robert W Corell, Lindsey Christensen, Noelle Eckley, Grete K Hovelsrud-Broda, Jeanne X Kasperson, Roger E Kasperson, Amy Luers, et al. Illustrating the coupled human–environment system for vulnerability analysis: three case studies. *Proceedings of the National Academy of Sciences*, 100(14):8080–8085, 2003. 23

[188] George E Uhlenbeck and Leonard S Ornstein. On the theory of the brownian motion. *Physical review*, 36(5):823, 1930. 32, 46

[189] Nicolaas Godfried Van Kampen. *Stochastic processes in physics and chemistry*, volume 1. Elsevier, 1992. 77

[190] Egbert H Van Nes and Marten Scheffer. Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *The American Naturalist*, 169(6): 738–747, 2007. 3, 39

[191] Egbert H van Nes, Babak MS Arani, Arie Staal, Bregje van der Bolt, Bernardo M Flores, Sebastian Bathiany, and Marten Scheffer. What do you mean,'tipping point'? *Trends in Ecology & Evolution*, 31(12):902–904, 2016. 1

[192] Detlef P Van Vuuren, Jae Edmonds, Mikiko Kainuma, Keywan Riahi, Allison Thomson, Kathy Hibbard, George C Hurtt, Tom Kram, Volker Krey, Jean-Francois Lamarque, et al. The representative concentration pathways: an overview. *Climatic change*, 109(1-2):5, 2011. 9, 10

[193] Adam J Vanbergen, Anahí Espíndola, and Marcelo A Aizen. Risks to pollinators and pollination from invasive alien species. *Nature ecology & evolution*, 2(1):16–25, 2018. 74

[194] Annelies J Veraart, Elisabeth J Faassen, Vasilis Dakos, Egbert H van Nes, Miquel Lürling, and Marten Scheffer. Recovery rates reflect distance to a tipping point in a living system. *Nature*, 481(7381):357–359, 2012. 30, 75

[195] Eric-Jan Wagenmakers and Simon Farrell. Aic model selection using akaike weights. *Psychonomic bulletin & review*, 11(1):192–196, 2004. 57, 70

[196] Rong Wang, John A Dearing, Peter G Langdon, Enlou Zhang, Xiangdong Yang, Vasilis Dakos, and Marten Scheffer. Flickering gives early warning signals of a critical transition to a eutrophic lake state. *Nature*, 492(7429):419–422, 2012. 31

[197] Colleen Webb. A complete classification of darwinian extinction in ecological interactions. *The American Naturalist*, 161(2):181–205, 2003. 66

[198] Elke U Weber. What shapes perceptions of climate change? *Wiley Interdisciplinary Reviews: Climate Change*, 1(3):332–342, 2010. 10, 13

[199] 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. 41, 93

[200] Peter Welch. The use of fast fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Transactions on audio and electroacoustics*, 15(2):70–73, 1967. 68, 69

[201] Kurt Wiesenfeld. Noisy precursors of nonlinear instabilities. *Journal of Statistical Physics*, 38(5):1071–1097, 1985. 46

[202] Eugene P Wigner. On the distribution of the roots of certain symmetric matrices. *Annals of Mathematics*, pages 325–327, 1958. 133, 134

[203] Mark S Williamson and Timothy M Lenton. Detection of bifurcations in noisy coupled systems from multiple time series. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(3):036407, 2015. 38, 39, 40, 41, 65

[204] C Wissel. A universal law of the characteristic return time near thresholds. *Oecologia*, 65(1):101–107, 1984. 30, 38

[205] Bradley K Woodworth, Nathaniel T Wheelwright, Amy E Newman, Michael Schaub, and D Ryan Norris. Winter temperatures limit population growth rate of a migratory songbird. *Nature communications*, 8:14812, 2017. 75

[206] N Wouters, V Dakos, M Edwards, MP Serafim, PJ Valayer, and HN Cabral. Evidencing a regime shift in the north sea using early-warning signals as indicators of critical transitions. *Estuarine, Coastal and Shelf Science*, 152:65–72, 2015. 31

[207] Qiang Yang, Mike S Fowler, Andrew L Jackson, and Ian Donohue. The predictability of ecological stability in a noisy world. *Nature Ecology & Evolution*, 3(2):251, 2019. 42

APPENDICES

Appendix A

Charting pathways to climate change mitigation in a coupled socio-climate model

A.1 Full model and parameter values

We couple an Earth system model (ESM) [117] with reduced ocean dynamics [133], to a dynamic model for social behaviour [88, 96, 91]. The full socio-climate model reads

$$\frac{dx}{dt} = \kappa x(1-x)(-\beta + f(T_f) + \delta(2x-1)), \tag{A.1}$$

$$\frac{dC_{\text{at}}}{dt} = \epsilon(t)\frac{1-x}{1-x_0} - P + R_{\text{veg}} + R_{\text{so}} - F_{\text{oc}}, \tag{A.2}$$

$$\frac{dC_{\text{oc}}}{dt} = F_{\text{oc}}, \tag{A.3}$$

$$\frac{dC_{\text{veg}}}{dt} = P - R_{\text{veg}} - L, \tag{A.4}$$

$$\frac{dC_{\text{so}}}{dt} = L - R_{\text{so}}, \tag{A.5}$$

$$c\frac{dT}{dt} = (F_d - \sigma T^4)a_E. \tag{A.6}$$

Climate variables are expressed as deviations from pre-industrial levels. Definitions of state variables and climate 'processes' are given in Table A.1, and baseline parameter values are provided in Table A.2. Functional forms for each process are outlined below.

Photosynthesis

Carbon uptake from the atmosphere via photosynthesis takes the following form

$$P(C_{\text{at}}, T) = k_p C_{\text{ve0}} k_{MM} \left(\frac{\text{pCO}_{2a} - k_c}{K_M + \text{pCO}_{2a} - k_c}\right)\left(\frac{(15+T)^2(25-T)}{5625}\right) \tag{A.7}$$

for $\text{pCO}_{2a} \geq k_c$ and $-15 \leq T \leq 25$, and zero otherwise. The mixing ratio of CO_2 in the atmosphere, pCO_{2a} is defined as the ratio of moles of CO_2 in the atmosphere to the total number of moles of molecules in the atmosphere k_a. Thus

$$\text{pCO}_{2a} = \frac{f_{\text{gtm}}(C_{\text{at}} + C_{\text{at0}})}{k_a} \tag{A.8}$$

where $f_{\text{gtm}} = 8.3259 \times 10^{13}$ is the conversion factor from gTC to moles of carbon and C_{at0} is initial level of CO_2 in the atmosphere. Note photosynthesis satisfies Michaelis-Menton kinetics in pCO_{2a} resembling increasing but saturating rates of photosynthesis as carbon in the atmosphere increases. The temperature term captures optimal photosynthesis at $T = 2$ (atmospheric temp. of 27°C) with declining rates for further increases in temperature.

Respiration

Plant respiration takes the form

$$R_{\text{veg}}(T, C_{\text{veg}}) = k_r C_{\text{veg}} k_A e^{-\frac{E_a}{R(T+T_0)}} \tag{A.9}$$

which increases with the amount of carbon present in the vegetation, and also with the temperature. This provides a positive feedback with increasing carbon levels. Soil respiration takes an analogous form:

$$R_{\text{so}}(T, C_{\text{so}}) = k_{sr} C_{\text{so}} k_B e^{-\frac{308.56}{T+T_0-227.13}}. \tag{A.10}$$

Turnover

There is an assumed constant fraction of plants dying in a given unit of time:

$$L(C_{\text{veg}}) = k_t C_{\text{veg}} \tag{A.11}$$

The stored carbon is then fed into the soil reservoir.

Ocean flux

Flux of CO_2 from the atmosphere to the ocean takes the form

$$F_{\text{oc}}(C_{\text{at}}, C_{\text{oc}}) = F_0 \chi \left(C_{\text{at}} - \zeta \frac{C_{\text{at0}}}{C_{\text{oc0}}} C_{\text{oc}} \right) \tag{A.12}$$

where χ is characteristic solubility of CO_2 in water and ζ is the evasion factor [133]. More complex ocean models couple these parameters to chemical dynamics within the ocean

itself [117], however we find good agreement when comparing this simplified climate model with the full Earth system model presented in [117].

Atmospheric Dynamics

We use the "grey-atmosphere approximation" as used in [117] to model atmospheric dynamics. This framework captures changes in global average surface temperature due to changes in albedo, solar flux, and the opacity of CO_2, H_2O_v and CH_4. The net downward flux of radiation absorbed at the planet's surface is given by

$$F_d = \frac{(1-A)S}{4}\left(1+\frac{3}{4}\tau\right), \tag{A.13}$$

where A is the surface albedo, S is the incoming solar flux and τ vertical opacity of the greenhouse atmosphere. Expressions for each opacity are given by

$$\tau(CO_2) = 1.73(pCO_2)^{0.263} \tag{A.14}$$

$$\tau(H_2O) = 0.0126\left(HP_0e^{-(L/RT)}\right)^{0.503} \tag{A.15}$$

$$\tau(CH_4) = 0.0231 \tag{A.16}$$

where pCO_2 is the mixing ratio of CO_2 in the atmosphere as defined earlier, H is the relative humidity, P_0 is the water vapor saturation constant, L is the latent heat per mole of water, and R is the molar gas constant.

Climate parameters

All climate parameters are borrowed from the Earth System model [117] except for those governing carbon transfer with the ocean, where we use a simplified framework [133]. We calibrate these ocean parameters to recover dynamics of the full Earth System model. The relative humidity H is calibrated to give a pre-industrial temperature of 288.15K. Parameter values along with upper and lower bounds are listed in Table A.2.

Social Parameters

For reference, social dynamics are modeled by

$$\frac{dx}{dt} = \kappa x(1-x)(-\beta + f(T_f) + \delta(2x-1)). \tag{A.17}$$

Justification for baseline parameters are provided below. The effect of varying key social parameters is detailed in the manuscript.

Social learning rate ($\boldsymbol{\kappa}$) determines the rate at which an alternate strategy propagates throughout a population, once the alternate strategy has a higher overall utility. Population level change in consensus can take decades (e.g. smoking as a health hazard) and so we select a value of κ accordingly. The time taken for consensus change is $O(1/\kappa)$ and so we use a baseline value of $\kappa = 0.05$ to correspond to change over a 20-year period. This is similar to previous studies modeling of evolving population consensus in socio-ecological systems [91, 96]. Upper/lower bounds of κ are taken to correspond population level consensus change from 5 / 50 years respectively.

Net cost of mitigation ($\boldsymbol{\beta}$) is set to one without loss of generality. Other payoff parameters are normalised to be relative to the net cost of mitigation.

Strength of social norms ($\boldsymbol{\delta}$) is a measure of the 'cost' that an individual incurs when acting against the majority opinion. In a population with complete consensus ($x = 1, 0$), the cost to adopt the alternate strategy is exactly δ. We assume this cost is on the same order as the net cost of mitigative practices (β) and so we investigate δ for the range $(0.5, 1.5)$.

Cost of warming ($\boldsymbol{f(T_f)}$) represents the perceived costs associated with a global temperature anomaly of T_f degrees Celsius. Despite the growing acceptance of climate change and its consequences, the rate of global CO_2 emissions per capita shows little sign of slowing down (Figure A.8). Temperature records report a current temperature anomaly of 1°C above pre-industrial values and CO_2 emissions per capita are hovering at approximately $1.3\,\mathrm{tC\,yr^{-1}}$, the highest rate in history. When modeling an individual's perceived cost of global warming at a temperature of T, we therefore use a sigmoidal function (Methods). This assumes small temperature anomalies (on the order of 1°C) receive low response, as is observed in empirical data (Figure A.8). At a specified threshold temperature, the response function increases in a non-linear fashion, capturing the expected non-linear increase in frequency of climate change disasters as temperature increases [118]. The cost

function eventually saturates for high enough temperature.

A.2 Additional figures and tables

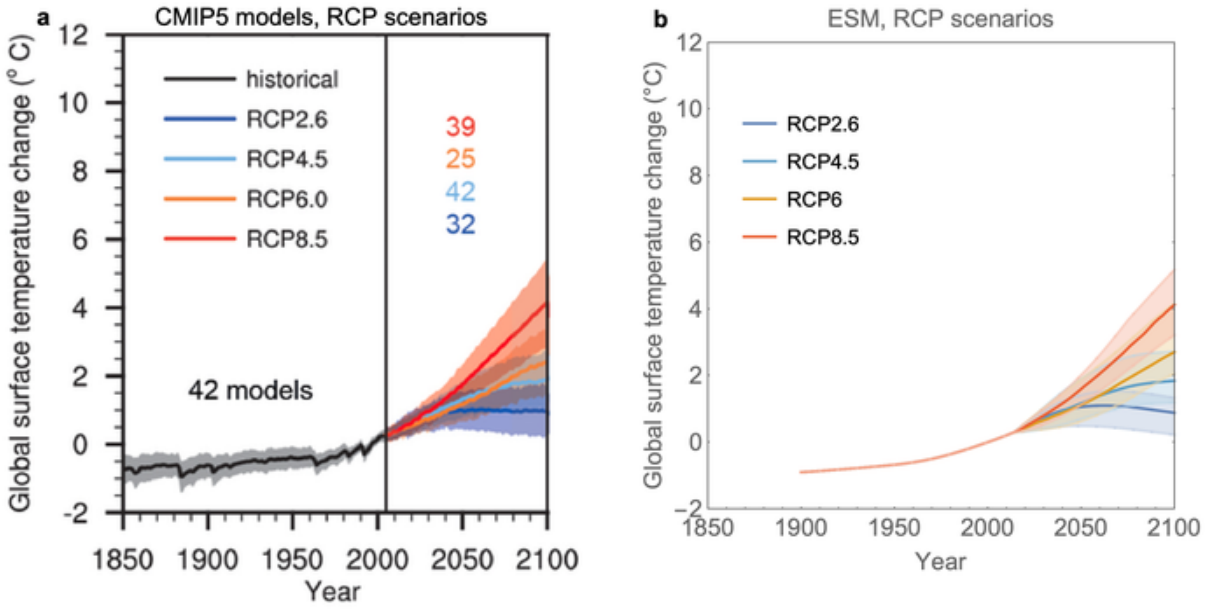

Figure A.1. Comparing temperature projections from the simple Earth system model with those of more complex climate models. a. Ensemble of simulations from the Coupled Model Intercomparison Project Phase 5 (CMIP5) using the Representative Concentration Pathways (RCPs) as emission scenarios (figure from the IPCC Fifth Assessment Report [179]. Displayed are 95% confidence intervals based on annual means. Numbers and their colour denote the number of models used for each RCP scenario. **b.** Ensemble of simulations from the simple Earth system model that we use in our socio-climate model with the same fixed emission trajectories. Parameters are drawn from triangular distributions with upper and lower bounds given in Table A.2.

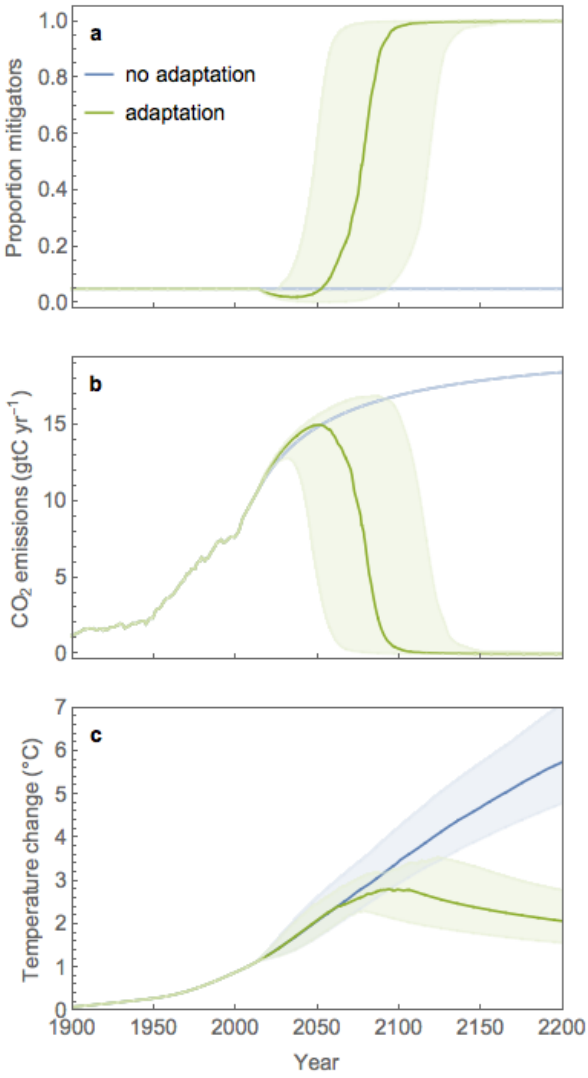

Figure A.2. Climate trends with and without adaptation to climate change. Removing adaptive behaviour from the model (by forcing the proportion of mitigators to remain at a constant, low value) results in saturating emissions and temperature increasing indefinitely (at least over the next two centuries). This is akin to the RCP8.5 scenario in the latest IPCC report (trajectory shown in Figure 1). Simulations above use parameter values drawn from triangular distributions with upper and lower bounds given in Table A.2.

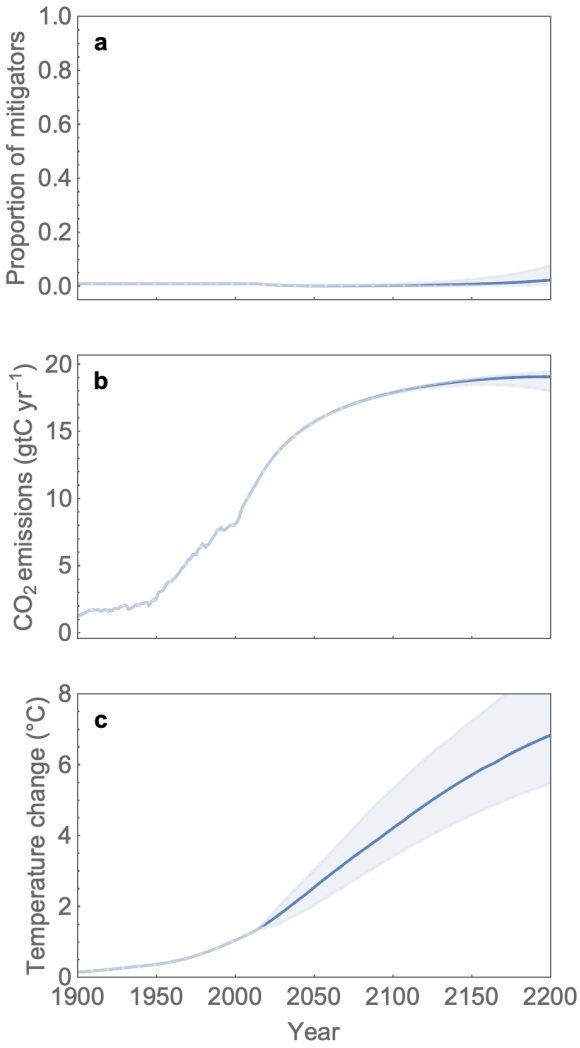

Figure A.3. Worst-case scenario of the coupled socio-climate model. Setting the social parameters to their bound that most favour non-mitigative behaviour, we get no spread of mitigative behaviour within the considered time-frame. This causes the temperature to increase in a manner similar to the RCP 8.5 scenario (Figure A.1). Fixed parameter values are $\kappa = 0.02$, $\beta = 1.5$, $\delta = 1.5$, $f_{\max} = 4$, $x_0 = 0.01$. All other parameter values are drawn from triangular distributions with upper and lower bounds given in Table A.2.

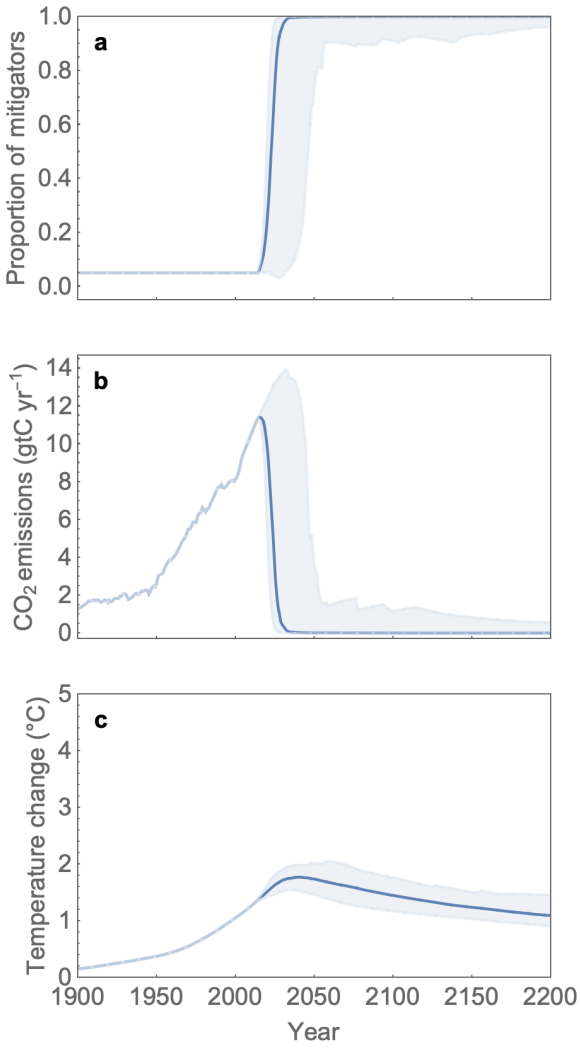

Figure A.4. Best-case scenario of the coupled socio-climate model. Setting the social parameters to their bound that least favour non-mitigative behaviour, we get very early spread of mitigative behaviour. This causes the temperature to evolve in a manner most similar to the RCP 2.6 scenario where temperature change stays below 2 degrees Celsius (Figure A.1). Fixed parameter values are $\kappa = 0.2$, $\beta = 0.5$, $\delta = 0.5$, $f_{\max} = 6$, $t_f = 50$. All other parameter values are drawn from triangular distributions with upper and lower bounds given in Table Table A.2.

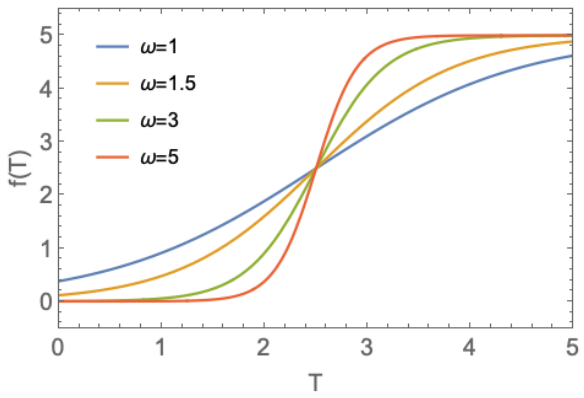

Figure A.5. Functional form for perceived costs associated with climate change. The incentive of individuals to mitigate is in part based on their perceived costs of climate change f at some projected temperature T. We adopt a sigmoidal response curve for $f(T)$ with variable curvature ω and horizontal shift T_c (the explicit form of $f(T)$ is provided in Methods). This form captures the expected non-linear increase in climate change impacts (cost) as temperature increases.

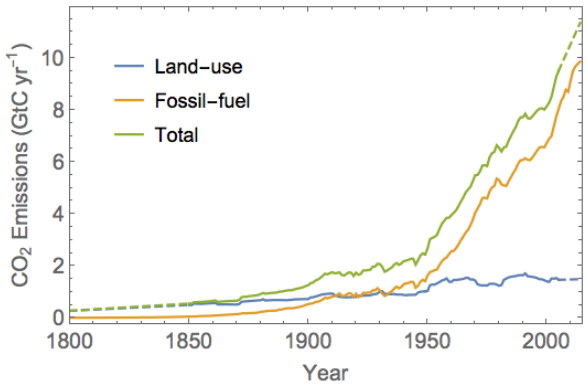

Figure A.6. Estimated historical CO_2 emissions. Data from the CDIAC on carbon emissions due to fossil-fuel burning and land-use changes, during the years 1800-2014. Land-use data is only available from 1850-2005 and so we linearly extrapolate (dashed line) to match the range of the fossil-fuel data. The sum of the emission trajectories is used to drive the model up to the year 2014, from which point the behavioural component of the socio-climate model is initiated.

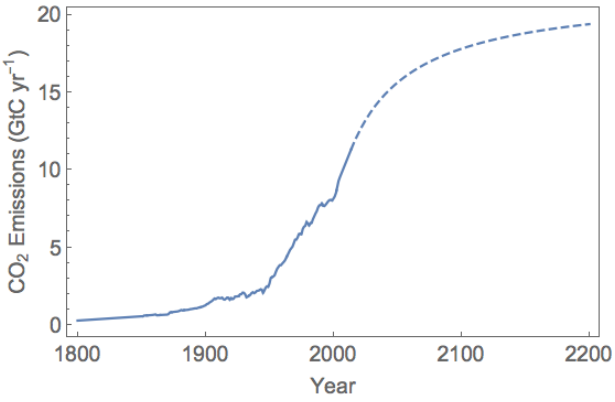

Figure A.7. Emissions in the absence of behavioral change, $\epsilon(t)$. In the socio-climate model, the factor $\epsilon(t)$ corresponds to the global CO_2 emissions should there be no change in human behavior. Pre 2014, $\epsilon(t)$ takes the historical emission trajectory (solid line) as human behaviour is not modeled here. Post 2014, $\epsilon(t)$ follows a saturating function (dashed line) to capture the saturation of global population size and energy needs. Details of this functional form can be found in the Methods section.

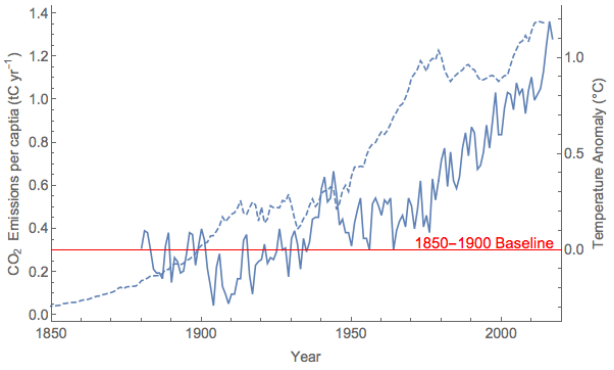

Figure A.8. Comparing per capita CO_2 emissions with global temperature changes. Temperature anomaly (solid line) above the 1850-1900 baseline value and CO_2 emissions per capita (dashed line) are shown for the 1850-2014. Data was obtained from the CDIAC data repository [17]. It is clear that despite exceeding a 1 degree temperature anomaly, the global emissions per capita of CO_2 show no obvious signs of decreasing. The current temperature anomaly is not yet high enough to spark a global decrease in emissions.

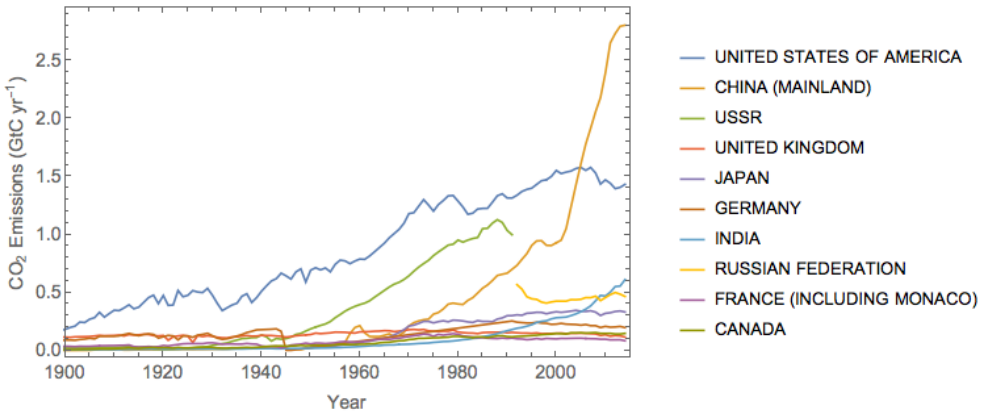

Figure A.9. CO_2 emissions by country. Total industrial CO2 emissions from the ten countries with the greatest cumulative output. Data was obtained from the CDIAC data repository [17]. Note the termination of the green curve marks the dissolution of the Soviet Union and the beginning of the yellow curve marks the formation of the Russian Federation.

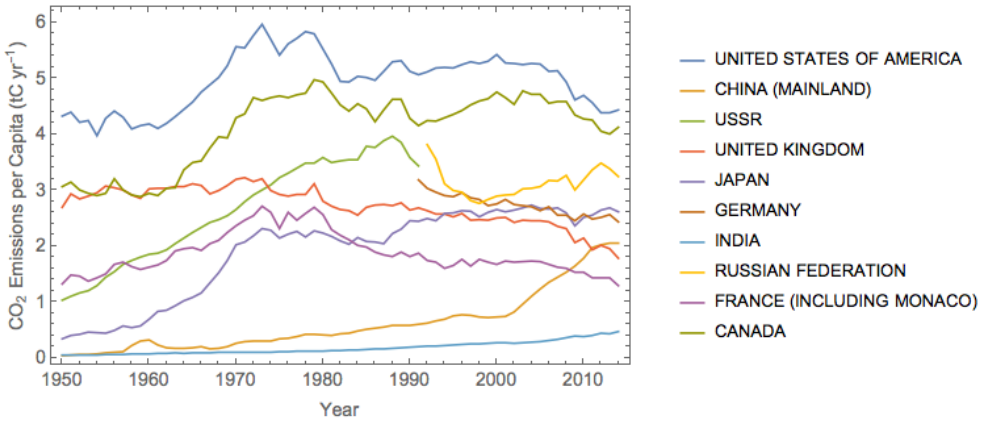

Figure A.10. CO_2 emissions per capita by country. Total industrial CO_2 emissions per capita from the ten countries with the greatest cumulative output. Data was obtained from the CDIAC data repository [17]. From 1950-1980 we observe strong increases in emissions per capita among most of these countries as is seen globally in Figure Figure A.8. Many of the countries peak in 1980 and follow slight downwards trends demonstrating a degree of country-level movement towards mitigation. Globally, this is not the case, however.

Variable / Process	Definition	Unit
x	proportion of mitigators in population	1
C_{at}	deviation of atmospheric CO2 (from pre-industrial 1800)	GtC
C_{oc}	deviation of CO2 in ocean	GtC
C_{veg}	deviation in CO2 in vegetation	GtC
C_{so}	deviation in CO2 in soil	GtC
T	deviation in temperature	K
$\epsilon(t)$	CO_2 emissions in absence of mitigation	GtC/yr
P	carbon uptake from photosynthesis	GtC/yr
R_{veg}	respiration from vegetation	GtC/yr
R_{so}	respiration from soil	GtC/yr
F_{oc}	flux of CO2 from atmosphere to ocean	GtC/yr

Table A.1. Climate model variables and dynamic processes.

Parameter	Definition	Baseline values / intervals	Unit	Source
C_{at0}	initial CO_2 in atmosphere	$(590, 596, 602)$	GtC	[94, 117]
C_{ao0}	initial CO_2 in ocean reservoir	$(1.4, 1.5, 1.6) \times 10^5$	GtC	[133]
C_{veg0}	initial CO_2 in vegetation reservoir	$(540, 550, 560)$	GtC	[94, 117]
C_{veg0}	initial CO_2 in soil reservoir	$(1480, 1500, 1520)$	GtC	[94, 117]
T_0	initial average atmospheric temperature	$(288, 288.15, 288.3)$	K	[117]
k_p	photosynthesis rate constant	$(0.175, 0.184, 0.193)$	yr^{-1}	[117, 94]
k_{MM}	photosynthesis normalising constant	1.478	1	[117]
k_c	photosynthesis compensation point	$(26, 29, 32) \times 10^{-6}$	1	[117, 76]
K_M	half-saturation point for photosynthesis	$(108, 120, 132) \times 10^{-6}$	1	[117]
k_a	mole volume of atmosphere	1.773×10^{20}	moles	[117, 186]
k_r	plant respiration constant	$(0.0828, 0.092, 0.1012)$	yr^{-1}	[117, 94]
k_A	plant respiration normalising constant	8.7039×10^9	1	[117]
E_a	plant respiration activation energy	$(54.63, 54.83, 55.03)$	J mol^{-1}	[117, 68]
k_{sr}	soil respiration rate constant	$(0.0303, 0.034, 0.037)$	yr^{-1}	[117, 94]
k_B	soil respiration normalising constant	157.072	1	[117]
k_t	turnover rate constant	$(0.0828, 0.092, 0.1012)$	yr^{-1}	[117, 94]
c	specific heat capacity of Earth's surface	$(4.22, 4.69, 5.16) \times 10^{23}$	JK^{-1}	[117, 94]
a_E	Earth's surface area	5.101×10^{14}	m^2	universal
σ	Stefan-Boltzman constant	5.67×10^{-8}	$\text{Wm}^{-2}\text{K}^{-4}$	universal
L	latent heat per mole of water	43 655	mol^{-1}	universal
R	molar gas constant	8.314	$\text{J mol}^{-1}\,\text{K}^{-1}$	universal
H	relative humidity	0.5915	1	calibrated
A	surface albedo	$(0.203, 0.225, 0.248)$	yr^{-1}	[117, 101]
S	solar flux	$(1231, 1368, 1504)$	Wm^{-2}	[117, 101]
$\tau(CH_4)$	methane opacity	$(0.0208, 0.0231, 0.0254)$	1	[117, 37]
P_0	water vapor saturation constant	$(1.26, 1.4, 1.54) \times 10^{11}$	Pa	[117, 134]
F_0	ocean flux rate constant	$(2.25, 2.5, 2.75) \times 10^{-2}$	yr^{-1}	[133]
χ	characteristic CO_2 solubility	$(0.2, 0.3, 0.4)$	1	calibrated
ζ	evasion factor	$(40, 50, 60)$	1	calibrated
κ	social learning rate	$(0.02, 0.05, 0.2)$	yr^{-1}	-
β	net cost of mitigation	$(0.5, 1, 1.5)$	1	-
δ	strength of social norms	$(0.5, 1, 1.5)$	1	-
$f_{\max}$	maximum of warming cost function $f(T)$	(4,5,6)	1	-
ω	nonlinearity of warming cost function $f(T)$	$(1, 3, 5)$	K^{-1}	-
T_c	critical temperature of $f(T)$	$(2.4, 2.5, 2.6)$	K	-
t_p	# previous years used for temperature projection	10	yr	-
t_f	# years ahead for temperature projection	$(0, 25, 50)$	yr	-
s	half-saturation time for $\epsilon(t)$ from 2014	$(30, 50, 70)$	yr	-
$\epsilon_{\max}$	maximum change in $\epsilon(t)$ from 2014	$(4.2, 7, 9.8)$	GtC yr^{-1}	-
x_0	initial proportion of mitigators	$(0.01, 0.05, 0.1)$	1	-

Table A.2. Climate model parameter values and definitions. Parameters values given as a tuple provide the lower bound, baseline, and upper bound values respectively.

Appendix B

Detecting and distinguishing tipping points using spectral early warning signals

B.1 Stability in high-dimensional systems

In this appendix, we show how results from random matrix theory can provide insight into the type of local bifurcations that are likely to occur in high-dimensional systems. This miniature study provides extra motivation for the study of EWS that are specific to bifurcation types. We begin by outlining how random matrix theory has been used in the past to draw connections between stability and complexity in complex systems, particularly ecological networks. We then make hypotheses about the types of local bifurcations that are likely to occur, given the type of interactions between the individual components of the system. Finally, we run simulations of random networks to test these hypotheses.

May's complexity-stability criterion

It was over 40 years ago, when May established a relationship between complexity and stability in food webs with a large number of species [123], using results from random matrix theory [202, 128]. Subsequently, random matrix theory has been used extensively to draw connections between complexity and stability in ecological networks [4, 131]. In

this context, complexity refers to a quantity weighted by connectance C (the probability that any pair of species interact), the average interaction strength between the species α, and the total number of species in the system n. May's result states that for large n, the system will almost certainly be stable provided

$$\alpha\sqrt{nC} < d, \tag{B.1}$$

where d is a coefficient that corresponds to the strength of self-regulatory effects (the stability of each species in isolation from the others). The system is almost certainly unstable if this condition does not hold. As a consequence, systems are likely to be unstable if they possess a large number of variables, are highly connected, or have weak self-regulatory effects. This result is in accordance with the empirical observation that species-rich ecosystems tend to have lower connectance than those with a smaller number of species [130].

The mathematics underlying this stability criterion comes from the field of random matrix theory, which has established remarkable theorems concerning the spectrum of eigenvalues for large, random matrices [202, 175, 184]. A random matrix is simply a matrix whose coefficients are drawn from a specific set probability distributions. Originally formalised in physics for studying statistical models of heavy nuclei [202], the field has grown rapidly, finding applications in e.g. data analysis, number theory and optimal control. Though how can random matrix theory bring insights to ecology?

An ecosystem can be modelled as a network, where each node corresponds to the dynamics of a single species, and connections capture the species interactions, which are in general non-linear. At equilibrium, the stability of the system may be obtained from the Jacobian matrix, whose elements correspond to the type and strength of interaction between each pair of species. Without loss of generality, we may assume that each of these elements has been drawn from some probability distribution. Random matrix theory can tell us the expected distribution of eigenvalues, given a probability distribution from which the matrix elements are drawn. This is useful, since these eigenvalues determine stability of the system, and to an extent, characterise possible bifurcations nearby.

The Circular Law

The result from random matrix theory that produced May's stability criterion (B.1) is Wigner's 'semicircle law' [202], which has now been proven in greater generality, and called

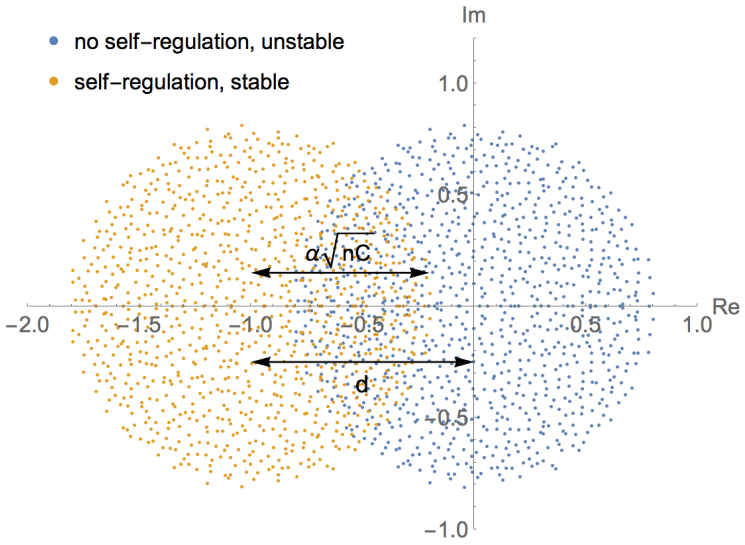

Figure B.1. Illustration of May's stability criterion. Shown are the eigenvalue spectra of two 1000×1000 matrices whose elements are, with probability C, drawn from a normal distribution with mean 0 and variance α^2 and otherwise set to 0. Without the assumption of self-regulating terms (blue) the corresponding system is unstable due to positive eigenvalues. With self-regulating terms (incorporated by replacing the diagonal elements of the matrix by $-d$), the corresponding system is stable since all eigenvalues are negative (orange). The radius of the spectrum is $\alpha\sqrt{nC}$ which must be less than d for the system to be stable (with high probability). This is the stability criterion that May derived. Parameters $d = -1$, $\alpha = 0.025$, $n = 1000$, $C = 1$.

the 'circular law' [184]. The circular law is stated simply as follows. Assuming each element of the an $n \times n$ matrix A is drawn from a probability distribution with mean zero and variance α^2, then the spectrum of eigenvalues of A converges to the uniform distribution over the disk with radius $\alpha\sqrt{n}$ centred at the origin. This property has become known as 'universality' [184]. Since ecosystems are typically sparsely connected, we can add to this and assume that each element of the matrix is assigned zero with probability $1 - C$ and is drawn from the given distribution with probability C (defining C as the connectance). The variance of the resulting distribution is reduced to $C\sigma^2$ and so using the property of universality, the eigenvalues of A converge to a disk of radius $\alpha\sqrt{nC}$ centred at the origin (Figure B.1, blue). Moreover, if one assumes that in isolation, each species exhibits stable dynamics (incorporated by setting the diagonal elements of A to $-d$), then the eigenvalue spectrum is is shifted to the left by d units (Figure B.1, orange). It is now clear that for

all eigenvalues in this system to be less than zero (with high probability) we require that the radius of the circle $\alpha\sqrt{nC}$ is less than d, which is the complexity-stability relationship that May put forward (B.1).

Interaction types and stability

Note that in deriving the stability criterion in (B.1), May considered large ecological networks in which all interactions (a_{ij}) are drawn from a single probability distribution, and thus each pair of species interacts with the same probability, and with a completely random type of interaction. Allesina and Tang [4] built upon this work by considering networks where the proportion of interaction types (e.g. mutualistic, cooperative, exploitative) could be specified. For example, imposing a predator prey interaction would necessitate interaction terms a_{ij} and a_{ji} having opposite signs, imposing a mutualistic interaction would necessitate a_{ij} and a_{ji} both being positive. We introduce their work here, since we find that the proportion of each interaction type in the system not only has a significant effect on stability, but also on the way in which stability may be lost, i.e. the type of bifurcation.

Allesina and Tang found that imposing interaction types on the system morphed the eigenvalue spectrum into an ellipse (Figure B.2). In particular, exploitative interactions (such as those found in predator-prey systems) narrowed the ellipse, reducing the size of the largest eigenvalue and therefore enhancing stability. Cooperative and competitive interactions widened the ellipse, increasing the size of the largest eigenvalue and therefore diminishing stability. This profound result demonstrates that random matrix theory has a lot to offer in understanding what makes complex systems stable. We suggest that random matrix theory can also give insight into the way in which a complex system destabilises, i.e. the type of bifurcation that it go through.

Random matrix theory and bifurcation type

Although much progress has been made on the stability of high-dimensional systems in theoretical and applied work, [131, 4, 49, 111] the description of ensuing dynamics following an instability is less clear. The loss in stability of a resident stable state can lead to all sorts of behaviour, including e.g. oscillations, chaos, or a transition to some faraway attractor. The immediate dynamics following this instability may be understood via the underlying

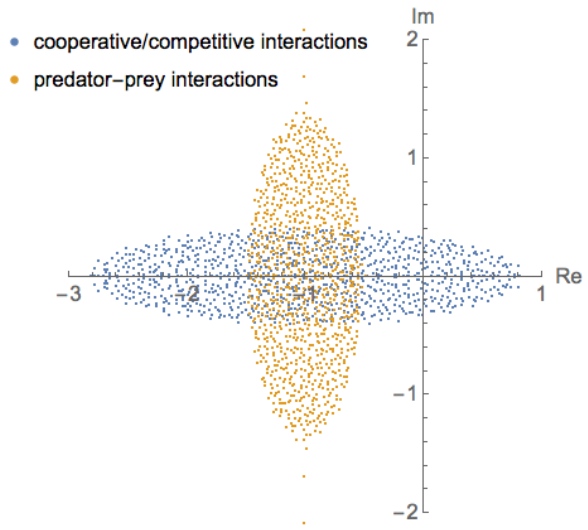

Figure B.2. Eigenvalue spectrum for matrices with imposed interaction types. Reconstruction of work by Allesina and Tang [4]. Shown are the eigenvalue spectra for two simulated 1000×1000 matrices. Blue dots correspond to eigenvalues of a matrix generated with half cooperative interactions and half competitive interactions. Orange dots correspond to eigenvalues of a matrix simulated solely with predator-prey interactions. This shows how predator-prey interactions are stabilising, whereas cooperative / competitive interactions are destabilising. See Allesina and Tang [4] for simulation methods and derived forms for the semi-major axis of each ellipse, which gives the stability threshold. Parameters $d = -1$, $\alpha = 0.05$, $n = 1000$, $C = 0.5$.

bifurcation. One can characterise the bifurcation to a certain extent by the behaviour of the destabilising eigenvalues (Figure 3.2).

We will consider continuous-time systems here. Therefore the two possible cases that result in the system destabilising are a real eigenvalue becoming positive (Fold/Transcritical/Pitchfork) or a complex-conjugate pair obtaining positive real part (Hopf). We investigate the extent to which these large random matrices resemble Hopf bifurcations when they are perturbed to a point of instability.

Notice that a significant portion of the eigenvalues in these spectra have a non-zero imaginary part (they lie off the real axis). In fact, there is a remarkable inverse power law relationship between the dimension of a random matrix and the expected proportion of real eigenvalues that it possesses [175] (Figure B.3). As a consequence, the proportion

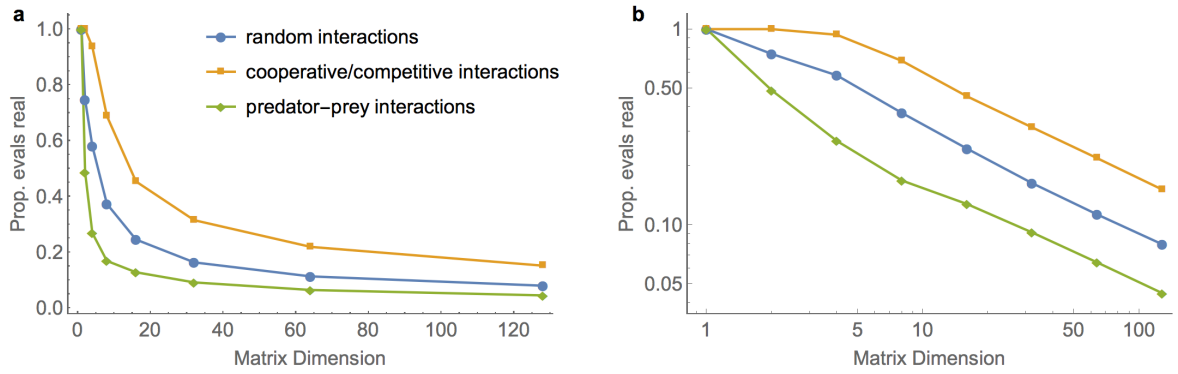

Figure B.3. Proportion of real eigenvalues rapidly decreases with system size. a. For select matrix dimensions, we simulate an ensemble of 1000 matrices for each of the interaction regimes. Points show the proportion of real eigenvalues amongst the entire ensemble. This proportion drops off much faster in predator-prey systems than for systems dominated by competitive and cooperative interactions. **b.** The same data on a log-log plot, demonstrating approximate power law behaviour for large dimensions. The power law exponent for the case of random interactions is approximately $\alpha = -0.54$, in agreement with early work by Sommers et al. [175]. Parameters $d = -1$, $\alpha = 0.05$, $C = 0.5$.

of real eigenvalues drops off rapidly as the system size increases. Moreover, the type of interactions in the system has a significant effect, with predator prey interactions resulting in a smaller proportion of real eigenvalues, compared to system made up of cooperative and competitive interactions. Due to these high proportions of complex eigenvalues, we hypothesise that high-dimensional systems, particularly those with predator-prey interactions, are more prone to Hopf bifurcations than systems with a small number of interacting components.

To test this hypothesis, we numerically simulate large random matrices with various proportions of interaction types using the methods of Allesina and Tang [4]. We then gradually reduce the self-regulating strength of a random selection of the species to force the system to destabilise. At the point where the system destabilises, we record whether this was due to a real eigenvalue or a complex conjugate pair. We do this for 1000 simulated community matrices and record the proportion of instabilities that correspond to a Hopf bifurcation. This is then done over a variety of system sizes and interaction types. Results

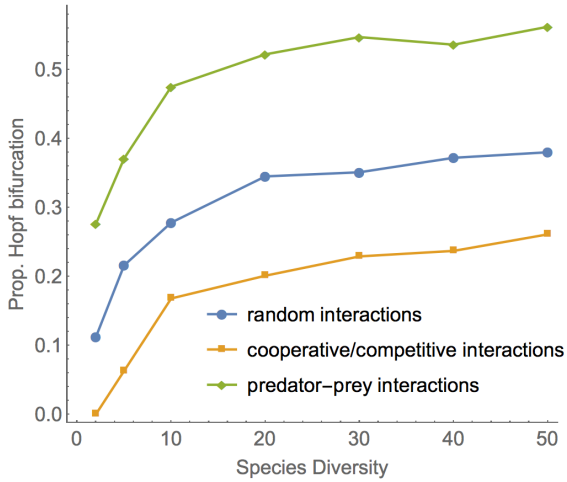

Figure B.4. Species diversity influences likelihood of Hopf bifurcation. For specific species diversity and interaction types, we simulated 1000 community matrices and incrementally perturbed them until they destabilised. The plot shows the proportion of these instabilities that manifested as Hopf bifurcations, showing the increasing trend with species diversity, and favouring predator-prey interactions among others.

are shown in Figure B.4. We find that the likelihood of the system encountering a Hopf bifurcation increases with species diversity (defined as the number of interacting species in the system) as expected. One might have expected a higher proportion in general for the high-dimensional systems due to the minute proportion of real eigenvalues. However, since the eigenvalue spectrum tends to an ellipse for large matrix dimensions, the eigenvalues with the largest real part will tend to lie on the real axis, and so with a reasonable probability will be the first eigenvalues to cause destabilisation. The simulations also demonstrate that predator-prey interactions promote Hopf bifurcations, which is in line with the observed oscillatory nature of predator-prey systems. Systems dominated by cooperative or competitive interactions are less conducive to oscillations.

In summary, this investigation has shown that Hopf bifurcations are an abundant feature of high-dimensional systems, particularly those with exploitative interactions between agents. Given that Hopf bifurcations are also candidates for critical slowing down, a metric that is able to distinguish them from other local bifurcations should be a useful addition to the current ecosystem forecasting tools.

B.2 Mathematical derivations

This appendix provides the mathematical derivations surrounding EWS in this chapter.

B.2.1 Analytical forms for EWS in continuous-time systems

Consider the generic continuous-time dynamical system depending on a single parameter

$$\dot{x} = f(x, \alpha), \quad x \in \mathbb{R}^n, \quad \alpha \in \mathbb{R}^1, \tag{B.2}$$

where f is a smooth function governing the deterministic dynamics of the system. We assume that the system sits in some stable equilibrium x_0, and therefore all eigenvalues of the Jacobian matrix about x_0 have negative real part. Upon varying α, a local bifurcation occurs at the point where an eigenvalue attains positive real part. There are generically only two possible ways this can occur [114]. Either a real eigenvalue becomes positive, or a complex conjugate pair attain positive real part (Figure 3.2). Both yield different bifurcations, and importantly, different EWS.

Case 1: A single real eigenvalue becomes positive. In the case where a real eigenvalue becomes positive, the system undergoes either a Fold, Transcritical or Pitchfork bifurcation, depending on higher order terms that are not captured by the eigenvalues. This unfortunately means that EWS using solely local stability properties of the system cannot distinguish these bifurcations in advance.

To derive generic EWS for these bifurcations, it is appropriate to consider them in their normal form, to which all other bifurcations of the same type are locally topologically equivalent [114]. Preceding each bifurcation, the linearised dynamics about the stable state of the system is identical, and given by

$$\frac{dx}{dt} = \lambda x, \tag{B.3}$$

where x represents the displacement from the equilibrium, and λ is the eigenvalue governing the stability of the system ($\lambda < 0$). Assuming the system is subject to small additive white noise, it obeys

$$\frac{dx}{dt} = \lambda x + \sigma \xi(t) \tag{B.4}$$

where $\xi(t)$ is a Gaussian white noise process and σ is the noise amplitude. Note that this is the one-dimensional OUp, for which EWS have already been derived in the chapter.

Case 2: A complex-conjugate pair of eigenvalues attain positive real part. In the case where a complex-conjugate pair of eigenvalues attain positive real part, the system either undergoes a super-critical (smooth) or a sub-critical (discontinuous) Hopf bifurcation, depending on higher terms that are not captured by the eigenvalues. Preceding each bifurcation in their normal form, the linearised dynamics about the stable state are given by

$$\frac{dx}{dt} = \mu x - \omega_0 y \tag{B.5}$$

$$\frac{dy}{dt} = \omega_0 x + \mu y, \tag{B.6}$$

where x and y represent displacements from the equilibrium in perpendicular directions, and the eigenvalues governing the stability of the system are now given by $\lambda = \mu \pm i\omega_0$. Note that a system must have at least two dimensions to undergo a Hopf bifurcation. To investigate EWS preceding these bifurcations, we include additive white noise to each of these components, to get the system

$$\frac{d\vec{x}}{dt} = J\vec{x} + B\vec{\xi}(t), \tag{B.7}$$

where $\vec{\xi}$ is a vector of Gaussian white noise processes (ξ_1, ξ_2), B is a matrix of noise amplitudes $\mathrm{diag}(\sigma_1, \sigma_2)$ and J is the Jacobian matrix given by

$$J = \begin{pmatrix} \mu & -\omega_0 \\ \omega_0 & \mu \end{pmatrix}. \tag{B.8}$$

The stochastic process in (B.7) is now a multi-variate OUp which, although more involved, can be subject to a similar treatment as the one-dimensional case. Integrating the process gives

$$\vec{x}(t) = e^{Jt}\vec{x}(0) + e^{Jt}\int_0^t e^{-Jt'} B\vec{\xi}(t')dt', \tag{B.9}$$

which has stationary dynamics given by

$$\vec{x}(t) = e^{Jt} \int_0^t e^{-Jt'} B\vec{\xi}(t')dt'. \tag{B.10}$$

The covariance matrix of a multi-dimensional process is defined as

$$\Sigma = \langle \vec{x}(t)\vec{x}^T(t) \rangle \tag{B.11}$$

which in this case gives

$$\Sigma = \int_0^t e^{J(t-t')} BB^T e^{J^T(t-t')} dt'. \tag{B.12}$$

In this derivation we used the property of independent white noise sources, $\langle \vec{\xi}(t)\vec{\xi}^T(t') \rangle = I\delta(t-t')$. Since the stochastic process is stationary, we have

$$\frac{d\Sigma}{dt} = 0, \tag{B.13}$$

and so taking the time derivative of (B.12) gives

$$J\Sigma + \Sigma J^T + BB^T = 0, \tag{B.14}$$

which is known as the continuous Lyapunov equation. In two dimensions, a closed form solution is known [73] and given by

$$\Sigma = \frac{-(J - (\operatorname{Tr} J)I)BB^T(J - (\operatorname{Tr} J)I)^T - \operatorname{Det} JBB^T}{2 \operatorname{Tr} J \operatorname{Det} J}, \tag{B.15}$$

which we can use to find the variance in system variables preceding a Hopf bifurcation. Using the Jacobian matrix in (B.8), the variance of each variable is given by

$$\operatorname{Var}(x) = \Sigma_{11} = \frac{\sigma_1^2}{2(-\mu)} + \frac{\omega_0^2(\sigma_2^2 - \sigma_1^2)}{4(-\mu)(\mu^2 + \omega_0^2)}, \tag{B.16}$$

$$\operatorname{Var}(y) = \Sigma_{22} = \frac{\sigma_2^2}{2(-\mu)} + \frac{\omega_0^2(\sigma_1^2 - \sigma_2^2)}{4(-\mu)(\mu^2 + \omega_0^2)}, \tag{B.17}$$

and the covariance term satisfies

$$\mathrm{Cov}(x, y) = \Sigma_{12} = \frac{\omega_0(\sigma_2^2 - \sigma_1^2)}{4(\mu^2 + \omega_0^2)}. \tag{B.18}$$

The autocovariance matrix is defined as

$$\phi(\tau) = \langle \vec{x}(t+\tau)\vec{x}^T(t) \rangle. \tag{B.19}$$

One can readily show using Eqn (B.7) and the statistical properties of white noise, that

$$\frac{\partial \phi}{\partial \tau} = J\phi, \tag{B.20}$$

which can be solved to give

$$\phi(\tau) = e^{J\tau}\Sigma \tag{B.21}$$

where $\Sigma = \phi(0)$. So one can directly compute the autocovariance matrix ϕ from the covariance matrix Σ. For the Jacobian in (B.8), this reduces to

$$\phi(\tau) = e^{\mu\tau} \begin{pmatrix} \cos(\omega_0\tau) & -\sin(\omega_0\tau) \\ \sin(\omega_0\tau) & \cos(\omega_0\tau) \end{pmatrix} \Sigma. \tag{B.22}$$

The power spectrum of each variable can be computed as before using the Wiener Khinchin Theorem. Since these expressions are long, we work with the leading order terms below under two different assumptions. (1) The noise amplitudes for each variable are similar, and (2) the system is very close to the bifurcation point, allowing us to assume that μ is small.

To leading order in noise amplitude difference

Here, we assume that the noise amplitudes σ_1 and σ_2 are similar, so $\sigma_2 = \sigma_1 + \epsilon$ where ϵ is small. Then to leading order, we have

$$\mathrm{Var}(x) = -\frac{\sigma_1^2}{2\mu} + O(\epsilon), \tag{B.23}$$

similar to the case of a single real eigenvalue. Note that all statistical expressions for the variable y can be obtained from those of x by interchanging σ_1 and σ_2. The autocorrelation coefficient $\rho_{xx} = \frac{\phi_{xx}(\tau)}{\phi_{xx}(0)}$ to leading order is

$$\rho_{xx}(\tau) = e^{\mu\tau}\cos(\omega_0\tau) + O(\epsilon). \tag{B.24}$$

From this expression, the power spectrum is computed by taking the Fourier transform,

$$S_{xx}(\omega) = -\frac{\sigma_1^2}{2\mu}\frac{1}{\pi}\int_0^\infty e^{\mu\tau}\cos(\omega_0\tau)\cos(\omega\tau)d\tau + O(\epsilon), \tag{B.25}$$

which can be evaluated to give

$$S_{xx}(\tau) = \frac{\sigma_1^2}{4\pi}\left(\frac{1}{(\omega+\omega_0)^2+\mu^2} + \frac{1}{(\omega-\omega_0)^2+\mu^2}\right) + O(\epsilon). \tag{B.26}$$

To leading order in proximity to bifurcation

To consider the case with dissimilar noise amplitudes, we can instead simplify matters by considering small μ, which assumes the system is close to the Hopf bifurcation. Taylor expanding about $\mu = 0$ gives variance

$$\mathrm{Var}(x) = \frac{\sigma_1^2+\sigma_2^2}{4(-\mu)} + O(\mu), \tag{B.27}$$

and autocorrelation

$$\rho_{xx}(\tau) = \cos(\omega\tau) + \mu\left(\tau\cos(\omega\tau) + \frac{\sigma_2^2-\sigma_1^2}{\sigma_1^2+\sigma_2^2}\frac{1}{\omega}\sin(\omega\tau)\right) + O(\mu^2). \tag{B.28}$$

Importantly, note that regardless of the noise amplitudes, the dominant term in the autocorrelation as the bifurcation is approached is $\cos(\omega\tau)$. This provides the distinguishing feature of a Hopf bifurcation, by taking positive or negative values depending on the lag time.

The $O(\mu)$ term in the power spectrum vanishes, and so can be computed to $O(\mu^2)$ as

$$S_{xx}(\omega) = \frac{1}{8\pi\omega_0}\left(\frac{2\sigma_1^2\omega_0 + (\sigma_1^2-\sigma_2^2)\omega}{\mu^2+(\omega+\omega_0)^2} + \frac{2\sigma_1^2\omega_0 + (\sigma_2^2-\sigma_1^2)\omega}{\mu^2+(\omega-\omega_0)^2}\right) + O(\mu^2), \tag{B.29}$$

which recovers recovers the 'noise-amplitude approximation' as $\sigma_1 \to \sigma_2$. We see that regardless of the noise amplitudes, the power spectrum is still maximised at $\omega = \pm\omega_0$ as

the bifurcation is approached ($\mu \to 0$), providing the characteristic feature of an upcoming Hopf bifurcation.

B.2.2 Analytical forms for EWS in discrete-time systems

Case 1: A real eigenvalue moves above one In the case where a real eigenvalue moves above one, the system undergoes either a Fold, Transcritical or Pitchfork bifurcation, analogous to the case of a positive real eigenvalue in a continuous system. Preceding each of these bifurcations in their normal form, the linearised dynamics satisfy

$$x_{t+1} = \lambda x_t, \tag{B.30}$$

where x_t denotes the displacement from the equilibrium, and λ is the eigenvalue governing the stability of the system ($|\lambda| < 1$). Assuming small additive white noise, the discrete then obeys

$$x_{t+1} = \lambda x_t + \sigma \epsilon_t \tag{B.31}$$

where ϵ_t is a normal random variable, and σ is the noise amplitude. This equation represents an AR(1) process which is the discrete-time analogue of the OUp, and has been studied thoroughly in the statistics literature [26].

The variance of this process can be computed by squaring both sides of (B.31) and taking the expectation, giving

$$\langle x_{t+1}^2 \rangle = \lambda^2 \langle x_t^2 \rangle + 2\lambda\sigma \langle \epsilon_t x_t \rangle + \sigma^2 \langle \epsilon_t^2 \rangle. \tag{B.32}$$

Given that the process is stationary, $\langle x_{t+1}^2 \rangle = \langle x_t^2 \rangle$. The white noise term is a normal random variable, independent of x_t and with unit variance, hence $\langle \epsilon_t x_t \rangle = 0$, and $\langle \epsilon_t^2 \rangle = 1$. The expression can then be simplified to obtain the variance,

$$\mathrm{Var}(x) = \langle x_t^2 \rangle = \frac{\sigma^2}{1 - \lambda^2}, \tag{B.33}$$

which diverges as the bifurcations are approached ($\lambda \to \pm 1$).

The lag-1 autocorrelation can be computed by multiplying (B.31) by x_t and taking expectations. This yields

$$\langle x_{t+1} x_t \rangle = \lambda \langle x_t^2 \rangle + \sigma \langle \epsilon_t x_t \rangle \tag{B.34}$$

which rearranges to

$$\rho_1 = \frac{\langle x_{t+1} x_t \rangle}{\langle x_t^2 \rangle} = \lambda. \tag{B.35}$$

When applied recursively, one obtains the autocorrelation function

$$\rho_\tau = \lambda^{|\tau|}. \tag{B.36}$$

From this we see that as $\lambda \to 1$ (Fold/Transcritical/Pitchfork), the autocorrelation increases monotonically to unity regardless of the lag-time. However, in the approach to the flip bifurcation, the autocorrelation increases or decreases depending on the lag time.

The power spectrum can be computed by taking the Discrete-Fourier-Transform of the autocovariance function

$$\phi_\tau = \langle x_t x_{t+\tau} \rangle = \frac{\sigma^2 \lambda^{|\tau|}}{1 - \lambda^2} \tag{B.37}$$

which gives

$$S(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} e^{-i\omega\tau} \phi_\tau \tag{B.38}$$

$$= \frac{1}{2\pi} \frac{\sigma^2}{1 - \lambda^2} \sum_{\tau=-\infty}^{\infty} e^{-i\omega\tau} \lambda^{|\tau|}. \tag{B.39}$$

Noting that the infinite sum can be split into two geometric series, one obtains the power spectrum as

$$S(\omega) = \frac{\sigma^2}{2\pi} \frac{1}{1 + \lambda^2 - 2\lambda \cos(\omega)}. \tag{B.40}$$

Case 2: A real eigenvalue moves below minus one In the case where a real eigenvalue moves below minus one, the system undergoes something entirely different - a Flip bifurcation, that is not possible in the continuous system. The analytical forms are the same as for the previous case, except the limiting behaviour preceding the bifurcation is obtained by taking $\lambda \to -1$.

Case 3: A complex-conjugate pair of eigenvalues attain a magnitude greater than one In this case, the system undergoes a Neimark-Sacker bifurcation[114], which

the discrete-time analogue of the Hopf bifurcation. Note this bifurcation is only possible in systems with dimension $n \geq 2$. The difference equation governing the linearised dynamics for this bifurcation is

$$x_{t+1} = r\cos\theta\, x_t - r\sin\theta\, y_t \tag{B.41}$$

$$y_{t+1} = r\sin\theta\, x_t + r\cos\theta\, y_t \tag{B.42}$$

where (x_t, y_t) is the (two-dimensional) displacement from equilibrium at time t, and $r \in (0, \infty)$ and $\theta \in (0, \pi)$ are parameters. With additive GWN, this system may be written

$$\vec{x}_{t+1} = J\vec{x}_t + B\vec{\epsilon}_t \tag{B.43}$$

where $\vec{\epsilon}_t$ is a vector of normal random variables (ϵ_1, ϵ_2), B is a matrix of noise amplitudes $\mathrm{diag}(a_1, a_2)$ and J is the Jacobian matrix given by

$$J = \begin{pmatrix} r\cos\theta & -r\sin\theta \\ r\sin\theta & r\cos\theta \end{pmatrix}. \tag{B.44}$$

The eigenvalues of J are a complex-conjugate pair given by $\lambda = re^{\pm i\theta}$, corresponding to case (e) in Figure 3.2. When $r = 1$, the eigenvalues reach the unit circle and the Neimark-Sacker bifurcation bifurcation occurs. How do the EWS behave?

The stochastic difference equation in (B.43) is a VAR(1) process (vector autoregressive process of order one). The variance of each variable can be obtained from the covariance matrix, defined as

$$\Sigma = \mathbb{E}(\vec{x}_{t+1}\vec{x}_{t+1}^T) \tag{B.45}$$

$$= \langle (J\vec{x}_t + B\vec{\epsilon}_t^T)(\vec{x}_t^T J^T + \vec{\epsilon}_t^T B^T) \rangle \tag{B.46}$$

$$= J\Sigma J^T + BB^T \tag{B.47}$$

where in the final step we have used properties of stationarity and independence of the noise. This is the discrete-time Lyapunov equation, which we can solve to find the covariance matrix. We solve this using symbolic computation in Mathematica. Expressions are long except for the case of equal noise amplitudes ($\sigma_1 = \sigma_2$), so we display those results here. However, differing noise amplitudes yields the same important features expected in the

EWS. In the case of equal noise amplitude, the variance in the state variable x is given by

$$\mathrm{Var}(x_t) = \frac{\sigma_1^2}{1-r^2}, \tag{B.48}$$

with the same scaling behaviour as in the one-dimensional case. We now compute the autocovariance matrix, which is defined as

$$\phi_\tau = \langle \vec{x}_{t+\tau} \vec{x}_t^T \rangle \tag{B.49}$$

which can be shown to satisfy $\phi_{\tau+1} = J\phi_\tau$ using the definition and equation (B.43). Solving this difference equation gives

$$\phi_\tau = J^\tau \phi_0 \tag{B.50}$$

$$= J^\tau \Sigma, \tag{B.51}$$

which, with the Jacobian in (B.44) becomes

$$\phi_\tau = r^\tau \begin{pmatrix} \cos(\theta\tau) & -\sin(\theta\tau) \\ \sin(\theta\tau) & \cos(\theta\tau) \end{pmatrix} \Sigma. \tag{B.52}$$

The autocorrelation for the variable x_t is then simply

$$\rho_\tau = r^\tau \cos(\theta\tau), \tag{B.53}$$

much like the autocorrelation preceding the Hopf bifurcation in the continuous-time case. We compute the power spectrum by taking the DFT of the autocovariance function for x_t

$$\phi_\tau = \frac{\sigma_1^2}{1-r^2} r^\tau \cos(\theta|\tau|) \tag{B.54}$$

where we have included the absolute value of τ in the exponent since the autocovariance

function is symmetric in τ. The power spectrum is then

$$S(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} e^{-i\omega\tau} \phi_\tau \tag{B.55}$$

$$= \frac{1}{2\pi} \frac{\sigma_1^2}{1-r^2} \sum_{\tau=-\infty}^{\infty} e^{-i\omega\tau} r^{|\tau|} \cos(\theta\tau). \tag{B.56}$$

which can be written as a geometric series* and computed as

$$S(\omega) = \frac{\sigma_1^2}{4\pi} \left(\frac{1}{1 + r^2 - 2r\cos(\omega - \theta)} + \frac{1}{1 + r^2 - 2r\cos(\omega + \theta)} \right). \tag{B.57}$$

B.2.3 Asymptotic behaviour of $S_{\max}$ up to local bifurcations

This section provides derivations of the asymptotic behaviour of $S_{\max}$ up to each local codimension-1 bifurcation. The derivations use the analytical forms of the power spectrum, as derived earlier and displayed in Table 4.1. The asymptotic behaviour is derived in terms of the distance to the bifurcation d, which we define as the smallest difference between the current eigenvalue (of the Jacobian matrix that governs the local dynamics of the system) and a location at which a bifurcation occurs.

Continuous-time: Fold, Transcritical, Pitchfork. These bifurcations occur when a real eigenvalue attains positive real part, and so $d = |\lambda|$, where $\lambda \in \mathbb{R}^{\leq 0}$. The power spectrum for the normal form of these bifurcations is

$$S(\omega) = \frac{\sigma^2}{2\pi} \left(\frac{1}{\omega^2 + \lambda^2} \right) \tag{B.58}$$

where σ is the white noise amplitude. The maximum value of this function occurs at $\omega = 0$, and is

$$S_{\max} = \frac{\sigma^2}{2\pi d^2} \sim \frac{1}{d^2} \quad \text{as} \quad d \to 0. \tag{B.59}$$

Continuous-time: Hopf. The Hopf bifurcation occurs when a complex-conjugate pair of eigenvalues $\lambda_\pm = \mu \pm i\omega$ attain positive real part. The distance to the bifurcation in this

*use the relation $\cos(\theta\tau) = (e^{i\theta\tau} + e^{-i\theta\tau})/2$

case is $d = |\mu|$. The power spectrum for the normal form of this bifurcation is

$$S(\omega) = \frac{\sigma^2}{4\pi}\left(\frac{1}{(\omega-\omega_0)^2+\mu^2} + \frac{1}{(\omega+\omega_0)^2+\mu^2}\right). \tag{B.60}$$

This function attains maximum value at $\omega = \omega_0$, which may be written

$$S_{\max} = \frac{\sigma^2}{4\pi d^2}\left(1 + \frac{d^2}{4\omega_0^2+d^2}\right). \tag{B.61}$$

As $d \to 0$, the bracketed term goes to 1 and we are left with the asymptotic behaviour

$$S_{\max} \sim \frac{1}{d^2} \quad \text{as} \quad d \to 0. \tag{B.62}$$

Discrete-time: Fold, Transcritical, Pitchfork. These bifurcations occur when a real eigenvalue reaches the value one from below, and so $d = 1 - \lambda$, where $\lambda \in \mathbb{R}$. The power spectrum for the normal forms of these bifurcations is

$$S(\omega) = \frac{\sigma^2}{2\pi}\left(\frac{1}{1+\lambda^2-2\lambda\cos(\omega)}\right), \quad \omega \in [-\pi, \pi]. \tag{B.63}$$

Provided that $\lambda > 0$, which is the region in which these bifurcations occur, the maximum value of this function occurs at $\omega = 0$ and is

$$S_{\max} = \frac{\sigma^2}{2\pi(1-\lambda)^2} = \frac{\sigma^2}{2\pi d^2} \sim \frac{1}{d^2} \quad \text{as} \quad d \to 0. \tag{B.64}$$

Discrete-time: Flip (period-doubling). This bifurcation occurs when a real eigenvalue reaches the value negative one from above, and so $d = \lambda + 1$. The power spectrum for the normal form of the Flip bifurcation is the same as for the previous bifurcations, however. The maximum of this power spectrum in the region $\lambda < 0$ is

$$S_{\max} = \frac{\sigma^2}{2\pi(1+\lambda)^2} = \frac{\sigma^2}{2\pi d^2} \sim \frac{1}{d^2} \quad \text{as} \quad d \to 0, \tag{B.65}$$

Discrete-time: Neimark-Sacker. The Neimark-Sacker bifurcation occurs when a complex conjugate pair of eigenvalues $\lambda = re^{\pm\omega_0}$ obtain unit magnitude, and so $d = 1 - r$. The

power spectrum for the normal form Neimark-Sacker bifurcation is

$$S(\omega) = \frac{\sigma^2}{4\pi}\left(\frac{1}{1+r^2-2r\cos(\omega-\omega_0)} + \frac{1}{1+r^2-2r\cos(\omega+\omega_0)}\right). \tag{B.66}$$

This takes a maximum value at $\omega = \omega_0$ may be written

$$S_{\max} = \frac{\sigma^2}{4\pi}\left(\frac{1}{(1-r)^2} + \frac{1}{1+r^2-2r\cos(2\omega_0)}\right) \tag{B.67}$$

$$= \frac{\sigma^2}{4\pi d^2}\left(1 + \frac{d^2}{d^2+2(1-d)(1-\cos(2\omega_0))}\right) \tag{B.68}$$

$$\sim \frac{1}{d^2} \quad \text{as} \quad d \to 0, \tag{B.69}$$

since the bracketed term goes to 1.

B.2.3.1 EWS in a system with multiplicative noise

Here we assume that multiplicative environmental noise is applied to the system, and show analytically how this can affect the EWS. We consider the one-dimensional system

$$\dot{x} = f(x) + \sigma g(x)\xi(t) \tag{B.70}$$

where f describes the deterministic dynamics of the system, σ is a noise amplitude factor, g captures how the noise amplitude varies with the system state, and $\xi(t)$ is a white noise process. In previous calculations that used additive white noise, we had $g(x) = 1$. Linearising the system about an equilibrium state x^*, we get

$$\dot{\eta} = f'(x^*)\eta + \sigma\left[g(x^*) + g'(x^*)\eta\right]\xi(t) + O(\eta^2) \tag{B.71}$$

where $\eta = x - x^*$. We now assume that η and σ are sufficiently small such that we may drop terms of order 2 and higher. This gives the approximation

$$\dot{\eta} = f'(x^*)\eta + \sigma g(x^*)\xi(t), \tag{B.72}$$

which is of the same form as with additive noise except for the additional $g(x^*)$ factor. Thus EWS expressions are the same, just with σ replaced by $\sigma g(x^*)$.

For example, if we consider environmental noise, which is often taken to have amplitude proportional to the size of the current state of the system [20], ($g(x) = x$), then we get EWS

$$\text{Var} = \frac{[\sigma x^*(\lambda)]^2}{2\lambda}, \quad S_{\max} = \frac{[\sigma x^*(\lambda)]^2}{2\pi\lambda^2} \tag{B.73}$$

where $x^*(\lambda)$ is the equilibrium value of the system (which in general depends on λ).

B.3 EWS in simulations of parametrised predator-prey model

In the study by Fussmann et al. [70], the authors parametrised a model to the output from the chemostat experiments. We run stochastic simulations of this model and compute EWS, to check agreement with theory. The model is given by

$$\frac{dN}{dt} = \delta(N_i - N) - F_C(N) \tag{B.74}$$

$$\frac{dC}{dt} = F_C(N)C - F_B(C)B/\epsilon - \delta C \tag{B.75}$$

$$\frac{dR}{dt} = F_B(C)R - (\delta + m + \lambda)R \tag{B.76}$$

$$\frac{dB}{dt} = F_B(C)R - (\delta + m)B \tag{B.77}$$

where

$$F_C(N) = \frac{b_C N}{K_C + N}, \quad F_B(C) = \frac{b_B C}{K_B + C}. \tag{B.78}$$

In the model N is the concentration of nitrogen (nutrient), C the concentration of *Chlorella vulgaris* (prey), R the concentration of reproducing *Brachionus calyciflorus* (predator), and B is the total concentration of *Brachionus calyciflorus*. Parameter values derived from experiments [70] are given in Table B.1. The conversion factors η_C and η_B are required to make the units of the model and experimental output consistent.

Under this parameter configuration, the model exhibits a Hopf bifurcation at $\delta = 0.009$ with a period of $T = 44.22$ days (H1), a Hopf bifurcation at $\delta = 0.151$ with a period of $T = 5.905$ days (H2), a Hopf bifurcation at $\delta = 0.970$ with a period of $T = 0.554$ days (H3), and a Transcritical bifurcation at $\delta = 1.427$ (T1). The corresponding bifurcation diagram is in Figure B.5a. We simulate the model at fixed dilution rates over the range $[0, 1.6]$ to investigate the behaviour of each EWS in the vicinity of the bifurcations. After a burn-in period of 500 time units to remove transient dynamics, we measure each EWS statistic over the following 4000 time units. The behaviour of the EWS considered is shown in Figure B.5b-f.

The variance of the system behaves as expected. As the dilution rate approaches each of the three Hopf bifurcations from the neighbouring equilibrium state, the variance of the system

Parameter	Definition	Value
δ	dilution rate	(.) per day
N_i	nitrogen inflow concentration	80 μmol/liter
b_C	maximum birth rate of *Chlorella*	3.3 per day
K_C	half-saturation constant of *Chlorella*	4.3 μmol/liter
b_B	maximum birth rate of *Brachiounus*	2.25 per day
K_B	half-saturation constant of *Brachionus*	15 μmol/liter
m	mortality of *Brachiounus*	0.055 per day
λ	decay of fecundity of *Brachionus*	0.4 per day
ϵ	assimilation efficiency of *Brachionus*	0.25
η_C	conversion factor μmol/lit to cells/ml of *Chlorella*	5×10^4 (cells/ml)(μmol/lit)$^{-1}$
η_B	conversion factor μmol/lit to females/ml of *Brachionus*	5 (females/ml)(μmol/lit)$^{-1}$

Table B.1. Predator-prey model parameters. Definition and value for each parameter in the predator-prey model given in Eqns. (B.74)-(B.77). Calibrated by Fussmann et al. [70] from chemostat experiments or obtained from published sources on *Chlorella* and *Brachionus*.

dynamics increases, consistent with theory. The variance is maximal in the oscillatory regime between H2 and H3 since oscillations are not filtered out by the smoothing of the time-series. Near the Transcritical bifurcation T1, there is an increase in variance from the right and the left, although this increase occurs relatively late (close to the bifurcation), and is a very slight change in magnitude.

Autocorrelation also behaves as expected, and we have displayed lag times of 3 and 6 as they exemplify the contrasting behaviour that autocorrelation can exhibit depending on the chosen lag time. Recall that the analytical form for the autocorrelation in the vicinity of a Hopf bifurcation has a multiplicative factor of $\cos(2\pi\tau/T)$ (Table 4.1) where τ is the lag time and T is the period of oscillations at the Hopf bifurcation. The other factor approaches one as the Hopf bifurcation is approached, just like the case for the Fold bifurcation. Since the period at H2 is approximately 6, this multiplicative factor will be negative for $\tau = 3$ (resulting in decreasing autocorrelation) and positive for $\tau = 6$ (resulting in increasing autocorrelation). Note that for certain lag times ($\tau = T/4, 3T/4$) this factor is zero, in which case autocorrelation does not serve as an EWS. In general, we do not know the period at a Hopf bifurcation in advance, and so choosing an appropriate lag time is difficult. The power spectrum is not hindered by this uncertainty, since it captures the equivalent information to autocorrelation at every lag time [26].

The peak in the power spectrum ($S_{\max}$) behaves very similarly to the variance. It as at least as good as Variance at signalling bifurcations. Preceding H2 it actually provides

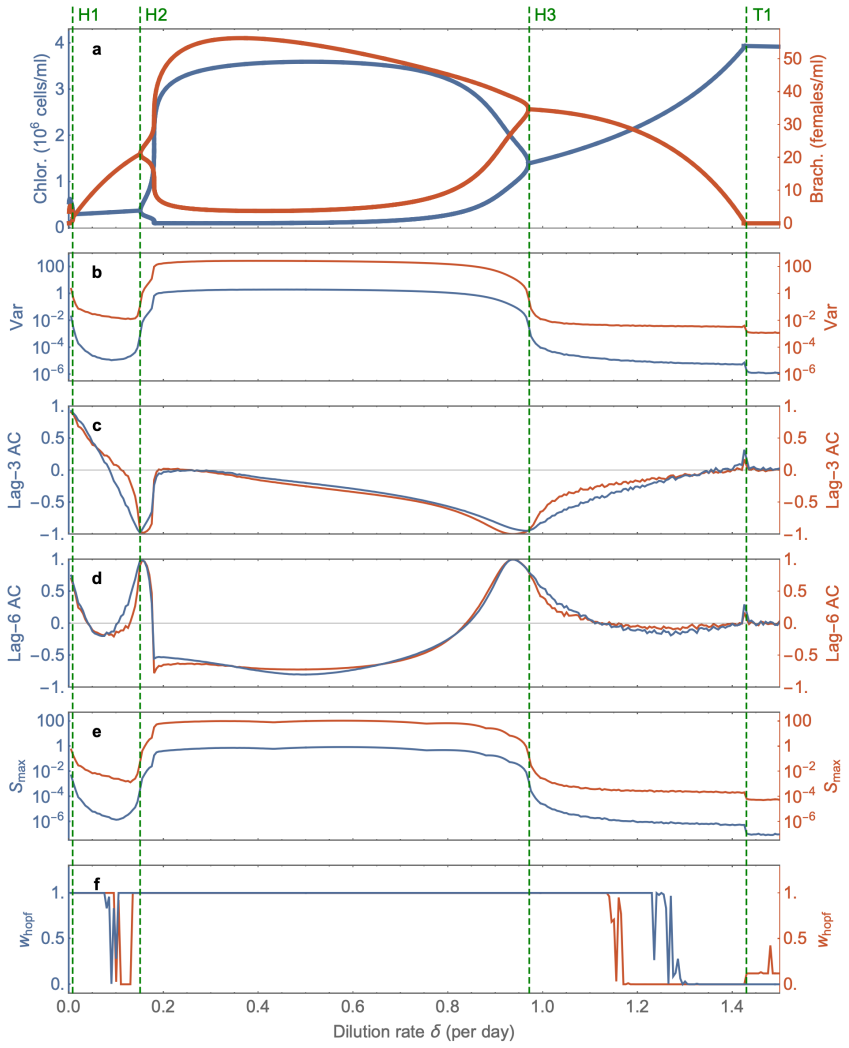

Figure B.5. EWS in stationary simulations of chemostat model. For fixed dilution rates spaced 0.005 apart, we simulate the model for 4500 time units allowing for a burn-in period of 500 time units to remove any transient behaviour. EWS are then computed on the remaining time-series data. **a.** Bifurcation diagram showing stable states and extreme points of limit cycles for the prey (*Chlorella*, blue) and the predator (*Brachionus*, red). **b.** Variance. **c-d.** Lag-3 and lag-6 autocorrelation. **e.** Peak in the power spectrum. **f.** Hopf AIC weight. Green dashed lines mark the Hopf bifurcations (H1, H2, H3) and the Transcritical bifurcation (T1). EWS are computed as in Methods.

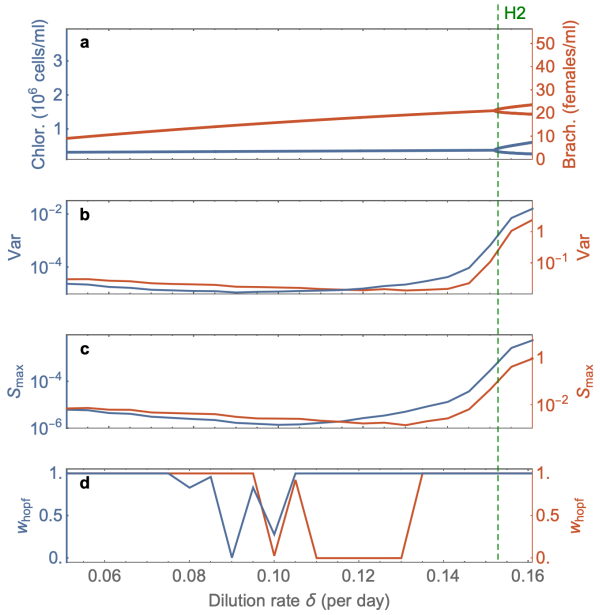

Figure B.6. EWS in stationary model simulations - zoom in on H2. As in Figure B.5 except zoomed in on the dynamics preceding H2. $S_{\max}$ does at least as well as variance in signalling the upcoming bifurcation, performing better than variance with the Brachionus time-series. The Hopf AIC weight registers the type of bifurcation for each species once $S_{\max}$ begins to increase.

earlier warning as can be seen in Figure B.6, which zooms in on the preceding dynamics.

The Hopf AIC weight (w_{hopf}) correctly signals which bifurcations are Hopf, and which are not. (Preceding H1-H3 it attains 1 the value 1, and preceding T1 it attains the value 0). The metric appears to be quite late in recognising a Hopf bifurcation in the Brachionus time-series preceding H2. Zooming in on this area (Figure B.6), we see that the signalling of w_{hopf} aligns with the moment when variance and $S_{\max}$ begin to increase.

B.4 Nonlinear optimisation procedure

As with any computationally feasible optimisation algorithm, the procedure can output a local minimum without reaching the global minimum. Moreover, the local minimum that the algorithm reaches is dependent on the initial parameter guess that it is given. As such, we sweep over a range of initial parameter guesses and store the best model fit (smallest squared error).

Model constraints For the fitting procedure, we impose some logical constraints on the parameters. This avoids the optimisation algorithm searching through irrelevant parameter combinations, thereby reducing computational time. The constraints we impose are as follows:

$$\sigma > 0, \tag{B.79}$$

$$\lambda, \mu < 0, \tag{B.80}$$

$$S_{\text{fold}}(\omega_{\max}) < \psi_f S_{\text{fold}}(0), \tag{B.81}$$

$$S_{\text{hopf}}(0) < \psi_h \, S_{\text{hopf}}(\omega_0). \tag{B.82}$$

Constraint (B.79) is imposed by the definition of the noise amplitude (a standard deviation of a distribution). We impose constraint (B.80), since we are only considering systems in a stable state (λ and μ correspond to the real part of the eigenvalues for the system's Jacobian matrix, which must be negative for the system to be stable). Constraint (B.81) asserts that the power at the edge of the frequency domain must be less than some proportion $\psi_f \in [0, 1]$ of the power at the zero-frequency. Note that this is guaranteed sufficiently close to a Fold bifurcation since $S_{\text{fold}}(0) \to \infty$ as $\lambda \to 0$, whereas $S_{\text{fold}}(\omega_{\max})$ is always finite. Our motivation for imposing this constraint is to better distinguish S_{fold} from the flat spectrum S_{null} which signifies no imminent bifurcation.

Constraint (B.82) asserts that the power at the zero-frequency is at most a proportion $\psi_h \in [0, 1]$ of the power at the intrinsic frequency of the Hopf bifurcation ω_0. This condition is always satisfied sufficiently close to a Hopf bifurcation since as $\mu \to 0$, $S_{\text{hopf}}(\omega_0) \to \infty$, whereas $S_{\text{hopf}}(0)$ always remains finite. Our motivation for this constraint is to distinguish S_{hopf} from S_{fold}, since without this constraint, S_{hopf} can fit a unimodal power spectrum just as well as S_{fold} can by choosing $\omega_0 = 0$, and a Hopf bifurcation may be erroneously flagged.

With the constraint, this power spectrum models fit only the two distinct behaviours that are expected preceding the Fold and Hopf bifurcations respectively.

The constraints (B.81) and (B.82) are algebraically equivalent to the parameter constraints

$$\omega_0^2 > \frac{\mu^2}{4\psi_h}\left(4 - 3\psi_h + \sqrt{\psi_h^2 - 16\psi_h + 16}\right), \tag{B.83}$$

$$\lambda^2 < \left(\frac{\psi_f}{1-\psi_f}\right)\omega_{\max}^2, \tag{B.84}$$

respectively, which are coded into the optimisation algorithm. We find that taking $\psi_f = 0.2$, $\psi_h = 0.5$ gives good performance.

Initial parameter guesses The analytical forms for the power spectra are fitted to the measured power spectrum using nonlinear optimisation procedures from the open-source Python package LMFIT [138]. Specifically, the Levenberg–Marquardt algorithm is used to minimise the least square error of the fit. As with any computationally feasible optimisation algorithm, the procedure can output a local minimum without reaching the global minimum. Moreover, the local minimum that the algorithm reaches is dependent on the initial parameter guess that it is given. As such, we sweep over a range of initial parameter guesses and store the best model fit (smallest squared error).

Appropriate initial parameter guesses are made based on three characteristic metrics of the power spectrum values $\{(\omega_i, S(\omega_i))\}_i$. This allows the algorithm to generalise to power spectra on significantly different scales. (For parameter guesses that are independent of the power spectrum, this would have been a problem). Consider the following metrics

$$S_{\max} = \max_i S(\omega_i) \tag{B.85}$$

$$\omega_{\text{dom}} = \arg\max_i S(\omega_i) \tag{B.86}$$

$$S_{\text{tot}} = \sum_i S(\omega_i)\Delta\omega \tag{B.87}$$

where $\Delta\omega$ is the difference between consecutive frequency values. Note that each of these metrics can be computed from the measured spectrum with very little computational effort. Now, we can find the equivalent metrics using the analytical forms to provide a functional relationship between the metrics and the parameter values. Then, given numerical values

for these metrics, we can compute the corresponding parameter values, and they will serve as our scale-independent initial parameter guesses. The mathematical forms for these initial parameter guesses are derived below.

For fitting S_{fold}, our baseline initial parameters $(\tilde{\sigma}, \tilde{\lambda})$ are set to satisfy

$$S_{\max} = \frac{\tilde{\sigma}^2}{2\pi\tilde{\lambda}^2}, \qquad S_{\text{tot}} = -\frac{\tilde{\sigma}^2}{2\tilde{\lambda}} \tag{B.88}$$

which are the analytical expressions for the peak in $S_{\text{fold}}(\omega)$ and its total variation respectively. This gives an initialisation of

$$\tilde{\sigma} = \sqrt{\frac{2}{\pi}}\frac{S_{\text{tot}}}{\sqrt{S_{\max}}}, \qquad \tilde{\lambda} = -\frac{S_{\text{tot}}}{\pi S_{\max}}. \tag{B.89}$$

For fitting S_{hopf}, our baseline initial parameters $(\tilde{\sigma}, \tilde{\mu}, \tilde{\omega}_0)$ are set to satisfy

$$S_{\max} = \frac{\tilde{\sigma}^2}{4\pi\tilde{\mu}^2}\left(1 + \frac{\tilde{\mu}^2}{\tilde{\mu}^2 + 4\tilde{\omega}_0^2}\right), \qquad S_{\text{tot}} = -\frac{\tilde{\sigma}^2}{2\tilde{\mu}}, \qquad \omega_{\text{dom}} = \tilde{\omega}_0 \tag{B.90}$$

which are the analytical expressions for the peak in S_{hopf}, its total variation, and dominant frequency respectively. This gives an initialisation of

$$\tilde{\sigma} = \sqrt{-2\tilde{\mu}S_{\text{tot}}}, \quad \tilde{\mu} = -\frac{1}{3\pi S_{\max}}\left(S_{\text{tot}} + \alpha^{\frac{1}{3}} + \frac{S_{\text{tot}}^2 - 12\pi^2\omega_{\text{dom}}^2 S_{\max}^2}{\alpha^{\frac{1}{3}}}\right), \quad \tilde{\omega}_0 = \omega_{\text{dom}} \tag{B.91}$$

where

$$\begin{aligned}\alpha = S_{\text{tot}}^3 &+ 9\pi^2\omega_{\text{dom}}^2 S_{\max}^2 S_{\text{tot}} \\ &+ 3\sqrt{3}\pi\sqrt{64\pi^4\omega_{\text{dom}}^6 S_{\max}^6 - 13\pi^2\omega_{\text{dom}}^4 S_{\max}^4 S_{\text{tot}}^2 + 2\omega_{\text{dom}}^2 S_{\max}^2 S_{\text{tot}}^4}\end{aligned} \tag{B.92}$$

Finally, for fitting S_{null} are baseline initial parameter $\tilde{\sigma}$ satisfies

$$S_{\text{tot}} = \tilde{\sigma}^2 \tag{B.93}$$

which is the analytical expression for the total variation of S_{null} (over the frequency domain $[-\pi, \pi]$). This simply gives us the baseline initialisation parameter of $\tilde{\sigma} = \sqrt{S_{\text{tot}}}$.

Once the baseline initialisation parameters are found, we can run the optimisation procedure to find the best fit. In cases where the power spectrum is non-typical, doing a sweep over a set of initialisation parameters increases the likelihood of finding the best fit. So we provide the option for the algorithm to sweep over $[0.5, 1, 1.5]$ times the baseline initialisation value for each parameter, and choose the best fit (lowest AIC score) Doing the sweep entails 3^3 initialisation settings for S_{hopf}, 3^2 settings for S_{fold} and 3 settings for S_{null}. Therefore, doing the sweep is 13 times more expensive (computationally) than doing no sweep at all.

B.5 Additional figures

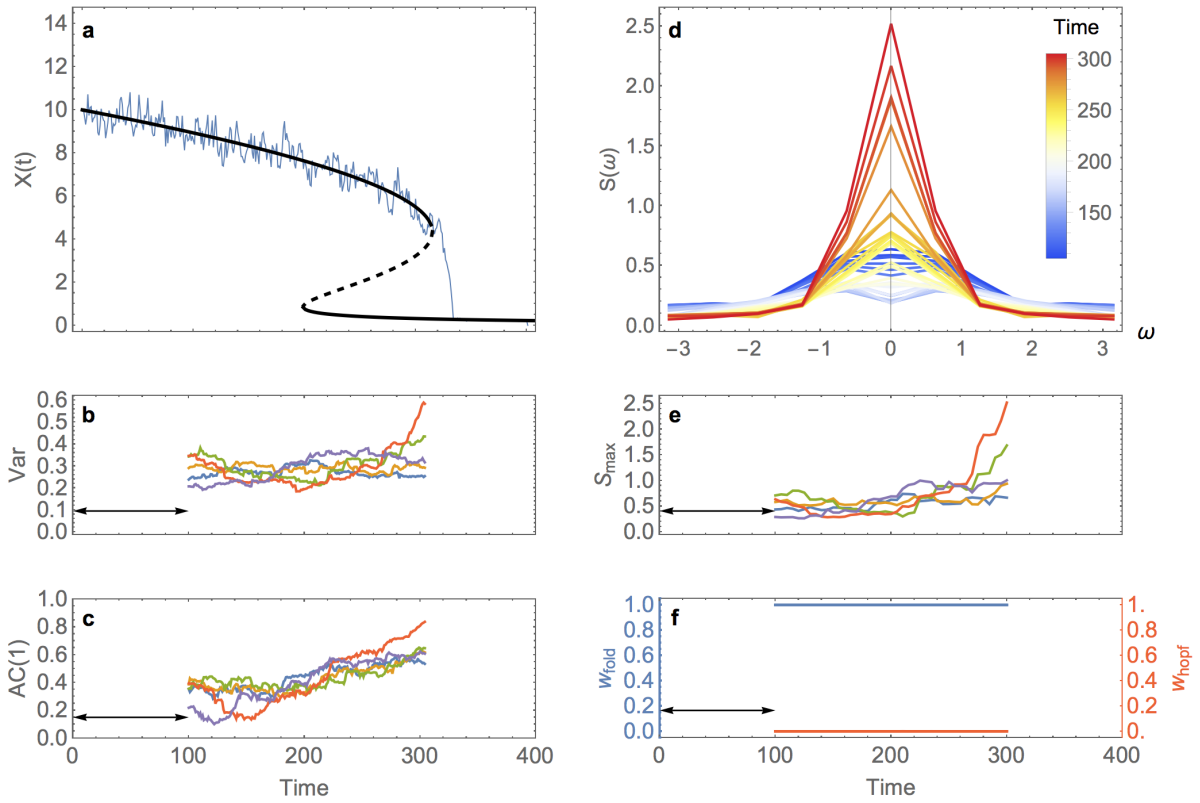

Figure B.7. EWS preceding the Fold bifurcation in the Ricker model. Individual realisations are shown for the Ricker model under the exploitation scenario. The growth rate is fixed ($r = 0.75$) whilst the harvesting rate (F) increases linearly across the interval $[0, 3]$. **a.** Time-series of a single realisation overlaid on bifurcation diagram. **b-c.** Variance and lag-1 autocorrelation computed over a rolling window of 5 realisations. **d.** Sample of power spectra at times leading up to the bifurcation for a single realisation. **e.** Evolution of peak in the power spectra ($S_{\max}$) for 5 realisations. **f.** AIC weights for the Fold (w_{fold}) and Hopf (w_{hopf}) power spectrum forms for 5 realisations. EWS are computed as in Methods. Fixed model parameters are $K = 10$, $h = 0.75$, $\sigma = 0.05$.

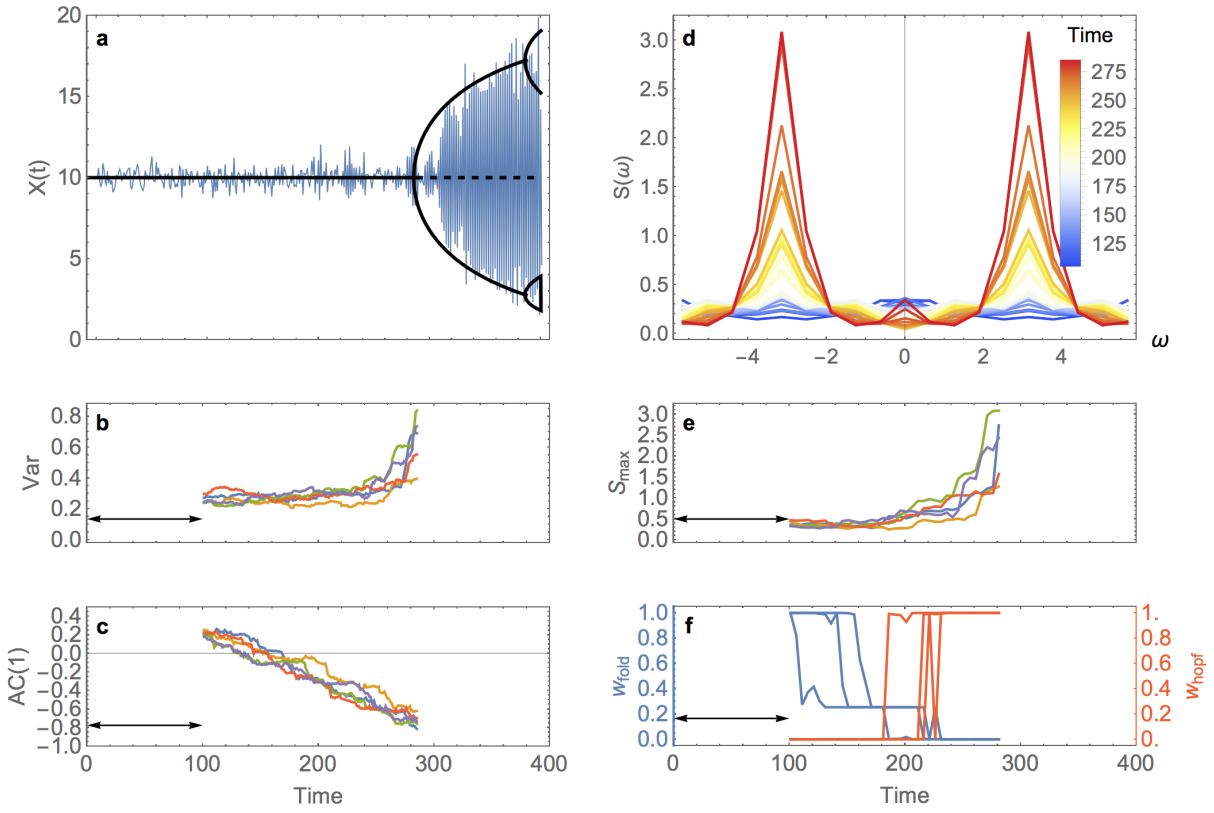

Figure B.8. EWS preceding the Flip bifurcation in the Ricker model. Individual realisations are shown for the Ricker model under the truncation scenario. The harvesting rate is fixed ($F = 0$) whilst the growth rate (r) increases linearly across the interval $[0.5, 2.6]$. **a.** Time-series of a single realisation overlaid on bifurcation diagram. **b-c.** Variance and lag-1 autocorrelation computed over a rolling window of 5 realisations. **d.** Sample of power spectra at times leading up to the bifurcation for a single realisation. **e.** Evolution of peak in the power spectra ($S_{\max}$) for 5 realisations. **f.** AIC weights for the Fold (w_{fold}) and Hopf (w_{hopf}) power spectrum forms for 5 realisations. EWS are computed as in Methods. Fixed model parameters are $K = 10$, $h = 0.75$, $\sigma = 0.05$.

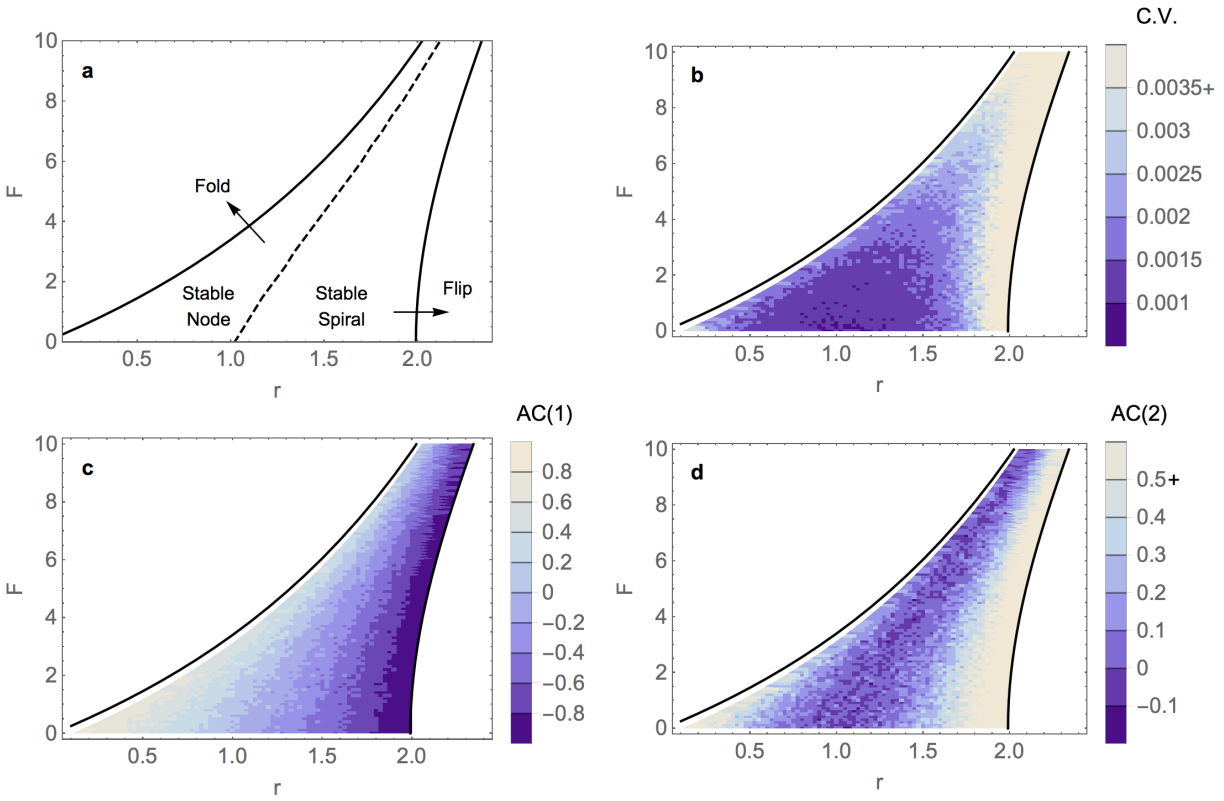

Figure B.9. Temporal EWS across the parameter plane of the Ricker model. For fixed values of (r, F) within the stable region of the bifurcation plane, we simulate the Ricker model and compute the EWS metrics. **a.** Regions of stability and bifurcation curves. Crossing the Flip/Fold bifurcation curve results in a transition to oscillations/an alternative stable state. Perturbations in the stable node/stable spiral regime result in a monotonic/oscillatory recovery trajectory. **b.** Coefficient of variation (C.V.) **c.** Lag-1 autocorrelation (AC(1)). **d.** Lag-2 autocorrelation (AC(2))

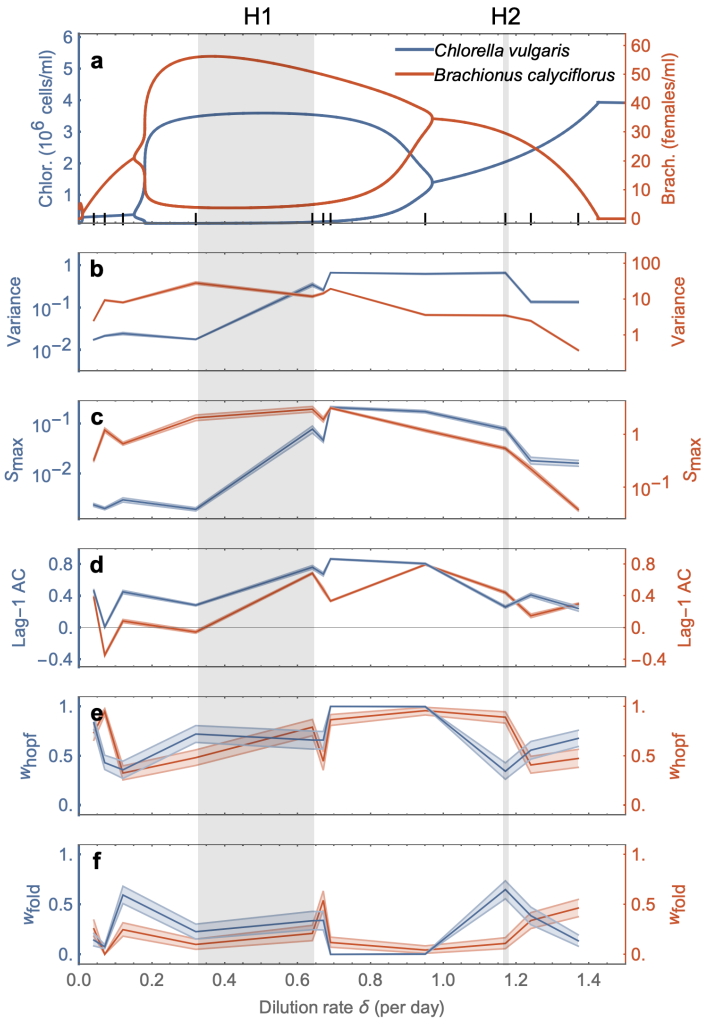

Figure B.10. Bootstrapped EWS by species in the empirical predator-prey system. For each chemostat time-series (with a fixed dilution rate δ), we computed 100 samples of bootstrapped EWS (Methods) and display the means with 95% confidence intervals. **a.** Bifurcation diagram of parametrised model [70] for Chlorella (blue) and Brachionus (red). Lines show stable states / limit cycle. Grey areas show regions where a Hopf bifurcation actually took place in experiments [70]. Bold ticks show the dilution rates used in experiments. **b.** Variance. **c.** Lag-1 autocorrelation. **d.** Peak in the power spectrum. **e-f.** AIC weights. Note that w_{null} can be obtained from the relation $w_{\text{fold}} + w_{\text{hopf}} + w_{\text{null}} = 1$. We computed EWS using a Hamming window of 80 days and a bootstrap block size of 40 days.

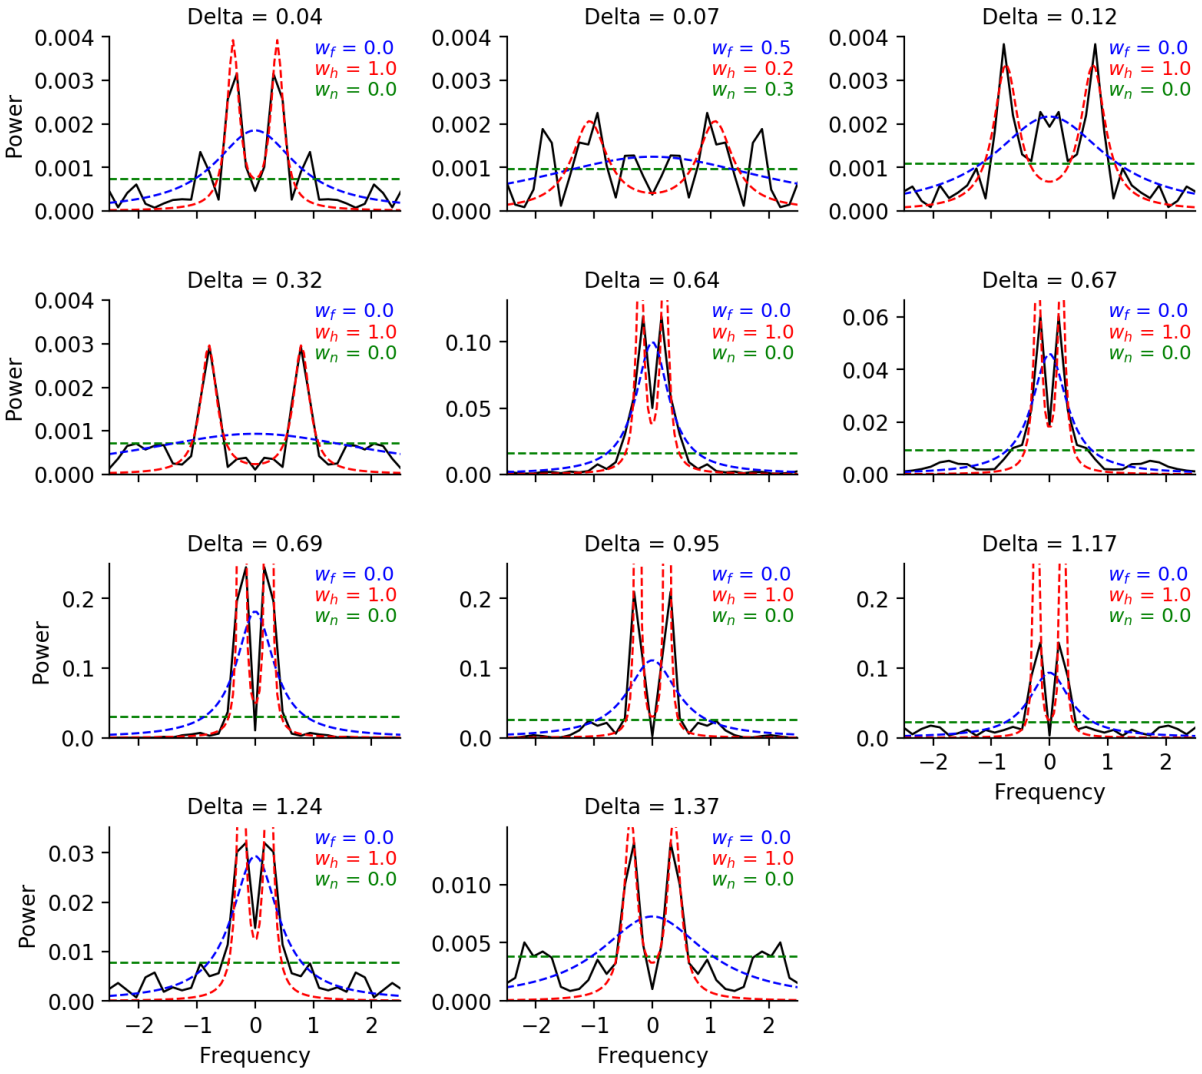

Figure B.11. Power spectra of the Chlorella time-series at each dilution rate (δ). The power spectrum (black) is computed using Welch's method with a Hamming window of length 40 days and a 50% offset. The analytical forms S_{fold} (blue), S_{hopf} (red) and S_{null} (green) are fit to the power spectrum to obtain the AIC weights w_f, w_h and w_n, which are an indication of which fit best explains the data.

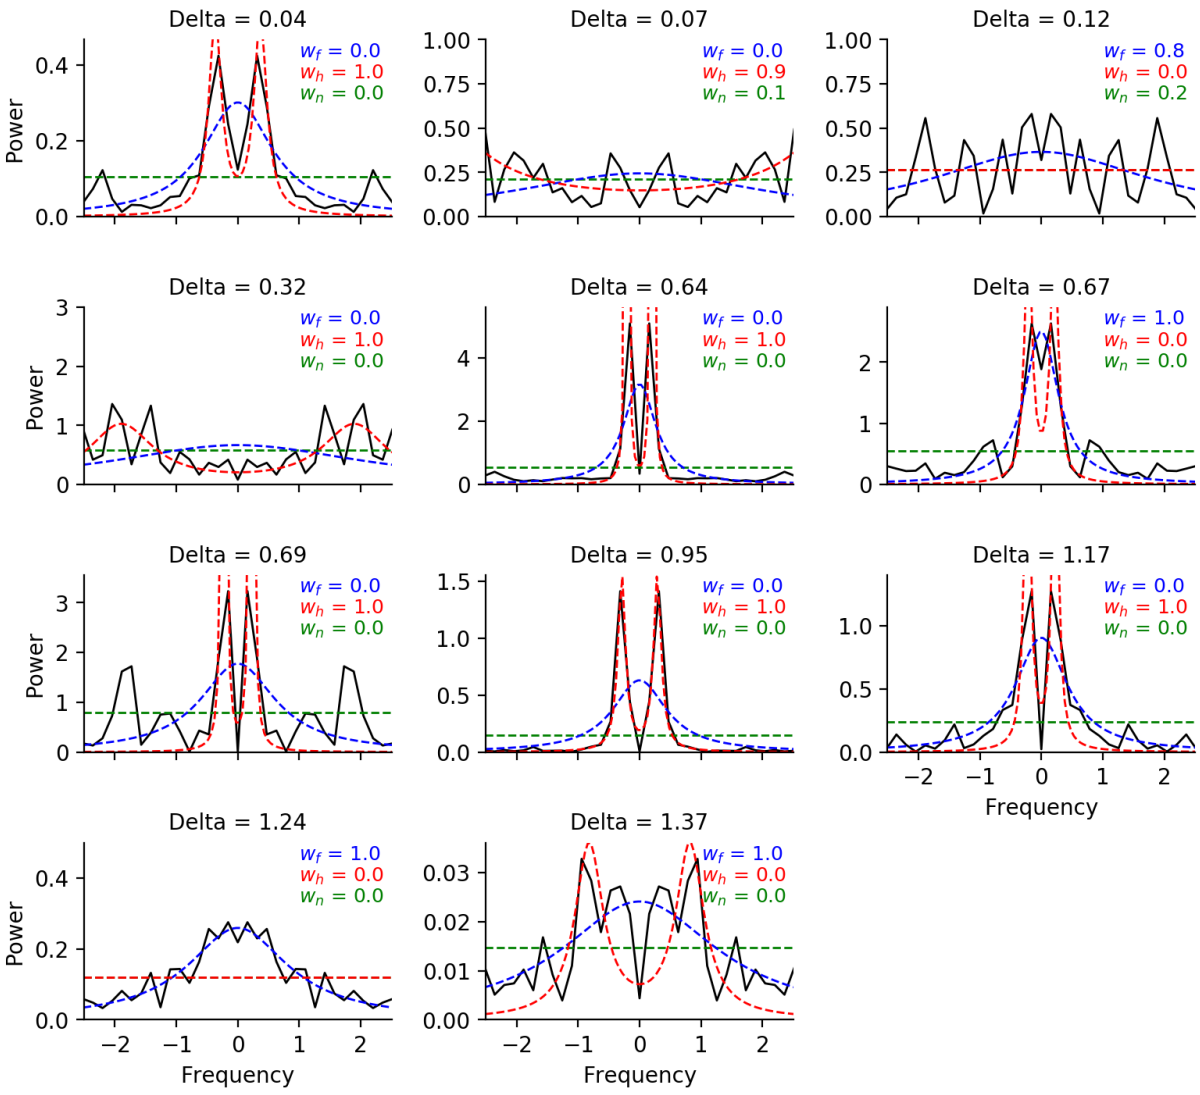

Figure B.12. Power spectra of the Brachiounus time-series at each dilution rate (δ). The power spectrum (black) is computed using Welch's method with a Hamming window of length 40 days and a 50% offset. The analytical forms S_{fold} (blue), S_{hopf} (red) and S_{null} (green) are fit to the power spectrum to obtain the AIC weights w_f, w_h and w_n, which are an indication of which fit best explains the data.

Appendix C

Early warning signals for population extinction in seasonal environments

C.1 EWS derivations in the standard Ricker model

A general non-seasonal population model

A general non-seasonal population model with demographic and environmental stochasticity may be written

$$N_{t+1} = f(N_t) + \sigma_d \sqrt{N_t} \epsilon_t^{(d)} + \sigma_e N_t \epsilon_t^{(e)}, \tag{C.1}$$

where $f(N_t)$ models the deterministic dynamics, σ_d and σ_e are noise amplitudes for the demographic and environmental noise processes respectively, and $\epsilon_t^{(d)}$ and $\epsilon_t^{(e)}$ correspond to white noise processes. For justification of these forms for environmental and demographic stochasticity, see e.g [20, 160].

Since the sum of two normal distributions is itself a normal distribution, Eqn. (C.1) can be rewritten

$$N_{t+1} = f(N_t) + \sigma(N_t)\epsilon_t \tag{C.2}$$

where ϵ_t is a normal random variable with zero mean and unit variance, and

$$\sigma(N_t) = \sqrt{\sigma_d^2 N_t + \sigma_e^2 N_t^2} \tag{C.3}$$

is the effective noise amplitude, which is state-dependent. Suppose the system is in an equilibrium state N^*. Denote $\xi_t = N_t - N^*$ as the residual dynamics about equilibrium (occurring due to noise). Then from Eqn (C.2), we get

$$N^* + \xi_{t+1} = f(N^* + \xi_t) + \sigma(N^* + \xi_t)\epsilon_t \tag{C.4}$$
$$= f(N^*) + f'(N^*)\xi_t + (\sigma(N^*) + \sigma'(N^*)\xi_t)\epsilon_t + O(\xi^2). \tag{C.5}$$

Neglecting quadratic and higher terms in residual dynamics and noise, we get

$$\xi_{t+1} = \lambda\xi_t + \sigma\epsilon_t \tag{C.6}$$

where $\lambda = f'(N^*)$ and $\sigma = \sigma(N^*)$ are functions of the equilibrium state N^*. Eqn (C.6) is an autoregressive process of order 1, from which the statistical metrics for EWS can be derived (e.g [30]). The variance of the system is given by

$$\text{Var} = \frac{\sigma^2}{1-\lambda^2}, \tag{C.7}$$

the lag-τ autocorrelation by

$$\rho(\tau) = \lambda^\tau, \tag{C.8}$$

and the power spectrum by

$$S(\omega) = \frac{\sigma^2}{2\pi}\left(\frac{1}{1+\lambda^2-2\lambda\cos(\omega)}\right). \tag{C.9}$$

The Ricker model

For the Ricker model, we have $f(N_t) = N_t e^{r-\alpha N_t}$, which has an equilibrium state of $N^* = r/\alpha$, and a local recovery rate of $\lambda = f'(N^*) = 1-r$. Inserting these terms in to the EWS expressions derived above gives a variance of

$$\text{Var} = \frac{\sigma^2}{r(2-r)}, \tag{C.10}$$

a lag-τ autocorrelation of

$$\rho(\tau) = (1-r)^\tau, \tag{C.11}$$

and a power spectrum of

$$S(\omega) = \frac{\sigma^2}{2\pi}\left(\frac{1}{1+(1-r)^2-2(1-r)\cos(\omega)}\right). \tag{C.12}$$

C.2 EWS derivations for the seasonal population model

A general seasonal population model

We consider the general form of

$$X_{t+1} = f(X_t, Y_t) + \sigma_b(Y_t)\epsilon_t^b, \tag{C.13}$$
$$Y_{t+1} = g(X_{t+1}) + \sigma_{nb}(X_{t+1})\epsilon_t^{nb}, \tag{C.14}$$

where X_t is the non-breeding population size and Y_t is the breeding population size of generation t. The functions f and g govern the deterministic dynamics in the breeding and non-breeding periods, respectively, as illustrated in Figure 5.2. We have assumed that dynamics in the breeding period depend on Y_t (via density dependent effects) and X_t via carry-over effects, which have been shown to be present in fruit fly populations [13].

Assume the system occupies a stable equilibrium (X^*, Y^*), which by definition satisfies

$$X^* = f(X^*, Y^*) \tag{C.15}$$
$$Y^* = g(X^*). \tag{C.16}$$

To examine the stochastic dynamics about the equilibrium state, we introduce the variables

$$\xi_t = X_t - X^* \tag{C.17}$$
$$\eta_t = Y_t - Y^*, \tag{C.18}$$

which correspond to deviations about the equilibrium state. We also use the shorthand notation

$$\phi_t^b = \sigma_b(Y^*)\epsilon_t^b \tag{C.19}$$
$$\phi_t^{nb} = \sigma_{nb}(X^*)\epsilon_t^{nb}, \tag{C.20}$$

for the noise terms at the equilibrium state. Substituting these terms into the model, taking Taylor expansions, and neglecting quadratic and higher terms, yields the process

$$\xi_{t+1} = \frac{\partial f}{\partial X}\xi_t + \frac{\partial f}{\partial Y}\eta_t + \phi_t^b \tag{C.21}$$

$$\eta_{t+1} = \frac{\partial g}{\partial X}\xi_{t+1} + \phi_t^{nb}, \tag{C.22}$$

where derivatives are evaluated at the equilibrium (X^*, Y^*). This can be rewritten as

$$\xi_{t+1} = a\xi_t + \gamma_t \tag{C.23}$$

$$\eta_{t+1} = b\xi_{t+1} + \delta_t, \tag{C.24}$$

where

$$a = \frac{\partial f}{\partial X} + \frac{\partial f}{\partial Y}\frac{\partial g}{\partial X}, \tag{C.25}$$

$$b = \frac{\partial g}{\partial X}, \tag{C.26}$$

$$\gamma_t = \frac{\partial f}{\partial Y}\phi_t^{nb} + \phi_t^b \tag{C.27}$$

$$\delta_t = \phi_t^{nb}. \tag{C.28}$$

Note that Eqn (C.23) is simply an AR(1) process again, and so statistical metrics can be easily derived (which correspond to EWS in X_t). The statistics of η_t can be derived from the fact that it is a linear function of ξ_t, given by Eqn (C.24). This system undergoes a bifurcation whenever $|a| = 1$ and is stable provided $|a| < 1$. The variance of ξ_t is given by

$$\mathrm{Var}(\xi_t) = \frac{\mathrm{Var}(\gamma_t)}{1-a^2} = \frac{\sigma_b^2 + f_Y^2\sigma_{nb}^2}{1-a^2} \tag{C.29}$$

showing how variance may increase depending on the relative speed at which $|a| \to 1$ and σ_b, $\sigma_{nb} \to 0$. The variance of η_t is given by

$$\mathrm{Var}(\eta_t) = g_X^2\mathrm{Var}(\xi_t) + \mathrm{Var}(\delta_t) = \frac{g_X^2(\sigma_b^2 + f_Y^2\sigma_{nb}^2)}{1-a^2} + \sigma_{nb}^2. \tag{C.30}$$

The lag-1 autocorrelation of ξ is simply

$$\rho_\xi(1) = a \tag{C.31}$$

which goes to 1 (-1) as the Transcritical (Flip) bifurcation is approached. The autocorrelation of η is

$$\rho_\eta(1) = \frac{\mathbb{E}\big[\eta_{t+1}\eta_t\big]}{\mathbb{E}\big[\eta_t^2\big]} = \frac{a\, g_X^2 \mathbb{E}\big[\xi_t^2\big] + g_X\, f_Y\, \sigma_{nb}^2}{g_X^2\, \mathbb{E}\big[\xi_t^2\big] + \sigma_{nb}^2}. \tag{C.32}$$

In the case where variance diverges at the bifurcation, $\mathbb{E}\big[\xi_t^2\big] \to \infty$ and so $\rho_\eta \sim a$ which goes to 1. Since there are two variables for this model, we can investigate the cross-correlation. We use the Pearson correlation coefficient which is given by

$$\rho = \frac{\mathbb{E}\big[\eta_t \xi_t\big]}{\sqrt{\mathbb{E}\big[\eta_t^2\big]\mathbb{E}\big[\xi_t^2\big]}} = \frac{a^2\mathbb{E}\big[\xi_t^2\big] + b^2\sigma_y^2 + \sigma_x^2}{\sqrt{\mathbb{E}\big[\xi_t^2\big](c^2\mathbb{E}\big[\xi_t^2\big] + \sigma_y^2)}}. \tag{C.33}$$

Asymptotically, we have

$$\rho \sim a^2 \quad \text{as} \quad |a| \to 1, \tag{C.34}$$

and so interestingly, the cross-correlation approaches one, regardless of the bifurcation that is being approached, and it does so quadratically with a.

C.3 Additional figures

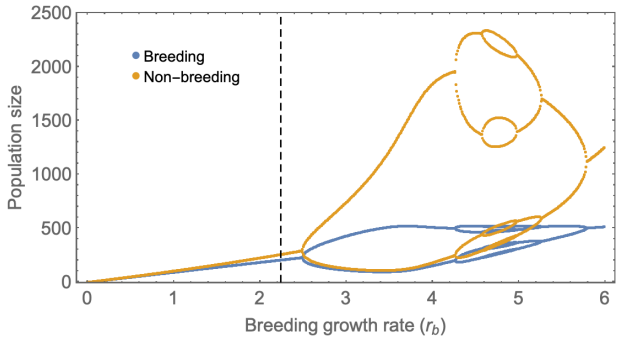

Figure C.1. Extended bifurcation diagram for the biseasonal Ricker model. As the breeding growth rate increases from low values, the system undergoes three Flip bifurcations which result in period-doubling. Upon further increase, the system undergoes three more Flip bifurcations that result in period-halving, to which the system returns to being at a fixed point. Note the region of chaos that was present in the standard Ricker model is no longer present. Parameters are at baseline values, given in Eqn. (5.4).

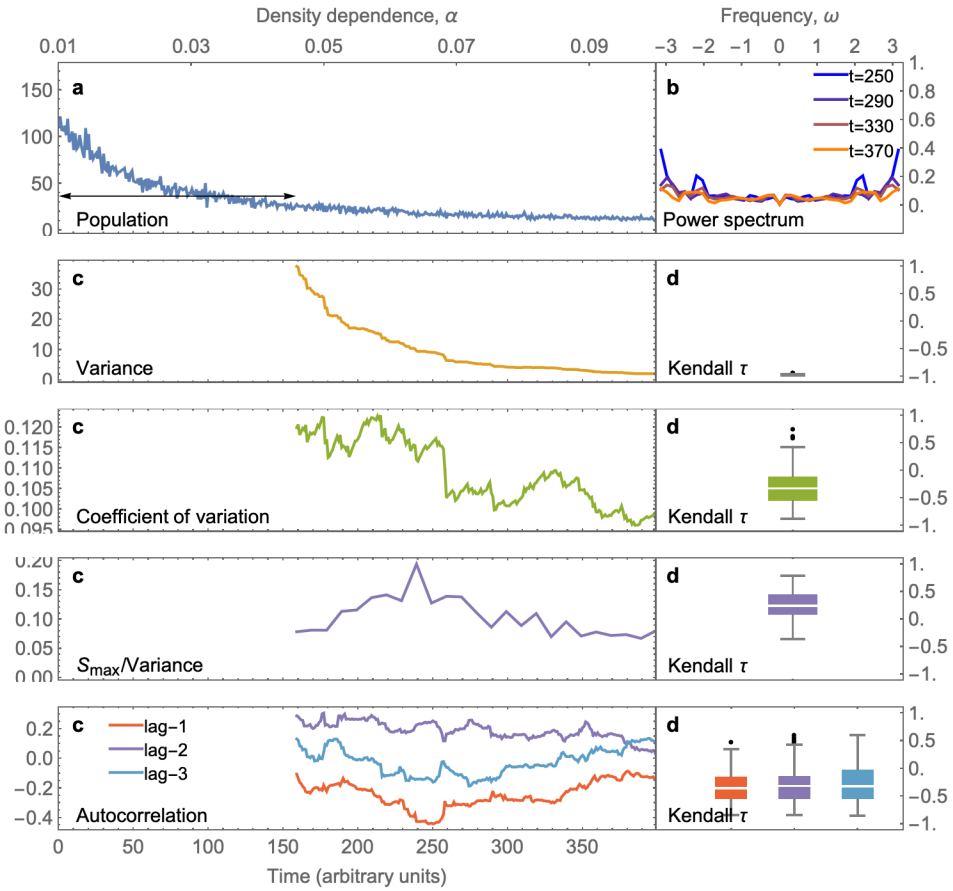

Figure C.2. EWS in the Ricker model with increasing density dependence. The parameter controlling density dependence (α) is increased linearly over the interval $[0.01, 0.1]$ during 400 time increments. The intrinsic growth rate (r) is kept fixed at 1.2. Note that the system does not undergo critical slowing down as expected (no bifurcation is approached in the model).

Appendix D

Software

D.1 A Python package: ewstools

During the course of this research, I have developed *ewstools*, a Python package for computing, analysing and visualising early warning signals (EWS) in time series data. All EWS computations in this thesis may be conducted using the relevant functions in this package. The objective for creating a package for this code was to provide a complementary code base to a currently well-used EWS package in R [55], with additional tools specific to my research, and using a platform for Python users.

The package is hosted on Github at `https://github.com/ThomasMBury/ewstools` where installation instructions, documentation, demos, and source code is available.