Academic Thesis: Declaration Of Authorship

I, Iain S. Weaver, declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

Macroscopic Principles for the Self-Organisation of Complex Ecoystems

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;

2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;

3. Where I have consulted the published work of others, this is always clearly attributed;

4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

5. I have acknowledged all main sources of help;

6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

7. Either none of this work has been published before submission, or parts of this work have been published as: [please list references below]:

Signed:

Date:

UNIVERSITY OF SOUTHAMPTON

Macroscopic Principles for the Self-Organisation of Complex Ecoystems

by

Iain S. Weaver

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

Faculty of Engineering, Science and Mathematics

School of Electronics and Computer Science

June 2015

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Iain S. Weaver

Many of the great challenges of our age are ultimately challenges of complexity. Anthropogenic climate change threatens to destabilise the ecological connections that span the biosphere as well as the delicate geopolitical functioning of our globalised civilisation. It seems a push towards developing an understanding of the seemingly unpredictable behaviour of such highly connected complex systems could hardly come at a more critical time. It is the spontaneous emergence of macroscopic phenomena from these complex systems that forms the core of this thesis.

Complexity generally results from the multiplicity of system components and is ubiquitous in natural systems. Developing an understanding of how macroscopic structures emerge from microscopic dynamics is central to furthering our understanding of complex systems. However, there exists a gulf between our capacity to formulate complex models and our ability to predict their behaviour. Additionally, large and complicated models require vast computational resources and will remain out of reach for many years, even when considering Moore's law, the observation that the increase of computing power is approximately exponential with time. This thesis contributes to building the bridge between model formulation and prediction with multiple directions of attack. We begin by analysing Watson and Lovelock's (1983) Daisyworld model by making explicit the relationship between the various model time scales, laying the foundation for our own conceptual ecosystem model with an arbitrarily diverse biota and a multidimensional environment. We make the important step of providing a framework to translate the model formulation and determine the collective behaviour of the biota.

Cellular automata and networks are methods for modelling discrete complex systems and are widely considered a staple of the complex systems simulator's toolbox. We go on to develop the analytical tools available to these models with a real space renormalisation of two dimensional non-equilibrium cellular automata along with demonstrating the utility of a master equation approach by analysing the properties of a new algorithm for growing complex networks. Lastly we investigate the utility of thermodynamic principles in predicting self-organisation in dissipative systems.

Contents

List of Figures

List of Tables

Acknowledgements

This thesis is submitted with thanks to my supervisors Dr James Dyke and Dr Adam Prugel-Bennett for hours of patience and discussion. This work was supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1).

Chapter 1

Introduction

1.1 Motivation

Interactions between simple entities can lead to unpredictable collective behaviour. These macroscopic phenomena can be said to emerge from the multiplicity of the components and are ubiquitous in natural systems, from ant colonies and avalanches to the Earth system as a whole, arguably the most complex system known (Colorni et al., 1991; Mackenzie and Mackenzie, 1998). It is home not only to a vast range of geochemical processes, but also a biosphere, the study of individual components of which comprises whole fields of research. The patchwork of vegetation across the surface of Earth emerges not merely from climatic variation but from local interactions between the biota and its environment (Neilson, 1995). However, there is hope; even when the behaviour of individuals is not well understood, it is commonly found that only key features on the microscopic level are relevant to important features on the large scale (Goldenfeld, 1992, Ch. 9). A simple example of this is social phenomena such as the behaviour of crowds of people or freeway traffic, although it is equally true for colonies of termites (Nagel and Schreckenberg, 1992; Mackenzie and Mackenzie, 1998). Disregarding microscopic complexities in favour of large populations of very simple autonomous agents enables us to replicate key emergent phenomena with surprising fidelity.

What can we say about the large scale behaviour of the Earth system then? Are environmental conditions static, intermittently stable or in constant flux? Certainly, over geological time, environmental conditions have remained well bounded for the existence of a widespread biosphere; global temperatures for example have ensured the prevalence of liquid water despite a number of planetary scale calamities such as snowball events (Hoffman et al., 1998). Since its inception as long as 3.7 billion years ago (Rosing, 1999), the elements of life have not merely persisted through growth and reproduction, but prospered, evolved, and diversified. They inhabit almost all environments across the globe from Arctic sea ice to the scalding heat of hydrothermal vents, where Wachtershauser (1990) argue the first metabolic processes originated (Horner et al., 1992). As well as flourishing, the emergence of life appears to have had profound effects on the environment, perhaps most significantly the emergence of photosynthesis, where oxygen is excreted as a by-product. This in spite of numerous pressures over a vast range of time scales to which life has been subjected; volcanic out-gassing and catastrophic impacts from asteroids exert an effectively instantaneous effect on aspects of the environment. In contrast, the variability of solar luminosity, which standard models suggest has increased from about two thirds of the modern value over geological time which can be considered sufficiently slow as to be neglected except over the very longest time scales (IPCC, 2007). In turn, terrestrial variables can be shown to vary over a similarly large range of time scales. The cryosphere, for example, incorporates responses ranging from the decadal in glacier extent to millennial in ice sheets (Goodison et al., 1999).

The global climate is seen to vary over this large range of timescales. Within an ice age, the broad state within which the Earth currently exists, glacial-interglacial transitions occur apparently cyclically, roughly every 100 thousand years. Decades of empirical study have tied a great deal of the variability to cyclical variations in the Earth's orbit (Imbrie et al., 1992). Ice core records shed significant light on the question of Earth system stability over these long times. As snow falls in regions of perpetual ice cover, the ice sheets of Greenland and Antarctica form layers with important clues and indicators of the climatic conditions year by year including temperature, precipitation and atmospheric composition. These records indicate that over the last four glacial-interglacial cycles, a period of roughly 400 thousand years, the temperature has fluctuated between two main bands which correspond to the warmer interglacial and cooler glacial Earth system states (Petit et al., 1999). This empirical approach yields extremely important insights into the long time and large scale behaviour of our planet, hinting at the existence of negative feedback and stabilising mechanisms at the heart of the Earth system. However it is not our approach in this thesis. We instead pursue mechanistic explanations both for the existence of abrupt transitions, such as the glacial-interglacial cycle and for the apparently long-term stabilising behaviour. To this end we begin from the roots of the Gaia hypothesis in Sec. 1.3 and press on to formulate very simple conceptual Earth system models, a reductionist approach which aims to distil key mechanisms responsible for the emergence of different phenomena. There is a wealth of literature related to this, particularly surrounding the Daisyworld paradigm (Wood et al., 2008). However, our approach is unique in that we begin by not by assuming stable states exist, but asking to what extent stabilising interactions are an emergent property of the Earth system, one featuring interactions between a diverse biota and a complex environment.

1.2 Target Audience

While there is no paucity of Earth system models, our conceptual approach provides unique insights and lays the foundation for a mechanistic description of a range of natural processes. All Earth system and ecosystem models are abstractions to some degree; modellers aim to incorporate the components needed to observe or predict a particular phenomenon over the time and spatial scales of interest, analogous to how the scales and features of a map are tailored to the needs of the user. This comparison fails to capture an important limitation of our modelling efforts; the ability to include all details and complexities far from our reach as it is limited not only by our incomplete knowledge, but also the amount of computing power available.

As an early example of a heavily abstracted model, Volterra (1926) study the fluctuations in the abundance of two species, a predator and its prey, in the absence of any spatial effects or other complexities. Modern efforts extend the range to very much more sophisticated modelling efforts; lake ecosystem models for example can contain

from dozens to hundreds of variables, while global climate models at the very limit of our ability to faithfully model the Earth system (Scheffer et al., 2001; Kanakidou et al., 2005). These more complicated models are typically parametrised by data and are driven by the urgency of predictive models for these potentially volatile systems. Simpler models such as Watson and Lovelock's (1983) Daisyworld do not attempt to provide utility in predicting the behaviour of any specific system. Rather, they attempt to isolate key features responsible for the emergence of a specific phenomenon. In doing so they are able to suggest a potential mechanism for such features to emerge in very much more complicated systems where it is otherwise difficult to attribute specific behaviour to any subsystem. Our approach is unique in that it begins from these very large abstractions in Chapter 2, and discards important limitations such as the low biotic diversity and environmental dimensionality which typically exist in conceptual models while retaining almost all tractability, generating a new model in Sec. 3.

The thesis aims to be of broad interdisciplinary interest. Researchers in the field of Earth systems science will find the work illuminating as we begin by disconnecting the classical Daisyworld model from concepts of Gaia and planetary homeostasis. This important step begins by identifying many of the less palatable assumptions in the Daisyworld model; the biota is small, the environment is simple and interactions are designed such that homeostasis emerges. In doing so we can show these assumptions to be unnecessary, and define a broader requirements for the emergence of stability and homeostasis in a coupled life-environment model. Our new model, explored in Chapter 3, will be intuitive to proponents of the Daisyworld paradigm while discarding many prescribed features which typically lead to a large proportion of the criticism of the model. These changes greatly broaden the horizons for study of such models to multidimensional homeostasis and importantly criticality and state transitions in complex systems, a problem of intense interest in contemporary economics, ecology, social science and perhaps most pressingly, climate science. Models predict numerous climatic and Earth systems to be approaching unprecedented critical points and in most cases, the end point of such a transition and its implications for the much wider, strongly coupled Earth system as a whole is entirely unclear.

Purely abiotic systems enjoy an equally rich plethora of emergent behaviour on all scales. We investigate the more physically grounded phenomenon of Rayleigh-Bénard convection in Chapter 7; a ubiquitous form of turbulence found on all scales from a heated pan of soup to the planetary atmosphere and mantle. We explore the utility of a principle which aims to forego resolution of microscopic system dynamics in favour of constrained global optimisation of some thermodynamic property, such as entropy production. The existence of such a principle would, if proven, revolutionise the way Earth-systems researchers model principally thermodynamic systems such as oceans and atmospheres; the aim of a recent book publication, where our contribution (Weaver

et al., 2013) sits alongside numerous contributions ranging from atmospheric physics to enzyme kinetics.

This work has been presented at the European Conference for Artificial Life along with publications in the Journal of Theoretical Biology and PLoS Computational Biology (Weaver and Dyke, 2013, 2012b; Dyke and Weaver, 2013). While Watson and Lovelock's (1983) Daisyworld did not originate in this field, a great deal of interest developed from the domain of artificial life; both explore notions of homeostasis and evolution, although the Daisyworld paradigm challenges artificial life researchers to consider not just organismic evolution, but rather the co-evolution of the biotic and abiotic elements of the system. Our work takes this extension further by way of generality, exploring the general properties of a biota adapting to a multidimensional environment, even under the influence of varying external perturbations, presenting new avenues and challenges for research into artificial life systems.

Throughout, we advocate a bottom-up approach, beginning with well understood microscopic components, and using a range of statistical methods along with simulation to extrapolate the interesting system-level behaviour. While the methods are largely standard within the domains of mathematics and physics, the applications are novel, and have found traction in the Journals Complexity and of Statistical Physics, where we have published a renormalisation method for one- and two-dimensional non-equilibrium cellular automata (Weaver and Prügel-Bennett, 2014, 2013).

1.3 Gaia hypothesis

Lovelock and Margulis (1974) propose the Gaia Hypothesis; that the persistence of environmental conditions conducive to life emerges from the organisation of the Earth system into self-regulating states, analogous to how organisms regulate their internal environment. The latter is uncontroversial; there are clear evolutionary advantages to organismic homoeostasis, maintenance of internal conditions is crucial to the continuation of the complex set of metabolic processes on which all organisms depend, and indeed all life exhibits some degree of internal regulation. In addition to this, certain large scale abiotic homeostatic mechanisms are generally appreciated, such as the weathering of silicate rocks which removes atmospheric CO_2, ultimately ending up in the ocean. Higher temperatures accelerate the process and serve to offset increases in global temperature to some extent (Walker et al., 1981; Berner, 1991). However, the sheer complexity of interactions between organisms and their environment greatly obfuscate apparently simple biotic homeostatic mechanisms and reveal complicated limitations. Schwartzman and Volk (1989) discuss the interactions of a range of biotic processes with the relatively simple silicate rock weathering, tentatively suggesting this chemical weathering is enhanced by life on Earth although there are a number of competing processes.

The Gaia hypothesis is one with many forms and variants, so it is important to be clear which version we advocate, or argue against. Lovelock's (1979) suggestion that the biosphere as a whole regulates its environment, homeostatic to the extent that it my be considered a *super-organism* is referred to as Geophysiological or Strong Gaia. This description is contentious for a number of reasons; to begin with, the notion of the Earth being organismic in the sense that it be alive is particularly unpalatable; for one thing, it is unable to reproduce and is not subject to natural selection, two staples of living systems (Margulis and Sagan, 2000). Opponents to the Gaia hypothesis often cite the existence of ice ages as evidence against the Gaia hypothesis. The reason for this is that ice ages can be shown to correspond to dramatic reduction of biomass and diversity and the shift to an ice age state can be said, in some sense, to be deeply unfavourable to life (Prentice et al., 2011; Tyrrell, 2013) (although the end of an ice age due to bio-geophysical processes sits well with Lovelock's view). This argument is in fact against another extreme form of the Gaia hypothesis called Optimising Gaia, whereby the biota is said to configure its own environment such that conditions are in some way optimal for life, such as by maximising global biomass (Lovelock, 1995; Ackland et al., 2003). The existence of some optima is clearly teleological and extremely hard to support scientifically. It can be argued that both of these hypotheses serve as useful metaphors for the Earth-system as a mechanism though neither represent testable hypotheses (Kirchner, 1989) and are not dealt with further in this thesis.

How then should we expect life to organise on the system level? Does it *conform* and simply co-evolve with its environment, or are there mechanisms by which the biota might collectively have a hand in regulation of key environmental factors? We have encountered Strong Gaia and now at the other end of the scale, Weak Gaia, or the Coevolution hypothesis (Kirchner, 1989, 2002). Here the biosphere is said to exert a strong force on the environment, exactly with the Gaia hypothesis although where Gaia suggests this feedback promotes robustness in the Earth system, coevolutionists make no claim about the eventual outcome of this coupling, and the apparent stability of the Earth system remains in question; there is no mechanism to account for the past, or predict future environmental conditions. Indeed the existence of this coupling was established before the Gaia hypothesis; biogeochemists already had an appreciation of the myriad of biologically mediated processes and their impact on the environment (Kirchner, 1989).

We seek knowledge and intuition of the Earth system, particularly the extent to which it can be viewed as an, at least, intermittently stable system. We do so through abstractions to very simple models from which salient Earth system features emerge, called conceptual ecosystem models, one of the earliest of which was Watson and Lovelock's (1983) Daisyworld which describes an Earth-like planet orbiting a star which is home to a very simple biota, consisting only of black and white daisies which completely cover the surface of the planet. The two daisy types are differentiated in the following ways;

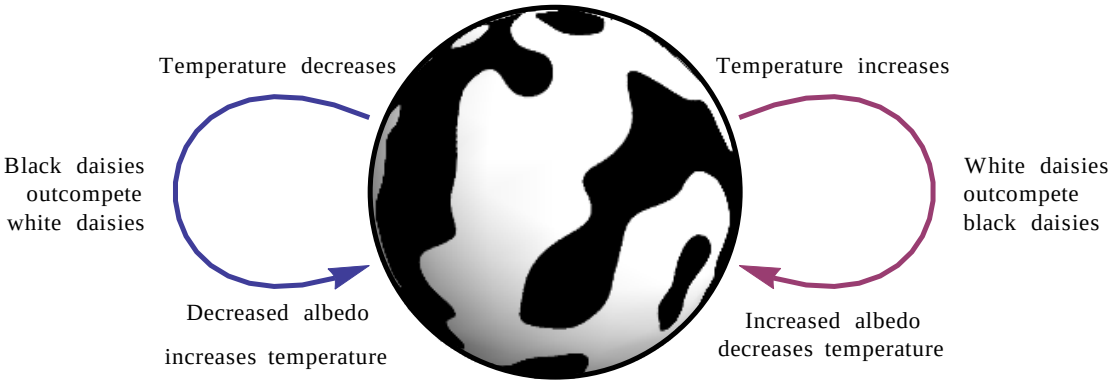

FIGURE 1.1: Watson and Lovelock's (1983) Daisyworld describes a planet covered in black and white daisies, designed such that a cybernetic feedback system emerges; increases and decreases of temperature are opposed by the proliferation of white and black daisies respectively.

black daisies have a small contribution to planetary albedo and therefore an increasing effect on the planetary temperature at radiative equilibrium, while white daisies have the opposite influence. The biota exists in a mixed state of black and white daisies at some temperature above which white daisies out compete black daisies and proliferate, and below which black daisies out compete white daisies. This cybernetic feedback system is special in that the configuration of daisy albedo and preferred climate leads to the negative feedback loop shown in Fig. 1.1; an increase in planetary temperature, such as by changes in the solar luminosity, leads to a relative increase in the population of white daisies, and increase in planetary albedo and an opposing change in temperature, regulating the global temperature in a reign-control state, so called as proliferation of one daisy time can only have either an increasing or decreasing influence on global temperature. Like reigns, they can pull but not push, and any single daisy type cannot produce stability; an antagonistic pair with opposing effects is needed.

Between the strong and weak Gaia hypotheses then lies a moderate theory where life-environment interactions lead to generally negative feedback processes. Why though should negative feedback, which promotes stability, dominate over positive feedback? Indeed there are many biotic processes which appear to have a destructive or destabilising influence on the environment (Kirchner, 2002). The reign-control system in Daisyworld exists by design and does not go so far as to suggest a mechanism by which the antagonistic species pair could emerge (aside from the fact that planets in other configurations would be barren, and we would not remark upon them at all, which is essentially the statement of the Anthropic principle). A large number of extensions to the Daisyworld model exist to begin to address this criticism. The key feature of homeostasis persists throughout many further studies including increases in daisy species diversity and inclusion of mutation and adaptation (Lovelock, 1992; Lenton and Lovelock, 2000; Wood and Coe, 2007). The original model exists only as weakly coupled differential equations, and other studies make the distribution of daisy species in space explicit

having little influence on the regulatory behaviour of the model, starting from von Bloh et al. (1997) and later being developed to study a range of pattern forming phenomena such as desertification (Ackland et al., 2003). In all instances, it can be argued that homeostasis is prescribed by the model formulation and while novel in the first instance, there remains no mechanistic explanation for the reign control system emerging. We address this point thoroughly in Chapter 3 by removing the prescriptive formulation of the model to one in tune with notions of complexity and diversity, randomly parametrising a large population to interact with a multidimensional environment while retaining the simplicity and tractability which made the original Daisyworld so compelling.

Such a model is ideally suited to press onwards to address a number of other contemporary questions. Criticality in Earth-system components is an ever more pressing issue. As pressure on the Earth system from humans, such as anthropogenic emissions, climb ever higher, the response of such systems to these sources of stress are unclear. To what extent do we expect complex systems to maintain robustness in the face of increasing perturbations? This question is of broad interest as complex systems are almost synonymous with systems which undergo abrupt transitions in much the same way as simple physical systems such as magnets. Below the critical temperature, the system is magnetised and is robust to the application of an opposing external magnetic field which can be thought of as an external perpetration or force. However, at some critical field strength, magnets of at least two spatial dimensions spontaneously align with an external field, a phenomena first proven by Peierls and Born (1936) and famously solved for a two dimensional square lattice arrangement by Onsager (1944). The position of this transition has received a great deal of attention and by now is generally well understood for the case of magnets, although failure to understand and predict critical behaviour of economic systems comes at the cost of millions of dollars. The cost of a poor understanding of transitions in social systems is unequivocally greater still and many climatic models indicate the price of ignorance to be ruinous. This thesis exposes its general model formulation to a range of multidimensional forcing in order to deepen our intuition of the critical behaviour of multidimensional complex systems. In doing so we are situated to examine the important notion of early-warning signals; identifiable system features that precede abrupt transitions, which have received much attention in climate science with varying degrees of success (Scheffer et al., 2009). While successful in some instances, there are a variety of methodologies for identifying early-warning signals but no systematic way of determining which is applicable to a specific system. Through minimal conceptual models, we may find appropriate early-warning signals and seek to reconcile key model features with those of real Earth systems.

1.4 Structure of thesis

Chapters 2-4 revolve around the questions of Earth-system stability, beginning with an investigation of Watson and Lovelock's (1983) classic Daisyworld model, and wading deeper by asking to what extent we expect parametrised Earth-system models to enjoy stability at all, probing deeper into the behaviour of such models when subjected to increasing environmental pressure, an important contemporary problem. Chapters 5-6 seek to extend this work by embedding complex systems in space such as design of cellular automata, or more generally network models, along with a discussion of thermodynamics and how complex phenomena emerge from purely abiotic processes, and potential macroscopic principles which explain their apparently spontaneous emergence. We focus on exploring the applications and utility of novel analytical techniques including the principle of Maximum Entropy Production, the renormalisation group and master equations to more abstract systems.

Chapter 2

The thesis begins by detailing Watson and Lovelock's (1983) Daisyworld, the pioneering conceptual ecosystem model we aim to understand, reduce and soon reimagine as a generic model of multidimensional complexity. Any investigation of feedback and regulation of the Earth system must be accompanied by consideration of time scales, particularly in determining limiting processes for the emergence of any regulatory dynamics. We begin with a great deal of simplification, enabled by separating time scales related to different model processes. In doing so, we are able to provide systematic analysis of the effects of relaxing this separation, finding the rate of change of external perturbation to impose limits on the regulatory power of the model, and non-zero heat capacity to result in temperature oscillations (Weaver and Dyke, 2012b). This work has been particularly important in informing subsequent investigations, conveniently highlighting the importance of relating model time scales in an appropriate and physical way.

Chapter 3

The extent to which a system as complex as the Earth, with vast spatio-temporal variability born of a myriad of interacting biotic and abiotic components, could be expected to enjoy any degree of stability at all is unclear, and somewhat contentious. Does this complexity promote fragility or robustness? Previous work in this direction typically revolves around the Daisyworld paradigm (Watson and Lovelock, 1983), encompassing a biota with limited variability and a simple environment consisting of few variables. This chapter begins by generalising this formulation of life-environment feedbacks, suggesting a minimal set of assumptions for the emergence of stationary, homeostatic behaviour. The new model shares many features with Daisyworld, with a number of notable differences. Firstly, the biota in Daisyworld is parametrised such that a reign-control state emerges, while

our model is built from a large number of randomly parametrised species. Secondly, while Daisyworld hinges on a single environmental variable, we generalise the model to any dimensionality. We find homeostasis to be a general property of this type of model and furthermore, this minimal model proves largely tractable in that the statistics of model fixed points may be calculated exactly from model parameters. Furthermore, by examining the behaviour of random Gaussian matrices we may classify these points, determining the distribution of stable, unstable, and those fixed points with oscillating behaviour.

Chapter 4

We proceed to examine in detail the nature of transitions between stable points in a multi-dimensional environment in Chapter 4. One-dimensional hysteresis loops are iconic phenomena and commonly used as a metaphor for transitions in very much higher-dimensional systems, while not taking into account the additional degrees of freedom which may result in cyclical, complex or chaotic behaviour, inaccessible to such a simple metaphor. On the other hand, complicated systems typically suffer from a lack of transparency and tractability. We extend the study of the simple model introduced in the previous chapter to explore critical transitions between the emergent network of fixed points in multiple dimensions. Our model is subjected to a range of multidimensional external forces, enabling us to identify phenomena beyond those accessible to simple, low dimensional, Daisyworld-like models, such as cusp bifurcations. Exactly as before, the simple formulation leaves the model tractable to a statistical analysis, and we are able to determine the distribution of transitions with external perturbing forces. Our main results can be proven to be highly general properties, as opposed to special instances of the model.

Chapter 5

A key challenge in developing a thorough understanding of the Earth system is not only in creating a framework consistent with a system of such spatio-temporal variability, but in developing the appropriate analytical tools with which to treat it. To this end, cellular automata receive a great deal of attention in the study of spatially embedded interacting systems as they are arguably the simplest spatial model possible. They have become a popular tool across a broad range of sciences and have found a plethora of applications from modelling free-way traffic flow to neuron activity. Indeed Daisyworld finds a number of spatial implementations; continuous one-dimensional by Adams et al. (2003), discrete two-dimensional by von Bloh et al. (1997) and even Ackland et al.'s (2003) continuous curved two-dimensional Daisyworld. However, the translation between our choice of cellular automata rules, the microscopic model dynamics, and the large scale behaviour they create is not simple which is itself a part of the charm; apparently simple local rules can lead to unpredictable, complex or even chaotic large scale phenomena. Renormalisation group theory seeks to bridge the gap by providing a framework to translate between the microscopic and large-scale dynamics (Goldenfeld, 1992,

Ch. 9). Chapter 5 contributes to this research by describing a renormalisation group approach to analysing the large scale behaviour of cellular automata, and applying it to both one and two dimensional models (a first for cellular automata which feature an absorbing state). We are able to demonstrate that when the renormalisation is not exact, its quality is improved by estimates of the stationary state distribution on the microscopic level, along with providing methodology to calculate the microstate distribution.

Chapter 6

Cellular automata are but one way of representing spatially embedded systems. Concepts of locality and adjacency are much richer, not confined to a regular lattice; in reality model elements may be spatially close, but not interact in any significant way. Many species share an environment but have no direct interactions Networks enable us to represent this type of system faithfully, and have been used as models of a diversity of complex systems from ecology to social sciences where spatial influence on the system dynamics is strong, but not necessarily simple. Indeed they offer an alternative form of Earth system models when compared to cellular automata models such as von Bloh et al.'s (1997) two-dimensional Daisyworld, and our model. Their generality arises from the range of interpretations of the role and relationship between edges and vertices. Grown networks, such as by iteratively adding edges and vertices to a graph, possess evidence of their history in various network properties. In Chapter 6 we make use of generating functions to better understand the effects of network growth on the degree and component size distributions along with degree correlation between connected vertices. We are able to pit the effects of history against the effects of preferential attachment and find qualitatively different behaviour in the degree correlation between connected vertices with increasing network connectivity.

Chapter 7

The Earth system and all encompassed processes work by degrading energy, transferring it from high to low densities such as from hot to cold regions in atmospheric circulation and precipitation. Circulation is driven by the temperature gradient between the equator and poles while the hydrological cycle is driven by evaporation of surface water, transferring energy to the cold atmosphere (Ito and Kleidon, 2005). Both processes are governed by surface heating from the sun; a driving force which maintains the requisite temperature gradients for these processes, and both mechanisms are, in isolation, sufficiently well understood that they can be described in simple terms as obeying thermodynamic laws. In contrast to equilibrium thermodynamics, there is a relative paucity of macroscopic governing principles applicable to far from equilibrium systems. The result of the diversity of interacting non-equilibrium processes on Earth appears extremely unpredictable, resulting in large uncertainty when predicting local phenomena such as weather in the short term, and the global climate in the long term. It has been suggested that

non-equilibrium dissipative systems may self-organise to optimise certain thermodynamic quantities. The principle of maximum entropy production proposes that the internal entropy production of a system is such a quantity (Paltridge, 1975). The potential role of life as some subsystem in a global thermodynamic engine did not go unnoticed; since its conception, the principle has even been compared directly to the Gaia hypothesis both as conflicting (Ackland, 2004) and complementary (Dyke, 2008) organisational principles. We investigate the utility of this contentious principle in predicting the organisation of a simple fluid system; the configuration of convection cells in a heated fluid. We describe a simple lattice gas model, which may be thought of as a generalised cellular automata, and use simulations to determine that the preferred convective states appear to maximise total heat flux, which is inverse to entropy production when the temperature gradient is large, a result which is in contradiction to the maximum entropy production principle.

1.5 Publications

Portions of this thesis have been previously published in the following, for which my main contribution is noted

Chapter 2

I S Weaver and J G Dyke. The importance of timescales for the emergence of environmental self-regulation. *Journal of Theoretical Biology*, 313(0):172–180, 2012b. ISSN 0022-5193

Original analysis of the Daisyworld model with explicitly separated timescales, along with analysis of relaxed constraints and numerical results.

Chapter 3

I S Weaver and J G Dyke. A novel approach to analysing fixed points in complex systems. In Thomas Gilbert, Markus Kirkilionis, and Gregoire Nicolis, editors, *Proceedings of the European Conference on Complex Systems 2012*, pages 523–533. Springer, 2012a

J G Dyke and I S Weaver. The emergence of environmental homeostasis in complex ecosystems. *PLoS computational biology*, 9(5):e1003050, May 2013. ISSN 1553-7358

Using the limit of a large biosphere to implement Dyke's (2010) 'Daisystat' model as a Gaussian process, and further providing numerical simulations and analytical treatment of the model.

Chapter 4

I S Weaver and J G Dyke. Tipping points in Complex Coupled Life-Environment

Systems. In *Advances in Artificial Life, ECAL 2013*, pages 387–394. MIT Press, September 2013. ISBN 9780262317092 The application of multidimensional forcing to Dyke's (2010) 'Daisystat' model, the following analytical treatment and identification of the critical slowing-down phenomena in the model, along with the simulated time series data and analysis.

Chapter 5

I S Weaver and A Prügel-Bennett. Renormalization group approach to 1D cellular automata with large updating neighborhoods. *Complexity*, (1), May 2014. ISSN 10762787

I S Weaver and A Prügel-Bennett. Renormalisation of two-dimensional cellular automata with an absorbing state. *Journal of Statistical Mechanics (in press)*, 2013

Largely independent construction of the renormalisation algorithm and numerical implementation.

Chapter 6

I S Weaver. Preferential attachment in randomly grown networks. *Journal of Statistical Physics (submitted)*, 2014

Almost entirely independent modification of the network growth algorithm, analytical work and numerical statistics.

Chapter 7

I S Weaver, J G Dyke, and K Oliver. Can the principle of Maximum Entropy Production be used to predict the steady states of a Rayleigh-Bérnard convective system? In Roderick C Dewar, Charley H Lineweaver, Robert K Niven, and Klaus Regenauer-Lieb, editors, *Beyond The Second Law: Entropy Production and Non-Equilibrium Systems*. Springer, 2013

Implementation of a random thermal flipping component in an FHPB model along with all numerical results and analysis.

Chapter 2

Timescales in Watson and Lovelock's Daisyworld

2.1 Introduction

To what extent the emergence and evolution of life on Earth has resulted in a planetary system that is in some respects homeostatic or self-regulating was the focus of the original Gaia Hypothesis (Lovelock and Margulis, 1974). In the original phrasing, it was suggested that this regulation may produce the optimal environmental conditions for life on Earth, such that some measure of activity or abundance is maximised; a form which has since become an extreme form of the idea, optimal Gaia. There are numerous objections one might raise to this idea, the dramatic reduction of biomass corresponding to an ice-age transition for example (Tyrrell, 2013). However, such arguments in no way undermine the weaker definition of the Gaia hypothesis, homeostatic Gaia, whereby we accept that life on Earth has a strong effect on environmental conditions, and hypothesise that the interactions have an overall stabilising, or homeostatic influence.

The antecedents of the Gaia hypothesis can be traced to the work of Vernadsky (1926) who popularised the term 'biosphere' and argued that one cannot understand the surface conditions of the Earth without factoring in life's effects. Lovelock's initial insight was to realise that any widespread biosphere should, in principle, be detectable from space as the effects of life would be to produce atmospheres far from equilibrium (Lovelock, 1965). A terrestrial example is the presence of both non-negligible amounts of oxygen and methane in the atmosphere. It is the continual biogenic production of methane that replaces the methane that would be rapidly lost as it reacts with oxygen (Cicerone and Oremland, 1988). Life, via unavoidable metabolic by-products, affects the movement and cycling of chemical species within the Earth system and in doing so has affected not only the Earth's atmosphere but also its oceans, crust and cryosphere. Lovelock speculated that the evolution of life and its interaction with different elements of the Earth system gave rise to a system in which the surface conditions of Earth had been regulated to within the bounds that would be required for widespread surface life such as maintaining surface liquid water over appreciable fractions of the globe which, while seemingly trivial, would not necessarily be the case in the absence of the diversity of biotic processes (Lovelock, 1979). The teleological implications of this were not lost on contemporary thinkers and the original hypothesis was controversial and subject to pointed critique (Dawkins, 1983; Doolittle, 1981). As more recent reviews show, the Gaia Hypothesis has matured over time and rather than being fundamentally incompatible with neo-Darwinism, geology or climatology, is now seen as being relevant to a range of subjects that investigate the Earth system (Lenton, 1998; Free and Barton, 2007). For example, the abiotic process of chemical weathering of silicate rocks has produced a negative feedback loop that has to a certain extent offset the increased luminosity of the sun over geological timescales (Walker et al., 1981; Berner, 1991). It has been proposed that land plants have an overall increasing effect on this process and so have altered the fixed points of atmospheric CO_2 (Schwartzman and Volk, 1989). Evidence for the

biological amplification of chemical weathering by land plants has been traced back to the Ordovician (Lenton et al., 2012).

In terms of quantitative models and simulations, much of the investigation into the plausibility of Gaian self-regulation has been progressed via developments and extension of the Daisyworld model (Watson and Lovelock, 1983). Daisyworld was originally intended as a mathematical proof of concept for planetary homeostasis that emerged from the interactions between simple life forms and their environment. Daisyworld is a grey, Earth sized planet that orbits a sun-like star. Much like our own sun, this star has increased in luminosity over geological timescales. On a lifeless planet, this increase in luminosity would produce a proportional increase in global temperature. However, Daisyworld is seeded with black and white daisies which prove to be important components in a planetary self-regulating system that leads to global temperature not increasing but remaining within comparatively narrow bands for a wide range of luminosity. It is assumed that there are sufficient nutrients and water for daisies to grow and the only environmental factor that determines their growth is the local temperature. As well as being affected by temperature, the changing coverage of black and white daisies affects the albedo of the planet and so the amount of energy from the star that is reflected back out into space and consequently global temperature. As luminosity increases, it is the proportional coverage of black and white daisies that changes while global temperature remains relatively stable.

The original model has since been significantly extended, developed and applied to new domains. One change most significantly of interest to ecologists is the addition of a more diverse or evolving biota. Dawkins (1983) was particularly interested in the apparent altruism in this model in the sense that the daisy species expend energy to produce pigments which ultimately lead to the regulatory mechanism. Lovelock (1992) introduces a third daisy type which appears grey and has identical radiative properties to a barren planet while enjoying a greater reproductive rate. However, this does not undermine the regulation as local differences in temperature still exist between the daisy types (Lansing et al., 1998). Indeed you can add an arbitrary number of differently shaded daisies to the model. The nature of the governing equations ensures competitive exclusion which only ever leads to the coexistence of two daisy species (Ackland, 2004), and the regulatory mechanism remains identical to the much more transparent two-species model which we will explore in detail.

An initial and ongoing criticism of Daisyworld, was that the model is too simple and too contrived to give us any useful information and insights into the actual Earth (Kirchner, 2002). While some of these criticisms were addressed at the time by Lenton and Wilkinson (2003) there has been relatively little concentration on the respective timescales of the different processes operating in Daisyworld before or since. In particular, while the luminosity of the star, which increases over geological timescales, can be approximated as remaining constant while daisy coverage changes, many perturbations of the Earth

system occur over much shorter timescales. The response of organisms and ecosystems to an impact of a sizeable meteorite into the Earth would not be adaptation but obliteration and less dramatic perturbations will occur over short timescales relative to biological organisms and ecosystems. Some of these may be related to geological processes such as earthquakes and volcanism (Zielinski et al., 1994). With regards to anthropogenic climate change, it is not simply the total amount of CO_2 that humans have emitted into the Earth's atmosphere but rather the rate at which this has happened and consequently the speed of the change in radiative forcing (IPCC, 2007). Other anthropogenic perturbations can occur over even faster timescales and affect different elements of the Earth system. For example humans' impacts on the nitrogen cycle have, post industrialisation, led to dramatic changes in the global nitrogen cycle (Vitousek et al., 1997) and such effects need to be considered alongside other processes such as land use change and deforestation. Models of self-regulating mechanisms need to be able to incorporate these processes or alternatively tell us when regulation will fail. They also need to consider how the different timescales operating within the Earth system interact to function as a coherent self-regulating mechanism and what are the bounds and limits for such interacting processes. As well as being subjected to different timescales of perturbation, the Earth system is comprised of a multitude of interacting sub-systems that operate over a wide range of temporal and spatial scales. The Earth's oceans have profound effects on climate and given that liquid water has approximately four times the mass heat capacity of air, interactions between oceans and atmosphere will feature lags, delays and other behaviour associated with the interactions of processes interacting over different timescales. This represents a fundamental challenge in understanding processes such as the El Niño/La Niña-Southern Oscillation (Battisti and Hirst, 1989).

It is appreciated that the biota will not only respond to changed climate, but that such changes could attenuate or amplify initial forcing. The development of a new generation of coupled general circulation and dynamic vegetation models predict land ecosystems and oceans act as net sinks for increased CO_2 until critical thresholds are reached at which point they will start to behave as net sources (Cox et al., 2000). Consequently, rather than environmental conditions remaining fixed while life adapts to a perturbation, the relaxation to a steady state is instead a moving target with biotic effects continually moving environmental conditions. Depending on the particular elements of the Earth system of interest, the oscillations and progress towards the steady state values may be more important than the steady state values themselves.

Time dependencies have been examined in Daisyworld previously (Zeng et al., 1990; Gregorio et al., 1992; Wood et al., 2006). However these have all proceeded on the basis that the timescales of certain processes have already been determined. In particular the increase in luminosity occurs over the very longest timescales. Other timescales such as the rates of daisy removal and establishment and the response of environmental variables (such as temperature) to changes insolation and albedo can remain implicit.

In this study we make explicit these timescales and explore the self-regulating mechanism as they are relaxed. Our motivation in doing so is to assess the robustness of self-regulation and consequently how it may be applied to the analysis of elements of the Earth system. To begin, we reintroduce Watson and Lovelock's (1983) Daisyworld model in Sec. 2.2, and identify the model timescales in Sec. 2.3. Here we show that by making the relevant timescales explicit and make comparisons with timescales of similar processes on Earth. We are able to eliminate the time dependence of global temperature and daisy coverage by instead solving for their steady state values. Following Hankin and Mitchell (2011), we eliminate the need to explicitly include bare habitable ground. However, rather than parametrising infiltration between daisy species, we further separate the timescales associated with changes in daisy coverage and begin analysis when the daisy growth process occurs on much shorter timescales than daisy removal. Solutions to this model in Sec. 2.4 highlight some aspects of the original Daisyworld model such as the decrease in temperature with increased solar forcing and hysteresis loops. These solutions represent limiting cases, and are exploited in Sec. 2.5 and Sec. 2.6 to relax the previous assumption of separated timescales and produce illuminating analysis of the model response to increased rates of external forcing and instantaneous shocks. Sec. 2.7 briefly investigates the role, if any, of space in the Daisyworld mechanism. The onset of temperature oscillations due to increased heat capacity and the impact of seed diffusion rates, and effects of local seeding are also explored and summarised in Sec. 2.8.

2.2 Model formulation

Watson and Lovelock's (1983) Daisyworld was originally intended as a mathematical proof of concept for planetary homeostasis that emerges from the interactions between simple life forms and their environment. Daisyworld is a grey, Earth sized planet that orbits a sun-like star which provides the planet with incident solar radiation L. Much like our own sun, this star has increased in luminosity over geological time scales. On a lifeless planet, this increase in luminosity would produce a proportional increase in global temperature. If the planet is modelled as a black body reflecting the fraction A of incident solar radiation, its steady state temperature T is found by the StefanBoltzmann law

$$C\frac{\mathrm{d}T}{\mathrm{d}t} = L(1-A) - \sigma T^4 = 0 \tag{2.1}$$

where C is the planetary heat capacity and σ is the StefanBoltzmann constant. However, Daisyworld is seeded with black and white daisies whose share of the planetary surface is denoted α_b and α_w respectively. These prove to be important components in a planetary self-regulation and lead to global temperature not increasing but remaining within comparatively narrow bands for a wide range of luminosity. It is assumed that there are sufficient nutrients and water for daisies to grow and the only environmental

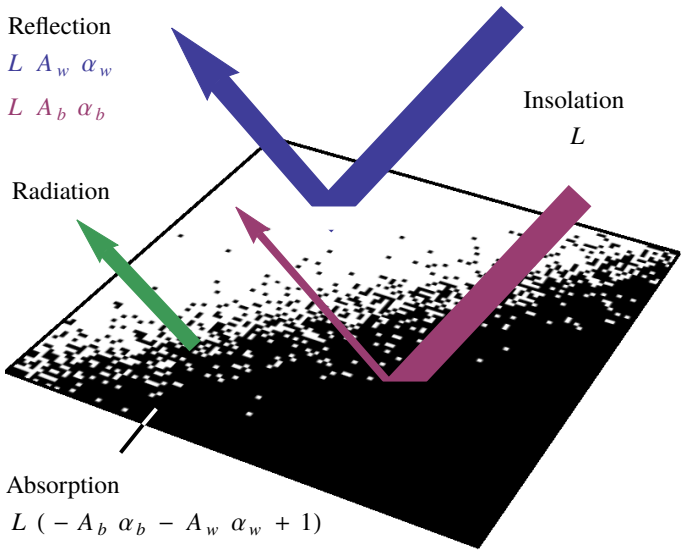

FIGURE 2.1: The surface of Daisyworld is inhibited by black and white daisies population share and albedo α and A respectively. Black daisies are more absorptive than white daisies such that $A_b < A_w$. The flux incident and reflected by the planet is indicated by arrows whose thickness illustrates the relative flux density. Fluxes interacting with the white and black daisies are coloured blue and red respectively.

factor that determines their growth is the local temperature. As well as being affected by temperature, the changing coverage of black and white daisies affects the albedo of the planet and consequently global temperature. The planetary albedo A is calculated from the black and white albedos and that of the unpopulated planetary surface, A_b, A_w and A_e respectively.

$$A = \alpha_b A_b + \alpha_w A_w + x A_e \tag{2.2}$$

where x is the fraction of the planetary surface not occupied by either daisy type

$$x = 1 - \alpha_b - \alpha_w \tag{2.3}$$

and A_e is chosen such that black daisies have a relatively low albedo and white daisies have a relatively high albedo, $A_b < A_e < A_w$. Fig. 2.1 shows the radiative balance of the planet.

Daisies proliferate into the bare regions of the planet x at a temperature dependant rate $\beta(T)$, and sites revert to bare ground at the natural death rate rate γ. The growth function β is in many senses arbitrary, providing it possesses a single local maximum at $T_{\text{opt.}}$, the optimal temperature for daisy proliferation. While analytically inconvenient, Watson and Lovelock's (1983) original model uses a parabolic function in the range of

habitability where the function would be positive, and zero elsewhere

$$\beta(T) = \begin{cases} \frac{(T-T_{\text{opt.}})^2}{\Delta T^2} & T_{\text{opt.}} - \Delta T \leq T \leq T_{\text{opt.}} + \Delta T \\ 0 & \text{otherwise} \end{cases} \tag{2.4}$$

where ΔT parametrises the width of the parabola by the points it intersects zero, $T_{\text{opt.}} \pm \Delta T$. The original Daisyworld model is in a sense *zero-dimensional* in that spatial distributions and non-uniformities are not explicitly resolved; the seeding, or spread of daisies therefore occurs globally. The change over time in the coverage of the daisies is written as a pair of differential equations

$$\frac{\mathrm{d}\alpha_b}{\mathrm{d}t} = \alpha_b\big(x\beta(T_b) - \gamma\big), \tag{2.5a}$$

$$\frac{\mathrm{d}\alpha_w}{\mathrm{d}t} = \alpha_w\big(x\beta(T_w) - \gamma\big) \tag{2.5b}$$

where T_b and T_w are the local temperatures of the black and white daisies respectively. This is an important feature as the minimal ingredient for self-regulation to emerge from this model is that the black and white daisies must be differentiated by their local temperatures. For finite diffusion rates, this property will emerge due to gradients in the absorption of heat between black and white daisies. Watson and Lovelock (1983) assume that all daisies of each type are subject to a uniform temperature field, neglecting any differences which would emerge from gradients in surface albedo. The local temperatures are given by

$$T_i = T + q\,(A - A_i) \tag{2.6}$$

which ensures the mean global temperature $T = \alpha_b T_b + \alpha_w T_w + x T_e$, where q parametrises the departure of local temperatures from the global mean. Given finite rates of diffusion of heat, the black daisies are slightly hotter and the white daisies are slightly colder than the global temperature. As q increases, the amount of insulation between the daisy fields increases and so the temperature difference between black and white daisy populations increases. Negative values for q, as well as being thermodynamically implausible, destroy homeostasis. In a spatial model, it is more sensible to model the diffusion of heat through Daisyworld explicitly, as opposed to this much more simple parametrisation. Indeed this has been done for one and two dimensional Daisyworlds where increasing diffusivity is shown to be analogous to decreasing the insulation q as expected (Adams et al., 2003; von Bloh et al., 1997; Ackland et al., 2003).

As luminosity increases, it is the proportional coverage of black and white daisies that changes while global temperature remains relatively stable. There are numerous ways to understand this behaviour, perhaps the most common being reign-control homeostasis, a term coined by Clynes (1969) which refers to asymmetrical model elements which combine to form dynamics which are in some sense symmetrical; the reigns which control the path of a horse individually guide the animal left or right, though their combination provides control in one dimension. The biotic elements of Daisyworld are asymmetrical

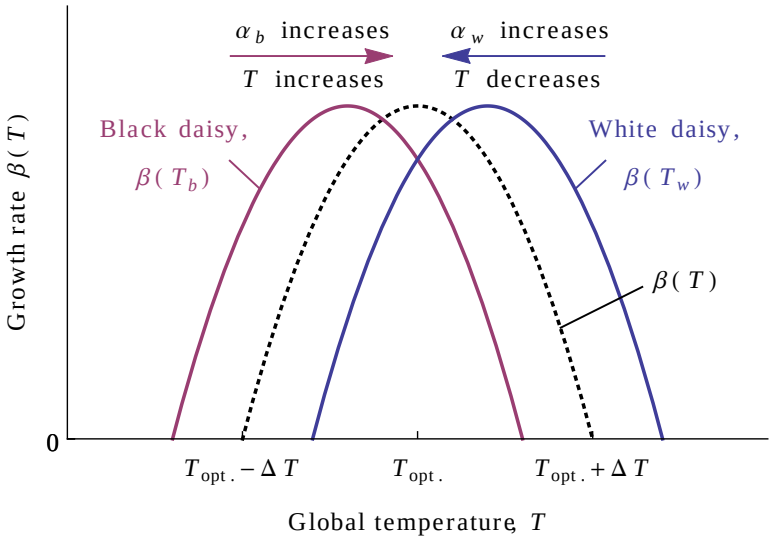

FIGURE 2.2: A reign control mechanism is established when the biotic elements with an increasing influence on their environment out-compete those with a decreasing effect at low temperatures, and vice versa (McDonald-Gibson et al., 2008). A point of contention, addressed in Chapter 3, is that randomly parametrised daisy species would not necessarily be seeded in this configuration. In the two dimensional parameter space, each daisy type may have a positive or negative influence on the planetary albedo, and a positive or negative effect on its temperature. If these traits were assigned at random, this negative feedback configuration therefore would only arise with probability $\frac{1}{4}$.

in that they only exert either an increasing or decreasing influence on the planetary temperature, while their combination leads to homeostasis. In Daisyworld, while daisies share the growth function β, their local temperatures differentiate them such that when the global temperature is below $T_{\text{opt.}}$, black daisies which have an increasing influence on the global temperature out-compete white daisies, and vice versa when the global temperature is above optimal for daisy growth. This effect is illustrated for the parabolic growth function in Fig. 2.2. This mechanism can also be seen in the absence of inter-species competition, such as Harvey (2004) which separates the daisy types into distinct 'beds' between which there is a heat flux, and therefore temperature gradient which differentiates the species sufficiently for the reign control mechanism to exist.

2.3 Separation of time scales

Assumptions of separated time scales, where certain processes occur very quickly in relation to other slow processes, are common in this type of modelling effort, if only because time scales are often unknown or obfuscated by very much more complicated and unresolved processes. Firstly, we demonstrate the utility of separating time scales in Watson and Lovelock's (1983) Daisyworld by making this assumption explicit and showing a range of simplifications to the model analysis which follow. The model incorporates four distinct time scales.

τ_β Daisy establishment

τ_H Mean temperature change

τ_γ Daisy removal

τ_L Insolation change

Watson and Lovelock (1983) assume zero heat capacity in Daisyworld. For a fixed luminosity, any change in the albedo of the planet produced an instantaneous change in temperature. For greater clarity, we refer to this convention as being in the limit $\tau_H \ll \tau_{\gamma,L}$. Changes in global temperature happen faster than any other element. Consequently, changes in the albedo due to γ and changes in insolation L occur sufficiently slowly that the model is always in a state of radiative equilibrium, and the global mean temperature, T, can be defined as the steady state temperature. The exact form of heating is commonly found to be unimportant to the fixed point dynamics of the model (Dyke, 2008). As such, we replace Eq. (2.1) with the linearised Steffan-Boltzmann equation to assist analysis

$$C\frac{\mathrm{d}T}{\mathrm{d}t} = L(1-A) - \sigma T_{\mathrm{eq}}^3 T \tag{2.7}$$

where T_{eq} is the steady state temperature of a bare planet ($x = 1$). For simplicity, we replace T, C and L with their dimensionless counterparts, normalised by the factor σT_{eq}^3

$$T = L(1-A). \tag{2.8}$$

At the other end of the spectrum, Watson & Lovelock assumed that changes to insolation occur so slowly that it may remain fixed while all other variables are numerically integrated to steady state. Therefore, $\tau_L \gg \tau_{\gamma,H}$. Now as well as reaching radiative equilibrium, the daisy turnover processes are able to achieve a steady state between adjustments in L. In this section we determine the nature of the steady state with these assumptions.

On Earth, the separation between abiotic forcing time scale τ_L and abiotic relaxation τ_H is reasonable. While there are a number of processes which may contribute to this abiotic relaxation, operating over a range of time scales, they are rapid compared to changes in external forcing. Due to its mass, ocean temperature lags behind changes in external forcing by $\sim 10^3$ years (Battisti and Hirst, 1989), while the response of the cryosphere range from decadal variance in glacier extent to millennial changes in ice sheets (Goodison et al., 1999). In contrast, abiotic forcing processes which influence the surface temperature can be shown to occur on the very longest timescales; fluctuations in the luminosity of the sun are sufficiently small to be neglected except over $\sim 10^9$ years, and the effects of secular cooling of the Earth is only felt over comparable time scales (Anderson, 1989).

On the other hand, the time scales associated with changes in the biota are less clear, and the extent to which the separation of time scales described here can be justified

is a limitation of the Daisyworld model. The response of the biosphere to changes in its abiotic environment range over ecological time, from tens to hundreds of years for different stages of ecological succession to evolutionary processes that can occur over geological time scales. As well as the distribution and abundance of species changing in response to environmental change, the frequency of alleles in biological populations can also change and so evolution can be an important process. The original Daisyworld featured life forms with fixed traits and no capacity for evolution. Subsequent studies relaxed this assumption and explored evolutionary dynamics with the conclusion that the behaviour of the model and its homeostatic behaviour is largely preserved (Lovelock, 1992; Lenton, 1998). For this section, our analysis is limited to the original Daisyworld model and so we do not include adaptation in our model.

2.4 Fixed points and stationary behaviour

With the exception of linearising the Boltzmann equation earlier, the following analysis proceeds along similar lines to Saunders (1994). Inspecting Eqs. (2.5a) and (2.5b) reveals the condition for a fixed point to be where both differential equations return zero

$$\alpha_b\big(x\beta(T_b) - \gamma\big) = 0, \tag{2.9a}$$

$$\alpha_w\big(x\beta(T_w) - \gamma\big) = 0 \tag{2.9b}$$

which leads to a number of trivial fixed points such as both species b and w extinct, $\alpha_b = \alpha_w = 0$ or a single extinction such as $\alpha_b = 0$ and $\alpha_w = 1 - \frac{\gamma}{\beta(T)}$. Initially, we are interested in solutions with $\alpha_b > 0$ and $\alpha_w > 0$ which clearly only occur for $\beta(T_b) = \beta(T_w)$. Here, our exact choice of β becomes relevant, although we may make some generalisations. If β is chosen to be a symmetric function with a single maximum at $T = T_{\text{opt.}}$, or q is sufficiently small that it can be approximated parabolic in the vicinity of $T_{\text{opt.}}$, we can write this condition as

$$T_{\text{opt.}} - T_w = T_b - T_{\text{opt.}}. \tag{2.10}$$

At this point, we replace the albedos A_b, A_e and A_w by their commonly chosen (and largely arbitrary) values of $\frac{1}{4}, \frac{1}{2}$ and $\frac{3}{4}$, and specify $T_{\text{opt.}} = \frac{1}{2}$. Rearranging Eq. (2.10) for $\alpha_{b,w}$ and substituting into the conditions for a fixed point, given by Eqs. (2.9a) and (2.9b) gives the black and white daisy populations at this fixed point as

$$\alpha_{b,w} = \frac{1}{2} - \frac{L-1}{L-q} - \frac{\gamma}{2\beta(T_{b,w})}. \tag{2.11}$$

This fixed point is of interest as it corresponds to a temperature $T \approx T_{\text{opt.}}$. As we will show, this model shares the essential features of previous Daisyworld models in which

this is an attractive fixed point. Consequently the model will self-regulate in that the temperature is maintained approximately constant over a range of insolation L.

Watson & Lovelock employ a fixed death rate which produces a steady turnover of daisies. We further simplify the model by assuming this turnover to occur on sufficiently long timescales compared to growth such that no significant fraction of bare ground exists at any time. This corresponds to the the limit of rapid daisy colonisation compared to removal, $\tau_\beta \ll \tau_\gamma$ or $\frac{\gamma}{\beta(T_b)} \to 0$ and ensures $x \to 0$. The model is then reduced to a two state model in which the surface of the planet is covered by a mixture of only black and white daisies with no bare ground. This limit is only sensible providing $\beta(T_{b,w}) > 0$ (which is surely the case where $T \approx T_{\text{opt.}}$). It is important to note that in doing so, we constrain our analysis away from some population effects which hinge on the relationship between establishment and removal. However, this reduction will be found to facilitate very simple, yet highly general analysis of a number of other effects. We have now identified and separated all four relevant timescales according to

$$\begin{gathered} \tau_{H,\beta} \ll \tau_\gamma \ll \tau_L \\ \text{fastest} \longrightarrow \text{slowest} \end{gathered}$$

The basin of attraction for this point must assume $0 < \alpha_b < 1$ which we denote with:

$$L^- \leq L \leq L^+, \tag{2.12a}$$

$$L^- = \frac{2+q}{3}, \tag{2.12b}$$

$$L^+ = 2 - q. \tag{2.12c}$$

Having determined the black daisy population as a function of L in Eq. (2.11), this result can be substituted into the definition of T in Eq. (2.8) along with the normalisation condition to yield the global mean temperature

$$T = \begin{cases} \frac{3}{4}L & L \leq L^- \\ \frac{L(1-q)}{2(L-q)} & L^- \leq L \leq L^+ \\ \frac{1}{4}L & L \geq L^+ \end{cases}. \tag{2.13}$$

The change in global temperature for a range of luminosity is shown in Fig. 2.3. Note how the decrease in temperature between L^- and L^+ is parametrised by q. The limit $\tau_\beta \ll \tau_\gamma$ is useful in providing a transparent analysis, but requires a different treatment if the function chosen for $\beta(T)$ is allowed to be zero. Indeed, the original Daisyworld model uses a parabola, cut off at the axis to ensure $\beta \geq 0$. For greater generality, let $\beta(T)$ be a function which is positive in the range $T^- < T < T^+$ but zero everywhere else. Following Ashby (1960) we call this the essential range, the range of environmental conditions that are essential for life. Outside this range, we have a lifeless planet with $T = A_e L$, while within the allowed temperature range, the model behaves exactly as

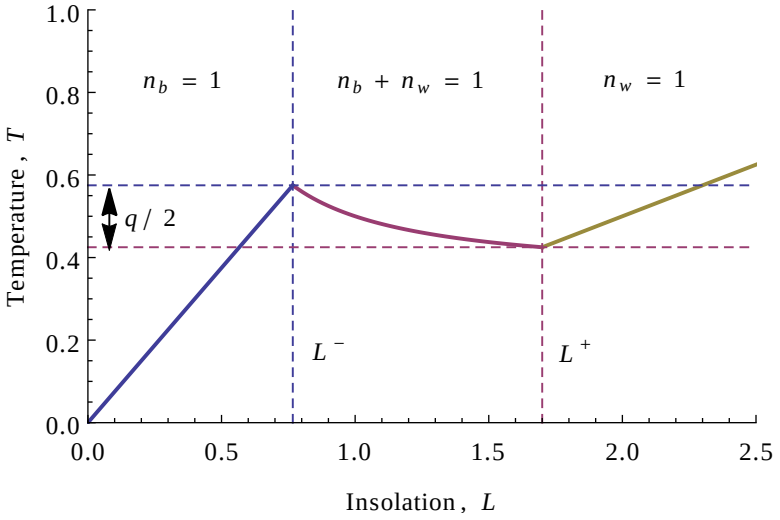

FIGURE 2.3: The geometry of self-regulation in Daisyworld for a growth function which always satisfies $\beta(T) > 0$, allowing the separation of timescales $\frac{\gamma}{\beta(T)} \to 0$. The onset, and width of the regulating region is determined by the limits of the effect of Daisies on the global temperature, L^- and L^+, given in Eqs. (2.12b) and (2.12c). Beyond these, no further changes can occur to planetary albedo. Outside the regulating range, $\beta(T_b) \neq \beta(T_w)$, leading only to trivial fixed points corresponding to extinction of one daisy type.

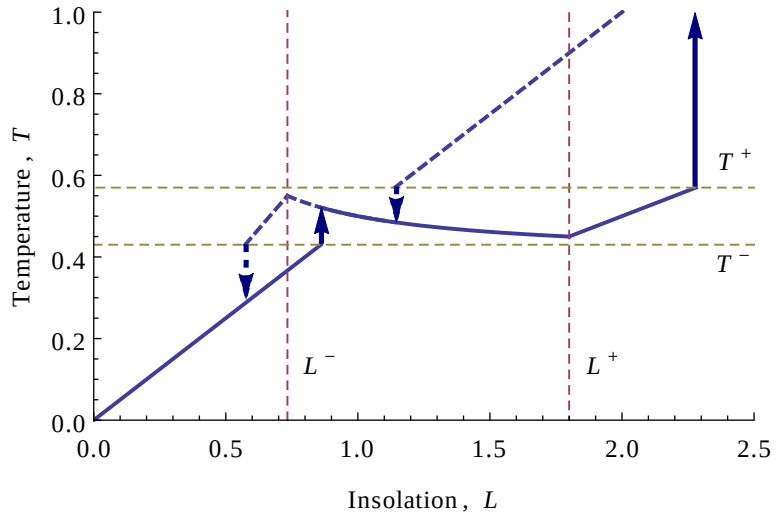

FIGURE 2.4: Global temperature, T, with increasing (solid) and decreasing (dashed) insolation, L. In contrast with Fig. 2.3, the growth function $\beta(T)$ is only non-zero in the range $T^- < T < T^+$. self-regulation can only emerge in this range, resulting in hysteresis. Once in a regulating state, the system is subject to the limits of regulation as before.

discussed. This leads to some interesting consequences, such as hysteresis loops shown in Fig. 2.4.

Next, we consider the robustness of the salient features of the model to departures from the simplest case of extreme separation, where the time scales above are sufficiently separated such that $\frac{\tau_H}{\tau_L}, \frac{\tau_\gamma}{\tau_H}, \frac{\tau_\beta}{\tau_\gamma} \to 0$.

2.5 Forcing time scales

We begin by challenging the model by comparing the time scales of changes in the biota, τ_γ and changes in the state of the environment, τ_H. So far we have assumed changes it insolation to occur over the very longest time scale in the system. i.e. τ_L is so large that insolation is fixed while daisy coverage and temperature change towards their stead states values. This is the same assumption as in Watson & Lovelock's original model and many subsequent studies. In this section we relax this assumption, represented schematically below.

$$\tau_{H,\beta} \ll \tau_\gamma \approx \tau_L$$
$$\text{fastest} \longrightarrow \text{slowest}$$

$\tau_{H,\beta}$ are the timescales for rates of heating and daisy establishment respectively. It is assumed that these occur so fast as to be instantaneous processes when compared to τ_γ, the rate of change of daisy removal and τ_L, the rate of change of insolation. While previously τ_γ and τ_L were taken to be separated, we explore the model's behaviour as the external driving on the system increases. We do this by keeping τ_H and τ_β fixed as the fastest timescales (they are sufficiently small that the model exists in a state of constant radiative equilibrium, and colonisation of bare ground is rapid) and evaluate the model's stability as the separation between τ_γ and τ_L decreases. We have already introduced changes in insolation and secular cooling as examples of driving forces, and observed that they occur only on the very longest timescales. However, it is simple to imagine planets where this is not the case, for example where the orbital period is sufficiently long and elliptical.

The first step in reducing the problem is to assume that the time evolution of the daisy coverage can be represented by a linear relaxation (or equivalently a Newtonian relaxation) towards the insolation dependant fixed point value, $\alpha^*_{b,w}(L)$, found with Eq. (2.11). This allows us to avoid making assumptions about the form of β and in particular sidestep issues arising from non-linearity in the daisies' response to temperature. We then introduce the daisy removal time scale, τ_γ, in order to examine how the evolution towards the fixed point changes as τ_γ increases and so daisy removal becomes slower with respect to the change in insolation.

$$\frac{\mathrm{d}\alpha_b}{\mathrm{d}t} = \frac{1}{\tau_\gamma}\big(\alpha^*_b(L) - \alpha_b\big), \tag{2.14a}$$
$$\frac{\mathrm{d}\alpha_w}{\mathrm{d}t} = -\frac{\mathrm{d}\alpha_b}{\mathrm{d}t}. \tag{2.14b}$$

As shown in Fig. 2.5, the behaviour of this linearised system is barely distinguishable from the solutions given in Eqs. (2.5a) and (2.5b). We are only able to formalise the

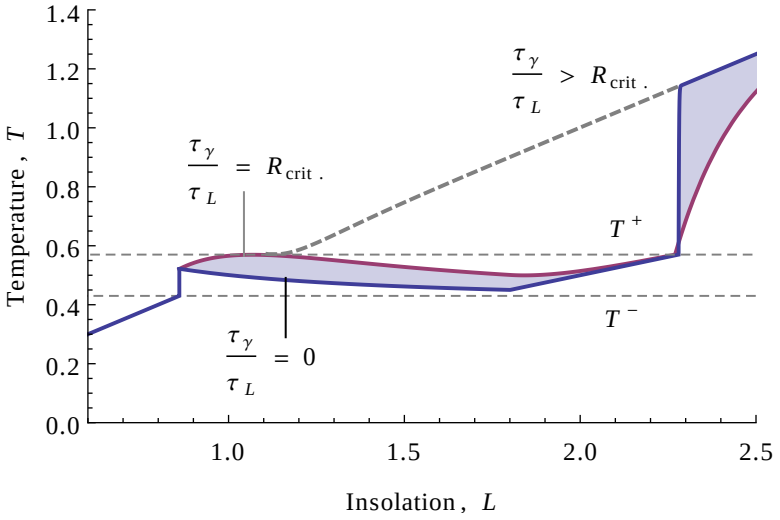

FIGURE 2.5: The onset of regulation is abrupt, and occurs at a critical value of $\frac{\tau_\gamma}{\tau_L}$. Above this limit, no regulation may occur while below this limit, regulating solutions are found in the shaded region, down to the limit where τ_γ and τ_L are fully separated. Dashed lines indicate the region where $\beta(T) > 0$, outside which only fixed points corresponding to extinction exist.

evolution of this system in terms of exponential integral functions which are not particularly meaningful to a reader, and no analytic solution exists. However, we find that the onset of homoeostasis is abrupt and occurs at the critical point $\frac{\tau_\gamma}{\tau_L} = R_{\text{crit}}$, where R_{crit} is some function of q, and the essential range. This threshold confines the system to either extinction, or self-regulation.

With a fixed and very fast establishment rate, increasing τ_γ slows down the daisies' responses to changing insolation. This is equivalent to insolation continuing to increase while the daisies are moving towards a fixed point. The rate of change of insolation can be so great as for the system to be driven out of the essential range without being able to oppose the change. The system is most sensitive to changes in τ_L ($R_{\text{crit}} \to 0$) at the limits of regulation where fixed point temperatures are close to the maximum of the growth function. As white daisy coverage approaches the maximum, there is less 'head room' to accommodate faster changes in insolation and so any increase away from the very longest of timescales for τ_L results in the collapse of homeostasis and extinction.

Adjustment of the timescales of certain processes, and its effects on the Daisyworld self-regulating mechanism has been considered implicitly in a number of previous studies. Luminosity timescales were relaxed in Dyke et al. (2007) and McDonald-Gibson et al. (2008) who developed zero-dimensional agent based versions of Daisyworld in which the population of daisies change during changes to external perturbations. self-regulation was observed under such conditions. Our analysis has shown that while it is not necessary for luminosity to remain fixed as the system evolves to a steady state, relaxing this assumption will inevitably reduce the region of luminosity over which the system will be stable. As Gregorio et al. (1992) and Wood et al. (2006) previously noted, the size of the

basin of attractor will change as luminosity changes. Such sensitivity is most pronounced at the limits of self-regulation, regions that are characterised by large hysteresis loops.

2.6 Model heat capacity

Next we compare the biotic time scales to those of external forcing, τ_L. In particular we examine how the model responds to sudden perturbations or shocks. A shock can be thought of as a fluctuation, driving a system away from it's current state on timescales much shorter than any dissipation can occur. Such fluctuations may originate within, or outside the system under consideration, and are not necessarily small. Planet Earth has been subjected to numerous such events of wildly differing magnitudes; events such as comet impacts, or changes in atmospheric composition by volcanic out-gassing or anthropogenic emission of greenhouse gasses occur on dramatically shorter timescales than any dissipation, resulting in an instantaneous perturbation to abiotic factors (IPCC, 2007). Similarly, life may be subject to similar perturbations, whereby the composition of the biota undergoes rapid change; forest fires and anthropogenic land use change for example occur on shorter timescales than reorganisation of the biosphere may otherwise occur.

In this section, we examine the role of timescales τ_H and τ_γ in dissipating such shocks, while assuming τ_L is sufficiently large that L may be held constant, and τ_β is small such that colonisation is rapid, as before. This can be represented schematically by evaluating the model's behaviour as the separation between τ_H and τ_γ changes.

$$\tau_\beta \ll \tau_H \approx \tau_\gamma \ll \tau_L$$
$$\text{fastest} \longrightarrow \text{slowest}$$

As in the previous section, we approximate the time evolution of $\alpha_{b,w}$ by a linear approach to the respective temperature-dependant fixed point values, but now we perturb the model temperature by a temperature fluctuation of size ΔT away from it's steady state value. By further approximating the parameter q to be small, and departures of T from it's fixed point value, T^* given by Eq. (2.13) to be small, we find the time evolution of α and T can be found in closed form, although the exact result is somewhat unwieldy. Fortunately we can gain useful insights into the behaviour of the model simply by studying the leading exponential

$$T(t) \approx T^* + \Delta T \exp\left(-\frac{t}{\tau_{\text{damp.}}}\right) \tag{2.15}$$

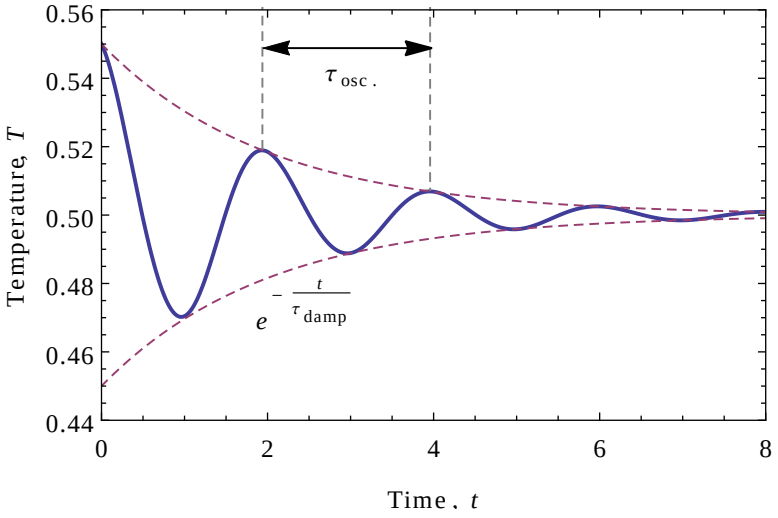

FIGURE 2.6: For sufficiently small $\frac{\tau_\gamma}{\tau_H}$, oscillations occur along with the exponential damping. Here they are shown for $\tau_H = 1$, $\tau_\gamma = 0.1$ for the case of $L = 1$.

where $\tau_{\text{damp.}}$ is the characteristic decay time of fluctuations in T, given by

$$\tau_{\text{damp.}} = \frac{2\tau_H}{1 + \sqrt{1 - 2L\frac{\tau_H}{\tau_\gamma}}}. \tag{2.16}$$

which may also be extracted very easily by way of linear-stability analysis about the known fixed points. Eq. (2.15) results in an asymptotic approach to the fixed point in T. This expression gives us two important insights; firstly, we have the exponent for exponential decay of fluctuations in T. The rate of decay of fluctuations is maximised by our original separation of timescales, where heating occurs on the shortest timescales, and $\tau_\gamma \gg \tau_H$. In this case, we have $\tau_{\text{damp.}} = \tau_H$ and the system relaxes back to the steady state value as quickly as possible. In the opposite limiting case of slow heating, the decay rate is minimised, and $\tau_{\text{damp.}} = 2\tau_H$.

Our second observation is that the sum of terms under the square root may be negative, resulting in a complex exponent and therefore oscillations. We can determine the limit of timescales which marks the onset of oscillations in T to be

$$\frac{\tau_\gamma}{\tau_H} < 2L. \tag{2.17}$$

Beyond this, damped oscillations occur with period $\tau_{\text{osc.}}$

$$\tau_{\text{osc.}} = \frac{4\pi\tau_H}{\sqrt{2L\frac{\tau_H}{\tau_\gamma} - 1}}. \tag{2.18}$$

This is illustrated by Fig. 2.6. This formulation of oscillations is consistent with Nevison et al.'s (1999) observation that oscillation period increases roughly linearly with the heating time scale, or equivalently the model heat capacity.

2.7 Maximum seed diffusion

Zero-dimensional Daisyworld assumes that diffusion of seeds is sufficiently rapid that the species are well mixed and only the abundance of a daisy type limits it's rate of growth. However, the biosphere is not well mixed, and local seeding clearly plays a significant role in the distribution of the biota. Ecological succession of newly habitable, or recently vacated ground occurs across a broad range of timescales, from r-selected pioneering grasses over decades, to K-selected plants and trees (MacArthur and Wilson, 1967). Introducing spatial factors such that daisies can only reproduce by seeding into neighbouring empty bare areas seeding begins to relax this assumption. The first spatially explicit Daisyworld was formulated by von Bloh et al. (1997) in which the original zero-dimensional model was translated onto a two-dimensional lattice in which each square or cell could be in an empty, black daisy or white daisy state. The update rules for each cell in this cellular automaton are probabilistic; an occupied cell will turn into an empty cell, with probability $\gamma\Delta t$, while occupied cells may seed empty cells in their seeding neighbourhood with probability $\beta(T)\Delta t$. In this cellular automata Daisyworld the diffusion of seeds from daisies to bare ground for establishment was limited to local interactions; a bare cell would be seeded by one of its occupied immediate neighbours.

Heat diffusion was modelled explicitly with the relevant transport equations, allowing the non-uniform temperature field to emerge, parametrised by the diffusion coefficient rather than the choice of q. von Bloh et al. (1997) make some progress in reconciling this model with the original Daisyworld model; a Taylor expansion is used to illustrate that the discretised Daisyworld shares key behaviour with the zero dimensional model under the homogeneous (or mean field) approximation. We demonstrate that the behaviour of our simplified model can be recovered from a two-state cellular automata; cells occupy either state b or w having albedo A_b or A_w respectively. To advance the cellular automata through time, we randomly select a single site where the daisy will die, and be replaced. The probability is uniform across all daisies, similar to the constant death rate γ used previously. The replacement rule we implement can be used to recover Eq. (2.11) under the mean field approximation is to update site i according to

$$P(b)_i = \frac{n_{b,i}\beta(T_b)}{n_{b,i}\beta(T_b) + n_{w,i}\beta(T_w)}, \tag{2.19a}$$

$$P(w)_i = 1 - P(b)_i \tag{2.19b}$$

where $P(b)_i$ and $P(w)_i$ are the local probabilities of a site being replaced by a black or white daisy respectively, and $n_{b,i}$ and $n_{w,i}$ are the numbers of black and white daisies in the neighbourhood of i, typically chosen to comprise only the four nearest adjacent cells. The separation of τ_γ and τ_β is satisfied as removal occurs sufficiently slowly that only one site is ever removed in an iteration, and is always occupied before the next removal. This cellular automata is implemented on a regular square lattice of 100×100

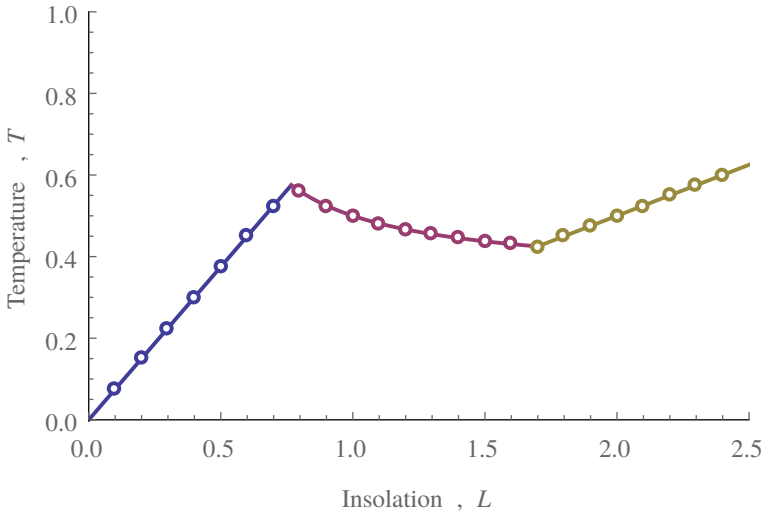

FIGURE 2.7: The results of simulating a 100×100 cellular automata Daisyworld using the rules given by Eq. (2.19b). Each data point is given 10^6 update steps to equilibrate, before taking an average temperature over the following 10^6 update steps. Temperature variance is at most the size of plot points, and indicates identical behaviour to the zero-dimensional Daisyworld equations which can be recovered from the cellular automata by mean field theory.

daisies and subjected to 10^6 replacement steps from an initially random field of black and white daisies. The steady state temperature of the cellular automaton is shown for a range of L in Fig. 2.7 along with the spatial distribution of daisy types at $L = 1$ shown in Fig. 2.8.

It is simple to show that regulation emerges from this cellular automaton under the mean field approximation by first approximating the neighbourhood of individual sites to be that of the global average; rather than local effects, sites are influenced only be a mean field (the classic example of an application of mean field theory is to the Ising model, explored in detail by (Stanley, 1987, Ch. 3.6)). In this instance we set the local populations, n, to the global mean populations, α, multiplied by the number of neighbouring lattice sites, z. Steady states then exist where we have no net change in the populations of black and white daisies, a condition which is expressed by

$$P(b)z\alpha_w = P(w)z\alpha_b. \tag{2.20}$$

Substituting in the normalisation condition gives

$$\alpha_b(1-\alpha_b)\beta(T_b) = \alpha_b(1-\alpha_b)\beta(T_w) \tag{2.21}$$

which only has non-trivial roots for $\beta(T_b) = \beta(T_w)$, a familiar observation made earlier in Eqs. (2.9a) and (2.9a), and can be solved under the same assumptions as previously. Lenton and van Oijen (2002) implement such a cellular automaton with highly non-linear versions of Eqs. (2.19a) and (2.19b). This alternative formulation does not yield regulating steady states of this nature, highlighting the importance of our choice of updating

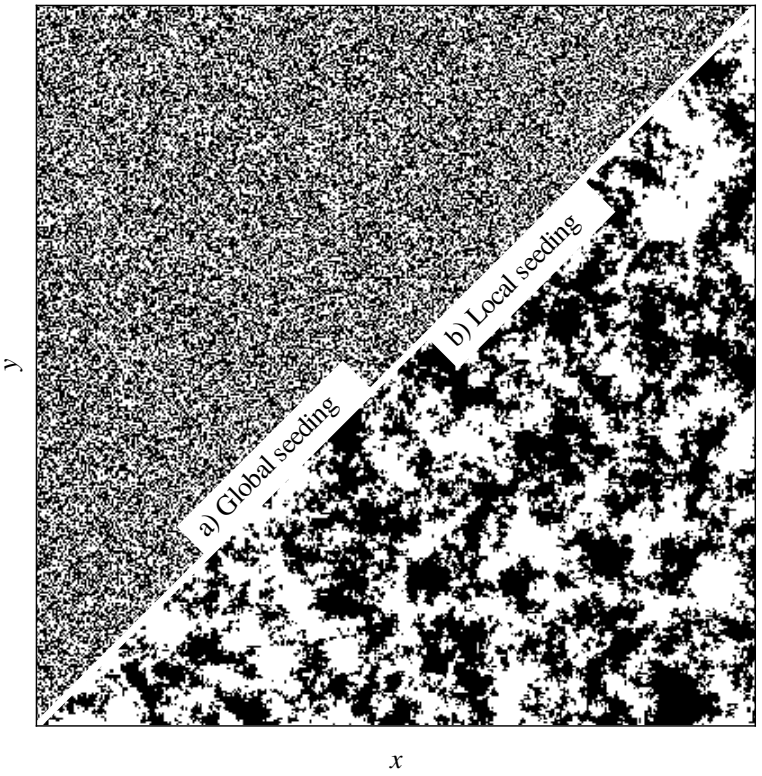

FIGURE 2.8: Comparison of the mean field picture where spatial correlations are minimised (a), and a simulated local seeding where strong spatial correlations are evident (b). Black and white daisies are regions represent the two daisy types and x and y are spatial dimensions. While a mean field approximation exactly predicts the stationary behaviour, it would fail in predicting spatial correlations.

rules. As predicted by the mean field equations, the spatially embedded Daisyworld shares identical stationary behaviour to Watson and Lovelock's (1983) zero-dimensional model.

From Fig. 2.8 we have seen that relaxing the assumption of maximum seed diffusion, by implementing a cellular automata with local replacement rules, does introduce spatial variations into Daisyworld. Mean fields often capture the large scale stationary behaviour of spatially embedded systems, though fail to anticipate finer details where spatial correlations are significant to the large scale behaviour. In this instance, Fig. 2.7 shows we correctly predict the homeostatic behaviour despite the emergence of seemingly long-range clustering of the daisy types produced by a cellular automata. The implication of this result is that either the long-range correlations reported by Wood et al. (2006) which would break the mean-field assumptions do not exist, or that while long-ranged correlations exist, they are irrelevant to the ability of the model to exhibit homeostasis. As Fig. 2.8 shows our mean field approximations do not capture the spatial correlations, which may be important for observing some strong spatial effects, such as desert formation, where the local temperature departs the essential range (Ackland et al., 2003). This point may only be addressed by the development of an exact solution

to the spatially embedded Daisyworld, incorporating the possibility of spatial correlations. Doing so may shed light on the role of local seeding and the resulting correlations in dissipative systems, such as the latent heat exchange field investigated by Baldocchi et al. (2005).

2.8 Discussion

Our analysis of a simplified version of the original Daisyworld model has demonstrated that some of the original assumptions relating to timescales can be relaxed and self-regulation is still observed. From Watson and Lovelock's (1983) original mathematical formulation of the zero-dimensional Daisyworld model, we have identified the four competing timescales. They are associated with changes in external forcing or insolation, equilibration of global temperature, and daisy turnover that includes establishment and removal. By making explicit a separation of timescales between external forcing, and internal processes, and further assuming establishment to occur much more rapidly than removal, the long time behaviour of the model can be expressed very simply. In this reduced model, it is simple to see that self-regulation may emerge, as well as to identify the limits over which self-regulation can occur. Indeed, the essence of the original model remains entirely intact. Of course, in applying this significant reduction, we constrain ourselves from investigating a number of population effects which may be introduced by slower rates of population establishment. However, this simplified model gave us keen insights as we depart from other simplifying assumptions. First, we relaxed the assumption of a separation between the external forcing and daisy turnover timescales. This is equivalent to increasing the rate of change of luminosity. Regulation is found to occur providing these timescales are related by a critical ratio, beyond which the self-organisation of daisy coverage is slow, and the population is driven to extinction before significant regulation occurs.

Next, the assumption that the global temperature equilibrates on much shorter timescales than daisy turnover was relaxed. This is equivalent to introducing heat capacity such that global temperatures evolve towards a steady state in response to changes in radiative forcing. Despite this generalisation, we found a closed form expression for the time evolution of global temperature when perturbed from its steady state value. As expected, the perturbation decays exponentially, though surprisingly, a more rapid response by the daisy population enables fluctuations to persist for significantly longer. At the threshold where decay time is exactly doubled, we find the onset of temperature oscillations.

Finally, we addressed Watson and Lovelock's (1983) assumption that there are no spatial effects in the daisy seeding process by introducing local seeding in a cellular automata version of the model. Through appropriate choice of probabilistic updating rules, we

can recover the exact same fixed points as in the previous case through a mean field approximation. While local seeding clearly results in very strong correlations which are not accounted for in a mean field approach, near perfect agreement with simulation indicates such correlations have no impact on the regulatory fixed points. Given that mean field approximations typically perform better with increasing dimensions, these results along with the rest of our analysis of time scales are promising in that they suggest how important dimensions of interaction of the very complex Earth system can be captured in relatively simple models that will be mathematically reducible and so able to provide important insights into the real world.

The Daisyworld model remains open to criticism. While black and white daisies organise into homeostatic configurations in this model, it is much simpler to design models where there is no self-regulation such as two daisy species with a greater albedo than a bare planet, or black daisies with a higher optimal temperature than white daisies. In this sense homeostasis has emerged by design and while Daisyworld variants exist with several daisy types, they are constrained by an inverse relationship between daisy absorptivity and optimal temperature. What reason is there to expect such configurations to occur on Earth, as opposed to their more numerous unstable counterparts? In Chapter 3 we explore the minimal ingredients for the emergence of homeostasis in model ecosystems, particularly the emergence of antagonistic species and the reign control mechanism.

Chapter 3

Generalisation to multidimensional environments

3.1 Introduction

The principle that environmental factors affect life is evident throughout the biosphere both regionally and globally, and throughout Earth's history (Gaston, 2000). Variables such as temperature and soil or atmospheric composition determine whether an organism can proliferate, and populations will develop to some resource limited level. Climatic shifts between, for example, greenhouse and ice house states are accompanied by mass extinction (Haywood, 2004). Indeed this principle is at the core of environmental niche modelling, where the distribution of species in the space of significant environmental factors, their realized-niche, is used to predict the spatial species distribution (Thomas et al., 2004). The influence of life on its environment however, and whether it is expected to have a stabilising, or destabilising effect is less clear. Lovelock and Margulis's (1974) original Gaia Hypothesis focused on the extent to which the emergence of life has promoted a self-regulating, or homeostatic system. In the previous chapter, we explored the properties of Daisyworld, a model in which we are able to design a biosphere such that homeostasis emerges with no mechanism more complicated than natural selection. In this chapter, we consider the ability of a general coupled life-environment system to generate stabilising feedback loops in increasingly complex environments under a minimal set of assumptions, without the need for such careful design.

The Gaia hypothesis proposes that life may self-organise into complex, self-regulating systems with maintain the habitability of their environment and at the higher level contribute to regulation of variables globally conducive for life to flourish. Controversy followed as it is unclear how regulatory mechanisms might emerge without the need for system-wide cooperation which would contradict the principle of natural selection at the species level (Doolittle, 1981). The Gaia hypothesis is discussed in detail along with its variants and implications in Sec. 1.3. We aim to address two short-comings of this model in detail. Firstly, we have shown in Sec. 2.4 that with small reductions, the model is extremely simple to treat but while useful in its transparency, the behaviour of such a one-dimensional dynamical system is extremely limited. Hysteresis loops are an example of an iconic phenomenon, and such a dynamic is well recognised in one dimension. Cyclical and chaotic behaviour however, are not accessible to low dimensional dynamical systems. It is therefore unclear how these features are expressed by high-dimensional, much more complex systems. The trade-off to this is that while more complicated many-body systems benefit from a much richer zoo of emergent behaviours, it typically comes at the expense of transparency, exchanging generality for a more faithful representation of a specific system.

Along with increasingly complex environments, the relevance of the reign-control mechanism to very much larger populations of diverse biotic elements in unclear. Daisyworld contains exactly two species which behave antagonistically in this way; black daisies have

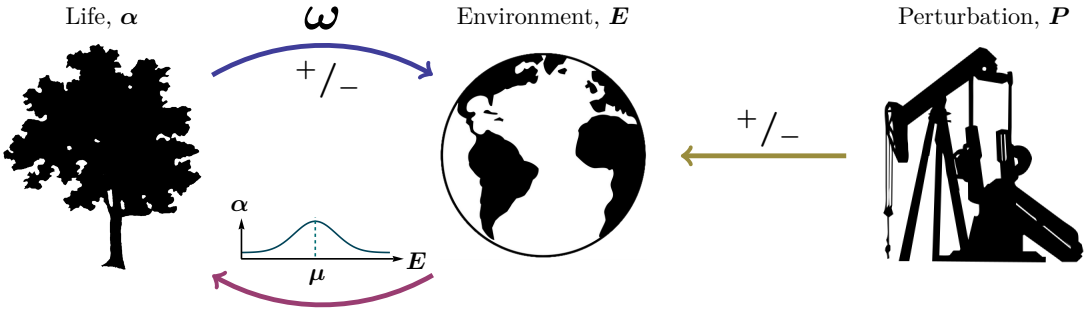

FIGURE 3.1: The abundance of a life elements, $\boldsymbol{\alpha}$, is some function of environmental conditions, $\boldsymbol{E}$, parametrised by $\boldsymbol{\mu}$. Environmental conditions are affected by the biota by some linear coupling $\boldsymbol{\omega}$ with additional contributions, $\boldsymbol{P}$, which are in some sense external to the life-environment system such as anthropogenic emissions. While concepts such as facilitation, ecosystems engineering and niche construction include assumptions that life can alter its environmental conditions in important ways, the notion that this system self-organises into a negative feedback loop is controversial.

a niche below the homeostatic fixed point and drive the environment in the positive direction, while white Daisies proliferate above the fixed point and drive the temperature in the negative direction (Harvey, 2004). However there is no reason to expect an ecosystem to organise in this way, and Daisyworld does not address the mechanism by which the pair of antagonistic species might emerge. While later work makes clear that the mechanism is not unique to a two-species model, a persistent theme of such models is that they are in some sense designed (Lenton, 1998).

Furthermore, the extent to which such a highly variable and complex system could be expected to enjoy any degree of stability at all remains an open and somewhat controversial question. Previous work in this direction typically revolves around this Daisyworld paradigm, encompassing a biota with limited variability and a simple environment consisting of few variables. The relevance of the reign-control mechanism to very much larger populations of diverse biotic elements in a multidimensional environment is unclear.

We begin by formulating the model, introducing our principle assumptions of life-environment feedback in Sec. 3.2. Some representative results are displayed by Sec. 3.3. High dimensional models require a novel computational method to produce whose implementation is given by Sec. 3.4. It turns out this interpretation is very useful in providing a thorough analysis which begins in Sec. 3.5. We first examine the density of fixed points with model parameters in Sec. 3.5.1 and further classify these into a range of fixed point types for a two dimensional model in Sec. 3.5.2. These results are discussed in detail in Sec. 3.6.

3.2 Model formulation

Watson and Lovelock's (1983) Daisyworld model describes a planet populated by two species of daisy; black and white daisies which are best adapted to low and high temperatures respectively. With an abundance of all other resources daisies need to reproduce, the planetary temperature is the limiting factor for their growth and is governed by radiative balance with the local star. While black daisies out-compete white at low temperatures, their colour means that the planet has a decreasing albedo as they proliferate, and therefore the steady-state temperature of the planet rises, and vis-versa for the white daisies. The reign-control state emerges as these antagonistic effects reach balance, and the planet becomes homeostatic in the sense that changes in solar luminosity are counteracted by adjustments in daisy populations, altering the planetary albedo and mitigating any increase in temperature one might expect. Chapter 2 describes and explores the model thoroughly, and one may see Wood et al. (2008) for an in depth review of numerous Daisyworld adaptations and extensions.

Dyke's (2010) Daisystat generalises Watson and Lovelock's (1983) model through the introduction of a large, diverse biota of K distinct, randomly parametrised elements. We make no assumption about the form the biota takes; be they individual organisms, species, ecotypes or populations, their abundance, metabolic activity or overall influence is denoted in the vector $\boldsymbol{\alpha}$ (Dyke, 2010). The biota lacks any form of self-interaction through processes such as competition or predation. While this is a component of the original Daisyworld model, where in the absence of environmental changes one type of daisy would go extinct in favour of the other even if it is marginally out competed, this is not a necessary ingredient to establish homeostatic behaviour.

These elements individually interact with an N-dimensional environment whose variables are denoted by the vector $\boldsymbol{E}$, in contrast to the single environmental dimension of Daisyworld. In essence, any coupling between biotic elements is mediated through the environment. These minimal ingredients, encapsulated in Fig. 3.1, have been shown to produce a range of interesting behaviours (Dyke and Weaver, 2013). Most importantly, the reign-control states found in Daisyworld emerge over a wide range of assumptions merely from the effects of a randomly assembled biota, without the need for parametrisation or tuning, a major criticism for the original Daisyworld mechanism. Notably this model lacks life-life and environment-environment feedbacks. Chemical weathering of silicate rocks is an example of a negative environmental feedback where increases in global temperature such as by increases in atmospheric CO_2 levels lead to increased silicate rock weathering, a process which captures CO_2. On the other hand, reduction of global ice and snow leads to increases in absorption of solar radiation and further increases in temperature, an example of positive feedback.

Our contribution to these fundamental questions is to identify the key components of Daisyworld and similar models which are responsible for the emergence of homeostasis,

and to circumvent the limitations of a low-dimensional environment, small population and need for parametrisation and prescription. Our proposed model is one of random couplings between two principle components, based on a minimal set of assumptions illustrated in Fig. 3.1. The model biota has a linear increasing or decreasing influence on the different dimensions of it's environment which in turn, has a non-linear influence on the distribution of species which make up the biota through a very simplistic niche idea.

Throughout this work we have used boldface notation to denote vectors and will reference individual elements with subscripted indices.

$$\boldsymbol{X} = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n. \end{bmatrix}$$

i) Environment affects life

The first component of Fig. 3.1 we discuss is the influence of the state of the environment $\boldsymbol{E}$ on the composition of the biota, represented by K distinct elements of vector $\boldsymbol{\alpha}$. We implement a very simple niche idea whereby each element of the biota has an optimal environment, a point in the space of environmental variables where its activity is maximised. The time evolution of each element of the biota is, in the simplest case, a linear relaxation towards some steady-state activity $\boldsymbol{\alpha}^*$ with characteristic time scale τ_α.

$$\tau_\alpha \frac{\mathrm{d}\boldsymbol{\alpha}(t)}{\mathrm{d}t} = \boldsymbol{\alpha}^*(\boldsymbol{E}, \boldsymbol{\mu}) - \boldsymbol{\alpha}(t) \tag{3.1}$$

where $\boldsymbol{\alpha}^*$ is the fixed point in the time evolution of $\boldsymbol{\alpha}$, dictated by the state of the environment and a niche parameter unique to the each biotic element $\boldsymbol{\mu}$, which determines the most advantageous environmental conditions for each biotic element. In the vicinity of $\boldsymbol{\mu}$, $\boldsymbol{\alpha}^*$ is maximised and decays by some function as the environment departs the niche. Niches are uniformly randomly distributed in the essential range of width σ_μ. For the purpose of this paper, the precise choice of $\boldsymbol{\alpha}^*$ can be shown to be arbitrary (Dyke and Weaver, 2013), and we choose a N-dimensional Gaussian for visualisation

$$\boldsymbol{\alpha}^*(\boldsymbol{E}, \boldsymbol{\mu}) = A \exp\left(-\frac{(\boldsymbol{\mu} - \boldsymbol{E})^\top (\boldsymbol{\mu} - \boldsymbol{E})}{2\sigma_E^2} \right) \tag{3.2}$$

where σ_E is the niche width, the characteristic width of the function and A parametrises the magnitude of the niche function at $\boldsymbol{E} = \boldsymbol{\mu}$, a redundant parameter which is set to unity for convenience. As the essential range becomes populated by a diverse biota, the mean activity increases with the size of the biota, K, while the deviations increase with

$\sqrt{K}$ such that in the limit of a very diverse biota of many unique elements, the total steady-state activity is approximately uniform over the essential range.

ii) Life affects environment

To be alive necessitates having some influence on the local environment through consumption or excretion of resources to maintain a metabolism although an organisms effect may extend in ways such as Watson and Lovelock's daisies which modify the planetary albedo through their colour. In the simplest case, the effects are linear and simply proportional to the activity of the biota, implemented by assigning each biotic element a unique influence on each aspect of its environment independently, stored in the matrix $\boldsymbol{\Omega}$ such that the summed effect on each environmental variable, $\boldsymbol{F}$ may be found simply from the matrix product

$$\boldsymbol{F}(t) = \boldsymbol{\Omega} \cdot \boldsymbol{\alpha}(t). \tag{3.3}$$

The effects in $\boldsymbol{\Omega}$ are chosen randomly with zero mean such that the model has no propensity for positive or negative feedback. Unlike the steady-state activity of a large population, the net effect of a population has zero mean and with appropriate scaling (discussed shortly) the variance in the surface is significant. The time evolution of the environmental variables is, in the simplest case, linearly driven by contributions from the biota $\boldsymbol{F}$ and external abiotic influences $\boldsymbol{P}$. That is to say, in the presence of even weak forcing we do not assume any fixed-point in $\boldsymbol{E}$ to exist in contrast to Daisyworld, where the environment always reaches radiative equilibrium.

$$\tau_E \frac{\mathrm{d}\boldsymbol{E}(t)}{\mathrm{d}t} = \boldsymbol{F}(t) + \boldsymbol{P}(t). \tag{3.4}$$

The linear time evolution of $\boldsymbol{\alpha}$ ensures this system has a $2N$-dimensional phase space (Weaver and Dyke, 2013) though this can be reduced to N-dimensions if the time scales of processes associated with changes in the biota, $\boldsymbol{\alpha}$ and the environment, $\boldsymbol{E}$ may be assumed separated, such that changes to the ecology occur on very much shorter time scales than those to the environment, or $\tau_\alpha \ll \tau_E$. In this instance, the activity of the biota may be replaced by its steady-state value $\boldsymbol{\alpha}^*$ and the model may be reduced to a single equation

$$\frac{\mathrm{d}\boldsymbol{E}(t)}{\mathrm{d}t} = \boldsymbol{\Omega} \cdot \boldsymbol{\alpha}^*(\boldsymbol{E}, \boldsymbol{\mu}) + \boldsymbol{P}(t). \tag{3.5}$$

It is useful to ensure that important model characteristics are invariant with the arbitrary model parameters, those which are neither random nor fundamental model components. This enables direct comparison of model behaviour over a range of parameters which might otherwise be obfuscated. We rescale the effect of the biota by introducing the normalisation constant

$$(\boldsymbol{\Omega} \cdot \boldsymbol{\alpha}) \to C^{-1}(\boldsymbol{\Omega} \cdot \boldsymbol{\alpha}) \tag{3.6}$$

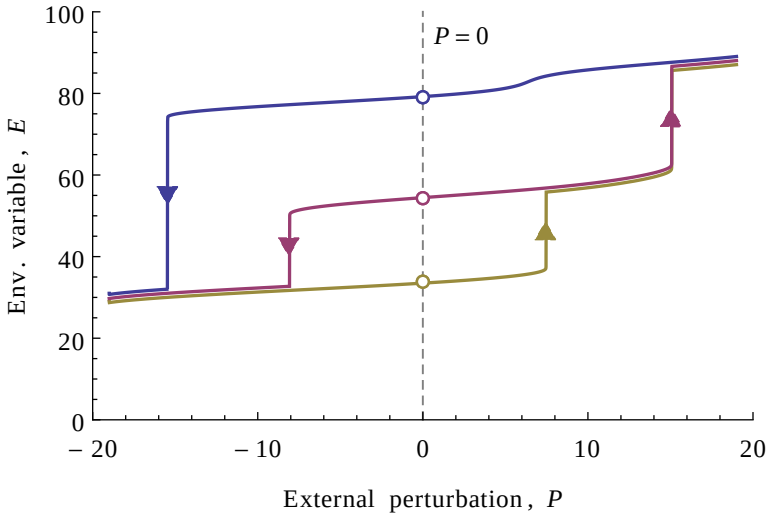

FIGURE 3.2: Hysteresis is caused by the existence of multiple homeostatic points for a given external perturbation (Dyke and Weaver, 2013). The yellow, red and blue stable points correspond to the coloured circles in the left figure. Recovery back to the red state after a transition to the blue state as a result of increasing $\boldsymbol{P}$ is only possible via a large decrease in $\boldsymbol{P}$ and a transition via the yellow state.

where we have used the constant C to normalize the total biotic effect, $\boldsymbol{F}(\boldsymbol{E})$ such that its variance is independent of our arbitrary choices; the distribution from which the random coupling matrix, $\boldsymbol{\Omega}$, is drawn, the fundamental niche width, σ_μ and the number of biotic components, K.

$$C^2 = 2\sqrt{\pi}K\sigma_E\sigma_\mu\sigma_\omega^2. \tag{3.7}$$

where σ_ω^2 is the variance of the distribution from which the random coupling matrix, $\boldsymbol{\Omega}$, is drawn.

3.3 Numerical Results

Despite it's very general formulation and random parametrisation, the model exhibits a diversity of homeostatic fixed points along with hysteresis in response to increasing perturbing force, illustrated for the case of a single environmental variable in Fig. 3.2. This is an iconic behaviour shared with Watson and Lovelocks Daisyworld model, except rather than homeostasis being designed, such points have emerged spontaneously from the biological complexity of the model and the simple feedbacks shown by Fig. 3.1. Fig. 3.3 shows how two species can be arranged in such a way that a single homeostatic fixed point emerges between their optimal temperatures. However, the diagram is much more interesting for a large number of species, shown in Fig. 3.3(b) which illustrates there is no need for design with a sufficiently diverse biota.

Along with a small, well designed biota, a criticism to the Daisyworld model is that it hinges on a single environmental variable, and therefore necessarily exhibits simple

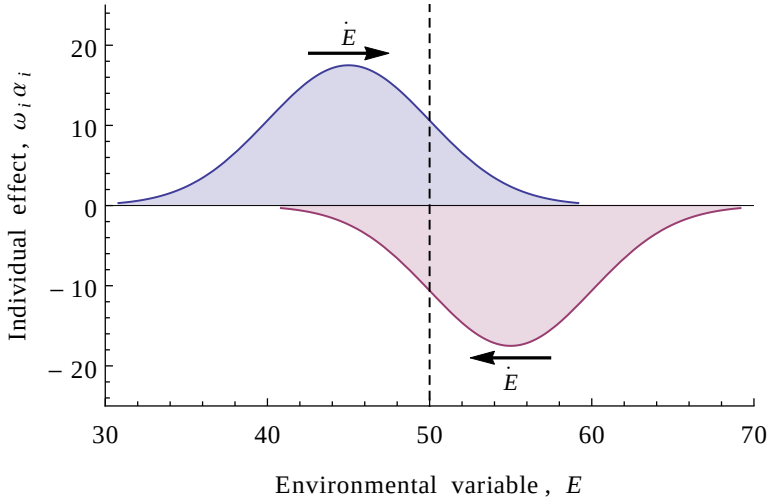

(a) Two species with increasing and decreasing influences on the environmental variable with niches at low and high values respectively.

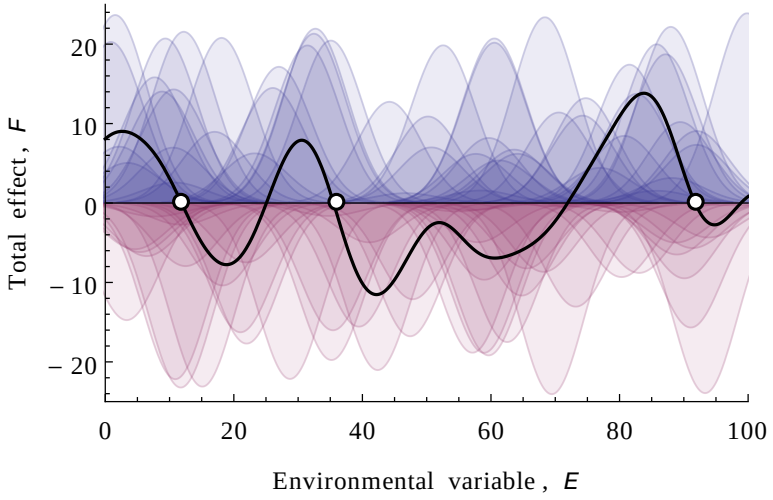

(b) One hundred species with randomly assigned niches and influences.

FIGURE 3.3: shows a rein control state. A biotic component that increases the environmental variable, E, counteracts the effects of a biotic component that decreases E. This results in E being regulated around values near the vertical dashed line. The probability of such a rein control pair being present in a population of two biotic components is $\frac{1}{4}$. As the number of biotic components is increased up to 100 in Fig. 3.3(b), a total effect F (solid black line) emerges as it is the sum of the individual biotic effects. Homeostatic stable points (denoted with circles) emerge whenever F undergoes a zero-crossing from left to right. These correspond to rein control homeostatic states (Dyke and Weaver, 2013).

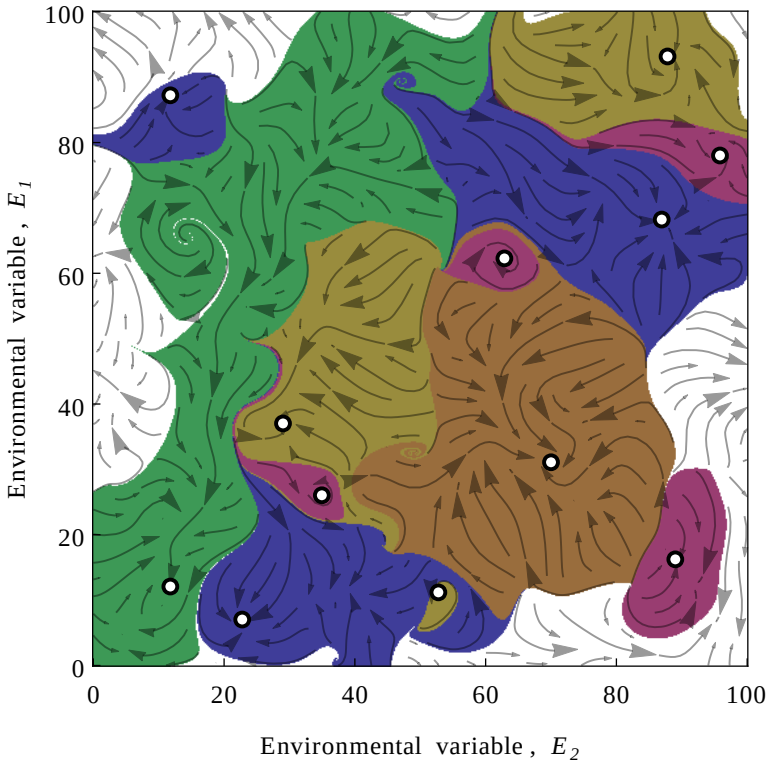

FIGURE 3.4: By a similar mechanism to that illustrated in Fig. 3.3(b), attractive and homeostatic fixed points emerge in environments of two or more dimensions. This figure shows the flow lines of the two environmental variables, $\boldsymbol{E}$, under the driving force of a random biota. Coloured basins of attraction show regions of initial conditions will follow stream lines until arrive at the same fixed points, indicated by open circles (Dyke and Weaver, 2013).

behaviour. On the other hand the Earth system could not possibly be reduced in this way (Kirchner, 2003). Initial intuitions may suggest that with increasing environmental variables the likelihood of finding a point stable in all dimensions simultaneously would vanish exponentially. However, transitions between neighbouring attractive points can be seen in higher dimensional systems. Fig. 3.4 directly examines the basins of attraction for a two dimensional model and illustrates the complicated structure of underlying basins of attraction. Two-dimensional transitions are achieved through forcing, or instantaneous displacement of the environment from one basin to another.

3.4 Gaussian processes

One may intuit that with very many biotic elements, there is a tendency towards uniformity in $\boldsymbol{F}$, reducing the likelihood of finding stationary points; the expected sum of random variables increases linearly with their number, while the variance increases with the root of the number, such that with large numbers of random variables, the variance

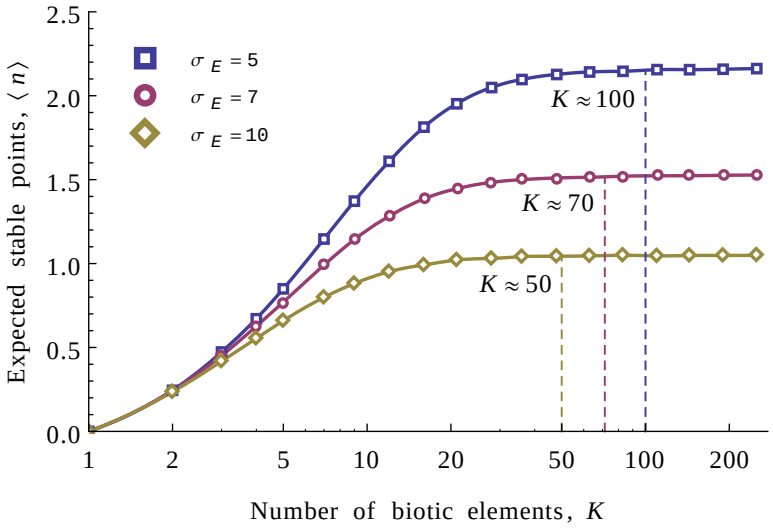

FIGURE 3.5: The expected number of fixed points is a function of the ratio of the niche function width σ_E to the essential range σ_μ (chosen to be 100 here), the environmental dimensions N (chosen as one for simplicity) and the diversity of the biota K which is shown to increase the expected number of fixed points in the model up to some saturation point where the biota can be said to be sufficiently diverse that the effect function $\boldsymbol{F}$ is well approximated by a Gaussian process (Dyke and Weaver, 2013).

tends towards negligibility. This is true except in the case of mean zero variables, ensured by $\langle \boldsymbol{\Omega} \rangle = \mathbf{0}$. Appropriate scaling of $\boldsymbol{F}$, provided by Eq. (3.6), guarantees there is no such tendency towards uniformity with increasing K.

Early intuitions were that increasingly complex systems enjoyed increased stability (Elton, 1958; Odum, 1971), a conclusion drawn from its corollary; that reduction in ecosystem complexity, such as by species invasion reducing biodiversity, appears to correspond to increases in system variability, reflected by fluctuations in species populations. The work of Gardner and Ashby (1970) and May (1972) contended this, showing that in a network interpretation of complexity, increased numbers and strength of connections ultimately led to instability. These historical perspectives of ecosystem stability we restrained by studying equilibrium population dynamics which are stationary in the absence of large perturbations. Later work has progressed our understanding by considering a broader concept of stability including permanency, resistance, resilience and inertia; similar, but subtly different concepts which introduce the concept of a stable, but not stationary system. McCann (2000) provide a brief review of the experimental work towards establishing a general relationship between complexity and stability.

Dyke et al. (2007) explore this contention in the Daisyworld model finding, as with our model, this is not the case; the independence of biotic elements, which are coupled only through their shared environment, ensures this instability does not emerge. As described in Fig. 3.3, very small numbers of biotic components require chance to generate the correct configurations for Daisyworld-like homoeostatic points to emerge. Fig. 3.3 shows this required configuration for a two species model. However, Fig. 3.5 shows

that increasing the biotic diversity K leads to an increase in the expected number of fixed points until it reaches saturation, where the net effects of the biota $\boldsymbol{F}$ is well approximated by a stationary Gaussian process. Computation of the total biotic force becomes expensive for large K, leading to exponentially increasing costs with increases in the environmental dimension N. To overcome this, we take this limit of $K \to \infty$, where the function $\boldsymbol{F}(\boldsymbol{E})$ may be represented by its covariance by the following method.

Firstly, we devise a spatial grid of n points in the space of environmental variables where we aim to sample the function $\boldsymbol{F}(\boldsymbol{E})$, and will interpolate between. The exact value of n should be chosen carefully as a trade-off between the quality of representation of $\boldsymbol{F}(\boldsymbol{E})$, and the computational cost (and memory requirement) of the matrix decomposition, which scales as $\mathcal{O}(n^2)$. At the very least, it should be chosen such that $n\,\sigma_E \geq \mu_R$. The spatial grid $\boldsymbol{\xi}$ and corresponding function samples, $\boldsymbol{Z}$, are illustrated by Fig. 3.6 and can be written as

$$\boldsymbol{\xi} = \begin{pmatrix} \boldsymbol{E}_1 \\ \boldsymbol{E}_2 \\ \vdots \\ \boldsymbol{E}_n \end{pmatrix} \qquad \boldsymbol{Z} = \begin{pmatrix} \boldsymbol{F}(\boldsymbol{E}_1) \\ \boldsymbol{F}(\boldsymbol{E}_2) \\ \vdots \\ \boldsymbol{F}(\boldsymbol{E}_n) \end{pmatrix} \tag{3.8}$$

Next, the covariance of $\boldsymbol{F}(\boldsymbol{E})$ between each element of $\boldsymbol{\xi}$ is computed in the covariance matrix, defined by

$$\mathbf{C} = \begin{pmatrix} k(\boldsymbol{E}_1, \boldsymbol{E}_1) & k(\boldsymbol{E}_2, \boldsymbol{E}_1) & \cdots & k(\boldsymbol{E}_n, \boldsymbol{E}_1) \\ k(\boldsymbol{E}_1, \boldsymbol{E}_2) & k(\boldsymbol{E}_2, \boldsymbol{E}_2) & \cdots & k(\boldsymbol{E}_n, \boldsymbol{E}_2) \\ \vdots & \vdots & \ddots & \vdots \\ k(\boldsymbol{E}_1, \boldsymbol{E}_n) & k(\boldsymbol{E}_2, \boldsymbol{E}_n) & \cdots & k(\boldsymbol{E}_n, \boldsymbol{E}_n) \end{pmatrix} \tag{3.9}$$

where $k(\boldsymbol{E}_i, \boldsymbol{E}_j)$ is the covariance of $\boldsymbol{F}$ between the points $\boldsymbol{E}_i$ and $\boldsymbol{E}_j$. The bulk of the computation is in the Cholesky decomposition of $\mathbf{C}$; finding the matrix $\mathbf{A}$ where

$$\mathbf{C} = \mathbf{A}^{\mathrm{T}}\mathbf{A}. \tag{3.10}$$

Importantly, while very costly to compute, calculation of $\mathbf{A}$ is deterministic and need only be carried out once for a specific grid in order to produce a large number of different functions $\boldsymbol{F}(\boldsymbol{E})$ for numerical validation. Finally, a matrix $\mathbf{W}$ of independent random values is generated corresponding to each of the n grid points and the N environmental variables,

$$\begin{matrix} & F_1 & F_2 & \dots & F_N & \\ \mathbf{W} = & \begin{pmatrix} w_{1,1} & w_{2,1} & \dots & w_{N,1} \\ w_{1,2} & w_{2,2} & \dots & w_{N,2} \\ \vdots & \vdots & \ddots & \vdots \\ w_{1,n} & w_{1,n} & \dots & w_{N,n} \end{pmatrix} & \begin{matrix} \boldsymbol{\xi}_1 \\ \boldsymbol{\xi}_2 \\ \vdots \\ \boldsymbol{\xi}_n \end{matrix} \end{matrix}, \tag{3.11}$$

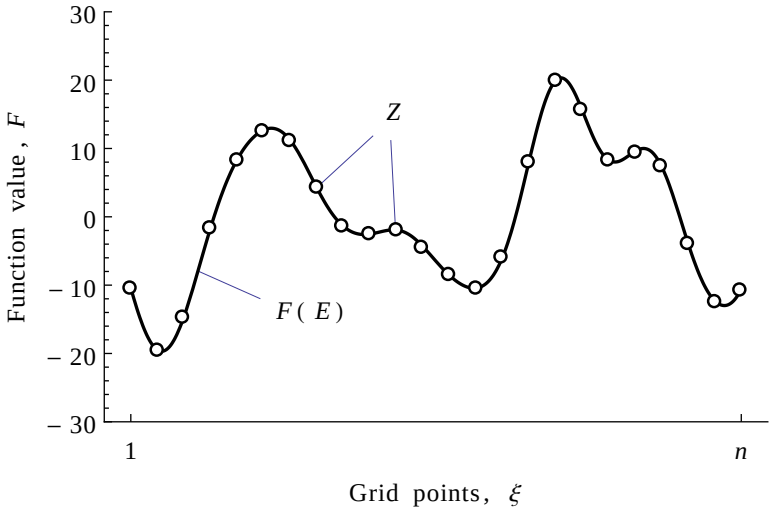

FIGURE 3.6: The vector $\boldsymbol{Z}$ contains function values corresponding to the n grid points, $\boldsymbol{\xi}$. These may then be simply interpolated to approximate the function.

and the sampling points $\boldsymbol{Z}$ are computed by the matrix product of $\mathbf{A}$ with $\mathbf{W}$. The function $\boldsymbol{F}(\boldsymbol{E})$ is found by basic interpolation between grid points, illustrated by Fig. 3.6.

$$\boldsymbol{Z} = \mathbf{A} \cdot \mathbf{W} \tag{3.12}$$

3.5 Generality of results

The mechanism responsible for these homeostatic fixed points is commonly referred to as reign-control. The term is as relevant here as it is to the two-species Daisyworld as our biota remains limited to exerting a monotonic increasing or decreasing effect on their environment. Fig. 3.3 illustrates a reign-control configuration in a Daisyworld-like model (similar to Fig. 2.2) while Fig. 3.3(b) shows three similar homeostatic fixed points emerging from a population of 100 randomly parametrised species.

In a two-species model such as Daisyworld, the emergence of two asymmetrical species required for a reign control state to exist is only one of four possibilities. Why, for example, should this be favoured over species which like it relatively hot and cold, while making it hotter and colder respectively? In such a case, the homeostatic fixed point is replaced by an unstable fixed point, and the system would not be homeostatic in any sense. One way to relax this assumption would be to allow the emergence of such destabilising daisies by introducing mutation in the albedo traits of daisies. The consequences of mutation in Daisyworld have been studied previously, for example Stocker (1995); Wood et al. (2006).

Our model consists of a random population where the niche and effects of individuals are uncorrelated such that there is no preference for positive or negative feedback. by formulation, significant parts of the population can be regarded as analogues of cold-loving, temperature-decreasing daisies and therefore with the capacity to disrupt homeostatic states. This model is shown in one dimension in Fig. 3.3(b) and shows a number of homeostatic fixed points, where a decrease in the environmental variable leads to an increasing effect from the biota and visa-versa.

It is important to establish the extent to which these observations are general with respect to both our model parameters, the biotic diversity K, environmental dimensions N and our implementation.

(1) Does the model behaviour change with increasing numbers of biotic components?

(2) How is it affected by increasing environmental complexity?

(3) To what extent is the choice of underlying functions important?

3.5.1 Density of fixed points

The emergence of fixed points in the environmental variable space is demonstrated by Fig. 3.3 and the purpose of this section is to determine their density as generally as possible. To begin with, we must characterise the summed effect function, $\boldsymbol{F}$, by deriving its covariance function in terms of the population of K biotic response functions $\boldsymbol{\alpha}(\boldsymbol{E})$ and their random coupling terms $\boldsymbol{\Omega}$. The following analysis is made much simpler under the approximation that the function $\boldsymbol{F}(\boldsymbol{E})$ is stationary such that the statistics of $\boldsymbol{F}(\boldsymbol{E})$ are uniform within the essential range, which requires us to assume assumes width of individual biotic functions to be small compared to the essential range, $R \gg \sigma_E$, and that the niche of our biotic elements are uniformly distributed. In reality, close to the edges of the essential range the amplitude of $\boldsymbol{F}(\boldsymbol{E})$ is reduced due to lower total biotic abundance. A consequence of this approximation is that the covariance of $F_i(\boldsymbol{E})$, $k_i(\boldsymbol{E}, \boldsymbol{E'})$, is also stationary, and depends only on the distance $|\boldsymbol{E} - \boldsymbol{E'}| = \Delta\boldsymbol{E}$

$$k_i(\Delta\boldsymbol{E}) = \mathbb{E}\left(\sum_{n,m=1}^{K} \omega_{i,n}\omega_{i,m}\alpha_n(\boldsymbol{E})\alpha_m(\boldsymbol{E}+\Delta\boldsymbol{E})\right). \tag{3.13}$$

At this point, we can exploit the absence of correlations first between individual biotic components, and then between the weights in $\boldsymbol{\Omega}$, and the biotic abundance. The first observation leads us to conclude the off-diagonal terms, where $i \neq j$, do not contribute to the covariance. The second enables us to separate the expectation values of $\boldsymbol{\omega}$ and α, giving

$$k_i(\Delta\boldsymbol{E}) = K\sigma_\omega^2\mathbb{E}\big(\alpha(\boldsymbol{E})\alpha(\boldsymbol{E}+\Delta\boldsymbol{E})\big). \tag{3.14}$$

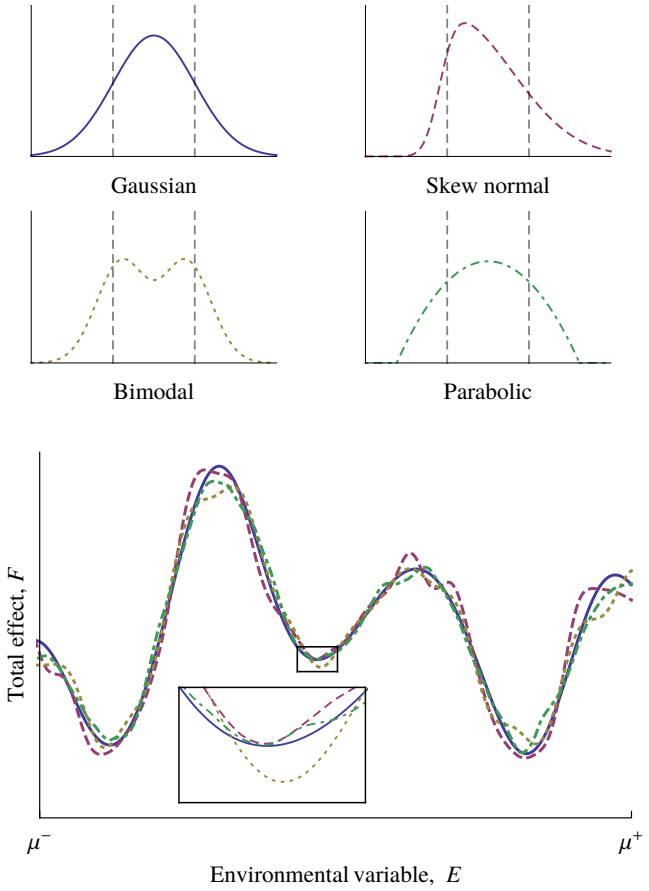

FIGURE 3.7: Four niche functions are shown with identical mean positions $\boldsymbol{\mu}$ and characteristic width σ_E. The total biotic force, $\boldsymbol{F}$, that the different niche functions produce in a population of 10^4 biotic functions is shown. Only the characteristic width and not the particular form of the niche function is important for the establishment of homeostatic states, and our choice of niche function is therefore arbitrary providing it has a well-defined width (Dyke and Weaver, 2013).

where σ_ω^2 is the variance of the random matrix $\boldsymbol{\Omega}$. The right side of this equation can be identified simply as the covariance of the individual biotic activity functions. This result shows that the covariance of the summed functions share the functional form of the individual functions of which it is comprised. The characteristic length is independent of the value of K and therefore the propensity for $\boldsymbol{F}$ to form homeostatic stable points is unaffected by the the biotic complexity of the model. Another consequence of this is to relegate our particular choice of the function α to arbitrariness since the only important factor is its characteristic width, illustrated for a range of possible choices in Fig. 3.7.

Next, we will use this result to derive an expression for the expected number of homeostatic fixed points within the essential range. We begin by stating that a model fixed point occurs in a small interval of $\boldsymbol{E}$ if the condition $\boldsymbol{F}(\boldsymbol{E}) - \boldsymbol{P} = 0$ is satisfied in the interval. Labelling the interval $\boldsymbol{\epsilon}$, this condition can be expressed

$$\big(F_i(\boldsymbol{E}) - P_i\big)\big(F_i(\boldsymbol{E}+\boldsymbol{\epsilon}) - P_i\big) < 0 \qquad \forall_i \tag{3.15}$$

and the expected number of such points in the unit interval, n_0, is found from a product of indicator functions of the form of Eq. (3.15)

$$n_0 = \mathbb{E}\left(\prod_{i=1}^{N} \frac{1}{\sqrt{\boldsymbol{\epsilon}^\top \boldsymbol{\epsilon}}} \exp\left(-\frac{P_i^2}{2\sigma_F^2}\right) \left[F_i(\boldsymbol{E})F_i(\boldsymbol{E}+\boldsymbol{\epsilon}) < 0\right]\right). \tag{3.16}$$

We have used square brackets $[\ldots]$ to denote the indicator function

$$\left[F_i(\boldsymbol{E})F_i(\boldsymbol{E}+\boldsymbol{\epsilon}) < 0\right] = \begin{cases} 0 & F_i(\boldsymbol{E}) \text{ and } F_i(\boldsymbol{E}+\boldsymbol{\epsilon}) \text{ have same sign} \\ 1 & F_i(\boldsymbol{E}) \text{ and } F_i(\boldsymbol{E}+\boldsymbol{\epsilon}) \text{ opposite sign} \end{cases} \tag{3.17}$$

The perturbation term has been removed from the indicator function in Eq. (3.16) by exploiting the Gaussian distribution of $\boldsymbol{F}(\boldsymbol{E})$. The expectation of an indicator function may be interpreted as the probability of its contents being true, and the product of several therefore gives the probability of many conditions being met simultaneously. Each term in the product may be treated independently due to the independence between the biotic effects on the different environmental variables F_i and F_j. The problem is therefore reduced to finding the value of the series of N expectation values. Expanding for small $|\boldsymbol{\epsilon}|$ gives

$$p = \mathbb{E}\Big(\left[F_i(\boldsymbol{E})^2 < -\epsilon_i F_i(\boldsymbol{E})F_i'(\boldsymbol{E})\right]\Big) \tag{3.18}$$

where $F_i'(\boldsymbol{E})$ is used to indicate the derivative of $F_i(\boldsymbol{E})$ in the $\hat{\boldsymbol{\epsilon}}$ direction

$$F_i'(\boldsymbol{E}) \equiv \nabla_{\hat{\boldsymbol{\epsilon}}} F_i(\boldsymbol{E}).$$

To find the expectation value of this indicator function, we need to know how $F_i(\boldsymbol{E})$ and $F_i'(\boldsymbol{E})$ are distributed. Rather than suffer any loss of generality, we make three important observations;

(1) at any point in $\boldsymbol{E}$ within the essential range, $F_i(\boldsymbol{E})$ is a sum of independent contributions from the biotic elements. Therefore, by the central limit theorem, each point follows a Gaussian distribution.

(2) this distribution has a mean of zero as previously stated. There is no tendency for positive or negative feedback between the biota and environment.

(3) $F_i(\boldsymbol{E})$ and $F_i'(\boldsymbol{E})$ are uncorrelated as a consequence of the independence of parameters $\boldsymbol{\mu}$ and $\boldsymbol{\omega}$.

The problem is now dramatically reduced, we need only find the variance of the Gaussian random variables $F_i(\boldsymbol{E})$ and $F_i'(\boldsymbol{E})$, labelled σ_F^2 and $\sigma_{F'}^2$ respectively. Having ensured normalisation of σ_F^2 by Eq. (3.6), we can write a similar expression for $F_i'(\boldsymbol{E})$, and

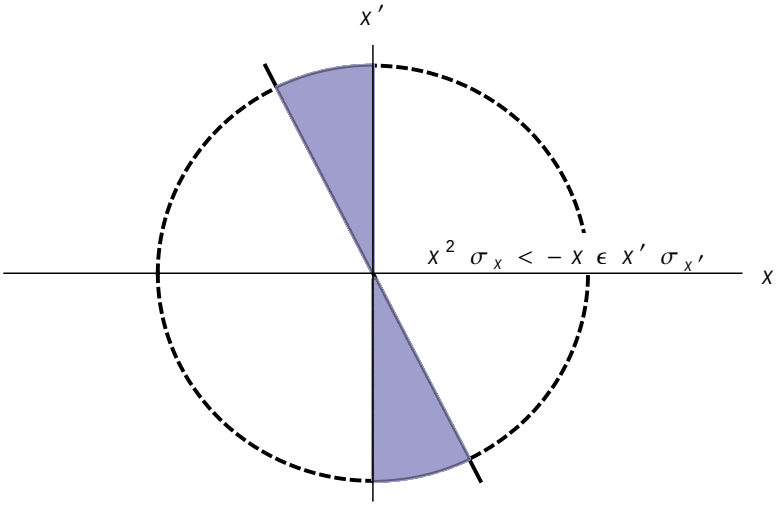

FIGURE 3.8: The solution to Eq. (3.22) may be found by determining the fraction of the unit circle which satisfies the spherically symmetric constraint indicated.

remove the directional derivative from the expectation value to give

$$\begin{aligned}\sigma_{F'}^2 &= \mathbb{E}\big(F'(\boldsymbol{E})F'(\boldsymbol{E}')\big)\Big|_{\boldsymbol{E}=\boldsymbol{E}'} \\ &= \nabla_{\boldsymbol{E},\hat{\boldsymbol{\epsilon}}}\nabla_{\boldsymbol{E}',\hat{\boldsymbol{\epsilon}}}k_i(\boldsymbol{E}-\boldsymbol{E}')\Big|_{\boldsymbol{E}=\boldsymbol{E}'} \\ &= -k_i''(\mathbf{0}) = \frac{1}{2\sigma_E^2}\end{aligned} \tag{3.19}$$

again, by normalisation of σ_F. Next, we substitute these into Eq. (3.18)

$$p = \iint \left[F_i(\boldsymbol{E})^2 < -\epsilon F_i(\boldsymbol{E})F_i'(\boldsymbol{E})\right]\psi(F_i;0,\sigma_F^2)\psi(F_i';0,\sigma_{F'}^2)\mathrm{d}F_i\mathrm{d}F_i' \tag{3.20}$$

where $\psi(x;\mu,\sigma^2)$ is the Gaussian probability density function centred at μ with variance σ^2

$$\psi(x;\mu,\sigma^2) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right). \tag{3.21}$$

After a change of variable, $\frac{F_i}{\sigma_F} \to x$ and $\frac{F_i'}{\sigma_{F'}} \to x'$, we can exploit spherical symmetry in x and x', illustrated in Fig. 3.8.

$$p = \iint \left[x^2\sigma_F < -\epsilon x\, x'\sigma_{F'}\right]\psi(x;0,1)\psi(x';0,1)\mathrm{d}x\mathrm{d}x' \tag{3.22}$$

We need only find the fraction of the $x\,x'$ plane for which $x^2\sigma_F < -\epsilon x\, x'\sigma_{F'}$, illustrated by Fig. 3.8. We therefore find the expectation of the indicator function to be

$$p = \frac{1}{\pi}\,\mathrm{atan}\left(\epsilon\frac{\sigma_{F'}}{\sigma_F}\right) \tag{3.23}$$

which can be expanded to first order for small ϵ, and substituted into Eq. (3.16) to give

$$n_0 = \exp\left(-\frac{\boldsymbol{P}^\top \boldsymbol{P}}{2}\right)\left(\frac{p}{\epsilon}\right)^N = \left(\frac{1}{\pi}\sqrt{\frac{1}{2\sigma_E^2}}\right)^N, \tag{3.24}$$

which is consistent with (Alder and Strassen, 1981, Theorem 4.1.1). Eq. (3.24) gives us the expected number of fixed points in a unit hyper-volume of the N-dimensional environmental variable space, assuming a large and diverse biota which maximises this value as previously illustrated by Fig. 3.5. This is a function only of the magnitude of external perturbing forces, $|\mathbf{P}|^2$ and the width of the biotic effect function σ_E^2.

3.5.2 Classification of fixed points

Following this we ask what fraction of these fixed points, where $\boldsymbol{F}(\boldsymbol{E}) = \mathbf{0}$ are stable and homeostatic in the face of external perturbation, and which fall into the diversity of other fixed point types. The distribution of fixed point types depends on the distribution of the eigenvalues $\boldsymbol{\lambda}$ computed from the Jacobian $\mathbf{J}(\boldsymbol{E})$ which, for a two-dimensional environment, may be written

$$\mathbf{J}(\boldsymbol{E}) = \begin{pmatrix} \frac{\mathrm{d}F_1(\boldsymbol{E})}{\mathrm{d}E_1} & \frac{\mathrm{d}F_1(\boldsymbol{E})}{\mathrm{d}E_2} \\ \frac{\mathrm{d}F_2(\boldsymbol{E})}{\mathrm{d}E_1} & \frac{\mathrm{d}F_2(\boldsymbol{E})}{\mathrm{d}E_2} \end{pmatrix} \tag{3.25}$$

Defining $\delta = \det(\mathbf{J}) = J_{1,1}J_{2,2} - J_{1,2}J_{2,1}$ and $\tau = \mathrm{tr}(\mathbf{J}) = J_{1,1} + J_{2,2}$, the six fixed point classifications are enumerated below, and illustrated by Fig. 3.9;

(1) $\delta < 0$
Real with one positive and one negative. The fixed point is a saddle.

(2) $\delta > 0,\ \tau < 0,\ \text{and}\ \tau^2 - 4\delta \geq 0$
Real and negative. The fixed point is a sink node.

(3) $\delta > 0,\ \tau > 0,\ \text{and}\ \tau^2 - 4\delta \geq 0$
Real and positive. The fixed point is a source node.

(4) $\delta > 0,\ \tau = 0$
Complex with zero real part. The fixed point is a center.

(5) $\tau < 0,\ \text{and}\ \tau^2 - 4\delta < 0$
Complex with negative real parts. The fixed point is a spiral sink.

(6) $\tau > 0,\ \text{and}\ \tau^2 - 4\delta < 0$
Complex with positive real parts. The fixed point is a spiral source.

Providing the real parts of all eigenvalues are negative, the fixed point is stable regardless of the imaginary parts. We can determine the relative probabilities of a fixed point falling

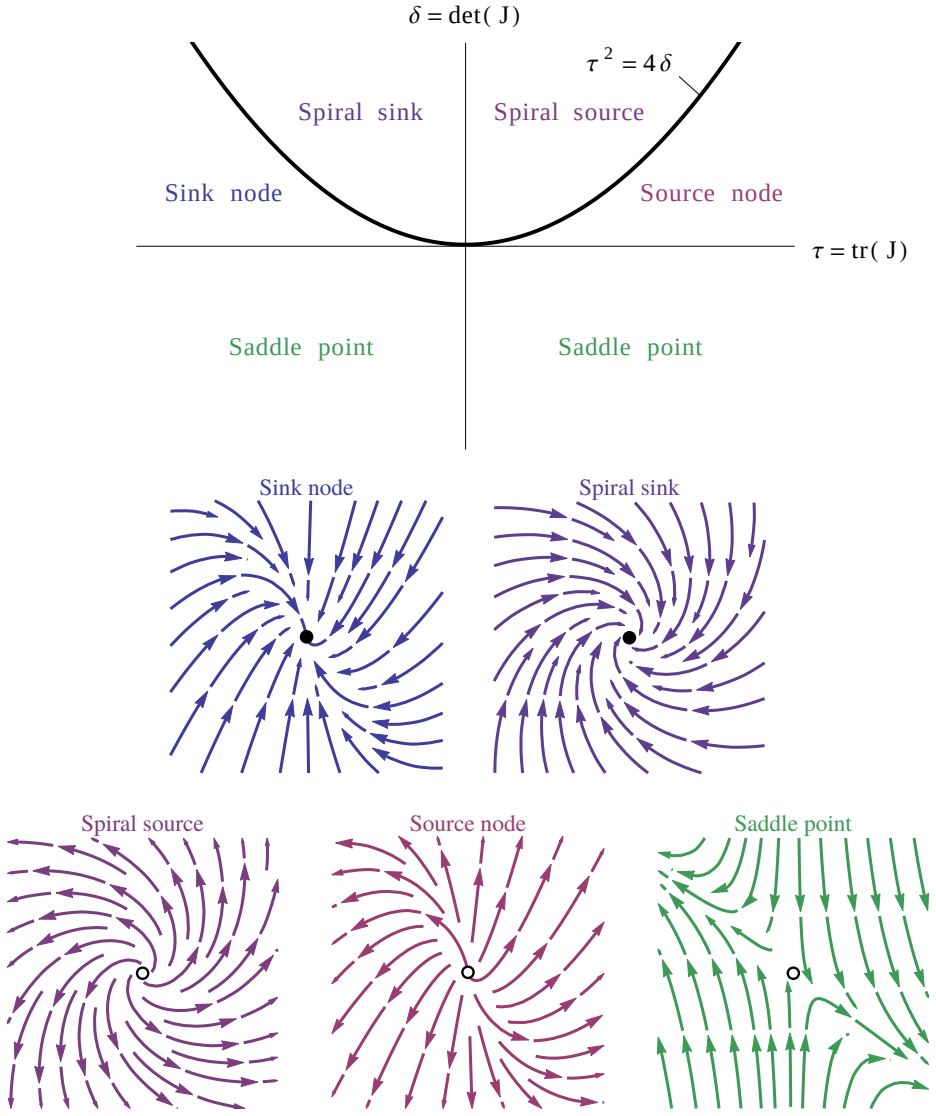

FIGURE 3.9: Model fixed points fall under six classifications. The relationship between the trace and determinant of the Jacobian $\mathbf{J}(\boldsymbol{E})$ determine whether eigenvalues are positive or negative and real or complex. These values govern the behaviour of the system in the vicinity of fixed points.

into any of these categories by first determining the distributions of δ and τ, then by integrating over the regions illustrated by Fig. 3.9.

The values in the Jacobian at any point are independent Gaussian distributed variables which are furthermore uncorrelated from the function value $\boldsymbol{F}(\boldsymbol{E})$. It is simple to show the trace of the Jacobian is also normally distributed with variance equal to the sum of those along the diagonal of the Jacobian.

$$\begin{aligned} X, Y &\sim \mathcal{N}\left(0, \sigma^2\right) \\ (Z = X + Y) &\sim \mathcal{N}\left(0, 2\sigma^2\right) \end{aligned} \tag{3.26}$$

where $\mathcal{N}\left(\mu, \sigma^2\right)$ is the normal distribution with probability density function

$$\psi(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right). \tag{3.27}$$

We can prove this relationship by introducing the moment generating function

$$\begin{aligned} M_Z(t) &= \mathbb{E}\left(e^{t(X+Y)}\right) \\ &= \mathbb{E}\left(e^{tX} e^{tY}\right) \\ &= \mathbb{E}\left(e^{tX}\right) \mathbb{E}\left(e^{tY}\right) \end{aligned} \tag{3.28}$$

due to the independence of the random variates X and Y. These expectations can be written and solved by

$$\begin{aligned} M_Z(t) &= \left(\int_{-\infty}^{\infty} \exp\left(tx - \frac{x^2}{2\sigma^2}\right) \mathrm{d}x\right)^2 \\ &= \left(\exp\left(\frac{\sigma^2 t^2}{2}\right)\right)^2 \\ &= \exp\left(\frac{2\sigma^2 t^2}{2}\right) \end{aligned} \tag{3.29}$$

where we have shown that the moment generating function for the sum of Gaussian variates with variance σ^2 is identical to that of a single Gaussian variate with variance $2\sigma^2$. It has been shown that if two distributions share the same moment-generating function, then they have identical distributions (Grimmett, 1986). Substituting our values of $\sigma = \frac{1}{\sigma_E}$ for the Jacobian elements, determined by Eq. (3.19) yields

$$\mathbf{J} \sim \mathcal{N}\left(0\ , \frac{1}{2\sigma_E^2}\right) \tag{3.30a}$$

$$(\tau = J_{1,1} + J_{2,2}) \sim \mathcal{N}\left(0\ , \frac{1}{\sigma_E^2}\right). \tag{3.30b}$$

The distribution of δ is more involved.

$$U, V, W, X \sim \mathcal{N}\left(0\ , \sigma^2\right) \tag{3.31a}$$

$$(Y = UV, WX) \sim \mathcal{K}\left(0\ , \sigma^2\right) \tag{3.31b}$$

$$(Z = UV + WX) \sim \mathcal{L}\left(0\ , \sigma^2\right) \tag{3.31c}$$

where $\mathcal{K}$ is a probability distribution defined by K_0 is the modified Bessel function of the second kind, and $\mathcal{L}$ is the Laplace double exponential function

$$\mathcal{K}(n,\sigma^2) \mapsto f(x;n,\sigma^2) = \frac{1}{\pi\sigma^2} K_n\left(\frac{\sqrt{x^2}}{\sigma^2}\right), \tag{3.32a}$$

$$\mathcal{L}(\mu,\sigma^2) \mapsto f(x;\mu,\sigma^2) = \frac{1}{2\sigma^2}\exp\left(\frac{\sqrt{(x-\mu)^2}}{\sigma^2}\right) \tag{3.32b}$$

where K_n is the modified Bessel function of the second kind. Again we can show this by way of moment generating functions. We begin by writing the distribution of UV, WX as

$$UV = \left(\frac{U+V}{2}\right)^2 - \left(\frac{U-V}{2}\right)^2. \tag{3.33}$$

Since we have already established $U+V$ to be normally distributed, the square follows the chi-square distributed. We have therefore shown that the distribution of the product of independent random variates is equivalent to the difference between pairs of independent Chi-squared variates. From here the moment generating function is much easier to find

$$\begin{aligned} M_{UV}(t) &= \mathbb{E}\left(\exp\left(t\left(\tfrac{U+V}{2}\right)^2 - t\left(\tfrac{U-V}{2}\right)^2\right)\right) \\ &= \mathbb{E}\left(\exp\left(t\left(\tfrac{U+V}{2}\right)^2\right)\right)\mathbb{E}\left(\exp\left(-t\left(\tfrac{U-V}{2}\right)^2\right)\right) \\ &= \int_{-\infty}^{\infty} e^{\frac{tx^2}{4}}\psi(x;0,2\sigma^2)\mathrm{d}x \int_{-\infty}^{\infty} e^{-\frac{tx^2}{4}}\psi(x;0,2\sigma^2)\mathrm{d}x \\ &= \frac{1}{\sqrt{1-\sigma^2 t}\,\sqrt{1+\sigma^2 t}} \\ &= \frac{1}{\sqrt{1-\sigma^4 t^2}} \end{aligned} \tag{3.34}$$

which can be shown to be the moment generating function of the Bessel function of Eq. (3.32b) (Craig, 1936).

$$\begin{aligned} M_Z(t) &= \mathbb{E}\left(e^{t(UV+WX)}\right) \\ &= \left(\frac{1}{\sqrt{1-\sigma^4 t^2}}\right)^2 \\ &= \frac{1}{1-\sigma^4 t^2} \end{aligned} \tag{3.35}$$

which is exactly the moment generating function of the Laplace function Eq. (3.32b) with $\mu = 0$. As before, we substitute our value of σ to yield the distribution of δ

$$(\delta = J_{1,1}J_{2,2} - J_{1,2}J_{2,1}) \sim \mathcal{L}\left(0, \frac{1}{\sigma_E^2}\right). \tag{3.36}$$

The univariate probability distributions Eqs. (3.30b) and (3.36) are not independent. It is simple to see that positive values of δ will correlate to large positive or negative values of τ. In order to generate the bivariate probability distribution of τ and δ, it is possible in principle to follow the usual method of moment generating functions, though in this instance it is far simpler to use a change of variable. As before, we begin by defining our independent variates

$$X, Y \sim \mathcal{N}\left(0\ , \sigma^2\right) \tag{3.37a}$$

$$Z \sim \mathcal{K}\left(0\ , \sigma^2\right) \tag{3.37b}$$

along with the covariates

$$(\tau \equiv T) = X + Y \tag{3.38a}$$

$$(\delta \equiv U) = XY - Z \tag{3.38b}$$

$$V = Z \tag{3.38c}$$

where we have introduced a sixth variable, V which ensures the mapping between our random variates is bijective, which will be important shortly. The aim is to transform the expectation of some function f to have the following form

$$\mathbb{E}\big(f(T,U,V)\big) = \iiint f(x+y, xy-z, z) p_X(x) p_Y(y) p_Z(z) \mathrm{d}x\mathrm{d}y\mathrm{d}z \tag{3.39a}$$

$$= \iiint f(t,u,v) g(t,u,v) \mathrm{d}t\mathrm{d}u\mathrm{d}v \tag{3.39b}$$

where $p_X(x)$, $p_Y(y)$ and $p_Z(z)$ are the probability densities of X, Y and Z, and $g(t,u,v)$ is the joint probability density of the random variables T, U and V, to be found. The redundant variable V can then be marginalised out by integration. We begin by the change of variable $(x,y,z) \to (t,u,v) = (x+y, xy-z, z)$ in Eq. (3.39a). We begin by inverting this mapping, which gives

$$\begin{aligned} x, y &= \frac{1}{2}\left(t \pm \sqrt{t^2 - 4(u+v)}\right) \\ z &= v \end{aligned} \tag{3.40}$$

and provides us with the domain $s^2 > 4(t+u)$. From this we compute the Jacobian of the coordinate transformation.

$$\begin{vmatrix} \frac{\mathrm{d}x}{\mathrm{d}t} & \frac{\mathrm{d}x}{\mathrm{d}u} & \frac{\mathrm{d}x}{\mathrm{d}v} \\ \frac{\mathrm{d}y}{\mathrm{d}t} & \frac{\mathrm{d}y}{\mathrm{d}u} & \frac{\mathrm{d}y}{\mathrm{d}v} \\ \frac{\mathrm{d}z}{\mathrm{d}t} & \frac{\mathrm{d}z}{\mathrm{d}u} & \frac{\mathrm{d}z}{\mathrm{d}v} \end{vmatrix} = \begin{vmatrix} \frac{1}{2}\left(1 + \frac{t}{k}\right) & -\frac{1}{k} & -\frac{1}{k} \\ \frac{1}{2}\left(1 - \frac{t}{k}\right) & \frac{1}{k} & \frac{1}{k} \\ 0 & 0 & 1 \end{vmatrix} \tag{3.41}$$

$$= \frac{1}{\sqrt{t^2 - 4(u+v)}} \tag{3.42}$$

where in Eq. (3.41) we have used $k = \sqrt{t^2 - 4(u+v)}$ for brevity. Substituting the results of Eqs. (3.40) and (3.42) into Eq. (3.39a) yields

$$\begin{aligned}\mathbb{E}\big(f(T,U,V)\big) = \iiint f(u,v,v)\ \psi\left(\frac{1}{2}\left(t+\sqrt{t^2-4(u+v)}\right),0,\sigma^2\right)\\ \psi\left(\frac{1}{2}\left(t-\sqrt{t^2-4(u+v)}\right),0,\sigma^2\right)\\ K_0\left(\frac{\sqrt{v^2}}{\sigma^2}\right)\frac{\big[t^2>4(u+v)\big]\mathrm{d}t\mathrm{d}u\mathrm{d}v}{\pi\sigma^2\sqrt{t^2-4(u+v)}}\end{aligned} \tag{3.43}$$

where we have favoured the use of an indicator function $[\ \cdots]$ to define the domain, as opposed to providing limits of integration. Recalling our definition of $g(t,u,v)$, the joint probability density of the random variables T, U and V from Eq. (3.39b)

$$\therefore\ g(t,u,v) = \begin{cases} \dfrac{1}{2\pi^2\sigma^4}\dfrac{\exp\left(-\dfrac{t^2-2(u+v)}{2\sigma^2}\right)K_0\left(\dfrac{\sqrt{u^2}}{\sigma^2}\right)}{\sqrt{t^2-4(u+v)}} & t^2>4(u+v)\\ 0 & t^2\le 4(u+v). \end{cases} \tag{3.44}$$

As mentioned, the variable v is redundant and has been included to ensure the mapping between our independent and codependent variates, $(x,y,z)\to(t,u,v)$ is bijective. This can now be marginalised out by integrating out v over the domain $t^2>4(u+v)$. The distributions determined to this point are summarised in Fig. 3.10(a).

Finally we can determine the share of the fixed points which fall into each category by integrating the joint probability distribution of τ and δ, Eq. (3.44), over the regions indicated by Fig. 3.9. As expected, the fraction $\frac{1}{4}$ of the fixed points given by Eq. (3.24) are stable. This can be determined for systems of any dimensionality, as fixed points belonging to each quadrant of Fig. 3.9 can be reached by inverting components of $\boldsymbol{F}(\boldsymbol{E})$. For example, the transformation $F_1\to -F_1$ does not change the properties of the system as $\boldsymbol{F}(\boldsymbol{E})$ is zero mean and symmetrical about $\boldsymbol{F}=\boldsymbol{0}$. However this transformation also causes $\delta\to-\delta$, so for example stable points are transformed into saddle points. By integrating the area above the curve shown in Fig. 3.9 we find the fraction $\frac{1}{\sqrt{2}}$ of stable fixed points in a two dimensional system to have oscillatory, or spiralling behaviour in their vicinity. Interestingly, the same result would be arrived at if we assumed eigenvalues to be uniformly distributed, which is in fact true for very large random matrices and is known as the circular law (Girko, 1985).

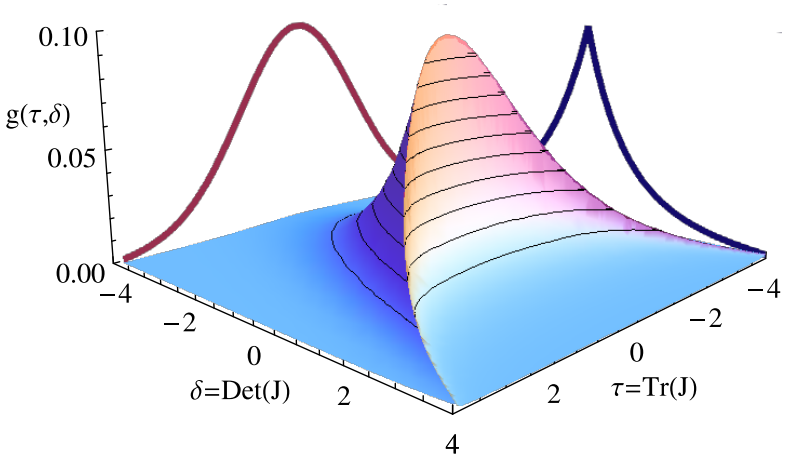

(a) The probability distribution of τ, given by Eq. (3.26), δ, given by Eq. (3.31c)

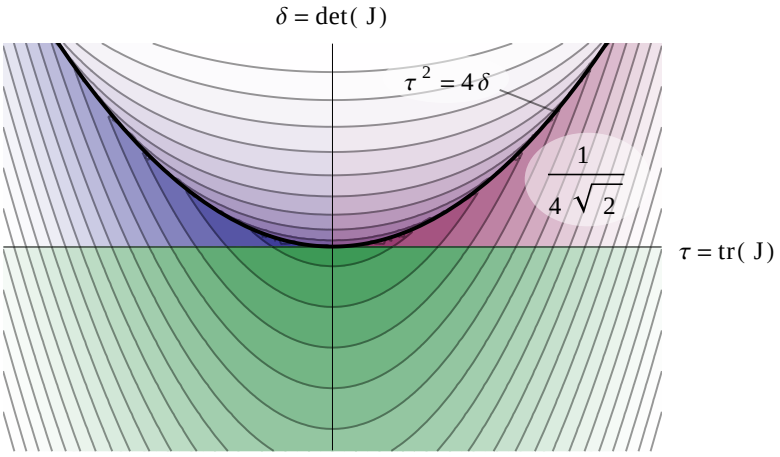

(b) The probability distribution is integrated to find the share of fixed points which fall into the categories of Fig. 3.9

FIGURE 3.10: We find each quadrant of the trace-determinant phase diagram to share $\frac{1}{4}$ of the integrated probability distribution, with the only non-trivial region being that partitioned by the curve and axis, $0 < \delta < \frac{\tau^2}{4}$. This region corresponds to fixed points with oscillatory behaviour and integrates to the fraction $\frac{1}{2\sqrt{2}}$.

3.6 Discussion

The relationship between complexity and stability is a well established topic in ecology (McCann, 2000). An important early paper reported an inverse relationship between diversity and stability in simple model ecosystems (May, 1972). As the number of linear connections between species increased, the probability that the ecosystem would be stable decreased. This finding was in part based on a study that found the stability of large systems underwent a catastrophic collapse at a certain level of connectivity (Gardner and Ashby, 1970). Our results can be interpreted in the light of these landmark studies. As the number of environmental variables, N, is increased we found an exponential increase in the number of stable states to be accompanied by a decreasing probability

that a randomly initialised system will remain with the essential range as it relaxes towards stable points. Increasing N corresponds to an increase in the number of biotic interactions as all such interactions are mediated via environmental variables. A more connected system is less stable in that respect. However, we found that the stability of the model increased with increased biodiversity up to a threshold, beyond which further increases in the numbers of biotic components had no effect on the expected number of stable points. How are we to understand this result in the light of May (1972)?

There are crucial differences between our model and that of May (1972). Importantly there are no inter-biotic competitive or trophic interaction; biotic elements can only interact through feedback with their shared environment. Homeostatic states emerged in Dyke et al. (2007); McDonald-Gibson et al. (2008) due to the presence of strong inter-species competition between a biota of finite size. This is not the case in our model where in the absence of competitive exclusion, a large number of independent biotic elements coexist. Another important distinction between our model and studies such as May (1972) is that we do not assume that interactions are linear or monotonic. That is, changes in one direction of the environmental vector do not always lead to increases in the activity of an element of the biota such that the fundamental niche of any biotic component is finite. For any environmental axis there will be conditions that are too low or too high for significant biotic activity. We have shown how the particular form of the biotic response function is irrelevant to the overall establishment of homeostasis. Providing they share an equal and well-defined characteristic width, the statistical behaviour of the model is independent of the choice of function, a result previously observed by Dyke and Harvey (2005).

The original Daisyworld was conceived as a proof of concept for a homeostatic biosphere. We have revisited this model with the aim to identify the smallest set of assumptions from which environmental homeostasis may emerge; life is affected by environmental conditions which in turn are affected by life. We have shown that self-regulation is a mechanism which may arise from a large population of random life elements, and explored this mechanism in a multi-dimensional environment. We have provided the results of numerical simulation in addition to a thorough analytical treatment, establishing our observations as fundamental emergent properties of this type of model, largely independent of parametrisation. Additionally we may go so far as to derive the distribution of fixed point behaviours, for example those which show oscillatory or saddle point dynamics. We note three important differences between the one-dimensional picture of hysteresis between bistable states; the magnitude of a regime-shift following a transition is highly variable, returning to a previous state after a transition may be impossible, and the density of viable states for a given perturbation is exponentially large with the number of environmental variables. These points prompt a number of questions when considering transitions in climatic or ecosystems; which aspects of the environment if

any are expected to undergo catastrophic changes? Can the previous state be recovered, and if so, which are the important dimensions of control? How many alternative states are consistent with external forces, and therefore stable? These questions will be the focus Chapter 4 while other questions regarding time dependence, competition and spatially distributions are discussed in Chapter 8, the concluding chapter of this thesis.

Chapter 4

Critical transitions in a general Earth-system model

4.1 Introduction

Along with key questions surrounding the stationary behaviour of Earth-system models, we explore the nature of critical transitions or regime shifts (Williams and Lenton, 2010). The idea of tipping points, where small changes to some aspect of an environment or ecosystem result in relatively abrupt, often catastrophic changes to the state of the environment or ecology, is approaching the fore when discussing Earth-systems. Increasingly the environment and its natural inhabitants are subjected to pressures both locally from processes such as farming and globally such as by changes to the global temperature or atmospheric composition. It is intuitive that such changes cannot continue without detecting some response from the Earth system, although the nature and mechanisms responsible for such changes remains a topic of contention. Will changes be smooth or abrupt, such as changes in the magnetisation of a ferromagnet with changes in temperature compared to spontaneously inverting poles in a strong magnetic field.

Our model is suited to offer intuition towards answering these important questions as even one dimensional environments produce a range of stable states and external forcing terms can be added to produce abrupt transitions between them. Higher dimensional environments have been shown to hold exponentially many such points and develop a diversity of intriguing dynamics (Dyke and Weaver, 2013). In this work we utilise our model, detailed in Sec. 3.2, by exploring its minimalist formulation in the context of time-varying external parameters, anticipating external forcing to result in characteristic transitions between the range of stable states whose behaviour and statistics we have explored exhaustively in Sec. 3.5 and Sec. 3.5.2. We begin by offering some preliminary results in Sec. 4.2, providing insights into the types of phenomena we can expect of our minimal ecosystem model. Sec. 4.3 provides an analytical treatment of the distribution of transition density with external perturbations. As a further step, we analyse the behaviour of the model in the vicinity of stable fixed points, particularly how this behaviour changes as the system approaches a transition in Sec. 4.4; a topic which has been gaining traction in the Earth-system context for several years now with an aim to identifying time series features which precede major transitions. The work concludes with an in-depth discussion in Sec. 4.5.

4.2 Results

The model formulation is presented in much more detail in Sec. 3.2. Here we summarise the model motivations, while presenting its mathematical representation in full. The model comprises of an environment with N variables, $\boldsymbol{E}$, and a biota of K variables, $\boldsymbol{\alpha}$ along with some perturbing term in each of the N environmental variables $\boldsymbol{P}$. The links between these model variables are illustrated in Fig. 4.1. Elements of the biota have a

positive or negative influence on each aspect of the environment which is linear with the abundance of a biotic element and parametrised by a $N \times K$ matrix $\boldsymbol{\Omega}$.

$$\boldsymbol{F}(t) = \boldsymbol{\Omega} \cdot \boldsymbol{\alpha}(t). \tag{4.1}$$

The environment in turn affects the biota but this connection is non-linear and based on the concept of niches; that each element of the biota has an ideal set of environmental conditions denoted $\boldsymbol{\mu}$ (Thomas et al., 2004). As the environment depart this point, the steady state abundance or activity of the biotic element decreases by some characteristic function which we choose to be Gaussian.

$$\boldsymbol{\alpha}^*(\boldsymbol{E}, \boldsymbol{\mu}) = \exp\left(-\frac{(\boldsymbol{\mu} - \boldsymbol{E})^\top (\boldsymbol{\mu} - \boldsymbol{E})}{2\sigma_E^2}\right). \tag{4.2}$$

While very simply expressed and intuitive, this is the least trivial aspect of the model; the choice of function is not obvious and as such it is important to see that this choice has no influence on the behaviour of a large model with a diverse biota, as shown in Sec. 3.5.1 (Dyke and Weaver, 2013). By allowing the biota to adjust relatively rapidly to this steady-state configuration $\boldsymbol{\alpha} * (\boldsymbol{E})$, we have shown we may replace the biotic effect with a Gaussian process, $\boldsymbol{F}(\boldsymbol{E})$; a stationary, smoothly varying function with a well defined covariance function. The dynamics of the model may be written

$$\frac{\mathrm{d}\boldsymbol{E}(t)}{\mathrm{d}t} = \boldsymbol{F}(\boldsymbol{E}) + \boldsymbol{P}(t). \tag{4.3}$$

where $\boldsymbol{P}(t)$ is a time varying contribution to the dynamics from external, or abiotic factors such as volcanic out gassing. This component is the focus of this chapter.

The model is characterised as producing a large number of fixed points. Sec. 3.3 shows the stationary behaviour of the model, finding fixed points arise with a calculable density. These points emerge from the diversity of the biota without the need for further tuning and scale exponentially with the complexity of the model environment. Fixed points are homeostatic in the sense that environmental conditions adjust only very slightly in the face of changes in perturbing force. Having already established the distribution of model fixed points with the magnitude of $\boldsymbol{P}$ to be Gaussian in Sec. 3.5.1, we move on to develop an understanding of the types of transitions we expect where fixed points annihilate, and the system reconfigures to some other stable point.

Hysteresis loops exist in a similar way to the one dimensional illustration Fig. 3.2, although additional dimensions of control open avenues for more interesting phenomena, such as irreversible transitions from environmental states which cannot be recovered by varying forcing in this one dimension. Applying perturbing forces in the two environmental dimensions reveals a greater depth of behaviour than a one dimensional model may exhibit. Fig. 4.2 shows the transition between two fixed points with perturbations

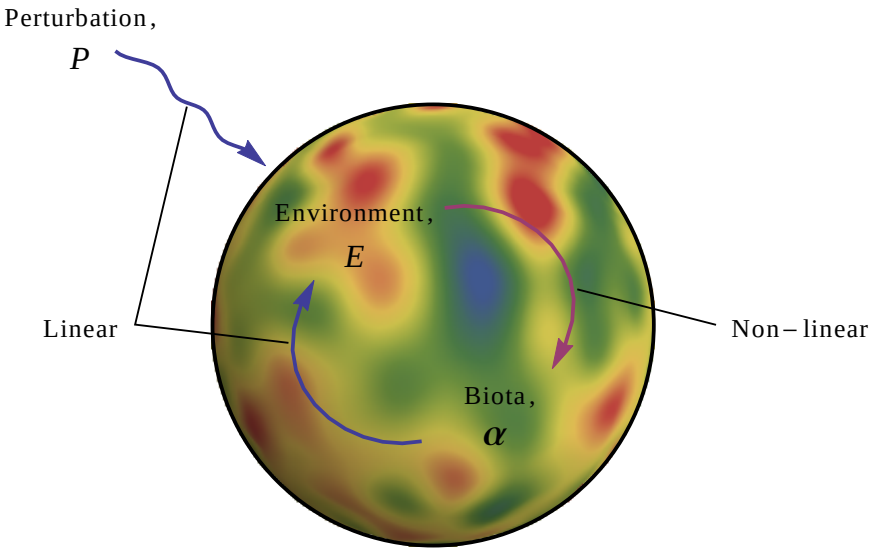

FIGURE 4.1: The abundance of a life elements, $\boldsymbol{\alpha}$, is coupled to the state of the environment, $\boldsymbol{E}$; Environmental conditions are affected by the biota by some linear coupling $\boldsymbol{\omega}$. There is an additional linear contribution from external influences, $\boldsymbol{P}$.

varied in two dimensions and is found to generate is a cusp bifurcation. These phenomena uniquely highlight the importance of understanding the major dimensions of control in a complex system as they describe a situation where one controlled variable changes the nature of a transition encountered by changes in another. Fig. 4.2 shows that changes in perturbing force P_2 cause changes in the state of the environmental variable E_1, though these changes can be smooth or discontinuous depending on the value of another perturbing effect P_1. This seemingly rare or designed phenomena has emerged simply from a randomly parametrised two dimensional model and highlights the types of structure and attractor one expects very much more complex Earth systems to produce.

4.3 Analysis of Transition Distributions

Precisely as before, we begin by determining the general properties of the model which may be established by rigorous exploration of the parameter space and simulation, and by simplification and an analytical treatment of key properties. In our case, simulation of models with $N > 4$ environmental variables are hugely expensive due to the exponentially increasing computation required with increasing environmental dimensionality. This limits the statistical rigour we may achieve, although the simplicity of the model by design, and the approximation of the system in the large K limit by a Gaussian process grants a great deal of utility when mining the model for analytically tractable properties. We were able to determine the distribution of stationary points, where the biota exists in a homeostatic configuration as function of external perturbation magnitude

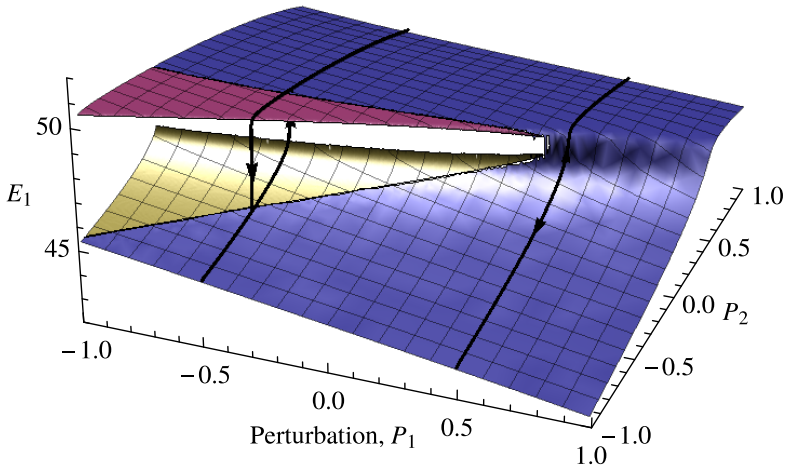

FIGURE 4.2: The shape of a single fixed point in a two-dimensional model with changes in perturbation P_1 and P_2 in the first and second environmental variables respectively. We can see that a transition is encountered by changing P_2 in the positive or negative direction. Where P_1 is positive, this is a smooth change while where P_2 is negative, the change is discontinuous, resulting in a bifurcation illustrated in Fig. 4.5.

in Sec. 3.5.1. Here we add to this the distribution of transitions between homeostatic states, and whether they are caused by increases in perturbation in the negative, or positive directions.

Proceeding along similar lines to previous exploration of the model (Dyke and Weaver, 2013), we can determine the density of fixed points though in this instance, it is interesting to express this as a function of the external perturbation. Simulation is in agreement with intuition that strongly forced environmental variables have a weaker propensity to form fixed points. Previously, we found the density of fixed points by asking the likelihood of the function $\boldsymbol{F}(\boldsymbol{E})$ crossing zero in all dimensions. The challenge here is similar, we begin by noting that the limits of homoeostasis, where a perturbing force is able to destabilize the environment, are dictated by local extrema in $\boldsymbol{F}(\boldsymbol{E})$ along some path. As a perturbation varies, the location of fixed points follows some random trajectory. Should we find a root in the derivative $\nabla_{\hat{\boldsymbol{\epsilon}}}\boldsymbol{F}(\boldsymbol{E})$ in *any* of the environmental variables the direction of this random trajectory then we encounter a local optima which limits the ability of the biota to oppose further increases or decreases in the perturbing force. Alternatively, such events may be described as points where unstable and stable fixed points converge, and annihilate.

We found the function value and its first derivative to be uncorrelated as a consequence of the independence of parameters $\boldsymbol{\mu}$ and $\boldsymbol{\omega}$ which greatly aided analysis. The same is true of the first and second derivatives. However, comparing the function to its second derivative finds correlation, that is to say the value of the function $\boldsymbol{F}(\boldsymbol{E})$ is not correlated to its gradient, but it is to the rate of change of gradient. Therefore, the density of transitions is, as expected, a function of the external perturbation. We have

determined how the first derivative is distributed in Eq. (3.19). Next we determine the distribution of the second derivative, and its covariance with $\boldsymbol{F}(\boldsymbol{E})$.

$$\boldsymbol{\Sigma}_i = \begin{pmatrix} \mathbb{E}\left(F_i\, F_i\right) & \mathbb{E}\left(F_i\, F_i''\right) \\ \mathbb{E}\left(F_i''\, F_i\right) & \mathbb{E}\left(F_i''\, F_i''\right) \end{pmatrix} \tag{4.4}$$

where we have used $\boldsymbol{F}''$ as shorthand for the second derivative of $\boldsymbol{F}$ to highlight the fact that the direction of the derivative does not matter due to isotropy in $\boldsymbol{F}$. It is intuitive that large positive or negative values of F_i may be close to extrema, and therefore correlate to a negative or positive value of F_i'' respectively. The diagonal terms of the matrix are simply the variances of $\boldsymbol{F}$ and $\boldsymbol{F}''$,

$$\mathbb{E}\left(F_i F_i\right) = 1 \tag{4.5a}$$

$$\mathbb{E}\left(F_i'' F_i''\right) = \frac{3}{4\sigma_E^4} \tag{4.5b}$$

while calculating the other terms is slightly more more involved.

$$\begin{aligned} \mathbb{E}\left(F_i F_i''\right) &= \frac{1}{\sigma_\mu^2}\mathbb{E}\left(\sum_j^K f_{i,j} \sum_k^K \left(\frac{(\boldsymbol{\mu}-\boldsymbol{E})^\top(\boldsymbol{\mu}-\boldsymbol{E})}{\sigma_E^2} - 1\right) f_{i,k}\right) \\ &= \frac{K}{\sigma_E^2}\mathbb{E}\left(f_i^2\left(\frac{(\boldsymbol{\mu}-\boldsymbol{E})^\top(\boldsymbol{\mu}-\boldsymbol{E})}{\sigma_E^2} - 1\right)\right) \end{aligned} \tag{4.6}$$

where we have taken only the diagonal terms as before. Evaluating the expectation value gives

$$\begin{aligned} \mathbb{E}\left(F_i F_i''\right) &= \frac{K}{C^2\sigma_E^2}\int \left(\frac{(\boldsymbol{\mu}-\boldsymbol{E})^\top(\boldsymbol{\mu}-\boldsymbol{E})}{\sigma_E^2} - 1\right) \exp\left(-\frac{(\boldsymbol{\mu}-\boldsymbol{E})^\top(\boldsymbol{\mu}-\boldsymbol{E})}{\sigma_E^2}\right)\left(\int \omega^2 \mathrm{d}\omega\right)\mathrm{d}\boldsymbol{\mu} \\ &= \frac{K\sigma_\omega^2}{C^2\sigma_E^2}\left(\int \frac{(\boldsymbol{\mu}-\boldsymbol{E})^\top(\boldsymbol{\mu}-\boldsymbol{E})}{\sigma_E^2} \exp\left(-\frac{(\boldsymbol{\mu}-\boldsymbol{E})^\top(\boldsymbol{\mu}-\boldsymbol{E})}{\sigma_E^2}\right)\mathrm{d}\boldsymbol{\mu} - \sqrt{\pi}\,\sigma_E\right) \\ &= \frac{K\sigma_\omega^2}{C^2\sigma_E^2}\left(\frac{1}{2}\sqrt{\pi}\,\sigma_E - \sqrt{\pi}\,\sigma_E\right) \\ &= -\frac{1}{2\sigma_E^2}\frac{K\sqrt{\pi}\sigma_E\sigma_\omega^2}{C^2} \\ &= -\frac{1}{2\sigma_E^2}. \end{aligned} \tag{4.7}$$

As expected, while $\boldsymbol{F}$ and $\boldsymbol{F}'$ were previously uncorrelated, they are clearly not independent as we find a negative correlation between $\boldsymbol{F}$ and $\boldsymbol{F}''$ as intuition suggests. The covariant Gaussian distribution has density

$$\rho_i(\boldsymbol{x}) = \frac{1}{2\pi\sqrt{\det(\boldsymbol{\Sigma})}} \exp\left(-\frac{1}{2}\,\boldsymbol{x}^\top \boldsymbol{\Sigma}_i^{-1}\boldsymbol{x}\right) \tag{4.8}$$

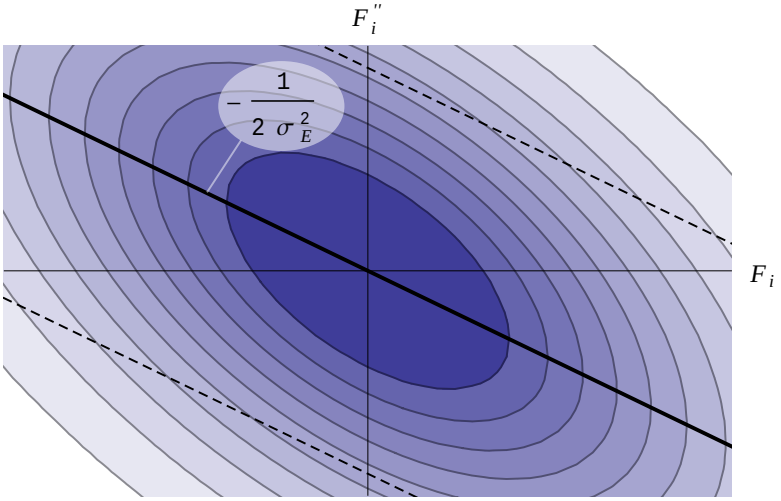

FIGURE 4.3: The distribution of the second derivative of $\boldsymbol{F}$ is Gaussian distributed shown here shaded where higher density is darker. The distribution variance $\frac{3}{4\sigma_E^4}$ is shown by the dashed line along with its mean $-\frac{1}{2\sigma_E^2}\boldsymbol{F}$ shown by the solid line. Large positive values of $\boldsymbol{F}$ are correlated with negative second derivatives and vice versa (in each dimension independently).

where $\boldsymbol{x} = \binom{F_i}{F_i''}$ and $\boldsymbol{\Sigma}$ is found by substituting Eqns. (4.5a), (4.5b) and (4.7) into Eq. (4.4)

$$\boldsymbol{\Sigma}_i = \begin{pmatrix} 1 & -\frac{1}{2\sigma_E^2} \\ -\frac{1}{2\sigma_E^2} & \frac{3}{4\sigma_E^4} \end{pmatrix}. \tag{4.9}$$

The probability density function Eq. (4.8) is shown in Fig. 4.3 and illustrates that the distribution of $\boldsymbol{F''}$ is Gaussian centred on $-\frac{1}{2\sigma_E^2}\boldsymbol{F}$ with variance $\frac{3}{4\sigma_E^4}$.

The rest of the analysis proceeds identically to Sec. 3.5.1, and so we pick up from Eq. (3.22) which expresses the probability, p, of encountering a local maxima in some small interval of the environmental variable space $\boldsymbol{\epsilon}$, which corresponds to an inability to exactly oppose further increases in forcing.

$$p = \iint \left[F_i'(\boldsymbol{E})^2 < -\epsilon F_i'(\boldsymbol{E}) F_i''(\boldsymbol{E}) \wedge F_i'' < 0\right] \psi\left(F_i'; 0, \sigma_{F'}^2\right) \psi\left(F_i''; -\frac{P_i}{2\sigma_E^2}, \sigma_{F''}^2\right) \mathrm{d}F_i' \mathrm{d}F_i'' \tag{4.10}$$

where we have previously established the variances

$$\sigma_{F'}^2 = \frac{1}{2\sigma_E^2} \qquad \text{and} \qquad \sigma_{F''}^2 = \frac{3}{4\sigma_E^4}.$$

Therefore the number density of extrema along the trajectory of a fixed point can be written

$$n_t = \frac{2p}{\epsilon \sigma_{F'}} \tag{4.11}$$

Unlike our previous application, there is no symmetry to exploit here though the integral is simple enough to perform numerically, shown in Fig. 4.4. Eq. (4.10) gives the

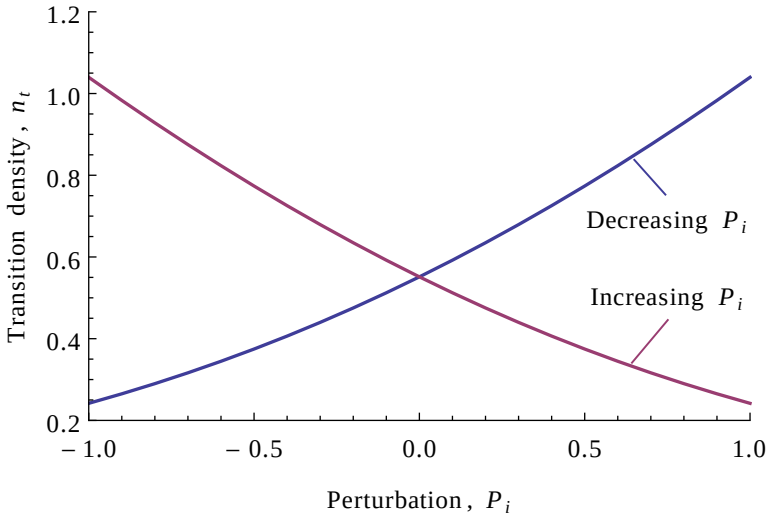

FIGURE 4.4: Transitions occur in the model where stable fixed points annihilate with unstable ones at extrema of $\boldsymbol{F}(\boldsymbol{E})$. The density of transitions per unit increase or decrease in P_i are given by integrating Eq. (4.10) numerically, shown here. Systems subjected to a positive perturbing force are increasingly likely to undergo a transition from further increases in perturbation, and vice versa.

probability of finding a maximum in one dimension in a small interval of $\boldsymbol{E}$ which is approximately related to a corresponding interval of $\boldsymbol{P}$ by a factor of $2\sigma_{F'}$.

4.4 Early Warnings in Complex Ecosystems

As systems approach a critical point they are at the cusp of a significant transition. Here, noise or very minor adjustments to perturbing forces can tip the system from one state to another; we say its susceptibility to these changes diverges and in practice, a stochastically perturbed system would see an increasing variance in its state variables as its susceptibility diverges (Carpenter and Brock, 2006). In addition to large susceptibility, the near-critical dynamics of such systems are slow such that fluctuations which do not cause transitions are reigned back to the steady-state very weakly. As before, the signal produced from a slow system would be detectable, but this time rather than increasing variance, one would detect longer trends due to slower relaxation (Kleinen et al., 2003). van Nes and Scheffer (2007) studies these phenomena in the context of ecosystem models, finding them to be reliable indicators for identification of near-critical ecosystems, those which are particularly vulnerable to reconfiguration or collapse.

This phenomena in the Earth-system context is no different from a term used by physicists in the context of phase transitions long before and is now called Critical slowing-down (CSD), referring to the long relaxation time of near-critical systems, those which are approaching phase, climatic or ecosystem transitions (Scheffer et al., 2009; Lenton, 2011; Sornette, 2002). As natural and anthropogenic pressures stress aspects of the

Earth system, it has been shown that certain time series find an increase in the auto-correlation coefficient precedes such transitions (Dakos et al. (2008) analyse eight ancient abrupt climate shifts). This signal suggests that the relaxation of the system towards its steady state slows as it approaches a regime shift.

Even as a simple dynamical system, we can show that our model is susceptible to this type of approach by computing the rate at which the model responds to fluctuations from its steady state in the vicinity of a transition. Fig. 4.5 illustrates the decay time of the system shown in Fig. 4.2 computed from the eigenvalues of the Jacobian of the system at the fixed points. Negative eigenvalues indicate that fixed points are stable, but additionally the inverse of the magnitude of eigenvalues provides us an estimate of the decay time of small fluctuations from the fixed point. Small negative eigenvalues correspond to a long decay time, indicating the system is in some sense slow, and recovery from small perturbations takes a long time. However it appears to posses no further information pertaining to the direction, magnitude or reversibility of the transition; clearly important and relevant questions when considering transitions in real systems. Fig. 4.5(a) shows that in both directions, the decay time diverges at the transition and we therefore expect to be able to pick out a CSD signal in this type of system. Fig. 4.5(b) also shows the decay time to become large, though strictly not divergent close to the point where the transition would occur for negative P_1. In a sense, this can be seen as misleading; in a real system, an investigator may expect an abrupt transition to some significantly different environmental state, though they find a fast but smooth variation.

We can consider the application of this to a simulated data set, beginning by generating a time series from a stochastically increasing perturbation, $\boldsymbol{P}(t)$. Up to this point, the difference between time scales of changes in the biota, environment and perturbing force have been large. If we relax this assumption, there is no clear function to choose for the relaxation of the composition of the biota into its steady state value. We proceed by assuming the simplest case of linear relaxation of populations towards the steady state value. This can be seen either as choosing a linear function simplest of a range of less trivial choices, or an assumption that the departure of the biota from its stationary state is sufficiently small that its return is approximately Newtonian.

$$\tau_\alpha \frac{\mathrm{d}\boldsymbol{\alpha}(t)}{\mathrm{d}t} = \boldsymbol{\alpha}^*(\boldsymbol{E}, \boldsymbol{\mu}) - \boldsymbol{\alpha}(t). \tag{4.12}$$

Such a case removes the need to resolve the individual populations $\boldsymbol{\alpha}$, reducing to a $2N$ dimensional system. Taking the product with $\boldsymbol{\Omega}$ gives

$$\begin{aligned} \tau_\alpha \boldsymbol{\Omega} \cdot \frac{\mathrm{d}\boldsymbol{\alpha}(t)}{\mathrm{d}t} &= \boldsymbol{\Omega} \cdot \boldsymbol{\alpha}^*(\boldsymbol{E}, \boldsymbol{\mu}) - \boldsymbol{\Omega} \cdot \boldsymbol{\alpha}(t) \\ \tau_\alpha \frac{\mathrm{d}\boldsymbol{F}(t)}{\mathrm{d}t} &= \boldsymbol{F}^*(\boldsymbol{E}) - \boldsymbol{F}(t) \end{aligned} \tag{4.13}$$

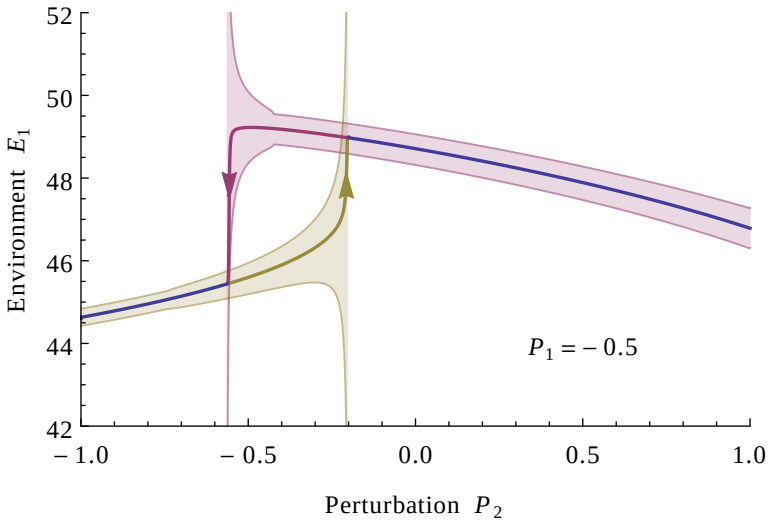

(a) Transition encountered with negative P_1.

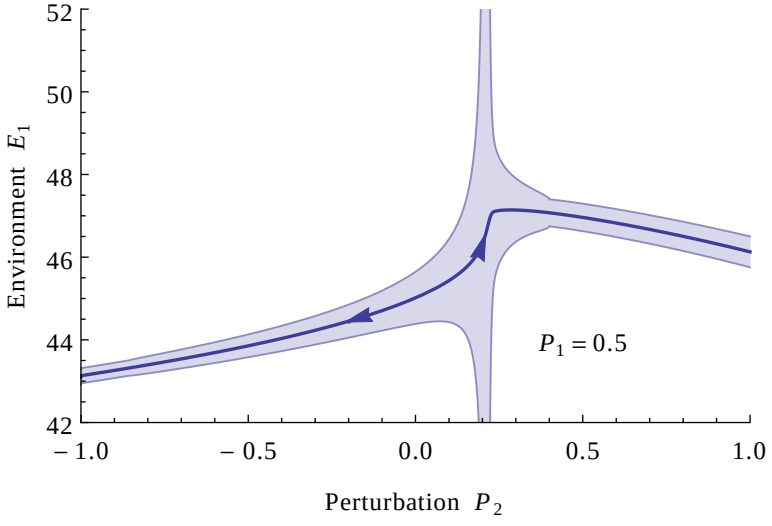

(b) Transition encountered with positive P_1.

FIGURE 4.5: Two cross-sections of Fig. 4.2 showing that the state of perturbation P_1 dictates the type of transition encountered by changes in P_2. Solid lines indicate the position of the fixed points, while shaded region widths are proportional to the decay time of small fluctuations from the fixed point. Fig. 4.5(a) shows a discontinuous transition caused by changes in P_2 where P_1 is positive while a smooth transition caused by changes in P_2 where P_1 is negative, illustrated by Fig. 4.5(b).

by the definition for $\boldsymbol{F}$ of Eq. (4.1), removing the need to keep track of individual populations. Sec. 2 has shown that the relationship between the time scales of changes in the biota and changes in external perturbation set limits on the ability of a system to self-regulate, and invite a range of new phenomena to emerge.

The perturbation is chosen to follow a Weiner process, where subsequent measurements are separated by small Gaussian distributed random numbers.

$$\big(P_1(t) - P_1(t + \Delta t)\big) \sim \mathcal{N}(\mu\Delta t, \sigma^2\Delta t) \tag{4.14}$$

where μ and σ^2 are the expected rate of increase in perturbing force and the variance of the increasing signal respectively. The perturbation drives the system through the transition shown in Fig. 4.5(a). Estimations the decay time of fluctuations within this time series may be found from the autocorrelation coefficient, exactly as Dakos et al. (2008), which can be summarised by the following steps:

i) Detrending

Detrending aims to remove long time trends in a time series, isolating the fast noise which is assumed to reflect the response of the system to small perturbations around a fixed point. Long time trends are represented by a smoothed signal, found by convolving the signal with, for example, a Gaussian kernel of width, σ (known as the bandwidth). For our purposes this process may be thought of as taking a weighted average, where the Gaussian kernel provides the weights, and may be formally written as

$$(f * g)(t) = \int_{-\infty}^{\infty} f(\tau)g(t - \tau)\mathrm{d}\tau \tag{4.15}$$

Where $f(t)$ is our signal, and $g(t)$ is a Gaussian with characteristic width σ (although commutativity ensures the distinction is unnecessary). The detrended signal, $f_{\mathrm{de.}}(t)$ is found simply by subtracting the smoothed signal from the time series.

$$f_{\mathrm{de.}}(t) = f(t) - (f * g)(t) \tag{4.16}$$

The convolution of a time series with a Gaussian is shown in Fig. 4.6(a).

ii) Auto-regression (AR)

Autoregressive models are commonly used to represent random processes in nature where the present state of a signal is correlated in some sense to its history. The simplest case possible is commonly referred to as an AR(1) model, where the number one indicates that we seek only a statistical relationship between temporally adjacent measurements

of the state of the signal

$$f_{\text{de.}}(t + \Delta t) = \alpha f_{\text{de.}}(t) + \epsilon(t), \tag{4.17}$$

where Δt is the time difference between subsequent signal measurements, or data points, α is the auto-regression parameter and $\boldsymbol{\epsilon}(t)$ is white noise, the residual between the time series data and AR(1) model. We aim to compute the value of α which best models our time series by fitting the model to minimize the sum of squares of our residual noise.

$$R^2 = \sum_{t=t_0}^{t_0+\lambda} \epsilon(t)^2 \tag{4.18}$$

The fitting parameter, α, quantifies the degree of correlation between subsequent points in the time series. For example, $\alpha \approx 0$ would indicate consecutive data points are minimally correlated. That is to say, the residuals cannot be decreased by assuming subsequent data points are in any way correlated. Increasing positive or negative values suggest positive or negative correlations between consecutive points.

The aim of a CSD analysis is to determine how the fitting parameter varies throughout a time series, and relate this to the response of the system. Large positive α indicates strong correlations and dynamics which are in some sense *slow* while the absence of correlations indicates the system dynamics are sufficiently fast that they are not resolved by the resolution of the time series. To this end, we choose the data range of Eq. (4.18), λ to be smaller than the entire data set. By fitting this model to successive *windows* along the full time series, we can observe how the correlation coefficient varies as we approach a transition.

iii) Decay time

Measurement of the AR(1) coefficient has been shown to give us intuition into the time scales of a system dynamics, though a little more effort is required to translate this into meaningful physical quantities which can then be related directly to our previous eigenvalue study of Fig. 4.5. We begin by assuming a small fluctuation, Δf, to decay exponentially towards their steady state value, f^*, at some characteristic rate, κ,

$$f(t) = f^* + \Delta f \exp\left(-\frac{t}{\kappa}\right). \tag{4.19}$$

This can be substituted into Eq. (4.17), and rearrange for the decay time κ.

$$\begin{aligned} \exp\left(-\frac{t}{\kappa}\right) - \alpha \exp\left(-\frac{t+\Delta t}{\kappa}\right) &= 0 \\ \exp\left(-\frac{\Delta t}{\kappa}\right) &= \alpha \\ \kappa &= -\frac{\Delta t}{\log(\alpha)}. \end{aligned} \tag{4.20}$$

The decay time can be extracted from Fig. 4.6(a) by this method, and is shown in Fig. 4.6(b).

iv) Kendall τ coefficient

The strength of an apparently increasing signal, such as the AR(1) coefficient shown in Fig. 4.6(b) can be quantified by the Kendall rank coefficient, τ, which measures the extent to which an observed trend represents a monotonic increase. It is found simply by the difference between the number of pairs of data points which indicate an increase over time (relative only to each other), $f(t_i) > f(t_j) \ \wedge \ t_i > t_j$, and those which show a decrease, called the number of concordant pairs, n_c, and discordant pairs, n_d, respectively. This is then normalised by the total number of pairs $\frac{1}{2}n(n-1)$.

$$\tau = \frac{n_c - n_d}{\frac{1}{2}n(n-1)}. \tag{4.21}$$

A signal dominated by noise will tend to have equal numbers of pairs increasing and decreasing over time, and therefore gives $\tau \approx 0$, while $\tau \approx \pm 1$ would indicate a monotonic increase or decrease. The significance of a given τ statistic can be estimated in the usual way; by determining the probability of arriving at a τ value of at least this size given random, uncorrelated data to produce a p-value. Fig. 4.6(b) yields a Kendall coefficient $\tau = 0.63$, establishing a positive trend with near certainty.

The time series of Fig. 4.6(a) shows a two dimensional model with a stochastically increasing forcing term in one environmental dimension. Fig. 4.5 analyses the Jacobian at the fixed points as this trajectory approaches a fixed point, and finds the largest eigenvalue to approach zero. This causes the decay time, that is the roughly exponential decay rate of small perturbations to the fixed point, diverges and the response of the system close to a perturbation is therefore expected to be *slow*, a recently popular suggestion for the behaviour of real Earth-systems in the vicinity of a critical transition (Lenton, 2011). Fig. 4.6(b) verifies this for the stochastically driven model by showing significant increases in the autocorrelation coefficient of a simulated time series, corresponding to a strongly increasing decay time.

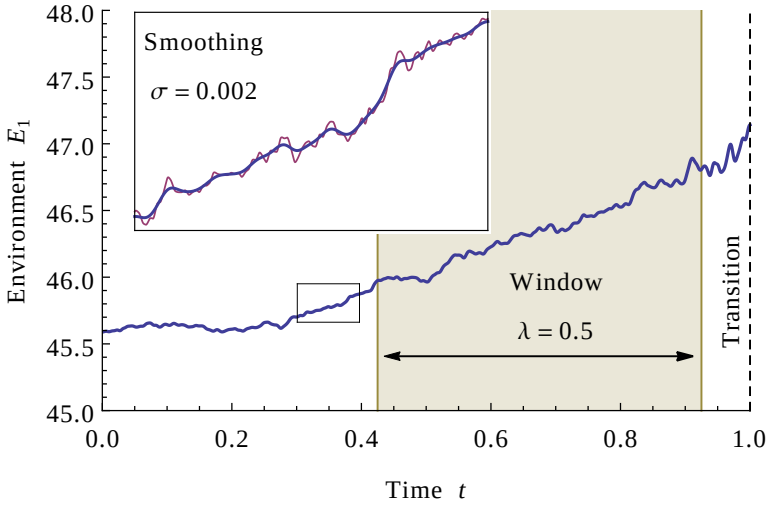

(a) A transition encountered by stochastically increasing forcing.

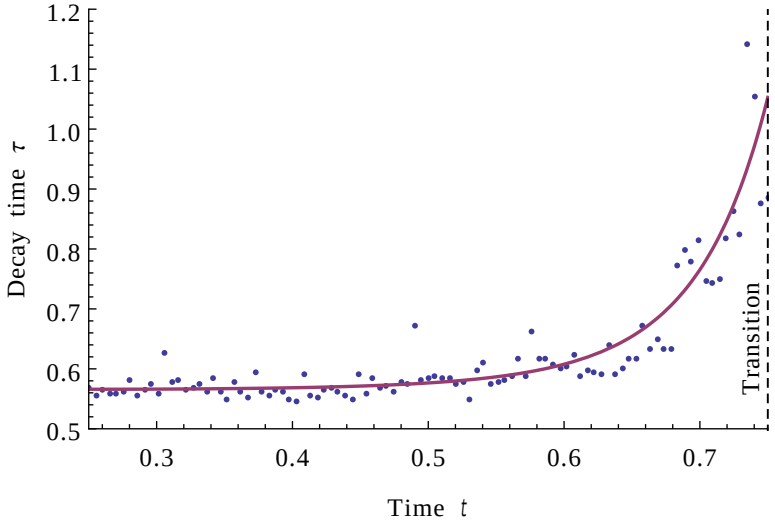

(b) The decay time is calculated from the AR(1) coefficients.

FIGURE 4.6: The same transition encountered by increasing perturbation in the positive direction in Fig. 4.5(a) is simulated with stochastically increasing forcing, modelled by a positive biased Wigner process which generates the time series of Fig. 4.6(a). An auto-regression process $\boldsymbol{E}(t+\Delta t) = \boldsymbol{E}(t)\alpha + \boldsymbol{\epsilon}(t)$ is fit to the time series and the decay time of small perturbation from the stationary points can be estimated by the AR(1) coefficient α shown in Fig. 4.6(b). The response of the system to perturbations slows as it approaches a transition which can be verified by the Kendall coefficient $\tau = 0.63$.

4.5 Discussion

Sec. 3.5 establishes that homeostasis emerges as a general property of this coupled life-environment system incorporating a diverse biota with individual increasing or decreasing effects on each environmental variable independently. In Sec. 3.5.1, we have shown that the density of fixed points is Gaussian with the magnitude of external perturbing forces; the expected number of stable biotic configurations is maximised when there is no perturbing force, and decreases with increasing perturbation magnitude. Sec. 3.5.2 goes on to classify the fixed points of a two environmental variable model, such as into those with stable, unstable and spiralling behaviour. We find their behaviour to be independent of the perturbing force and their distribution to be analytically tractable.

In a similar fashion, this chapter finds a diversity of transition types emerge from the model. We observe irreversibility in the sense that forcing applied to one environmental variable may destabilise a fixed point in such a way that the fixed point cannot be recovered by variations in perturbation to one or any dimensions. Additionally, the emergence of cusp bifurcations, where changes to the state of the environment caused by one forcing dimension may be smooth or abrupt depending on the state of another, highlights the novelty of our model implemented in a multidimensional environment. We have gone on to derive the density of transitions encountered by increases or decreases in perturbing forces. Transitions which occur at local extrema of the effects of the model biota in the space of environmental variables. Such points indicate where there can be no further increase in the effect of the biota to offset further increases in external forcing, and thus the model undergoes a transition. We have found that as the magnitude of the perturbation is increased, the model is always increasingly likely to encounter a transition due to further increases in perturbation. In other words, a strongly perturbed model is relatively unlikely to encounter a transition when the perturbation is decreased, rather than increased.

The relative fragility of a stressed life-environment system to further increases in stress to decreases is intuitive, although it is not clear that real Earth systems behave this way. However, there is evidence of critical slowing down in a range of time series data which undergo abrupt transitions. By examining our model, it is simple to see that the leading eigenvalue of the system Jacobian diverges in the vicinity of transitions. This is a result of transitions occurring at local extrema in the biotic effect function where the gradient approaches zero. Here, eigenvalues tend towards zero, and the decay time of small fluctuations from the fixed point diverge. This can be detected in simulated model time series using the usual method of detrended fluctuation analysis (Dakos et al., 2008). The AR(1) coefficient, which measures the degree of correlation between subsequent points in a time series, appears to diverge in the vicinity of the transition which is consistent with a diverging decay time.

Up to this point, the focus of this thesis has been on the Earth system and a thorough investigation of our abstract ecosystem model. There are many extensions one could propose to bring this model in line with much of the preceding research which focuses on cellular automata implementations, building ecosystems from networks, and thermodynamic consequences of homeostasis. These possibilities are discussed in detail in the closing chapter of this thesis. However, before for these models to be informative to the same extent as our current model, methods must be created, developed and refined to enable a thorough, general and transparent analysis to be performed. The following three chapters instead focus on establishing, and developing the techniques required to analyse this type of abstract model when advanced to spatial or network forms, along with promising thermodynamic principles for treating non-equilibrium systems.

Chapter 5

Renormalisation of cellular automata with absorbing states

5.1 Introduction

Watson and Lovelock's (1983) Daisyworld is commonly referred to as a zero dimensional model; there are no spatial effects which would limit the range of daisy seeding or heat diffusion. Embedding this model in space and introducing these effects has been a popular avenue of research and typically proceeds in one of two possible directions. Adams et al. (2003) and Adams and Carr (2003) model a planet with longitudinal uniformity by implementing a one-dimensional Daisyworld, subject to a greater insolation close to the centre (equator) than the ends (poles). The implementation is simply a spatial form of the differential equations encountered in Sec. 2.2, modified to allow diffusion of heat. The alternative is to study the interaction between the spatial distribution of daisies and the temperature field by using local seeding dynamics which we introduced in Sec. 2.7 such as von Bloh et al. (1997, 1999) along with Ackland et al. (2003) who additionally uses a non-uniform insolation. These studies employ cellular automata (CA), a staple component of many fields of computational research into spatially embedded interacting systems.

Cellular automata consist of a coarse grid of cells with both internal and local updating procedures; the future state of a point in the CA is determined by its state, and the state of surrounding cells in some small neighbourhood by discrete deterministic or probabilistic updating steps. They have become a popular tool across a broad range of sciences and have found a plethora of applications from modelling free-way traffic flow to neuron activity (Nagel and Schreckenberg, 1992; Ilachinski, 2001). Cellular automata can be divided into two main types. Deterministic cellular automata have transition matrices where transition probabilities are all $\in \{0, 1\}$. This means that at each update there is only one possible outcome, and the time evolution of the CA is deterministic. Elementary cellular automata (ECA) are the simplest deterministic models as they are binary, having only two possible states for each cell, denoted 0 or 1. Probabilistic cellular automata (PCA) have a range of possible outcomes at each update whose relative likelihood is given by the transition matrix with entries $\in [0, 1]$.

A wealth of research surrounds the study of equilibrium probabilistic cellular automata, most famously the Ising model, a ferromagnetic system. The study of equilibrium systems has a rich history and there are a large number of tools available to their analysis. However there is a relative paucity of work related to non-equilibrium CA in two-dimensions or more. Such models are often governed by dynamics determined by the author's intuition and attempts to characterise key microscopic mechanisms responsible for large-scale emergence and rarely have properties to exploit such as some well defined free-energy. There are no physical laws that must be obeyed on the large scale to inform our analysis and we are left only with our local updating probabilities. This class of model pervades almost all areas of science, from species distribution modelling to generating music (Balzter et al., 1998; Burraston and Edmonds, 2005) and indeed

spatial Daisyworld models (von Bloh et al., 1997, 1999; Ackland et al., 2003). In this chapter we aim to formulate a general framework to tackle such systems, and assess its accuracy and utility against simple examples of cellular automata.

Our previous example of a two-dimensional Daisyworld CA in Sec. 2.7 behaves identically to the zero dimensional Daisyworld when the heat diffusion is large and the temperature field is uniform. However, even in this simple case the Daisy distribution is non-trivial and the strong tendency for daisies of the same colour to cluster leads to large connected clusters which follow a power-law distribution; a class of particularly interesting behaviour commonly described as scale-free such that the large scale, or coarse grained dynamics are identical to the microscopic resulting in nested, or fractal structures, characterised by a power-law distribution and spatial scaling exponent. This type of system garnered a great deal of attention as it became clear that many systems of wildly disparate microscopic origins gave rise to critical behaviour dominated by power-law statistics. In its steady state, Daisyworld behaves identically to a model of directed percolation which will be discussed shortly, although the transition probabilities of Daisyworld are determined endogenously and the fractal appearance emerges without need for further tuning.

While it is tempting to suggest so, not all scale-free behaviours belong to the same class of systems, sharing some form of universal dynamics. There turn out to be a range of inequivalent methods for modelling these systems with distinct properties. For example, forest fires appear to obey a power-law frequency distribution of magnitude, though there are at least two competing descriptions (Hunt, 2008). One path taken is percolation theory which studies the clustering of random graphs and will be the focus of this work. The applicability of this, along with the wealth of research surrounding it, to a range of power-law behaviours is discussed by Hunt (2009). Statistical physics aims to relate the micro and macroscopic worlds and is perhaps at its most transparent in renormalisation theory which considers the change in the physical description of a system under some blocking, or scale transformation. For this reason, Licata (2010) describes the method aptly as a 'mathematical zoom lens' (for an introduction to renormalisation theory, see (Goldenfeld, 1992, Ch. 9)).

In this chapter we construct a real-space renormalisation method, and demonstrate its utility with application to a binary PCA with an absorbing state in one- and two-dimensions, models which represent a $1+1$ and $2+1$ dimensional directed percolation such as that shown in Fig. 5.1. This is a natural progression from the work of Edlund and Nilsson Jacobi (2010) who present a concise and systematic approach to renormalisation of such one-dimensional CA, which itself builds on previous work labelled as application of a dynamically driven renormalisation group (Tomé and de Oliveira, 1997; De Oliveira and Satulovsky, 1997). It is known that in applying some scale transformation to a CA, we do not guarantee the existence of an exact coarse grained dynamics, one which exactly incorporates all possible microscopic transitions. In such instances it is natural

to choose the best fit by way of least-squares. The important contribution of Edlund and Nilsson Jacobi (2010) has been to show that what previously was referred to as the dynamically driven renormalisation group amounts to weighting this least-squares fit by the steady-state probability distribution such that error is concentrated on the transitions between relatively uncommon states.

We begin by briefly reintroducing the novel matrix representation of PCA transitions and states as used by Edlund and Nilsson Jacobi (2010) in Sec. 5.2. The corresponding real-space renormalisation algorithm is sufficiently general to treat PCA with any dimensionality or alphabet and is detailed in Sec. 5.3. As with previous work, coarse-grained degrees of freedom are insufficient to provide exact renormalisations, and therefore the best-fit renormalisation is found by the method of least-squares, which may be weighted by estimates of the stationary probability distribution of states whose computation is detailed in Sec. 5.4, along with an interpretation of the least squares fitting in terms of the dynamically driven renormalisation group (DDRG) in Sec. 5.5. The results of the renormalisation are given in Sec. 5.6 and Sec. 5.7 where we demonstrate the renormalisation flow in three dimensions along with the critical transition probabilities for steady-state probability distribution estimates using mean-field and nearest-neighbour correlations.

5.2 Cellular automata

We consider two-state PCA where lattice sites are denoted by either $\sigma \in \{\square, \blacksquare\}$. Throughout this work, notation will be provided for the two-dimensional renormalisation where the method remains general since the reduction to a single dimension is trivial. We denote the state of $n \times n$ blocks of lattice sites by the notation $\boldsymbol{\sigma}^n \in \{\square, \blacksquare\}^n$ whose probability distribution is held in the state vector $\boldsymbol{s}^n$ of length 2^{n^2}.

Transition probabilities of single site σ are governed by the $k \times k$ block of sites in the previous interval and the transition matrix $\mathbf{P}$, which encompasses the microscopic dynamics of the PCA. $\mathbf{P}$ is a $2 \times 2^{k^2}$ transition matrix with elements $P(\sigma|\boldsymbol{\sigma}^k)$ giving the transition probability from a state $\boldsymbol{\sigma}^k \in \{\square, \blacksquare\}^k$ to the state $\sigma \in \{\square, \blacksquare\}$. In our notation so far, we can compute the single site probability distribution by applying the transition matrix

$$\boldsymbol{s}(t+1) = \mathbf{P}\boldsymbol{s}^k(t). \tag{5.1}$$

In order to update larger blocks of cells, we define a general transition probability matrix $\mathbf{P}_n$ which gives the probability vector $\boldsymbol{s}^n$ from its neighbourhood, $\boldsymbol{s}^{k+n-1}$. For example, the transition matrix for a $k = 2$ PCA which maps blocks of cells of size $n+1$ to n is

$$\mathbf{P}_n\left(\boldsymbol{\sigma}^n \middle|\, \boldsymbol{\sigma}^{n+1}\right) = \prod_{i,j=1}^{N} \mathbf{P}\left(\sigma_{i,j}^n \middle| \begin{matrix} \sigma_{i,j}^{n+1} & \sigma_{i+1,j}^{n+1} \\ \sigma_{i,j+1}^{n+1} & \sigma_{i+1,j+1}^{n+1} \end{matrix}\right). \tag{5.2}$$

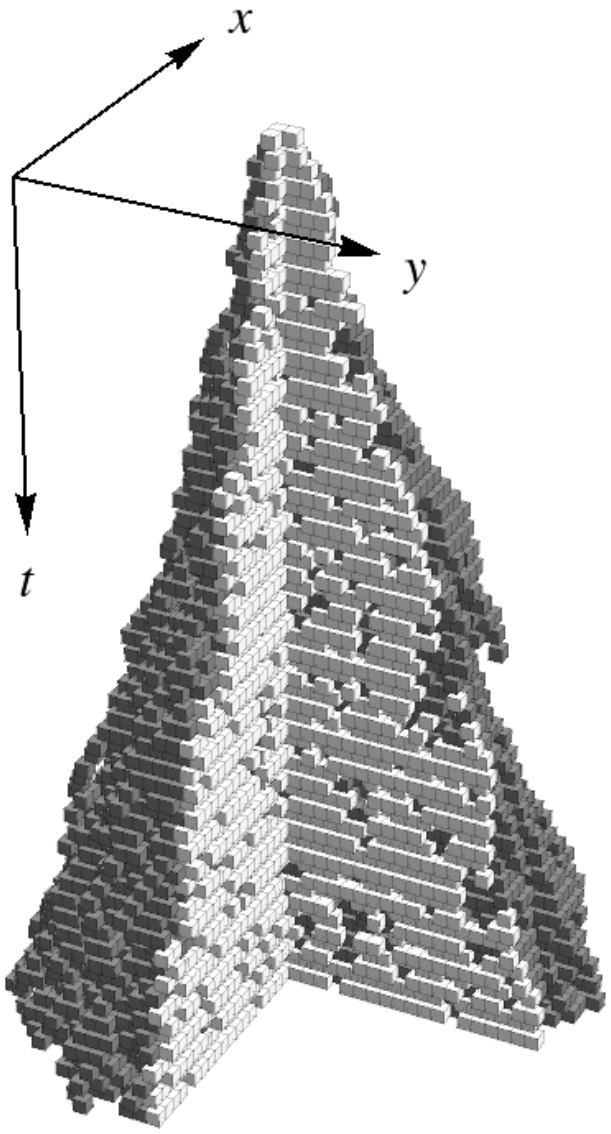

FIGURE 5.1: A model of directed percolation initialised from a single seed in $2+1$ dimensions shown here above the critical point, where clusters grow and percolate through time. The lighter region shows a cross-section, indicating the cluster is not compact but contains holes.

In this work, we study PCA with an absorbing state, where the transitions from the vacuum state (all sites □) to any other state do not occur. In a binary PCA, such a model is expected to fall into the percolation universality class whose behaviour is illustrated by Fig. 5.1. As such, the large scale critical behaviour of these models is postulated to be universal, governed by a fixed point in the renormalisation flow which is attractive in the critical plane, precisely as illustrated for the case of one-dimensional PCA by Weaver and Prügel-Bennett (2014).

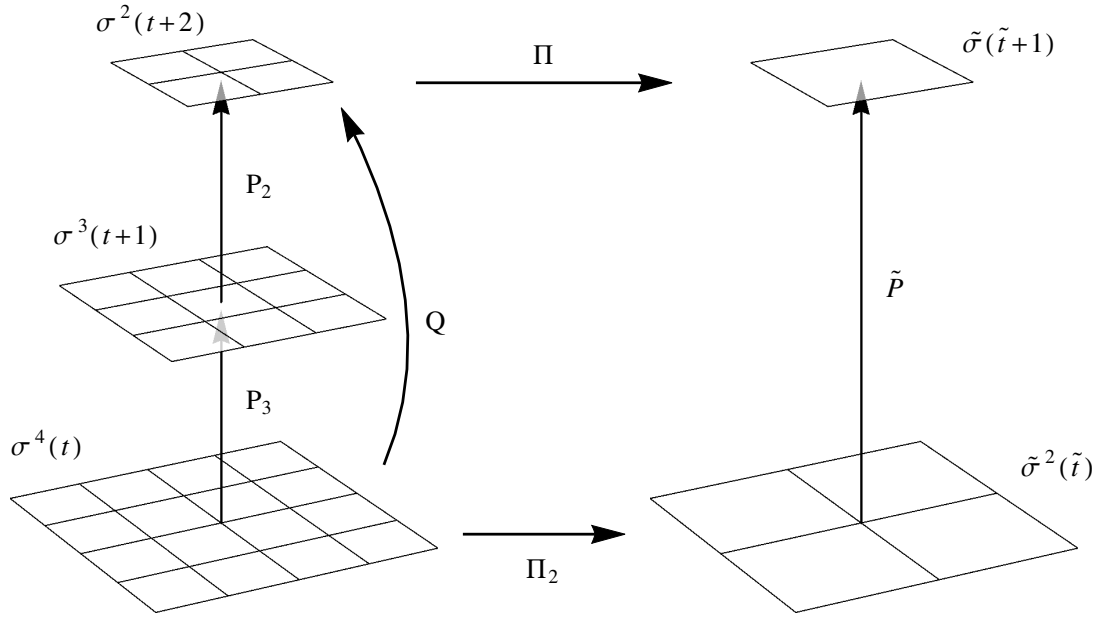

FIGURE 5.2: The renormalisation algorithm for the $k = 2$ two-dimensional lattice. In the case of an exact renormalisation, the above matrices must commute; carrying out two time steps through $\mathbf{Q} = \mathbf{P}_2\mathbf{P}_3$ followed by coarse graining by $\mathbf{\Pi}$ should be identical to beginning with the projection $\mathbf{\Pi}_2$, and applying the appropriate coarse-grained dynamics $\tilde{\mathbf{P}}$.

5.3 Renormalisation algorithm

The scale transformation begins with the coarse gaining projection $\mathbf{\Pi}$ which transforms the probability vector of blocks of N cells, $\boldsymbol{s}^N$, onto a single re-normalised cell $\tilde{s}$

$$\tilde{s} = \mathbf{\Pi}\ \boldsymbol{s}^N \tag{5.3}$$

where $\mathbf{\Pi}$ is a $2 \times 2^{N^2}$ matrix with elements $\Pi(\tilde{\sigma}|\boldsymbol{\sigma}^N)$ giving the probability of projecting the state $\boldsymbol{\sigma}^N \in \{\square, \blacksquare\}^N$ onto $\tilde{\sigma} \in \{\square, \blacksquare\}$. In this work, we apply only the smallest trivial coarse graining $N = 2$. This spatial renormalisation is accompanied by a stroboscopic renormalisation in time where we advance N updates at a time, $\tilde{t} = Nt$. Two updates can be carried out simply by taking the product of transition matrices

$$\mathbf{Q} = \mathbf{P}_2\mathbf{P}_{1+k}. \tag{5.4}$$

A coarse-grained dynamics $\tilde{\mathbf{P}}$ which, if it exists, satisfies the condition

$$\mathbf{\Pi}\ \mathbf{Q} = \tilde{\mathbf{P}}\ \mathbf{\Pi}_k. \tag{5.5}$$

where $\mathbf{\Pi}_k$ coarse grains $k \times k$ blocks of cells. Eq. (5.5) amounts to a statement that the coarse grained dynamics applied to the coarse grained projection of $\mathbf{s}^{2k}$ should be consistent with the microscopic dynamics, illustrated by Fig. 5.2.

5.4 Stationary probability distribution

When an exact solution to Eq. (5.5) does not exist, there is no macroscopic transition matrix which faithfully represents all allowed microscopic transition probabilities. In this instance we find the the most representative renormalisation by minimising the summed square error between $\mathbf{P}$ and $\tilde{\mathbf{P}}$. This is achieved by finding the least-squares solution to Eq. (5.5). Previous work has shown the quality of the resulting renormalisation to be improved by incorporating the dynamics of the PCA, $\mathbf{P}$, referred to as the dynamically driven renormalisation group (Vespignani et al., 1996). In essence, the dynamics are utilised by taking a least-squares solution weighted by the stationary probability distribution of the $2k \times 2k$ cells at the microscopic level, shown in Fig. 5.2. The steady-state probability distribution satisfies the condition

$$\mathbf{P}_m \boldsymbol{s}_{\mathrm{st}}^{m+k-1} = \boldsymbol{s}_{\mathrm{st}}^{m} \tag{5.6}$$

where $\boldsymbol{s}_{\mathrm{st}}^{m}$ is the steady-state probability distribution of blocks of $m \times m$ lattice sites from which the larger distribution, $\boldsymbol{s}_{\mathrm{st}}^{m+k-1}$, must be extrapolated in a consistent way. For example, the first improvement which can be made over a mean field approximation considers only nearest-neighbour lattice correlations, illustrated in Fig. 5.3. Consistency with a $m = 2$ distribution $\boldsymbol{s}^2$ requires that the distribution $\boldsymbol{s}^3$ satisfies the constraint

$$\sum_{\sigma_1,\sigma_2,\sigma_3,\sigma_6,\sigma_9} s^3 \begin{pmatrix} \sigma_1 & \sigma_2 & \sigma_3 \\ \sigma_4 & \sigma_5 & \sigma_6 \\ \sigma_7 & \sigma_8 & \sigma_9 \end{pmatrix} = s^2 \begin{pmatrix} \sigma_4 & \sigma_5 \\ \sigma_7 & \sigma_8 \end{pmatrix} \tag{5.7}$$

for all configurations of $\boldsymbol{\sigma}^2$, a total of $2^4 = 16$ constraints of which 12 are unique (not related by lattice symmetry). Terms which are summed out are highlighted for clarity. These constraints may be written more conveniently as

$$\mathbf{C}\boldsymbol{s} = \boldsymbol{d} \tag{5.8}$$

where $\mathbf{C}$ is a sparse matrix with non-zero elements corresponding to the sum in the LHS of Eq. (5.7) and $\boldsymbol{d}$ is a vector with entries corresponding to the constraints in the RHS of Eq. (5.7). Extrapolation of any one-dimensional probability distribution is trivial as the constraints have a unique solution which may be found by applying Bayes' theorem

$$s(\{x_1, \cdots, x_K\}) = \frac{\prod_{i=1}^{K-M+1} s(\{x_i, \cdots, x_{M+i-1}\})}{\prod_{i=1}^{K-M} s(\{x_{i+1}, \cdots, x_{M+i-1}\})}. \tag{5.9}$$

However in two dimensions, there is no unique solution. We therefore search for the maximum entropy solution; the distribution which maximises the information entropy while satisfying any constrains. The rationale for this choice is that it minimizes the amount of prior information built into the distribution alongside our constraints (Jaynes,

1957). The Lagrangian for this problem contains three components; the entropy of the distribution $\boldsymbol{\sigma}$, the constraints of Eq. (5.8) multiplied by one set of Lagrange multipliers $\boldsymbol{\lambda}$ and the normalisation condition multiplied by an additional Lagrange multiplier μ.

$$\begin{aligned}\mathcal{L}(\boldsymbol{\lambda}, \mu, s(\boldsymbol{\sigma})) = &- \sum_{\boldsymbol{\sigma}\in\boldsymbol{\sigma}'} D(\boldsymbol{\sigma}) s(\boldsymbol{\sigma}) \log(s(\boldsymbol{\sigma})) \\ &+ \sum_i \lambda_i \left(\sum_{\boldsymbol{\sigma}\in\boldsymbol{\sigma}'} s(\boldsymbol{\sigma}) C_i(\boldsymbol{\sigma}) - d_i \right) \\ &+ \mu \left(\sum_{\boldsymbol{\sigma}\in\boldsymbol{\sigma}'} D(\boldsymbol{\sigma}) s(\boldsymbol{\sigma}) - 1 \right) \end{aligned} \tag{5.10}$$

where $s(\boldsymbol{\sigma})$ and $D(\boldsymbol{\sigma})$ are the probability of configuration $\boldsymbol{\sigma}$ and its degeneracy (the number of equivalent constraints determined by the lattice symmetry) respectively, $\boldsymbol{C}(\boldsymbol{\sigma})$ and $\boldsymbol{d}$ contain the constraints and $\boldsymbol{\lambda}$ and μ are Lagrange multipliers, and sums are taken only over the unique configurations, $\boldsymbol{\sigma}'$, not related by lattice symmetry. We will now show that minimizing the Lagrangian provides the maximum entropy distribution which satisfies our constraints. Taking the derivative wrt. $s(\boldsymbol{\sigma})$

$$\frac{\mathrm{d}\mathcal{L}}{\mathrm{d}s(\boldsymbol{\sigma})} = -D(\boldsymbol{\sigma}) \log(s(\boldsymbol{\sigma})) - D(\boldsymbol{\sigma}) + \boldsymbol{C}(\boldsymbol{\sigma})\boldsymbol{\lambda} + \mu s(\boldsymbol{\sigma}) = 0 \tag{5.11}$$

Rearranging for $s(\boldsymbol{\sigma})$ gives the maximum entropy distribution as a function of the Lagrange multipliers and constraints

$$s(\boldsymbol{\sigma}) = \exp\left(\frac{\boldsymbol{C}(\boldsymbol{\sigma})\boldsymbol{\lambda}}{D(\boldsymbol{\sigma})} + \mu - 1 \right). \tag{5.12}$$

The normalisation parameter μ can be replaced simply by dividing through by the partition function, $Z(\boldsymbol{\lambda})$

$$s(\boldsymbol{\sigma}) = \frac{1}{Z(\boldsymbol{\lambda})} \exp\left(\frac{\boldsymbol{C}(\boldsymbol{\sigma})\boldsymbol{\lambda}}{D(\boldsymbol{\sigma})} \right) \tag{5.13}$$

where

$$Z(\boldsymbol{\lambda}) = \sum_{\boldsymbol{\sigma}\in\boldsymbol{\sigma}'} D(\boldsymbol{\sigma}) \exp\left(\frac{\boldsymbol{C}(\boldsymbol{\sigma})\boldsymbol{\lambda}}{D(\boldsymbol{\sigma})} \right). \tag{5.14}$$

Substituting Eq. (5.13) into Eq. (5.10) gives the dual-objective

$$\begin{aligned}\mathcal{L}(\boldsymbol{\lambda}) &= - \sum_{\boldsymbol{\sigma}\in\boldsymbol{\sigma}'} D(\boldsymbol{\sigma}) s(\boldsymbol{\sigma}) \left(\frac{\boldsymbol{C}(\boldsymbol{\sigma})\boldsymbol{\lambda}}{D(\boldsymbol{\sigma})} - \log(Z(\boldsymbol{\lambda})) \right) \\ &\quad + \sum_{\boldsymbol{\sigma}\in\boldsymbol{\sigma}'} s(\boldsymbol{\sigma}) \boldsymbol{C}(\boldsymbol{\sigma})\boldsymbol{\lambda} - \boldsymbol{d}\boldsymbol{\lambda} \\ &= \log\left(Z(\boldsymbol{\lambda})\right) - \boldsymbol{d}\boldsymbol{\lambda} \end{aligned} \tag{5.15}$$

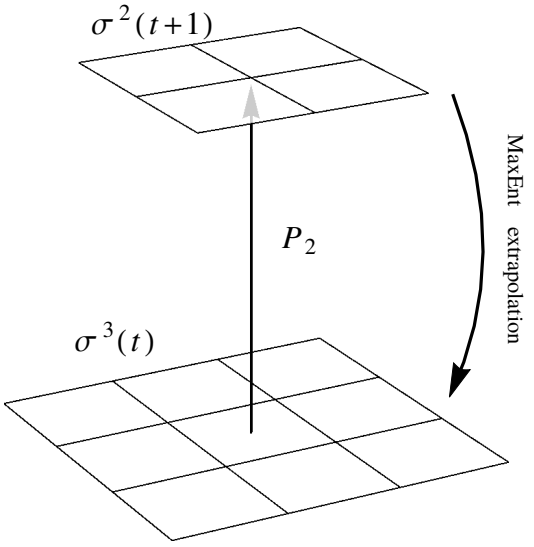

FIGURE 5.3: A summary of the algorithm to compute the steady-state distribution σ^2_{st} which incorporates nearest-neighbour correlations. The solution in unchanged by an extrapolation $\text{ext}(\sigma^2_{\text{st}} \to \sigma^3_{\text{st}})$ followed by the transition matrix $\mathbf{P}_3$, the solution of Eq. (5.6).

whose minimum can be shown to satisfy the constraints

$$\begin{aligned}\frac{\mathrm{d}\mathcal{L}}{\mathrm{d}\lambda_i} &= \frac{1}{Z(\boldsymbol{\lambda})}\frac{\mathrm{d}Z}{\mathrm{d}\lambda_i} - d_i \\ &= \frac{1}{Z(\boldsymbol{\lambda})}\sum_{\boldsymbol{\sigma}\in\boldsymbol{\sigma}'} C_i(\boldsymbol{\sigma})\exp\left(\frac{\boldsymbol{C}(\boldsymbol{\sigma})\boldsymbol{\lambda}}{D(\boldsymbol{\sigma})}\right) - d_i \\ &= \sum_{\boldsymbol{\sigma}\in\boldsymbol{\sigma}'} C_i(\boldsymbol{\sigma})s(\boldsymbol{\sigma}) - d_i = 0. \end{aligned} \tag{5.16}$$

With the ability to extrapolate small blocks of cells to larger blocks we can solve Eq. (5.6) either through iteration or using normal root-finding methods. As with previous studies, we expect the accuracy of the renormalisation to improve with the quality of our estimate of the steady-state probability distribution $\boldsymbol{s}_{\text{st}}$ which can be improved by accounting for longer range correlations in Eq. (5.6). Increasing the range, m, has been shown to give more representative weightings for the renormalisation and provides improvements in estimates of physical quantities such as the critical scaling exponent (Weaver and Prügel-Bennett, 2014). As expected, increasing m in our two-dimensional implementation is very much limited by a computational ceiling, but also $m > 3$ results in Eq. (5.8) becoming rank-deficient and extrapolation cannot be performed in a consistent way as it can for $m \leq 3$. We limit the study to the cases of $m = 1$ mean-field and $m = 2$ nearest-neighbour correlations along with the unweighted $m = 0$ case.

5.5 Least-squares

Having determined the stationary-distribution we seek a weighted least-squares solution to Eq. (5.5), achieved by defining the square error weighted by the probability vector $\boldsymbol{s}_{\mathrm{st}}^{2k}$

$$\mathcal{L}(\tilde{\mathbf{P}}) = \mathrm{Tr}\left(\boldsymbol{\Pi}\mathbf{Q} - \tilde{\mathbf{P}}\boldsymbol{\Pi}_k\right)\mathbf{D}\left(\boldsymbol{\Pi}\mathbf{Q} - \tilde{\mathbf{P}}\boldsymbol{\Pi}_k\right)^{\top} \tag{5.17}$$

where $\mathbf{D}$ is a diagonal matrix with the stationary probability distribution $\boldsymbol{s}_{\mathrm{st}}^{2k}$ along the diagonal. The solution is found where the derivative with respect to $\tilde{\mathbf{P}}$ goes to zero

$$\frac{\mathrm{d}\mathcal{L}}{\mathrm{d}\tilde{\mathbf{P}}} = 2\tilde{\mathbf{P}}\boldsymbol{\Pi}_k\mathbf{D}\boldsymbol{\Pi}_{k\top} - 2\boldsymbol{\Pi}\mathbf{Q}\mathbf{D}\boldsymbol{\Pi}_{k\top} = 0. \tag{5.18}$$

Solving for $\tilde{\mathbf{P}}$ we find

$$\begin{aligned}\tilde{\mathbf{P}} &= \boldsymbol{\Pi}\mathbf{Q}\mathbf{D}\boldsymbol{\Pi}_{k\top}\left(\boldsymbol{\Pi}_k\mathbf{D}\boldsymbol{\Pi}_{k\top}\right)^{-1} \\ &= \boldsymbol{\Pi}\mathbf{Q}\mathbf{D}^{\frac{1}{2}}\left(\boldsymbol{\Pi}_k\mathbf{D}^{\frac{1}{2}}\right)^{\dagger}\end{aligned} \tag{5.19}$$

where $\mathbf{A}^{\dagger}$ denotes the pseudo inverse $\mathbf{A}^{\top}(\mathbf{A}\mathbf{A}^{\top})^{-1}$.

5.6 Renormalisation of one-dimensional CA

Renormalisation flow refers to the change of the renormalised transition probabilities through the space of possible values across successive coarse-graining. The coarse graining projection, $\boldsymbol{\Pi}$, for models of percolation must ensure the vacuum state contains no particles in order to preserve it as an absorbing state. This can be achieved by using

$$\boldsymbol{\Pi} = \begin{array}{c|cccc} & \square\square & \square\blacksquare & \blacksquare\square & \blacksquare\blacksquare \\ \hline \square & 1 & 0 & 0 & 0 \\ \blacksquare & 0 & 1 & 1 & 1 \end{array} \tag{5.20}$$

which maps blocks of cells to the vacuum state, $\square$, if any only if they contain no particles, $\blacksquare$.

To visualise the renormalisation group flow we first consider the simplest case, $k = 2$ PCA, whose transition matrix $\mathbf{P}$ is given by Eq. (5.29). This is a convenient choice as the renormalisation has only three parameters and we can plot the flows in 3-dimensions. $k > 2$ is not so trivially visualised as the transition matrix is generalised to PCA with larger updating neighbourhoods by the exponentially larger set of parameters $p_1, \ldots, p_{2^k-1}$.

$$\mathbf{P} = \begin{pmatrix} & \square\square & \square\blacksquare & \blacksquare\square & \blacksquare\blacksquare \\ \square & 1 & 1-p_1 & 1-p_2 & 1-p_3 \\ \blacksquare & 0 & p_1 & p_2 & p_3 \end{pmatrix} \tag{5.21}$$

Domany-Kinzel PCA have the additional constraint of being symmetrical, such that they lie on the $p_1 = p_2$ plane. A previous renormalisation study of this plane is provided by Tomé and de Oliveira (1997). More generally, symmetrical CA with $k > 2$ are classified as models of the contact process (Liggett, 1985, Ch. 6). Providing the coarse-graining projection is also symmetric, the renomalisation flow is across the surface of the plane as shown in Fig. 5.4(a). Renormalisation flow is into one of the two stable fixed points corresponding to the vacuum and dense states, separated by the critical line. Flow along this line is towards a third, unstable fixed point; the critical point.

Compact directed percolation (CDP) adds a second absorbing state, ensured by $p_3 = 1$. These conditions ensure that connected clusters are compact; they contain no holes. In an update, a cluster of particles will increase or decrease in size with simply derived probabilities and an analytical approach to the problem is similarly straightforward Domany and Kinzel (1984); Essam (1999). Unlike the Domany-Kinzel plane, renormalisation flow is not always in the $p_3 = 1$ plane; the projection given by Eq. (5.20) does not preserve the dense absorbing state. The components of renormalisation flow across the plane are shown in Fig. 5.4(b). This plane usefully highlights qualitative differences between approximations with $M < 2$ where cell correlations are not resolved, and $M \geq 2$; the line $p_2 = 1 - p - 1$ defines a set of CA with parity $(0 \to 1)$ symmetry, and separates flow between the dense and vacuum states. It is comprised of points which become fixed when correlations are resolved in our estimate of the steady-state distribution.

The final plane of interest is the critical plane, which divides renormalisation flow between the fixed points corresponding to the dense and vacuum states. The unique feature of this plane is that while in the previous two planes, flow is repelled from some critical line, renormalisation in the critical plane is attracted to a single critical point. The implications of this are profound; PCA which lie in the critical plane may be differentiated microscopically, but share identical large-scale dynamics, including properties such as their spatial scaling characteristics. This phenomena is responsible for a wide range of systems without common microscopic dynamics sharing many characteristics at criticality. We find it illustrative to show the intersection of the critical plane with the Domany-Kinzel, and CDP planes, given by Fig. 5.5. We find only one fixed point in the renormalisation flow of $2 \leq k \leq 5$ which supports the DP hypothesis which would require PCA with any k to share large scale critical behaviour.

As illustrated by Fig. 5.5, PCA in the critical plane are universal in that they share the large-scale behaviour of the critical point of the renormalisation transformation. As we approach the critical point, the behaviour of certain physical quantities, such as the

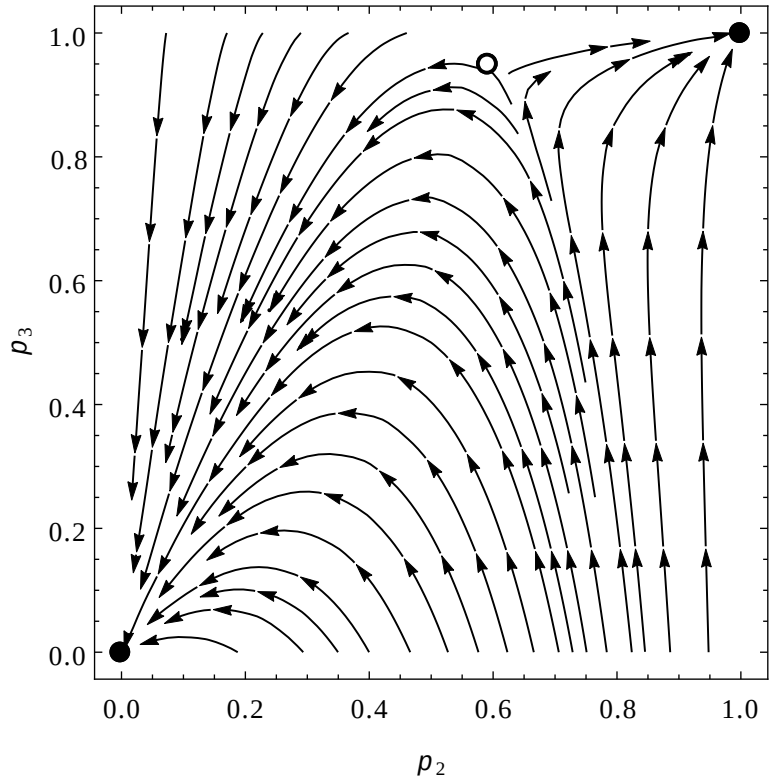

(a) The Domany-Kinzel plane, $p_1 = p_2$

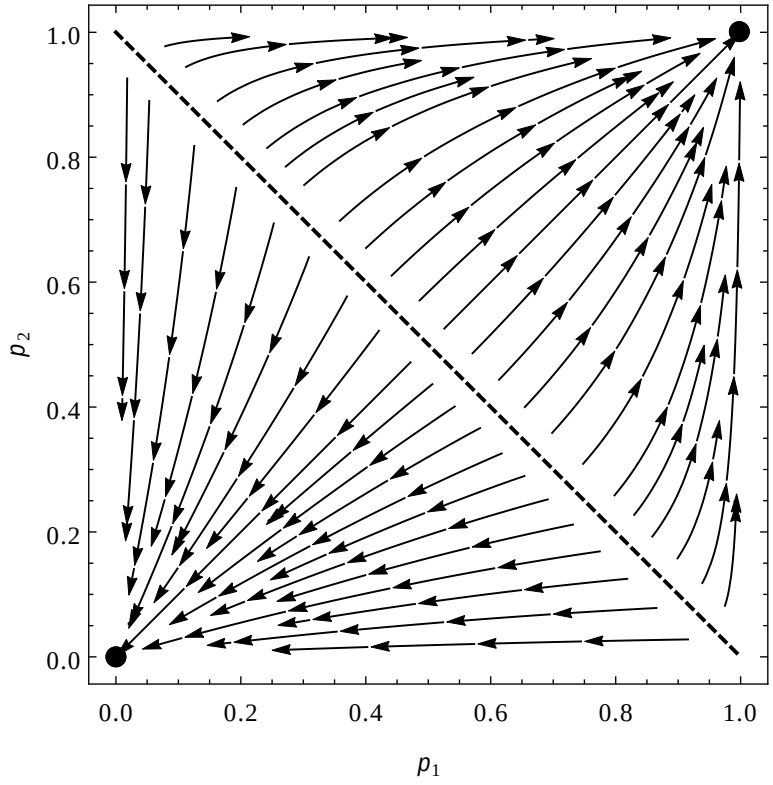

(b) The Compact-Directed Percolation plane, $p_3 = 1$

FIGURE 5.4: The flow of the $k = 2$ PCA, given by Eq. (5.29) through successive renormalisations. Solid points indicate stable fixed points while the open point indicates the critical point of the renormalisation; an unstable fixed point where successive scale transformations do not alter the dynamics. Additionally, the dashed line of Fig. 5.4(b) indicates a set of unstable fixed points.

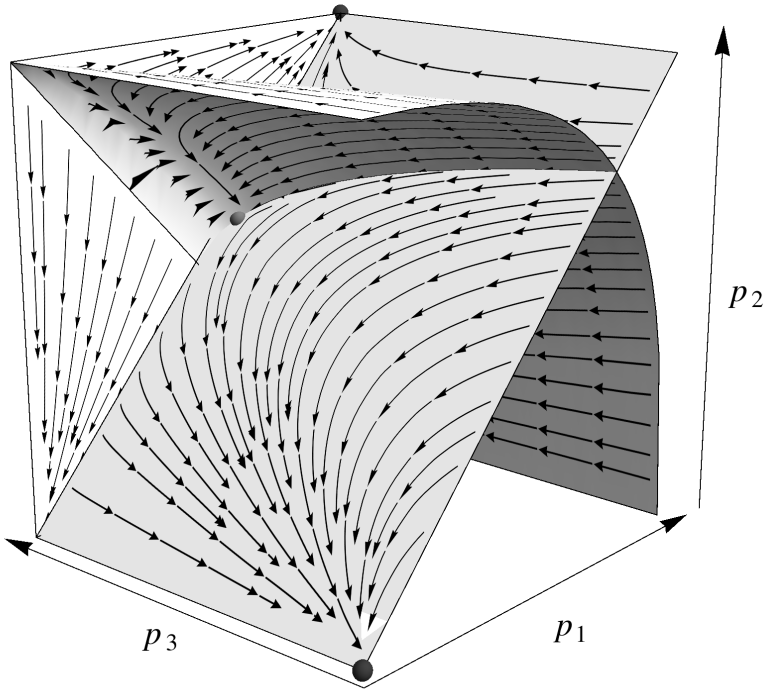

FIGURE 5.5: The intersection of the Domany-Kinzel and CDP planes with the critical plane. Renormalisation flow in the critical plane is towards the fixed point of the renormalisation group. The plane therefore describes a set of PCA which differ in their microscopic rules, but share identical large-scale behaviour. In this sense, they differ only in their irrelevant parameters which vanish on the large scale. The two darker points mark stable fixed points at the corners of the cube corresponding to the vacuum and dense states. The lighter point marks the unstable critical point.

size of percolating clusters, are governed by critical exponents. With successive renormalisation, we expect the spatial correlation length, ξ, to decrease on our renormalised length-scale by a factor N.

$$\tilde{\xi} = \frac{\xi}{N}. \tag{5.22}$$

Clearly, if clusters have finite size on the microscopic scale, they will vanish with repeated renormalisation. However, as we approach the critical plane, the spatial correlation length diverges as

$$\xi \propto |\mathbf{P} - \mathbf{P}^*|^{-\nu}. \tag{5.23}$$

Linearising the renormalisation transformation in the vicinity of the fixed point

$$\tilde{\mathbf{P}} \approx \mathbf{P}^* + \mathbf{J}(\mathbf{P}^*)(\mathbf{P} - \mathbf{P}^*) + \mathcal{O}(\mathbf{P} - \mathbf{P}^*)^2 \tag{5.24}$$

where $\mathbf{J}(\mathbf{P}^*)$ is the Jacobian of the renormalisation around the fixed point. On the large scale, after many iterations of Eq. (5.24), the largest eigenvalue of $\mathbf{J}(\mathbf{P}^*)$, λ, dominates and the increase in the distance $|\mathbf{P} - \mathbf{P}^*|$ in a single coarse graining on the large scale

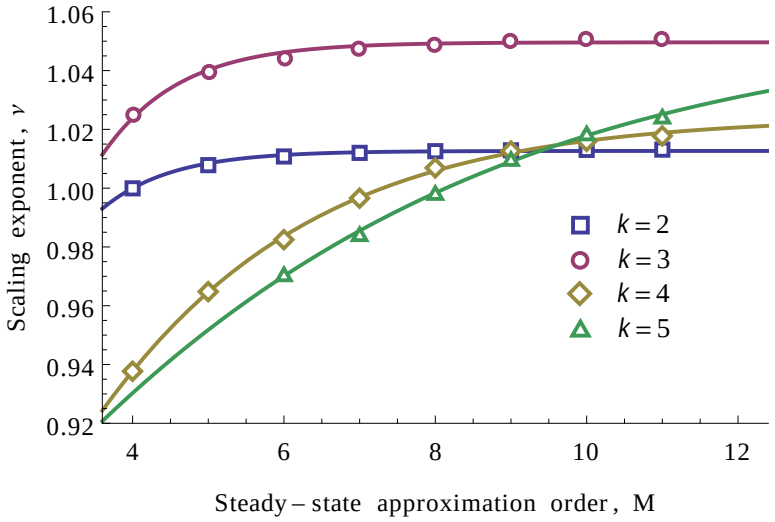

FIGURE 5.6: Calculation of the critical exponent from Eq. (5.27) around the fixed point of the renormalisation, along with the best-fit curve $\nu = \nu_\infty - ae^{-bM}$. Estimates of the critical exponent of the percolation universality class are refined with increasing M, the range of correlations considered in estimating the model steady-state probability distribution. Importantly, the limit $M \to \infty$ does not converge to the experimental value of 1.097 (Hinrichsen, 2006) for all k.

k	ν
2	1.013
3	1.050
4	1.024
5	1.056
expt.	*1.097* (Hinrichsen, 2006)

TABLE 5.1: Estimates of the critical exponent of the directed percolation universality class for a range of k in the limit $M \to \infty$, along with the experimental value. Odd and even values of k appear to approach the experimental value differently

may be written to first order

$$|\tilde{\mathbf{P}} - \mathbf{P}^*| \approx \lambda\,|\mathbf{P} - \mathbf{P}^*|\,. \tag{5.25}$$

From Eq. (5.23), the change in the mean cluster size is

$$\tilde{\xi} = \xi\lambda^{-\nu}. \tag{5.26}$$

Equating Eq. (5.22) to Eq. (5.26) and solving for ν gives

$$\nu = \frac{\log N}{\log \lambda}. \tag{5.27}$$

Calculation of the critical exponent is shown for a range of k and M in Fig. 5.6 along with estimates of the large M limit in Table 5.1. Agreement with the experimental

value $\nu = 1.0969$ (Hinrichsen, 2006) is improved by two factors. Firstly, estimates of the critical exponent are improved by producing more accurate estimates of the steady-state probability vector $\mathbf{s}_{\text{st.}}^{2k}$. By increasing M, the range of correlations considered in this approximation, we refine estimates of the critical exponent, as shown in Fig. 5.6. Furthermore, increasing the rule-space of the CA by increasing k also provides improvements but with two important caveats; odd and even k appear to converge separately, and improvement is only found for sufficiently large M. Our most refined estimate of the critical exponent is found from $k = 5$ extrapolated to $M \to \infty$ (by fitting to the function $\nu = \nu_\infty - ae^{-bM}$) which gives $\nu = 1.056$, within 4% of the experimental value.

While we expect large k to converge to the measured critical exponent, this convergence appears to be slow, and is in practice made difficult by two factors. Firstly, the computation required in the calculation of the steady-state distribution increases exponentially with k. Additionally, Fig. 5.6 shows that increasing k results in a higher value of M required to see an asymptote. As calculation of the steady-state distribution is also exponential in M, the combination of these factors severely limits further refinement of the estimate.

5.7 Renormalisation of two-dimensional CA

While the method we have presented is entirely general with respect to the size of the updating neighbourhood, k, both the computational and memory costs scale super-exponentially, and quickly become prohibitive. For this reason we limit ourselves to study symmetrical $k = 2$ transition matrices. This contains only five renormalisation parameters not related by lattice symmetry, which we label

$$\mathbf{P} = \begin{array}{c} \\ \square \\ \blacksquare \end{array} \begin{array}{c} \begin{array}{cccccc} \boxplus & \boxplus & \boxplus & \boxplus & \boxplus & \boxplus \end{array} \\ \begin{pmatrix} 1 & 1-p_1 & 1-p_2 & 1-p_2' & 1-p_3 & 1-p_4 \\ 0 & p_1 & p_2 & p_2' & p_3 & p_4 \end{pmatrix} \end{array}. \tag{5.28}$$

A coarse graining projection must be chosen which preserves the vacuum state in order that it be absorbing on all scales, that is to say it must never project a block of cells which contain a ■ onto a □. Therefore, the relevant projection is

$$\mathbf{\Pi} = \begin{array}{c} \\ \square \\ \blacksquare \end{array} \begin{array}{c} \begin{array}{cccccc} \boxplus & \boxplus & \boxplus & \boxplus & \boxplus & \boxplus \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} \end{array}. \tag{5.29}$$

To visualise the flow of the five renormalisation parameters p_1, p_2, p_2', p_3 and p_4 we plot the hyper-surface given by assuming a totalistic transition matrix $p_2 = p_2'$ with the transition probability p_4 taking its critical value in Fig. 5.7. The location of the critical point of the renormalisation group, the transition matrix which is invariant with our

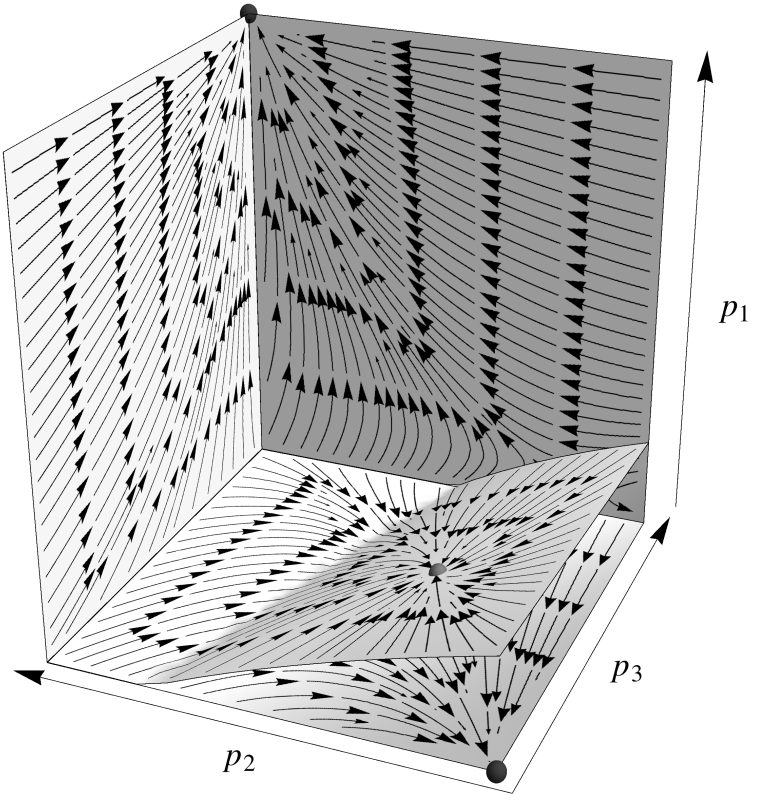

FIGURE 5.7: The unweighted renormalisation flow of the across the surface of four intersecting planes with $p_2 = p_2'$ and $p_4 = 0.4502$. The curved plane is a projection of the critical plane and contains the critical point, illustrating that the point is attractive in this plane. The two darker points mark stable fixed points at the corners of the cube corresponding to the vacuum and dense states. The lighter point marks the unstable critical point.

Approximation	m	p_1	p_2	p_2'	p_3	p_4
Unweighted	0	0.1473	0.2664	0.2714	0.3677	0.4502
Mean-field	1	0.2703	0.4619	0.4667	0.6025	0.7031
Nearest-neighbour	2	0.2650	0.5150	0.4600	0.6888	0.8379

TABLE 5.2: The critical transition probabilities estimated by our renormalisation group method where the renormalisation is unweighted, and weighted by approximations of the steady-state probability distribution considering only mean-field and nearest-neighbour cell-correlations.

scale transformation, is shown for a range of steady-state approximations in Table 5.2. Seemingly a result of our poor estimates of the steady state spin distribution, our estimates of the critical scaling exponent are extremely unsatisfying; 1.72, 3.98 and 4.86 for our unweighted, mean-field and nearest-neighbour schemes respectively to be compared to the numerical estimate of 0.73 (Hinrichsen, 2006).

5.8 Conclusion

We have shown that the renormalisation algorithm used by Edlund and Nilsson Jacobi (2010) is sufficiently general to yield meaningful results when applied to two-dimensional PCA. To provide the dynamically driven component, an estimate of the stationary probability distribution of states, requires a unique extrapolation of a probability distribution over a small block of cells to a large one. In one-dimension, Bayes' theorem is sufficient. We find that refining estimates of the steady-state probability distribution used to weight the renormalisation does not cause calculations of the critical exponent to converge on the experimental value. Rather, PCA with larger k have many more renormalisation parameters and are found to lead to significant improvements in estimates of the critical exponent of the DP universality class.

In two-dimensions there is a hyper-surface of consistent extrapolated parameters and we employ the principle of maximum entropy to choose the configuration with the highest entropy, constrained by the requirement that our smaller probability distribution be recoverable. This has proven effective for $m = 2$ (to incorporate nearest-neighbour correlations into the steady-state distribution) though the method encounters unfortunate complications for the cases of $m = 3$ and $m \geq 4$. For $m = 3$, a steady-state including next-nearest-neighbour correlations, there is a transition to a vacuum steady-state which is in disagreement with the dynamics of the renormalisation. A critical point could not be found when weighted with this approximation as it appears to lie within the range of renormalisation parameters in which the steady-state is expected to be a vacuum. On the other hand, in the case of $m \geq 4$ there is no consistent way to extrapolate the probability distribution consistently; the constraint matrix is rank deficient.

Unlike previous work relating to one-dimensional CA, an absorbing state is not necessary for the existence of large-scale structures. In one-dimension, if large structures are destroyed with a finite probability, there can be no structure on the very largest scales, and therefore only trivial fixed points in the corresponding renormalisation group. Two-dimensional systems do not have this quirk as the presence of a hole in a connected cluster does not destroy it, and indeed they may arise whether or not there exists an absorbing state. However, we find that applying a coarse graining projection where parity symmetry is preserved by probabilistic blocking, does not reveal any interesting fixed points in the renormalisation, rather there is a tendency towards unstructured behaviour on the large scale (i.e.: $p \to \frac{1}{2}$).

Previous work concentrated on determining the scaling behaviour at the critical point, estimated from the leading eigenvalues of the Jacobian at the critical point of the renormalisation (Weaver and Prügel-Bennett, 2014). The quality of this approximation has been shown to be improved by increasing the range of interaction, the size of updating neighbourhood k, and by refining estimates of the steady-state distribution used to weight the renormalisation. For the one dimensional case, these refinements appear to

yield significant improvement. In the two-dimensional case we are severely limited by the computational complexity of the problem and are only able to produce very coarse estimates of the scaling exponent at the critical point, and unable to comment on the improvement achieved by following the process through for larger updating neighbourhoods, or more accurate steady-state representations.

Chapter 6

Preferential attachment in randomly grown networks

6.1 Introduction

Having model elements interact in a plane is the simplest spatial realisation of discrete dynamical systems you may choose. Points are constrained to interact with a small and predictable set of neighbours in the same way that atoms may transmit energy through a solid lattice. Such a process may be generalised onto a more interesting or irregular structure called a network, where connections between interacting components are no longer strictly spatially local. Daisyworld and similar models have found implementations in a range of network topologies which allow daisies to seed over long range (Punithan et al., 2011). In Sec. 2.7 we demonstrated that a mean-field approximation is sufficient to determine the stationary behaviour of a simple spatial Daisyworld, so these studies are more interested in the pattern formation which results from Daisyworld's local seeding rules, and the impact of introducing long-range connections on the spatial species distribution which remains largely impenetrable to any general analysis.

Along with this, networks enjoy very broad interest in several fields. An appreciation that long-range interactions exist in human society, for example, is critical to understanding the spread of epidemics, a phenomena which occurs under a much broader range of parameters on a lattice (Newman, 2002b). Natural systems such as biochemical reaction networks and food webs are very naturally suited to a network representation and importantly they do not have a static topography, experiencing changes in ecological or even evolutionary time scales (Yodzis, 1981; Jeong et al., 2000). In this chapter we study the general properties of networks which are grown iteratively in time with a simple growth algorithm which is shown to lead to an interesting range of network types.

Of particular interest historically has been the apparent ubiquity of power-law degree distributions, and what came to be called scale-free networks, in a vast range of natural systems. In this regime of network, we find extremely well connected vertices, far more than could exist if connectivity were Gaussian. The vast disparity of connectivity suggests there is no characteristic scale, and there is no typical number of connections coming from a single vertex; the variance of vertex connectivity has diverged. Examples include networks of scientific collaborators (Newman, 2001) and Hollywood co-stardom (Barabási, 1999) along with transport networks such as airways (Guimerà et al., 2005) and roads (Kalapala et al., 2006). Further examples of scale-free networks, and power-laws elsewhere in the natural world and human society along with in-depth discussion can be found in reviews by Mitzenmacher (2004) and Newman (2005).

Barabási (1999) began to answer the question of the origin of these commonalities, demonstrating that though a process of preferential attachment whereby newly added vertices are connected to existing vertices with probability proportional to their number of connections, the scale-free degree distribution may emerge without any further mechanism. In contrast, Callaway et al. (2001) introduces a minimal model of network

growth in the absence of preferential attachment. Along with a number of interesting mathematical properties, they note that the model history results in older vertices tending to be more highly connected, purely due to having existed longer than younger vertices. Furthermore, these tend to be connected amongst themselves, introducing a positive degree correlation.

We aim to reintroduce the randomly grown network of Callaway et al. (2001) as a special instance of a more general algorithm for random growth where graph vertices are added one at a time along with new edges which may attach preferentially to varying extents via a linear attachment kernel. Mathematically, this introduces complications, though much of the tractability of the original model is retained. However, a number of key observations, particularly the associative mixing are apparently disrupted. Beginning with a brief description of the model in Sec. 6.2, we determine the degree distribution of this type of network in Sec. 6.3, finding power-law tails result from any tendency for preferential attachment. Sec. 6.4 investigates the point at which the giant component of such a network emerges; that is, the point at which the largest connected component occupies a finite fraction of the complete network, as opposed to being constrained to finite size clusters. Building on this result, Sec. 6.5 finds the corresponding discontinuity in mean connected component size as the giant component emerges. Finally, Sec. 6.5 explores the effects of the opposing forces of history and preferential attachment. The history of the network allows old vertices more opportunities to become interconnected, and naturally leads to a positive degree correlation coefficient. However, the hub structure of Fig. 6.1 indicates highly connected nodes are overwhelmingly connected to vertices with no addition connections, which would lead to a low degree of correlation. These results are discussed in the context of growing complex networks in Sec. 6.7.

6.2 Model

We begin from a single vertex and iteratively add a new vertex along with a random number of edges from some distribution with rate δ. While Callaway et al. (2001) consider only $\delta \leq 1$, this can in principle be very much larger. New edges join a random pair of existing vertices with probability proportional to their weight, made up of contributions from their existing connections k and a fixed constant m. k provides the preference for adding connections to already well connected nodes while m offsets this by providing a chance to connect randomly. In our terms, the probability $P_{i,k}$ that vertex i is linked to vertex j by a newly added link is given by

$$P_{i,j} = \frac{(k_i + m)(k_j + m)}{\left(\sum_{n=1}^{t}(k_n + m)\right)^2} \tag{6.1}$$

where k is the vertex degree, and m parametrises the preference for the new edge to join vertices with a high degree. This model differs significantly from other models of

preferential attachment in that networks produced are generally sparse except for high δ and there is no distinction made between the existing and newly added vertices (as opposed to the fully connected network of Barabási (1999), where newly added vertices are always connected).

The analysis which follows is significantly simplified by considering the model case where both ends of a new link have the same bias, m, towards already connected vertices. The extension to allow the two ends of new edges distinct biases is trivial and displays the same qualitative behaviour. At each stage, we compare analytical progress to network properties extracted from a numerical simulation of this type of graph.

6.3 Degree distribution

We begin by deriving the steady-state degree distribution for this type of grown graph. The master equation approach gives the rate of change in number of vertices with degree k, D_k. The special case of $D_{0,t}$ is simple since we add isolated vertices at a rate of 1, and find they are connected at rate

$$\frac{\mathrm{d}D_0}{\mathrm{d}t} = 1 - \frac{2\delta m}{t(m+2\delta)} D_0. \tag{6.2}$$

This is very simply translated into a rate equation to give the expected change in the population D_k in the discrete time interval t, $t+1$

$$\mathbb{E}\big(D_{0,t+1} - D_{0,t}\big) = 1 - \frac{2\delta m}{t(m+2\delta)} D_{0,t}. \tag{6.3}$$

Similarly, the same formulation is applied more generally to higher degree vertices. The change in $D_{k,t}$ is the difference between the expected number of vertices with degree $k-1$ which gain an edge, and those of degree k which gain an edge.

$$\mathbb{E}\big(D_{k,t+1} - D_{k,t}\big) = \frac{2\delta(k+m-1)}{t(m+2\delta)} D_{k-1,t} - \frac{2\delta(k+m)}{t(m+2\delta)} D_{k,t}. \tag{6.4}$$

Numerical simulation can be used and finds that for sufficiently large t, the frequency distribution D_k increases linearly with simulation time (Callaway et al., 2001). As such, we assume the graph grows to a steady state where D_k is related to the steady state degree distribution, d_k, by

$$D_{k,t} = d_k t. \tag{6.5}$$

This expression can be shown to be appropriately normalised since as stated, model time t is exactly equal to the number of vertices, $\sum_k D_k$. We seek a solution to Eq. (6.4) of

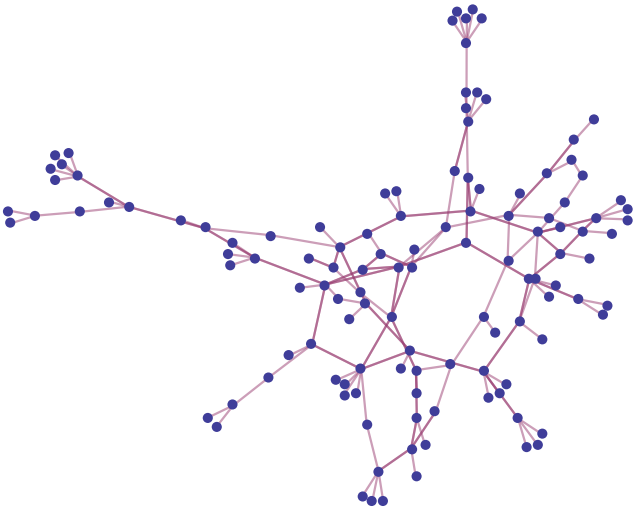

(a) $m \to \infty$

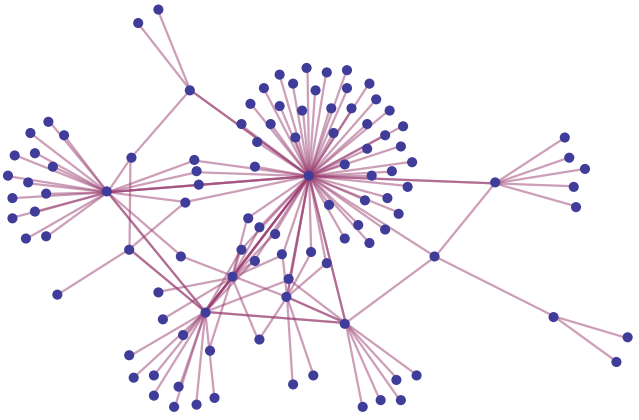

(b) $m = 1$

FIGURE 6.1: A sample from the core of a graph grown by our algorithm for the three cases considered with $\delta = \frac{1}{2}$. The images show that finite values for m produces *hub and spoke* type structures, indicative of power-law degree distributions expected from preferential attachment rules.

this form by substituting Eq. (6.5) into Eq. (6.4)

$$d_k = \frac{k+m-1}{\frac{m}{2\delta}+k+m+1} d_{k-1}. \tag{6.6}$$

Similarly for Eq. (6.3)

$$d_0 = \frac{m+2\delta}{m+2\delta m+2\delta}. \tag{6.7}$$

Solving eq. (6.6) with Eq. (6.7) as a lower bound where $k=1$ ($d_{k-1}=d_0$) gives

$$d_k = \frac{(m+2\delta)\,(m)_k}{(m+2\delta m+2\delta)\left(\frac{m}{2\delta}+m+2\right)_k} \tag{6.8}$$

where we have used Pochhammer notation, defined here as

$$(m)_k = \frac{\Gamma(k+m)}{\Gamma(m)}.$$

In the limit of $m \to \infty$, this simplifies significantly.

$$d_k = \frac{(2\delta)^k}{(2\delta+1)^{k+1}}.$$

This result gives the degree distribution for case (a) where edges connect random vertices without preference, and is identical to the result for a randomly grown graph (Callaway et al., 2001). Unsurprisingly, the graph degree distribution has a power-law tail. The exponent can be found in the large k limit using Stirling's approximation. We find

$$d_k \sim k^{-(\frac{m}{2\delta}+2)} \tag{6.9}$$

except in the case where $m \gg 2\delta$, where d_k tends towards an exponential distribution, illustrated in Fig. 6.2. Interestingly, strong power-law behaviour ($\gamma < 3$) is only found where $2\delta > m$. We can say that a highly skewed degree distribution is only found for sufficiently high δ; the sum weight of vertices due to their edges must be greater than the sum weight of vertices for simply existing in the graph in order to see significant preferential attachment. This effect becomes important later.

6.4 Giant component

We know to expect the case of $m \to \infty$ to undergo a phase transition across which the expected component size jumps discontinuously (Callaway et al., 2001). We begin to uncover the corresponding behaviour of finite m by deriving an expression for the size of the giant connected component, if any.

The number of connected components of size x, N_x can similarly be derived by a master equation approach. The expected change in the number of connected components of size

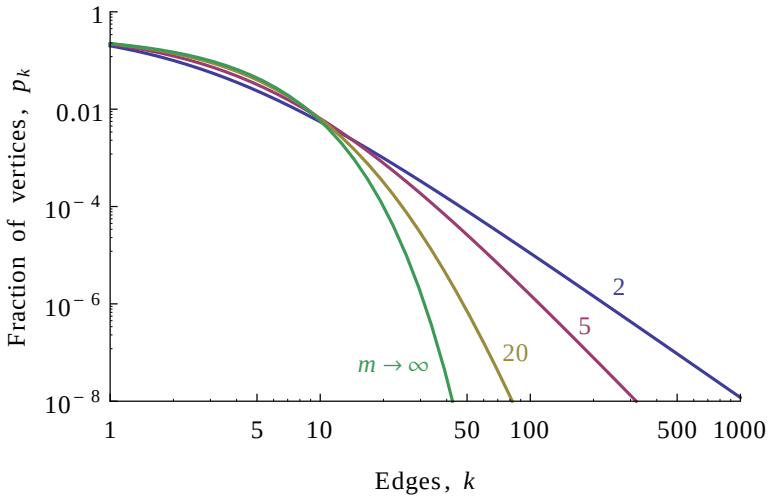

FIGURE 6.2: Edges added to a random graph connect randomly, but vertices are weighted by $m + k$ where k is the connectivity of a vertex and m parametrises the preference for highly connected nodes. Finite m leads to preferential attachment and a power-law tail while the limit $m \to \infty$ produces an exponential degree distribution (Callaway et al., 2001).

x at time t, $N_{x,t}$ has two separate contributions. Firstly, from the likelihood of joining different combinations of components which when connected form a component of size x, and also from the likelihood of joining a component of size x to any other component. At first, an issue appears to be that we only expect the sum of the weighting within a component to increase linearly with component size for $m \to \infty$, where individuals are weighted equally and the number of internal links is irrelevant. The problem is significantly simplified by assuming connected clusters not to contain any cycles or self-connections which is found to be precise in the limit $\frac{m}{2\delta} \to \infty$. Indeed the reality is found to be only marginally different from this. Even in the strong power-law domain, the total expected number of connections within a component increases linearly with the component size x, and we expect the qualitative behaviour of the network to be unchanged by assuming $2(x-1)$ edges in finite size connected components. It is simple to determine and include this linear relationship in the analysis, though it serves mainly to obfuscate the main results and we opt to proceed using the 'no-cycles' approximation.

We begin precisely as before, by writing down the master equation for the change in population of connected components of size x at time t, $C_{x,t}$. This has two components;

(1) Connecting a link to a member of a component of size x to anything

(2) Adding a link between two components of size y and $x - y$

$$\mathbb{E}\big(C_{x,t+1} - C_{x,t}\big) = \delta \sum_{y=1}^{x-1} \tfrac{(m+2)y-2}{t(m+2\delta)} C_{y,t} \tfrac{(m+2)(x-y)-2}{t(m+2\delta)} C_{x-y,t} - \tfrac{2\delta[(m+2)x-2]}{t(m+2\delta)} C_{x,t}, \tag{6.10}$$

which bounded at $x = 1$ by

$$\mathbb{E}\big(C_{1,t+1} - C_{1,t}\big) = 1 - \frac{2\delta m}{t(m+2\delta)} C_{1,t}. \tag{6.11}$$

Precisely as before, assuming a steady-state distribution c_x exists provides a set of equations and boundary conditions which can be solved iteratively.

$$c_x = -\frac{2\delta[(m+2)x-2]}{(m+2\delta)} c_x + \delta \sum_{y=1}^{x-1} \left(\frac{(m+2)y-2}{(m+2\delta)} c_y \right) \left(\frac{(m+2)(x-y)-2}{(m+2\delta)} c_{x-y} \right). \tag{6.12}$$

Similarly for Eq. (6.11)

$$c_1 = \frac{m+2\delta}{m+2\delta m+2\delta}. \tag{6.13}$$

Taking the summation $\sum_x x c_x$ gives the total fraction of vertices which belong to connected components of finite size; the complement to the fraction of the network incorporated into a giant component

$$S = 1 - \sum_x x c_x. \tag{6.14}$$

Eq. (6.13) can be solved by iteration from the boundary of Eq. (6.13), although a very large number of terms are required for accuracy, particularly in a strong power-law regime. Accuracy is more easily attained at a lower computational cost by employing generating functions. We define the generating function for the probability distribution of connected component sizes, $g(z)$ as

$$g(z) = \sum_x c_x \exp(xz).$$

Recall that the distribution c_x is not normalised. The sum $g(0)$ gives the ratio of the number of connected components to the number of vertices, such that $g(0) = 1$ only when $\delta = 0$. We are interested in $g(0)$ along with its first derivative,

$$g'(z) = \sum_x x c_x \exp(xz).$$

The interpretation of $x a_x$ is the likelihood of a randomly selected vertex belonging to a connected component of size k. The sum $g'(0)$ is therefore unity when no giant connected component exists (all components have finite size). Above the phase transition, the

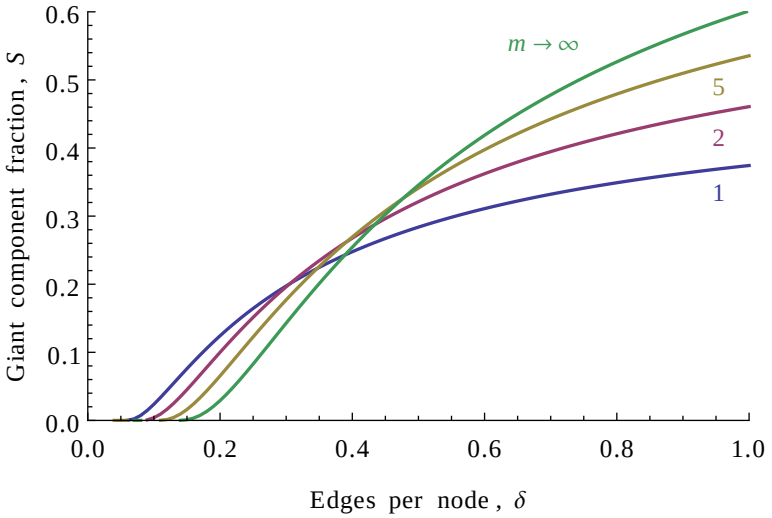

FIGURE 6.3: The fraction of the network connected to the giant component increases abruptly, but not discontinuously at some critical value of the parameter δ which determines the ratio of edges to vertices in the random graph. Unsurprisingly, smaller m produces strongly skewed degree distributions and the edges are more strongly concentrated amongst fewer vertices, resulting in a more rapidly emerging but slowly growing giant component.

fraction of the graph occupied by the giant component, S, can be extracted by

$$S = 1 - g'(0). \tag{6.15}$$

Multiplying Eq. (6.13) by $\exp(xz)$, summing over x and substituting Eq. (6.13) for the case where $x = 1$ gives

$$\begin{aligned} &2\delta(m+2\delta)(m+2)g'(z) = \\ &\quad \delta(m+2)^2 g'(z)^2 + 4\delta g(z)^2 - 4\delta(m+2)g(z)g'(z) \\ &\quad -(m+2\delta)(m-2\delta)g(z) + (m+2\delta)^2. \end{aligned} \tag{6.16}$$

Solving this equation with the initial condition $(z_0, g(z_0)) = (\exp(z_0), \exp(z_0)c_1)$ for large negative values of z_0. It is found that $\ln 10^{-6}$ provides sufficient accuracy. The giant component size is shown in Fig. 6.3. As we might expect, the stronger preference for vertices of higher degree causes a giant connected component to emerge at smaller values of δ compared with the random model. The reason for this is that the limited edges will tend to be concentrated over a smaller fraction of vertices, producing a giant connected component. However, its growth is inhibited by an increasing number of internal connections which do not increase its absolute size.

6.5 Mean Component Size

We can find exactly where the giant component emerges by, in addition to $g'(0)$, deriving $g''(0)$ defined as

$$g''(z) = \sum_{x=1}^{\infty} x^2 a_x \exp(xz)$$

useful for its relation to the expected component size $\langle c \rangle$ by

$$\langle c \rangle = \frac{\sum_{x=1}^{\infty} x^2 a_x}{\sum_{x=1}^{\infty} x a_x} = \frac{g''(0)}{g'(0)}.$$

The position of the phase transition can be determined by examining the form of $g(0)$ in the $g'(0) = 1$ regime, that is to say the range of δ for which no giant connected component exists. Substituting this into Eq. (6.16) at $z = 0$, and solving for $g(0)$ gives

$$g(0) = \frac{m^2 - 4\delta^2 + 4\delta(m+2) \pm (m+2\delta)^2}{8\delta}. \tag{6.17}$$

Since we know that as $\delta \to 0$ the ratio $g(0)$ approaches unity, we choose the negative signed solution, which contracts Eq. (6.17) to $g(0) = 1 - \delta$, a consequence of our approximation that finite components contain no cycles. Differentiating Eq. (6.16) and applying L'Hôpital's rule in the limit $z \to 0$ with $g(0) = 1 - \delta$ and $g'(0) = 1$ gives

$$\begin{aligned} g''(0) &= \frac{m^2 + 4\delta(3m + \delta + 4)}{4\delta(m+2)^2} \\ &\pm \frac{(m+2\delta)\sqrt{m^2 + 4\delta^2 - 4\delta m(2m+3)}}{4\delta(m+2)^2} \end{aligned} \tag{6.18}$$

which is real for $\delta < \delta_c$ where

$$\delta_c = \frac{1}{2}\left(3m + 2m^2 - 2\sqrt{2m^2 + 3m^3 + m^4}\right). \tag{6.19}$$

We choose the solution with the negative sign since we know that for $\delta = 0$, each vertex belongs to a component of size one, such that $\langle s \rangle = 1$. δ_c indicates the position of the emergence of a giant connected component. Expanding this in the limit $m \to \infty$ gives $\delta_c = \frac{1}{8}$, exactly as with the random grown graph (Callaway et al., 2001). Eq. (6.18) contains a singularity at $\delta = \delta_c$, past which $g(0) > 1 - \delta$ as edges are added to the giant component. In the absence of any closed form expression of $g(0)$ from Eq. (6.16) for the general case of $g'(0) < 1$, we solve for $g''(0)$ numerically with the initial condition $\big(z_0, g(z_0), g'(z_0)\big) = \big(\exp(z_0), \exp(z_0)a_1, \exp(z_0)a_1\big)$ for large negative values of z_0. Results are displayed in Fig. 6.4.

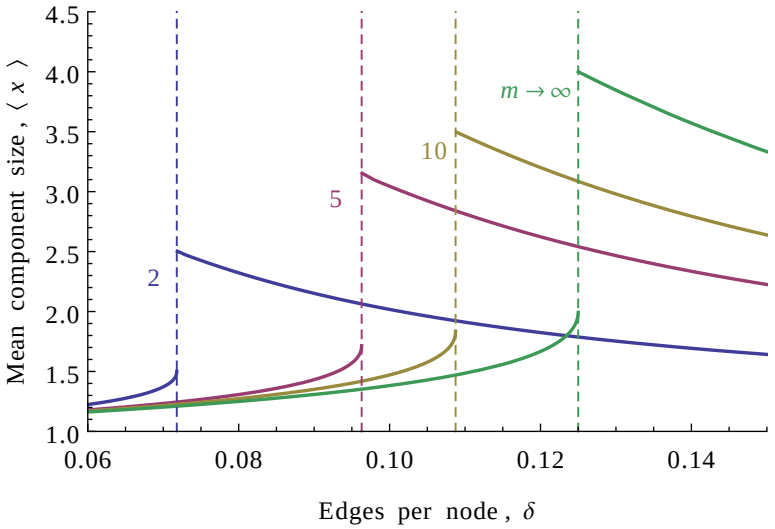

(a) $m \to \infty$

FIGURE 6.4: The mean size of finite connected components for a range of network parameters. Interestingly, the emergence of the giant component is accompanied by an abrupt change in the expected component size which increases asymptotically to $\delta_c = \frac{1}{8}$ in the limit of $m \to \infty$ (Callaway et al., 2001). Dashed lines indicate the transition by marking the singularity of Eq. (6.18), Eq. (6.19).

6.6 Degree Correlation

Differences can also be seen by examination of associative mixing in our model (Barabási, 1999; Callaway et al., 2001; Newman, 2002a). The explanation for positive associative mixing in the randomly grown graph is that older vertices, introduced at small t will not only accumulate more connections, but they will be more likely to be find connections amongst themselves in comparison to younger vertices introduced at larger t. The result is that vertices of high degree are likely to find connections to other vertices of high degree, resulting in associative mixing. However in power law degree distributions, highly connected vertices act as hubs for vertices of low degree.

The number of edges connecting vertices with degree k to l at time t, is defined as $E_{kl,t}$ (as introduced by Krapivsky et al. (2001)). This term has three contributions;

(1) $E_{kl,t}$ is increased when a vertex of degree $k-1$ is already connected to a vertex of degree l and receives another connection.

(2) $E_{kl,t}$ is increased by unconnected vertices of degree $k-1$ and $l-1$ becoming connected

(3) $E_{kl,t}$ is decreased when vertices of degree k and l are connected, and either receives an additional connection

The master equation resulting from these terms is

$$\begin{aligned}\mathbb{E}\big(E_{j,k,t+1} - E_{j,k,t}\big) = \\ 2\delta\left(\tfrac{d_{j-1,t}(j+m-1)}{(m+2\delta)t}\tfrac{E_{j-1,k,t}}{d_{j-1,t}} + \tfrac{d_{k-1,t}(k+m-1)}{(m+2\delta)t}\tfrac{E_{j,k-1,t}}{d_{k-1,t}}\right) \\ +2\delta\left(\tfrac{k+m-1}{(m+2\delta)t}d_{j-1,t}\tfrac{k+m-1}{(m+2\delta)t}d_{k-1,t}\right) \\ -2\delta\left(\tfrac{d_{j,t}(j+m)}{(m+2\delta)t}\tfrac{E_{j,k,t}}{d_{j,t}} + \tfrac{d_{k,t}(k+m)}{(m+2\delta)t}\tfrac{E_{j,k,t}}{d_{k,t}}\right).\end{aligned} \tag{6.20}$$

By adding that $E_{0,k,t} = E_{j,0,t} = 0$ for all j and k (vertices with no connections cannot have neighbours), Eq. (6.20) is sufficient for a boundary condition. In the steady state, we assume

$$E_{j,k,t} = 2\delta t e_{j,k}$$

The matrix $e_{j,k}$ gives the distribution of the degree of vertices at each end of a randomly chosen edge, normalised such that $\sum_j e_{j,k} = p_k$. The factor of 2 arises from the symmetry of $E_{j,k,t}$ under the interchange of j and k. Summation over j and k will double-count the total number of edges, δt.

Substituting this into Eq. (6.20), then solving for e_{kl} gives

$$\begin{aligned}e_{j,k} &= \frac{2\delta\big((j+m-1)e_{j-1,k} + (k+m-1)e_{j,k-1}\big)}{m+2\delta(j+k+2m+1)} \\ &\quad + \frac{(j+m-1)(k+m-1)p_{j-1}p_{k-1}}{(m+2\delta)\,(m+2\delta(j+k+2m+1))}.\end{aligned} \tag{6.21}$$

We translate this matrix into the degree correlation coefficient, ρ, following Callaway et al. (2001).

$$\rho = \frac{c}{\sigma^2} \tag{6.22}$$

where σ^2 and c are the variance of the vertex degree distribution at either end of a random edge, and the covariance of the degrees of vertices at the ends of a random edge respectively, and defined as

$$\sigma^2 = \frac{\sum_k (k-\mu)^2 k p_k}{\sum_k k p_k} \tag{6.23a}$$

$$c = \sum_{j,k}(j-\mu)(k-\mu)e_{j,k} \tag{6.23b}$$

$$\mu = \frac{\sum_k k^2 p_k}{\sum_k k p_k}. \tag{6.23c}$$

Here, the skewed degree distribution has an impact. Examining the long tail of the degree distribution given in Eq. (6.9) we can see that the variance of the distribution diverges when the exponent of the tail $\frac{m}{2\delta} \leq 1$. Here μ diverges and the degree correlation is exactly zero.

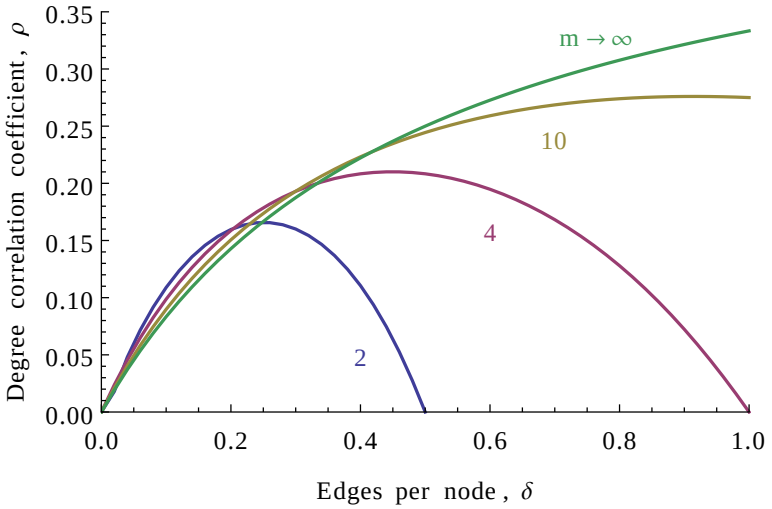

FIGURE 6.5: The correlation coefficient ρ calculated using Eq. (6.22). As the variance of p_k diverges for $\frac{m}{2\delta} \leq 1$, the correlation coefficient decays to zero. The limit $m \to \infty$ therefore corresponds to a monotonically increasing value of ρ.

Eq. (6.21) can be approached by multiplying by $jkh(y,z)$ and taking the summation over j and k, where

$$h(y,z) = \sum_{j,k} e_{j,k} \exp(ky + jz), \tag{6.24}$$

and subsequently solved along the line $y = z$. From this result, Eq. (6.23b) can found by first rewriting as

$$\begin{aligned} c &= \sum_{j,k} jke_{j,k} - 2\mu \sum_{j,k} je_{j,k} + \mu^2 \sum_{j,k} e_{j,k} \\ &= \partial_y \partial_z h(y,z)|_{y,z=0} - \mu^2. \end{aligned} \tag{6.25}$$

Fig. 6.5 shows some contrast between the networks grown with some tendency for preferential attachment and the randomly grown case with $m \to \infty$. While all instances have a historical tendency for older vertices both to have higher connectivity and be more likely to be connected, we find that this does not imply $\rho > 0$. For any finite m, ρ increases to a maximum at $\delta = \delta_0$, where δ_0 is some increasing function of m, before decaying with increasing δ. However for any finite graph, we find that finite size allows the model to maintain a positive correlation even for highly skewed degree distributions. At the point $m = 2\delta$ the variance in degree distribution diverges, and degree correlation is exactly zero. The limiting case is unbiased growth, $m \to \infty$, where the variance of p_k is finite for all finite δ, and positive associative mixing approaches its maxima asymptotically.

6.7 Discussion

We have introduced a general form of a randomly grown graph where vertices are added along with new edges, either end of which may connect with a degree of bias towards high degree vertices. Qualitatively similar properties are found in the case where both ends have the same propensity for preferential connection. This extension retains a great deal of the tractability of the special case of Callaway et al. (2001), where there is no tendency for preferential attachment, and reveals that a number of properties observed are unique to this case.

Firstly we have shown that, as expected, the addition of preferential attachment in randomly grown networks introduces skewed degree distributions with power-law tails. We find that even weak preferential attachment accelerates the emergence of a giant connected component, and correspondingly a discontinuity in the mean size of finite connected components. This is in agreement with first intuitions; the introduction of preferential attachment encourages the existence of a giant component at small δ compared to the unbiased case by disproportionately concentrating edges between older well-connected vertices.

Finally the addition of preferential attachment finds degree correlation to increase up to some critical value after which the diverging variance of vertex degree begins to dominate. Degree correlation decays to zero for increasingly well connected graphs, with increasingly skewed degree distributions. This is opposed to the asymptotically increasing degree correlation found from unbiased random growth. Even large simulations indicate this convergence to be slow, as the behaviour is dominated by rare, highly connected vertices indicative of scale-free networks. A similar result is obtained by Krapivsky et al. (2001) where scale-free networks are shown to be able to possess associative mixing to some extent. This result is due to the variance of degree distribution in scale free networks increasing very much faster than the covariance of of the degrees of connected vertices.

Chapter 7

Maximum entropy production principle in convection cell formation

7.1 Introduction

The Earth system considered as a whole is a complex, dissipative system, ensured by the constant solar insolation providing energy for the persistence of non-equilibrium phenomena on Earth. This has proven problematic as, in contrast to equilibrium thermodynamics, it is unclear which thermodynamic properties, if any, govern non-equilibrium systems. It is well understood that equilibrium systems minimise their free energy, that is energy which can be converted to do work, such as from temperature gradients. There have been a number of proposed equivalent principles for application to dissipative systems of which the Maximum Entropy Production (MEP) principle has gained significant traction (Ozawa and Ohmura, 1997; Ito and Kleidon, 2005). MEP states that dissipative systems are characterised by a steady state in which the rate of thermodynamic entropy production is maximised and a number of examples have demonstrated the utility of the MEP principle, such as the prediction of atmospheric heat transport from simple considerations (Paltridge, 1975; Lorenz, 2006; Shimokawa and Ozawa, 2002) and geological process within the Earth (Lorenz, 2002). Indeed some implementations of Daisyworld models explicitly model heat flux across the temperature gradient established by the difference between high albedo white daisies and the relatively absorptive black daisies and have investigated the impact of optimising entropy production on the regulatory mechanism (Pujol, 2002; Dyke, 2008) along with assessing MEP when compared to alternative optimisation principles (Ackland et al., 2003).

MEP studies typically proceed on the basis of formulating a model of the system in question, and using constrained optimisation to produce an MEP state. The resultant model output is then compared to empirical data. While the success of these studies are both fascinating and very promising, there is a significant paucity of experimental or simulated validation which would otherwise shed light on the types of system for which this approach may be successful (Ozawa, 2003; Martyushev and Seleznev, 2006), and while in progress, a theoretical grounding for MEP is yet to be established (Dewar, 2005, 2009). Here we present a series of simulations of Rayleigh-Bènard convective systems, produced in order to evaluate the applicability of MEP and other competing principles. In particular, we ask if the formation of convective cells within the system correlate to its entropy production in a way that would allow certain bulk properties of the system to be predicted without recourse to simulation of the internal kinetics.

Convection cells occur in a vast range of systems across an enormous range of scales, from currents in the core of early stars to a heated pan of soup (Getling, 1998). While these phenomena may appear as simple curiosities, they have a grand impact on global temperatures, producing cells in the atmosphere. In this work, we do not interest ourselves with the specifics of scale and the fluid properties, rather we are interested in the driving principles behind convection cells as a generic example of self-organisation. One type of convection cell is a Rayleigh-Bénard convection cell. They occur where fluids

with positive thermal expansion coefficient are heated from below. Warm parcels of fluid experience an upwards buoyancy force from the induced vertical density gradient. At the surface, they are able to deposit, or radiate heat, eventually becoming more dense than the rising fluid below. Being unable to descend here, currents form outward from the up-welling point until the inverse density gradient is encountered, and the parcel descends. Through this mechanism, a fluid subject to a temperature gradient is able to self-organise to a dissipative state, transporting energy from the hot to cold reservoir. The internal entropy production is easily computed for this phenomena, and the associated fluid mechanics are irrefutably complex, making an ideal case study for the investigation of extremum principles for non-equilibrium systems. The broad spectra of systems subject to this phenomena can hardly be overstated, an observation which caused Getling to remark "Convection due to non-uniform heating is, without overstatement, the most widespread type of fluid motion in the Universe" (Getling, 1998).

Rather than prescribe system-level properties such as heat flux or rates of convection, we allow these to emerge via the interaction of the simulated particles within the system. We implement a lattice gas cellular automata (LGCA) as our approach for modelling our convecting fluid system (Vichniac, 1984). Furthermore, in order to be able to produce statistically significant results, a novel computational method was developed that allowed large scale simulations to be conducted. We introduce the LGCA class of models in more depth, particularly our choice of lattice gas in Sec. 7.2 along with the necessary modifications to produce convection cells. Results are presented and discussed in detail in Sec. 7.3 and Sec. 7.4.

7.2 A Model of Rayleigh-Bénard Convection

We propose an experiment in which a fluid rests between two large thermal reservoirs, separated from the fluid by a thin insulating layer of negligible heat-capacity. The bottom reservoir is relatively hot, having a fixed temperature T_{h}^{*} while the top is cold with temperature T_{c}^{*}. This imposes a vertical temperature gradient in the fluid. The temperature difference across the fluid itself is at most $T_{\mathrm{h}}^{*} - T_{\mathrm{c}}^{*}$ but this vertical gradient is eroded by an upward buoyancy-forced convective heat flux within the fluid. A smaller vertical temperature gradient within the fluid can be created if the insulation at the boundaries is increased. In a steady state the heat flux into the fluid, Q_{in}, from the hot reservoir is equal to the flux out of the fluid, Q_{out}, into the cold reservoir. This arrangement is shown in Fig. 7.1 and similar systems have been implemented in laboratory experiments (Jellinek and Lenardic, 2009).

The challenge is to develop a simulation of a convective system that captures the important processes over sufficient temporal and spatial scales whilst remaining computationally feasible as well as intuitive. With a number of modifications (detailed in the

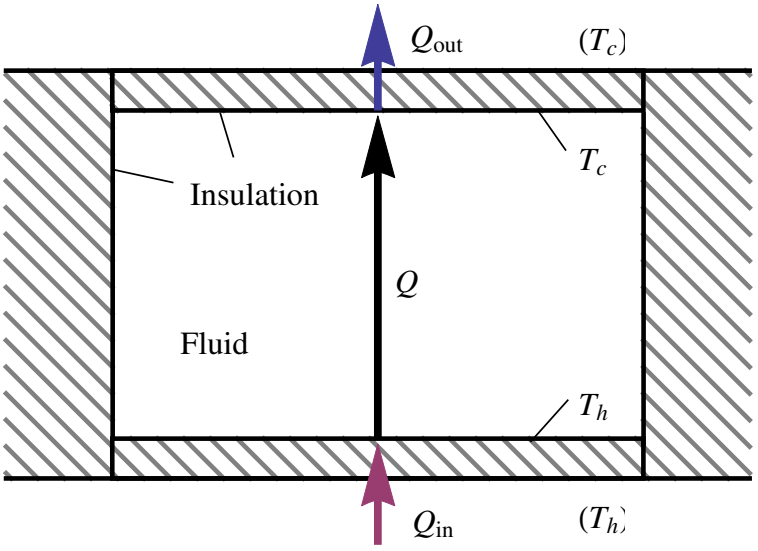

FIGURE 7.1: A fluid is subjected to a vertical temperature gradient by a pair of thermal reservoirs at constant temperatures T_{h}^* and T_{c}^*. The sides of the fluid are well insulated from the surroundings, while the boundaries between the fluid and thermal reservoirs have a small amount of insulation which can be tuned to adjust the temperature gradient established across the fluid.

following sections) LGCA represent an ideal trade-off in terms of efficiency and fidelity.

7.2.1 Lattice Gas Cellular Automata

LGCA consist of discrete particles traversing a regular lattice, characterised both by the choice of lattice vectors, that is the vector describing the edges attached to each vertex, $\vec{c}_i$ and collision rules. While in terms of implementation of the model we talk of particles, we are not constrained to modelling on the molecular scale. Indeed from this abstraction, the equations of motion of a continuous fluid can be extracted by application of Chapman-Enskog theory (Frisch, 1987). A commonality shared by all lattice gas models is that they advance in alternating collision and propagation steps:

Collision

Particles at a vertex participate in mass and momentum conserving collisions. Some configurations allow for scattering collisions in which particles redistribute their momenta. These collisions provide viscosity and enable us to recover the Navier-Stokes equations.

Propagation

Between collisions, particles travel along lattice vectors to neighbouring vertices separated by the characteristic lattice spacing ΔL. Particles which would travel across a solid surface have their momentum state inverted, known as a no-slip boundary due to the result of zero net velocity at a boundary.

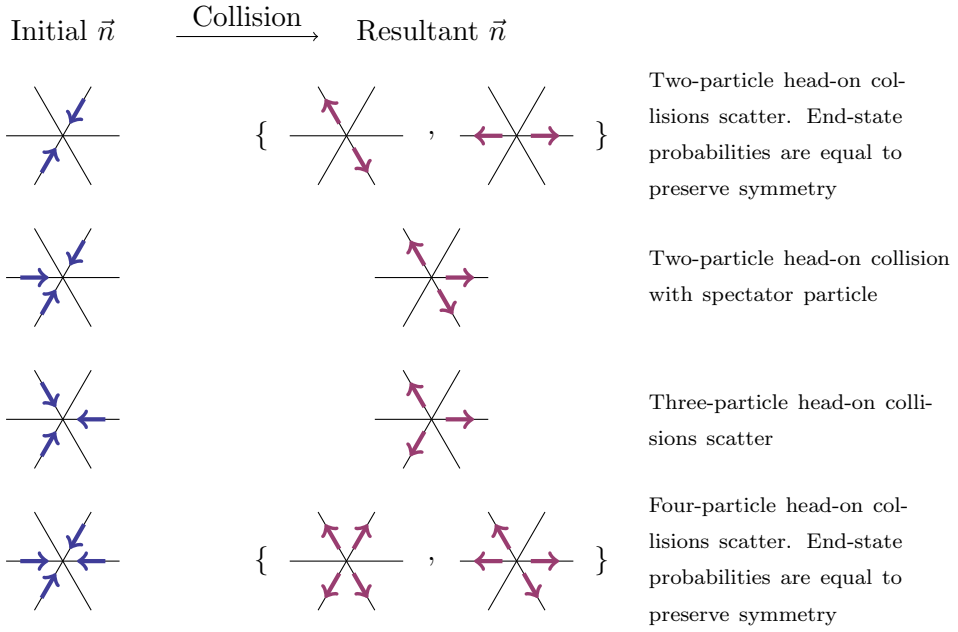

FIGURE 7.2: Explicit definitions of the FHP lattice gas collision rules which change the momentum states at a vertex. The FHP lattice has six momentum states at each vertex, where arrows indicate the presence of a particle in momentum state n_i and the vector $\vec{n}$ holds all the momentum state occupancies for a vertex. Two- and four-particle head-on collisions are probabilistic, and the resulting states occur with equal probability.

The original LGCA was developed in 1973 used a simple 2D square lattice and is commonly named the HPP lattice gas after its authors (Hardy, 1973; Hardy et al., 1976). These early LGCA suffered greatly from anisotropy, resulting in square vortices. A similar method was implemented later in 1986 on a hexagonal lattice called the FHP lattice gas, again named after the authors (Frisch et al., 1986). The increased symmetry of this lattice ensures it does not suffer so strongly from anisotropy. A thorough account of the zoo of lattice gas models can be found in Wolf-Gladrow (2005).

As a hexagonal lattice with six lattice vectors, each vertex in the FHP lattice has six momentum states, whose occupancy is denoted by the binary variable n_i with $i = 1 \dots 6$. The state $n_i = 1$ corresponds to an occupied vertex while 0 represents a hole. Almost all collisions are trivial since there is only one mass and momentum conserving outcome to a collision. The exceptions to this are mainly zero net-momentum collisions which have multiple possible outcomes. The choice proves extremely important to the model dynamics and scattering (momentum-state altering) collisions are chosen. The collision rules are summarised in Fig. 7.2 to highlight this distinction. The final element of the model is the propagation step. Between collision steps, momentum states travel along lattice vectors to participate in collisions at a neighbouring vertex.

7.2.2 FHP Buoyancy Modification

The traditional FHP model does not have a defined temperature since for most hydrodynamical problems thermal anomalies are assumed small, and temperature is therefore not important (Landau and Lifshitz, 1987). However, the temperature gradient serves as the driving force for convection for which Schaffenberger et al. (2001) modify the usual FHP model by introducing an effective particle temperature, along with an associated buoyancy force modelled by spontaneous momentum-state flipping. As well as a momentum-state, particles have a temperature state θ_i which can take the value ± 1 for a hot or cold particle respectively. The temperature state is a passive scalar; it follows particles passively during the propagation step. Hot particles can be thought of as being less dense than the surrounding medium and experience a positive buoyancy force. Cold particles are relatively dense and experience a negative buoyancy force. The inclusion of θ states modifies the lattice gas in the following ways:

Collision

In the collision step the total particle thermal energy, $\sum_i \theta_i$, is conserved. However, the temperature states are randomly distributed amongst the participating particles. The thermal diffusion coefficient then becomes a function of density only.

Thermal flipping

A force can be simulated by spontaneous flipping of particle momentum states, resulting in a change of net-momentum Frisch (1987). The buoyancy force is simulated by inverting particle y-momentum, without altering x-momentum. Particles in the negative y-momentum state with $\theta = 1$ flip into the corresponding positive y-momentum state with some small probability, here labelled by ω. The thermal flipping rule is illustrated in Fig. 7.3. While flipping causes spontaneous changes in local momentum, symmetry of the thermal flipping rule with θ ensures that on the average, global momentum is conserved.

Propagation

As a passive scalar, θ is propagated between lattice points along with particles. Additionally particles reflected from a boundary set θ to ± 1 depending on the simulated temperature of the boundary.

7.2.3 A Lattice Boltzmann Model

The simulation results this model required coarse-graining and long time averages to be taken over a very large lattice (Schaffenberger et al., 2001). Even then, they exhibit large amounts of noise, and rapid transitions can be seen between different convective states for parameter values where, say, period-1 (where the lattice is occupied by a single

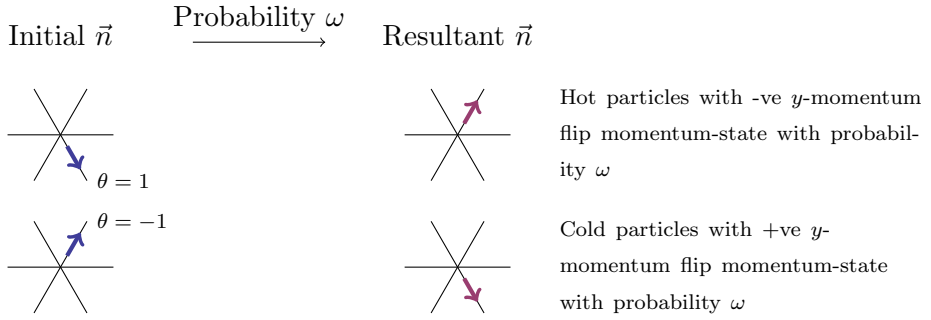

FIGURE 7.3: Explicit definitions of FHP buoyancy rules. Hot particles have positive buoyancy and flip to positive y-momentum states without altering their x-momentum, and vice versa for negatively buoyant cold particles. This can be thought of as equivalent to applying a force in $\hat{y}$ where the strength of the force is determined by the rate of such transitions.

convection cell) and period-$\frac{1}{2}$ (where two cells of half the size occupy the lattice) are similarly favourable. While the effects of such fluctuations may be important in real systems, we are interested in the thermodynamic properties of steady-state convection cells, a very difficult parameter to extract from this type of simulation.

Fortunately, the Boltzmann transport equation grants us reprieve from many of these problems. The very simple and intuitive lattice gas rules can be transformed from a binary model of lattice site occupations n_i to a continuum model of mean momentum-state occupancies N_i. Through some careful manipulation a continuum model can be made to approximate the real lattice gas dynamics with the advantage of much smaller lattices and removing the need for large space and time averages. In the following steps, we write $N_i(\vec{x}, t)$ to be the mean occupancy of momentum-state i of the vertex at $\vec{x}$ at simulation time t. The collision and propagation steps can be written as

$$N_i(\vec{x} + \vec{c}_i \Delta t, t + \Delta t) = N_i(\vec{x}, t) + \Delta_i \tag{7.1}$$

where $\vec{c}_i \Delta t$ is a vector of length ΔL, the lattice spacing, and therefore the left side of this equation can be thought of as the propagation step, advancing mean momentum occupancy along lattice vectors. The right side can be thought of as the collision step, and includes the collision term, Δ_i. This operator adjusts the mean occupancy propagated between lattice sites according to the collision rules. It includes all the information in Fig. 7.2. Rather than enumerating all possible momentum-state altering collisions here, we can call on the BGK approximation (Bhatnagar et al., 1954). The large and highly non-linear collision operator can be replaced by a less computationally intensive linear relaxation towards $N_i^{(\mathrm{eq})}$, the local mass and momentum conserving equilibrium with characteristic time scale τ_1

$$\Delta_i = -\frac{1}{\tau_1} \left[N_i(\vec{x}, t) - N_i^{(\mathrm{eq})}(\vec{x}, t) \right]. \tag{7.2}$$

Frish et al. Frisch (1987) proved the equilibrium momentum-state occupancies $N_i^{(\mathrm{eq})}$ of a vertex with a given density ρ and momentum $\vec{v}$, given by

$$\rho = \sum_{i=1}^{6} n_i(\vec{x}, t) \qquad \vec{v} = \sum_{i=1}^{6} N_i(\vec{x}, t)\vec{c}_i \tag{7.3}$$

obeys Fermi-Dirac statistics, since the particles are subject to the exclusion principle; each momentum state may be occupied by at most one particle, and therefore $0 \leq N \leq 1$. The proof itself is somewhat involved (Wolf-Gladrow, 2005, p. 64-67), so here we present only the result;

$$N_i^{(\mathrm{eq})}(\rho, \vec{v}) = \frac{1}{1 + e^{(h + \vec{c}_i \cdot \vec{q})}} \tag{7.4}$$

where h and $\vec{q}$ are the Lagrange multipliers associated with our conserved quantities. The full derivation of these Lagrange multipliers is cumbersome (Wolf-Gladrow, 2005, p. 248-251) and again we give only the result, which can be substituted into the linearised collision operator of Eq. (7.2)

$$N_i^{(\mathrm{eq})}(\rho, \vec{v}) = d + d(d-1)q_1 \vec{c} \cdot \vec{v} + \frac{1}{2} d(d-1)(2d-1) q_1^2 c_{i,\alpha}^2 v_\alpha^2 + d(d-1)h_2 \vec{v}^2 \tag{7.5}$$

where the sub-scripted α notation represents a sum over the x and y components, and

$$d = \frac{\rho}{6}, \quad q_1 = \frac{2}{d-1} \quad \text{and} \quad h_2 = \frac{1-2d}{(d-1)^2}. \tag{7.6}$$

In a similar way, we re-implement the buoyancy force as linear relaxation towards thermal equilibrium occupancy, $N_i^{(\mathrm{eq}')}$. This necessitates the addition of another time scale, τ_2 which parametrises the strength of this buoyancy force in a similar way to ω in the discrete model of Schaffenberger et al. (2001). Momentum states with $c_{i,y} = 0$ are unchanged, while the rest are relaxed towards

$$N_i^{(\mathrm{eq}')}(\vec{n}, \theta) = \frac{\mathrm{sign}(c_{i,y}) + \theta \Sigma_i^{(\mathrm{pair})} + \sqrt{\theta^2 (\Sigma_i^{(\mathrm{pair})} - 2)\Sigma_i^{(\mathrm{pair})} + 1}}{2\theta} \tag{7.7}$$

where $\mathrm{sign}(c_{i,y})$ denotes the sign of the y-momentum of momentum state i, and $\Sigma_i^{(\mathrm{pair})}$ is the sum of mean-occupancy of corresponding momentum states between which flipping can occur. Our completed collision operator is found by summing the relaxations towards momentum-conserving equilibrium, and this thermal equilibrium occupancy.

$$\Delta_i = -\frac{1}{\tau_1}\left[N_i(\vec{x}, t) - N_i^{(\mathrm{eq})}(\vec{x}, t)\right] - \frac{1}{\tau_2}\left[N_i(\vec{x}, t) - N_i^{(\mathrm{eq}')}(\vec{x}, t)\right]. \tag{7.8}$$

7.2.4 Boundary Conditions

To impose a temperature gradient across the lattice gas, we heat the top and bottom boundaries at different rates, as described in Sec. 7.2. A Newtonian-relaxation boundary

condition refers to a forcing scheme where the strength of forcing is linearly related to boundary temperature by some characteristic relaxation time scale λ. Here, the rate of heating or cooling of the boundaries is proportional to the difference in temperature between a surface, T, and its temperature in the absence of any additional dissipation, T^*. This boundary condition has traditionally been used in climate research as the simplest available representation of atmosphere-ocean heat exchange (Scott et al., 1999; Stommel, 1961), where the relaxation time scale λ is typically treated as an uncertain but temporally constant parameter encapsulating the effect of a number of small-scale radiative and convective processes. Within the context of our experiment, λ is governed by the thermal resistance of the insulating material at the upper and lower boundaries of the fluid, as illustrated in Fig. 7.1. Increasing the thickness of this material increases λ. Such schemes can be used to model sensible heat transfer from thermostats, or in the case of small temperature anomalies ($T^* - T \ll T^*$) it can be used to approximate flux driven forcing by linearising the Stefan-Boltzmann law. In this scheme, the rate of change of the boundary temperatures are given by;

$$\frac{\mathrm{d}T_h}{\mathrm{d}t} = \frac{1}{\lambda}\left(T_h^* - T_h\right) - Q, \qquad \frac{\mathrm{d}T_c}{\mathrm{d}t} = \frac{1}{\lambda}\left(T_c^* - T_c\right) + Q \tag{7.9}$$

where T_h^* and T_c^* are the steady-state temperatures of the lower and upper boundaries respectively. Along with λ these account for our control parameters. The remaining variables, the rate of heat exchange between the boundaries Q, and their temperatures T_h and T_c, emerge from the organisation of the fluid.

7.2.5 Defining Entropy Production

The entropy production of a closed system is simply found by summation of the various heat fluxes divided by the temperature at which the exchange takes place. The entropy production of the entire system S_{closed} in this instance is the sum of contributions from the heat flux into the fluid, Q_{in}, dissipation by the fluid, Q, and heat flux into the cold reservoir, Q_{out};

$$S_{\text{closed}} = Q_{\text{in}}\left(\frac{1}{T_h} - \frac{1}{T_h^*}\right) + Q\left(\frac{1}{T_c} - \frac{1}{T_h}\right) + Q_{\text{out}}\left(\frac{1}{T_c^*} - \frac{1}{T_c}\right) \tag{7.10}$$

where in the steady state $Q_{\text{in}} = Q = Q_{\text{out}}$, and Eq. (7.10) can be reduced to

$$S_{\text{closed}} = Q\left(\frac{1}{T_c^*} - \frac{1}{T_h^*}\right). \tag{7.11}$$

Due to our assumption that the thermal reservoirs are large and that their temperatures are approximately constant Eq. (7.11) indicates that entropy production and heat flux are linearly related by our largely arbitrary choices of reference temperatures for a closed system. In such a closed system, there would be no way to distinguish a principle of

maximum heat flux from maximum entropy production. However, previous studies have highlighted that this measure of entropy production does not yield MEP states consistent with observation (Essex, 1984). Agreement with empirical data is found by considering only entropy production by system components with sufficient complexity to self-organise. In the case of our model, conduction of heat from the thermal reservoirs through the insulating material is simple, while the fluid layer offers many degrees of freedom to adopt a range of possible dissipative structures, such as convection cells. As Virgo (2010) discusses in detail, no established theory exists to account for which sources of entropy production should be considered for MEPP to be successful, though we follow previous applications of the principle and will consider only the internal entropy production S due to the internal heat flux Q

$$S = Q\left(\frac{1}{T_c} - \frac{1}{T_h}\right). \tag{7.12}$$

We may begin to examine how the internal entropy production varies with heat flux by solving Eq. (7.9) for Q in their steady state and substituting into Eq. (7.12)

$$S = Q\left(\frac{1}{\lambda Q + T_c^*} + \frac{1}{\lambda Q - T_h^*}\right). \tag{7.13}$$

Further substituting $\Delta T = T_h^* - T_c^*$ and expanding to first order for small temperature anomalies $\Delta T \ll T$ gives

$$s = q(1 - 2q) \tag{7.14}$$

where s and q are our non-dimensionalised entropy production and heat flux rates respectively,

$$s = \frac{\lambda T^2}{(\Delta T)^2} S, \qquad q = \frac{\lambda}{\Delta T} Q. \tag{7.15}$$

This result illustrates that the entropy production is a parabolic function of the non-dimensionalised heat flux q in the limit of small temperature anomalies. This is caused by what Virgo (2010) describes as negative-feedback boundary-conditions; Eq. (7.12) is a trade-off between maximisation of the heat flux Q and the temperature gradient. The function has a single maximum at $q_{\text{crit.}} = \frac{1}{4}$ which is easily determined from Eq. (7.14) and illustrated by Fig. 7.4.

We define a rapid-forcing regime to be one where λ is sufficiently small that relaxation of the boundaries to their steady-state temperature occurs on much shorter time scales than the dissipation of heat between them, and $q < q_{\text{crit.}}$. Here, MEP is indistinguishable from maximising heat flux q, while in a slow-forcing regime, where λ is large and $q > q_{\text{crit.}}$, entropy production and heat flux are anti-correlated and the maximisation principles can be easily distinguished.

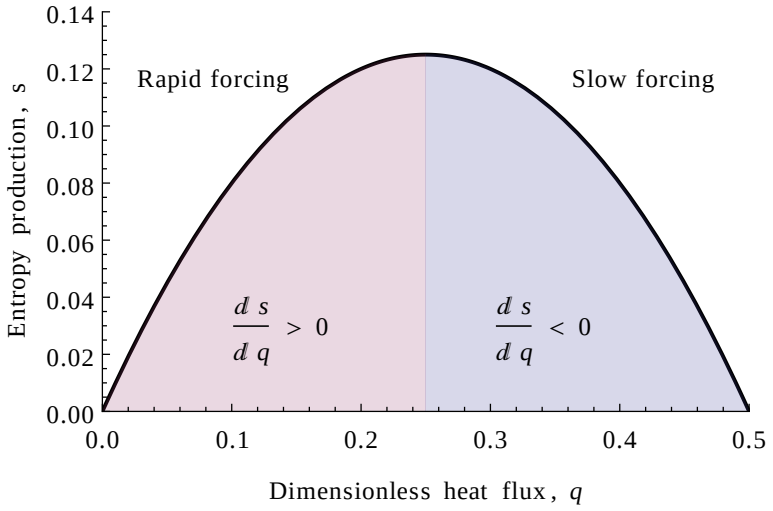

FIGURE 7.4: Non-dimensionalised entropy production s as a function of heat flux q for our Newtonian relaxation boundary condition in the limit $\Delta T = T_h^* - T_c^* \ll T_c^*$. In the rapid-forcing regime, maximising entropy production and maximising heat flux are equivalent, while in the slow-forcing regime they are easily distinguished.

7.3 Model Results

We present results from two sets of simulations, one with rapid forcing, one with slow forcing. For both sets we find a relationship between the vertical heat flux, q, and the stability of the emergent convection cells, where the system favours configurations that produce the greatest flux of heat from hot to cold reservoir. This phenomena is responsible for producing hysteresis loops. With rapid forcing, this maximum rate of heat flux is equivalent to maximum entropy production. With slow forcing, the maximum rate of heat flux is equivalent to minimum entropy production. This section examines these observations in greater detail. For all simulated results, the model uses mean particle density $\bar{\rho} = 1.2$, upper and lower boundary steady-state temperatures $T_c^* = 273\text{K}$, $T_h^* = 293\text{K}$ and lattice height $L_y = 100\frac{\sqrt{3}}{2}\Delta L$. Finally Eq. (7.8) requires the characteristic model relaxation times for momentum and thermal diffusion which are fixed at $\tau_1 = \tau_2 = 20$.

We examine the rate of heat flux and corresponding entropy production of a range of convection cell configurations outside the lattice dimensions in which they might typically occur. To begin with, we initialise a small cavity and evolve the model to a steady-state convection such as that shown in Fig. 7.5. From here, we increase the lattice width, L_x, some small fraction, injecting new particles to maintain the gas density at a constant. Since the gas is already in a steady-state of a specific period, we expect it will be attracted to a steady-state of the same period in the new, slightly larger cavity. This process can be used to expand and contract the lattice for a range of convection configurations.

7.3.1 Meta-stability and Hysteresis

A snapshot of two simulations are displayed in Fig. 7.5, highlighting the ability of the cavity to support a range of convection cell configurations. By slow expansion, we are able to maintain prescribed convective configurations over a large range of model dimensions. Additionally, perturbation experiments can be carried out (simply achieved by applying small, global mass- and momentum-conserving fluctuations to lattice momentum states), and find transitions to higher q states are strongly favoured. This can be seen clearly where the fluctuations caused by expanding the model lattice produce clean hysteresis loops, illustrated in Fig. 7.6 for the case of a transition between period-1 and period-$\frac{1}{2}$ convection cells. This phenomena is illustrated more generally in Fig. 7.7(a). From this, we infer that maximal q states are the most stable.

7.3.2 Heat Flux and Maximum Entropy Production

MEP is but one of a plethora of suggested extremum principles, albeit one which has experienced a wide range of successes. Still the challenge remains of being able to identify the appropriate principle to apply to a given system a priori. In part, this is down to the difficulty in disentangling the competing principles for specific systems. Here, we attempt to disambiguate between a principle of maximum heat flux, and MEP by tuning λ such that dissipation through the fluid is more rapid than forcing, entering the slow-forcing regime of Fig. 7.4 where the two principles are mutually exclusive.

Rapid-forcing Regime

In the rapid-forcing regime, the thermal forcing on the model boundaries occurs on much shorter time scales than the dissipation of the established gradient. This exists for small λ where forcing is sufficiently strong that a principle of MEP and maximum heat flux would be indistinguishable, as illustrated in Fig. 7.4. In our model, $\lambda \approx 20$ results in $q < q_{\text{crit.}}$, shown in Fig. 7.7.

Slow-forcing Regime

In contrast, $\lambda \approx 100$ finds $q > q_{\text{crit.}}$, where MEP and maximum heat flux are mutually exclusive. The results of simulation are shown in Fig. 7.8. The heat flux curves of Fig. 7.8(a) are characteristically identical to those in the strong-forcing case, Fig. 7.7(a), while the corresponding entropy production curves are inverted. Here, maximum heat flux is equivalent to a minimum entropy production principle and we conclude that the model as described appears to be represented by a principle of maximum heat flux, rather than of MEP considering only internal entropy production.

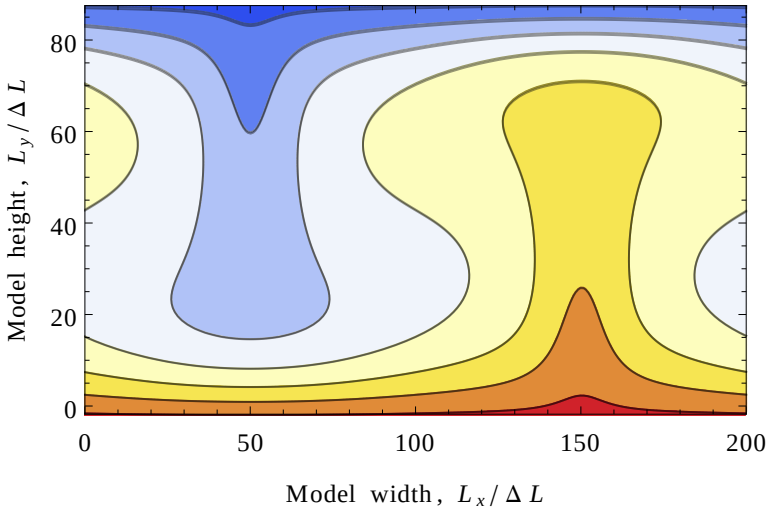

(a) Period-1 convection cell

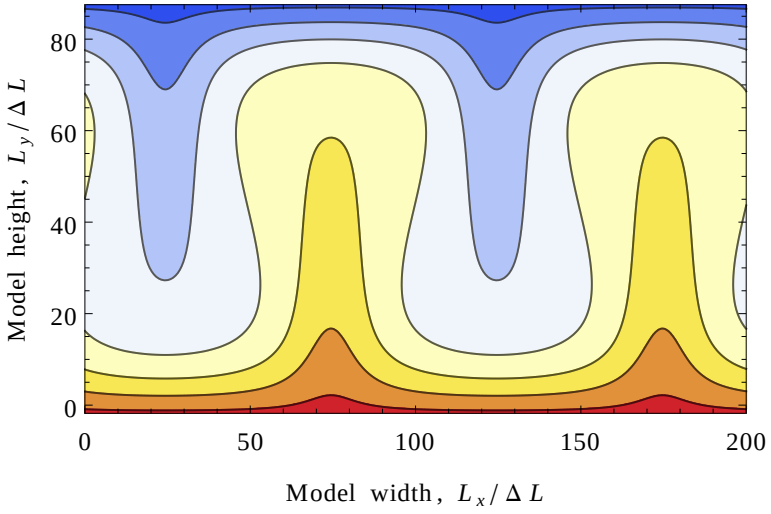

(b) Period-$\frac{1}{2}$ convection cells

FIGURE 7.5: Convection cell configurations produced by the modified lattice Boltzmann model with $\tau_1 = \tau_2 = 20$, $\frac{L_x}{\Delta L} = 200$. Contours show isotherms and are shaded light to dark with decreasing temperature.

7.4 Discussion

Through a principled and reductionist approach, we have implemented the modified FHP lattice gas of Schaffenberger et al. (2001) as a lattice Boltzmann model, and mapped the heat flux and entropy production characteristics of a range of convection cell periods for a range of model dimensions. We find that more efficient dissipative convection configurations are universally more stable resulting in hysteresis loops between convection regimes. It is found that in the regime of rapid forcing, assuming the system self-organises to maximise heat flux or maximum entropy production would lead to accurate predictions of the systems behaviour. In regimes of slow forcing the two assumptions are diametrically opposed: the system continues to self-organise into states of maximum

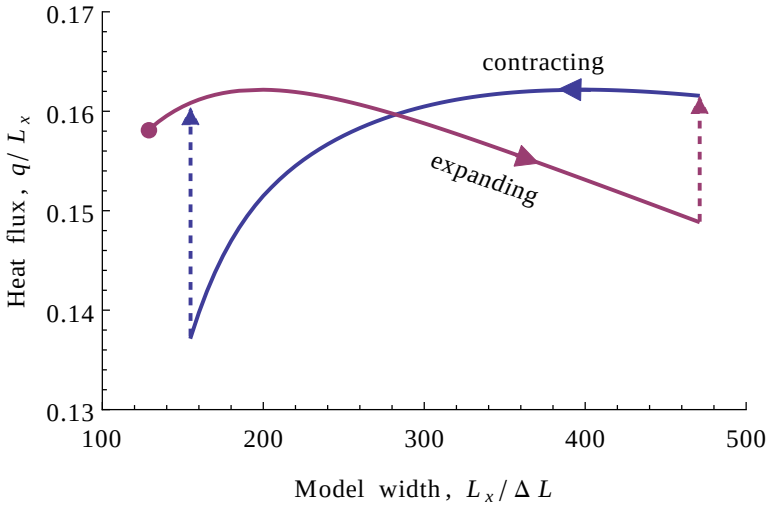

FIGURE 7.6: A hysteresis loop formed from expansion and contraction model with $\lambda = 20$. Expanding a period-1 convection cell results in a metastable convection regime. Fluctuations caused by the slow expansion cause a phase transition to a more efficient period-$\frac{1}{2}$ configuration.

heat flux which are associated with minimum entropy production. Consequently, having MEP emerge within a system is not sufficient evidence for us to conclude that the principle of MEP will accurately predict steady states for that system over a range of forcing or assumptions.

How are we to understand this result in the light of previous studies that have accurately reproduced the steady states of non-equilibrium systems by assuming them to be those characterised by maximum entropy production?

It is important to note that, in our experiment, the heat transport from hot to cold is an emergent property. This differs from from previous studies where heat flux has been manually adjusted to maximise entropy production (Ozawa and Ohmura, 1997; Paltridge, 1975; Lorenz et al., 2001). We are able to reconstruct the parabolic entropy production function via altering the rate of thermal forcing acting on the hot and cold reservoirs. Next we need to revisit the formulation of entropy production. We have considered an open system; only the aspect of the model which exhibits self-organisation, the lattice gas between the heated surfaces, contributes to the entropy production. However for a closed system, it has been shown that inclusion of the entropy production in heating the surfaces results in the entropy production S being a linear function of the heat flux Q. In this case principles of maximum heat flux and entropy production would be equivalent for any forcing rates. This has proven an important caveat in previous applications of MEPP. As an example, for some simple climate box models, including entropy production in absorption and emission of radiation leads to the conclusion that maximum entropy production occurs where diffusion rates are infinite (Essex, 1984).

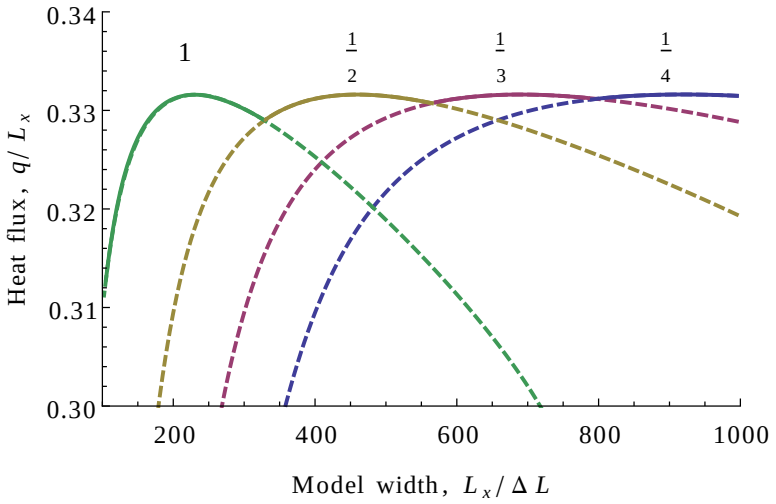

(a) Heat flux

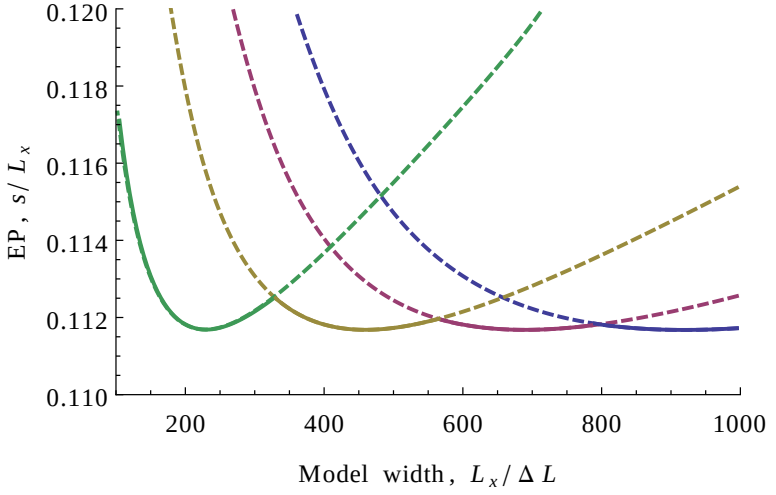

(b) Entropy production

FIGURE 7.7: Specific heat flux (Fig. 7.7(a)) and entropy production (Fig. 7.7(b)) with model width L_x for $\lambda = 20$. Labels show the period of the emerging convection cells in terms of L_x. Dashed lines indicate metastable states, where transitions to higher q states are favoured. Crossover points between different maximally efficient convective regimes are in exactly the same position in both cases. Assumptions of maximum heat flux or MEP accurately represent the system's behaviour. It is impossible to distinguish between a maximum heat flux and MEP principle, consistent with the expectation of the rapid-forcing regime.

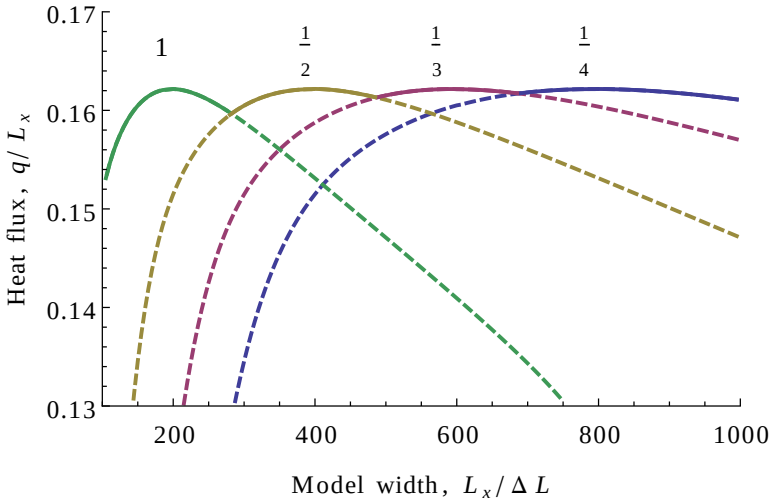

(a) Heat flux

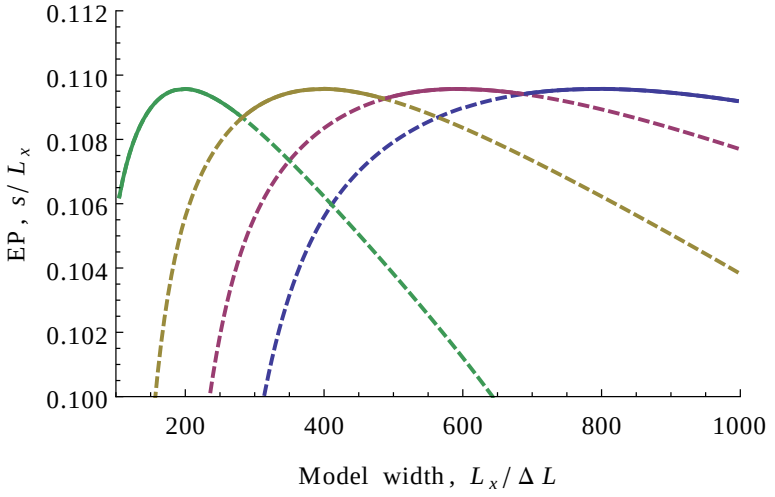

(b) Entropy production

FIGURE 7.8: Specific heat flux (Fig. 7.8(a)) and entropy production (Fig. 7.8(b)) with model width L_x for $\lambda = 100$. Labels show the period of the emerging convection cells in terms of L_x. Here we can see that MEPP predictions would select inefficient convection cells, while we find the model to favour maximally efficient configurations.

Agreement with observations requires us to only consider the entropy produced by dissipation within the Earth system where we assume there to be sufficient degrees of freedom to realise a range of dissipative states.

Finally we need to consider the epistemological basis of the principle of MEP. As has recently been argued, the principle of MEP may not be understood as a physical law, but rather as an information theoretic procedure or algorithm (Dewar, 2009; Dyke, 2010). It is when predictions deviate from observations that we gain information about those systems and processes we are studying.

Chapter 8

Conclusion

8.1 Chapter 2

This thesis has opened the discussion of the Gaia hypothesis with Watson and Lovelock's (1983) initial foray into conceptual ecosystem modelling; Daisyworld, a minimalistic model of two distinct species interacting with an environment so simple that only the global temperature influences the proliferation of the biota. The contribution of this model was to show that maintaining stable and favourable environmental conditions for life to proliferate requires nothing more complicated than interspecies competition providing the traits of the biosphere are correctly designed. Specifically, the model uses species with an increasing and decreasing influence on the global temperature, called black and white daisies. It is important that black daisies out compete their white counterparts when the temperature is relatively low. This configuration leads to a reign control state; while having an individually increasing or decreasing influence on global temperature, the competition between species leads the system to stable configurations.

By separating the model time scales associated with the removal and proliferation of daisies, changes in global temperature and changes in solar illumination on the Daisyworld planet we have shown that the existence of a homeostatic fixed point emerges trivially, and leaves us poised to relax the assumption of disparate time scales. Firstly, allowing the daisy coverage and global temperature to vary on similar time scales we introduce oscillating behaviour in the vicinity of the model fixed point with a period related to these time scales. Additionally, while previous studies fix the changes in solar illumination to the longest time scales, we find that relaxing this constraint imposes limits on the ability of the model to exhibit homeostasis; should the biota respond too slowly to changes in solar forcing, homeostasis abruptly disappears. Finally, we give the equations for a spatial implementation of Daisyworld in the form of cellular automata updating rules and show that in a mean field approximation, where the neighbours of each cell are assumed to be of the global average distribution, Watson and Lovelock's (1983) equations are exactly recoverable, and the behaviour of the model is unchanged. It is noteworthy however that the mean field approximation does not accurately reflect the distribution of black and white daisies which show strong clustering and that a method to resolve such spatial inhomogeneities is introduced in Chapter 5.

8.2 Chapter 3

Daisyworld is a useful proof of concept and without it much of the following research could not have happened. It does however have a number of short comings. In this chapter we focused on two main points which are the single environmental variable which strongly limits the potential for complex behaviour, and the apparent need to design the biosphere in order that homeostasis emerges rather than extinction. We address the first point by incorporating multiple environmental variables, and the second by allowing the

model to encompass an arbitrarily large amount of diversity in the biosphere by randomly parametrising not only the preferred environmental conditions of the biotic elements, but also their influence on each dimension of their environment. This enables configurations of species which, in the context of Daisyworld, would have a destabilising influence on the environment. Having developed an abstract biosphere model, we discovered a plethora of complex behaviour including an abundance of homeostatic fixed points connected in complicated attractor landscapes. This result is important as in this model, homeostatic fixed points have emerged not by design, but from the interactions between a diverse and randomly parametrised biota; a particularly relevant result in the context of the mounting evidence for a multiplicity of stable states in real Earth systems (Scheffer, 1990; Scheffer and Jeppesen, 1998; Van Geest et al., 2007).

We go on to prove the generality of this result using novel analytical methods, deriving an expression for the expected number of fixed points. Furthermore, we can apply variable forcing terms to the environmental dimensions, akin to the solar insolation term of the original Daisyworld model, and derive the distribution of model fixed points with external perturbing force. As expected, the number of homeostatic configurations decreases with the magnitude of perturbing forces. Furthermore we study the behaviour of the model in the vicinity of fixed points, determining the relative frequency of fixed points with spiralling, or oscillatory dynamics in a two environmental variable model.

A natural direction to extend the model is to provide a framework for a spatially distributed biota with a local environmental configuration in a similar vein to von Bloh et al.'s (1997) spatial implementation of the original Daisyworld model. In such a model, both environmental conditions and the composition of the biota are able to diffuse by some process, and the governing equations are extended to incorporate this; the time evolution of the biota and that of the environment, given by Eqs. (3.1) and (3.5) respectively, become

$$\tau_\alpha \frac{\mathrm{d}\boldsymbol{\alpha}(\boldsymbol{x},t)}{\mathrm{d}t} = \boldsymbol{\alpha}^*(\boldsymbol{E},\boldsymbol{\mu}) - \boldsymbol{\alpha}(\boldsymbol{x},t) + D_\alpha \nabla^2 \boldsymbol{\alpha}(\boldsymbol{x},\boldsymbol{t}) \tag{8.1}$$

$$\tau_E \frac{\mathrm{d}\boldsymbol{E}(\boldsymbol{x},t)}{\mathrm{d}t} = \boldsymbol{F}(\boldsymbol{x},t) + \boldsymbol{P}(t) + D_\tau \nabla^2 \boldsymbol{E}(\boldsymbol{x},t) \tag{8.2}$$

where $\boldsymbol{x}$ denotes our spatial vector, and $D_{\alpha,E}$ parametrises the diffusivity of the biotic composition and environmental variables. For the sake of simplicity, this expression causes all biotic elements to diffuse equally, although in principle an index may be included to generalise the process further.

In this work we have consistently used the limit of a rapidly adjusting biota, where $\tau_\alpha \ll \tau_E$. This is relatively easily implemented in a spatial model; the steady-state distribution of the biota may be computed between updates in the distribution of environmental variables ensuring separated time scales. Fig. 8.1 shows this model is able

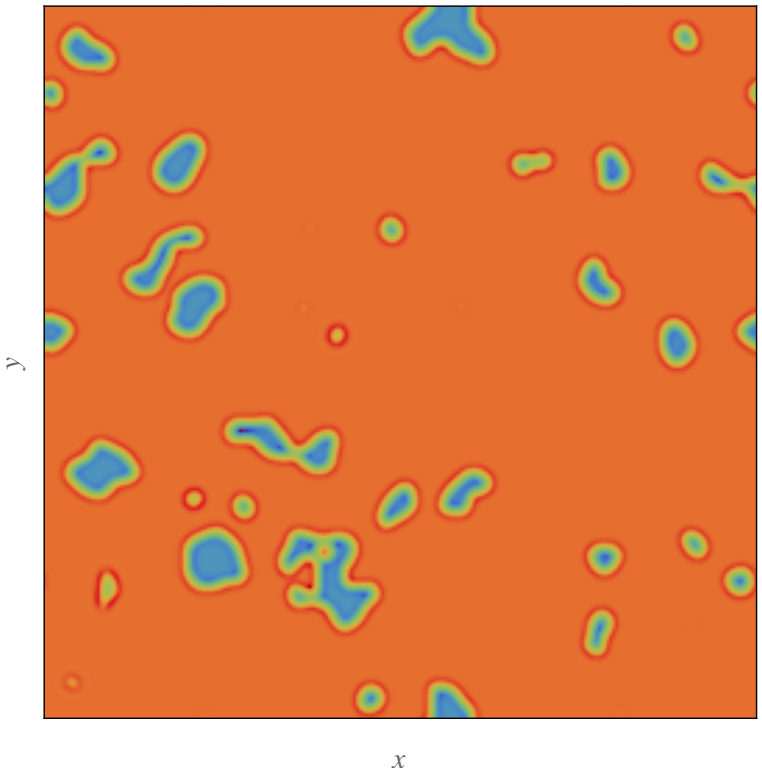

FIGURE 8.1: A spatial implementation of our model with one environmental variable where x and y are the two spatial coordinates. Shading is red to blue for high and low environmental values respectively. This model allows for multiple fixed points to be expressed simultaneously. Environmental converge to fixed points which are locally stable, leading to stationary distributions which are inhomogeneous.

to generate inhomogeneous environments, that is unlike our model, the environment is able to simultaneously occupy distinct, locally-stable fixed points.

8.3 Chapter 4

In Chapter 3, we developed a conceptual ecosystem model based on a minimal set of assumptions, a diverse biota of independent elements and their interaction with a multidimensional environment. We went on to study the stationary behaviour of the model by deriving the distribution and characteristics of model fixed points. In this chapter we proceeded to subject the model to perturbations in different environmental dimensions finding new phenomena to emerge such as cusp bifurcations where a transition between fixed points caused by varying perturbing forces in one environmental variable may be discontinuous or smooth depending on the forcing on other environmental dimensions. Precisely as before, the complex behaviour of the model belies a large amount of tractability which we utilise to derive a number of analytical results. We compare the likelihood of encountering a transition between fixed points with a unit increase or decrease in perturbation magnitude, finding that large positive or negative perturbations

correspond to a relatively high probability of encountering a transition due to further increases in forcing magnitude compared to decreases. In other words, a large positive perturbing force is more likely to cause the system to transition to a new state due to further increases in perturbation than decreases; an intuitive but potentially important result for the management of stressed ecosystems.

Our model describes an Earth system which, in the absence of changes in external forcing factors, approaches a fixed point where it would stay indefinitely. In reality the Earth system has been in a state of constant flux, undergoing smooth and abrupt changes, not necessarily always predicated by changes in abiotic forcing parameters. Organismic evolution appears to have been responsible for dramatic climatic shifts; as organisms evolve to utilise new environmental resources, they significantly change their concentrations and therefore the composition of the environment. As we have seen numerous times through the course of this thesis, non-linear feedbacks can very simply translate significant changes to large-scale, and catastrophic transitions.

One simple implementation of this idea to Daisyworld was carried out by Robertson and Robinson (1998) and allows daisy species to adapt their optimal temperature such that it converges on the current model temperature. The result is to introduce a new model behaviour where if adaptation occurs sufficiently quickly compared to changes in the external forcing, which determines the time scale of change in the model climate, the biota is able to persist in a roughly static configuration for arbitrarily long times. In our model, it is unclear what effect introducing this convergence would have. Species may for example converge their initially random niche position $\boldsymbol{\mu}$ on the current value $\boldsymbol{E}$ at a rate proportional to their abundance; potential members of the biota which are not abundant, or are dormant do not have the possibility of adapting to the current climate.

Williams and Lenton (2007) on the other hand have a more physically grounded model to capture the effects of an evolving biota on their environment. Their Flask model consists of a diverse microbial community in a vat of fluid with an inlet and outlet for various chemicals which, in the absence of a biota, would reach some steady state. Allowing microbes to adapt their required inputs and outputs produces steady-states in the chemical composition of the model which are periodically disrupted by the emergence of rebels, organisms which utilise the environment more efficiently to reproduce, but in doing so generate destabilising quantities of other chemicals. Disruptions may drive the flask to different steady-states, or even lead to internally generated ecosystem extinction. There is no sense of individual efficiency, or unexploited resources in our model. Instead, elements of the biota could be freed to evolve their effect on the environment by modifying the effect matrix $\boldsymbol{\Omega}$

$$\boldsymbol{\Omega}(t) = \boldsymbol{\Omega}(0) + \boldsymbol{\Xi}(t) \tag{8.3}$$

where $\Xi(t)$ is a matrix whose elements are independent random processes such as a Wigner function. The effect function would now evolve through time according to the statistics of $\Xi(t)$, such as that illustrated by Fig. 8.2(a). In response to these variations, it is unclear whether fixed points will vary smoothly, or if we expect they can annihilate over time and cause major transitions without the application of external perturbing forces. The answer is a combination and is shown by Fig. 8.2(b). Fixed point environmental conditions wander smoothly and randomly over time, with the only major changes corresponding to abrupt transitions, where fixed points annihilate and the system reconfigures to another stable point. The picture is very similar to an evolutionary system which displays punctuated equilibrium; long periods of slow change disrupted by intermittent, spontaneous catastrophes where major aspects of the biota and environment undergo rapid change. The existence of the phenomena was initially contentious when Gould and Eldredge (1977) conjectured biological evolution to occur in bursts, separated by relatively long periods of low activity, and has since been accepted as a staple of evolutionary theory (Gould and Eldredge, 2000).

8.4 Chapter 5

Very early on in their development, it was noticed that probabilistic cellular automata (PCA) possess useful features for modelling spatial ecological processes (Hogeweg, 1988). However, while equilibrium PCA have received a great deal of attention for their applications to physical phenomena such as ferromagnets and critical fluids (Brush, 1967), the much broader class of non-equilibrium PCA have received relatively little attention, and as such the methods for predicting, understanding and analysing their large scale behaviour are somewhat undeveloped. This chapter has contributed to the development of such methods by introducing a renormalisation algorithm originally implemented by Edlund and Nilsson Jacobi (2010), although our representation retains sufficient generality that it can be applied without modification to PCA of any dimensionality.

The renormalisation applies a scale transformation, coarse graining the PCA in space and time by a constant factor. It then seeks the coarse grained model dynamics which are consistent with the microscopic dynamics. Iterating this procedure tells us about the large scale model behaviour based only on the known microscopic dynamics and our choice of coarse graining procedure. We find an exact renormalisation is almost certainly not possible, and only exists in some special cases. Instead, we can produce estimates of large scale behaviour by minimising the squared error between our coarse grained and microscopic dynamics. Our estimates are improved by weighting this minimisation by the stationary probability distribution of states, a process which has previously been referred to as the dynamically driven renormalisation group (Vespignani et al., 1996). In essence, we concentrate our errors on transitions from states which are relatively improbable on the microscopic level, in favour of better representing common transitions.

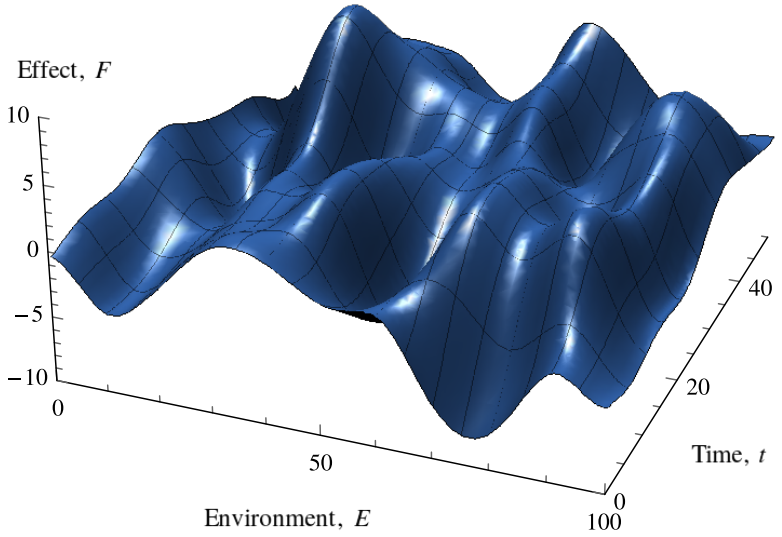

(a) The summed effect of a time varying biota.

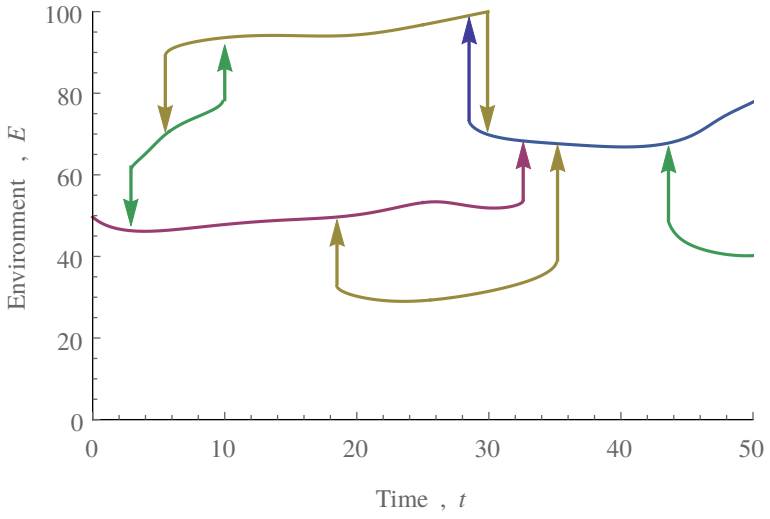

(b) Changes in the environmental variable with a time varying biota.

FIGURE 8.2: A biota which evolves its influence on the environment randomly through time causes the biotic effect to vary, shown in Fig. 8.2(a) for a single environmental variable. The model dynamics are governed by the position of fixed points in this function which wander randomly over time, with major changes corresponding to fixed points becoming unstable, shown in Fig. 8.2(b). Transitions in both directions are shown as the arrow of time in this case is arbitrary and the statistics of the biota are unchanged.

Edlund and Nilsson Jacobi (2010) demonstrated that better estimates of the stationary probability distribution in general produce better estimates of large scale behaviour while we go on to show the approximation to be further improved by renormalising PCA with larger updating neighbourhoods. Additionally, we provide the first application of this type of renormalisation to a two dimensional PCA along with a method to estimate the stationary probability distribution of states for such two dimensional PCA. Nearest neighbour correlations are included by a method of maximum entropy extrapolation (Jaynes, 1982).

There remains much more ground to cover before the method can be applied to more complicated PCA. The algorithm is relatively simple to implement on massively parallel architecture which would enable significant improvements, although the time scaling of inverting large matrices has proven to be a major limiting factor in refining the accuracy of our estimates.

8.5 Chapter 6

Not all interacting systems are well suited to spatial embedding. Locality is one of many types of interaction and multiple elements may be local to one another without interacting in any significant way. Similarly, components may be separated in space though strongly coupled. We can describe more complex system of interaction with networks which link pairs of interacting components explicitly. The analysis of these systems, Network theory, began in principle as long as 250 years ago although its applications have only come to the fore in the past two decades (Shields, 2012). This chapter combines the preferential attachment dynamics of Barabási (1999) with the network growth algorithm of Callaway et al. (2001) and find the combined model to contain similarities with and certain qualitative differences from its parents. As with Callaway et al. (2001), we are able to derive the degree distribution, connected component size distribution and degree correlation coefficient which emerge from a network grown by iteratively adding vertices and edges. The new graph generates a power law tail in the vertex degree distribution, an anticipated effect of preferential attachment although this has consequences for the degree correlation coefficient since the variance of this distribution diverges with increasing connectivity, the degree correlation tends to zero.

Networks are often used to model ecosystems, even Watson and Lovelock's (1983) Daisyworld has found network implementations, investigating the influence of long-range Daisy seeding on the regulation and daisy distribution (Yodzis, 1981; Jeong et al., 2000; Punithan et al., 2011). These networks tend to be very much more complicated than the likes of Barabási (1999), Callaway et al. (2001) and Weaver (2014). As such it remains to be seen what generalities, if any, may be drawn by applying network theory to abstractions of these ecosystem models. In the case of our ecosystem model, introduced

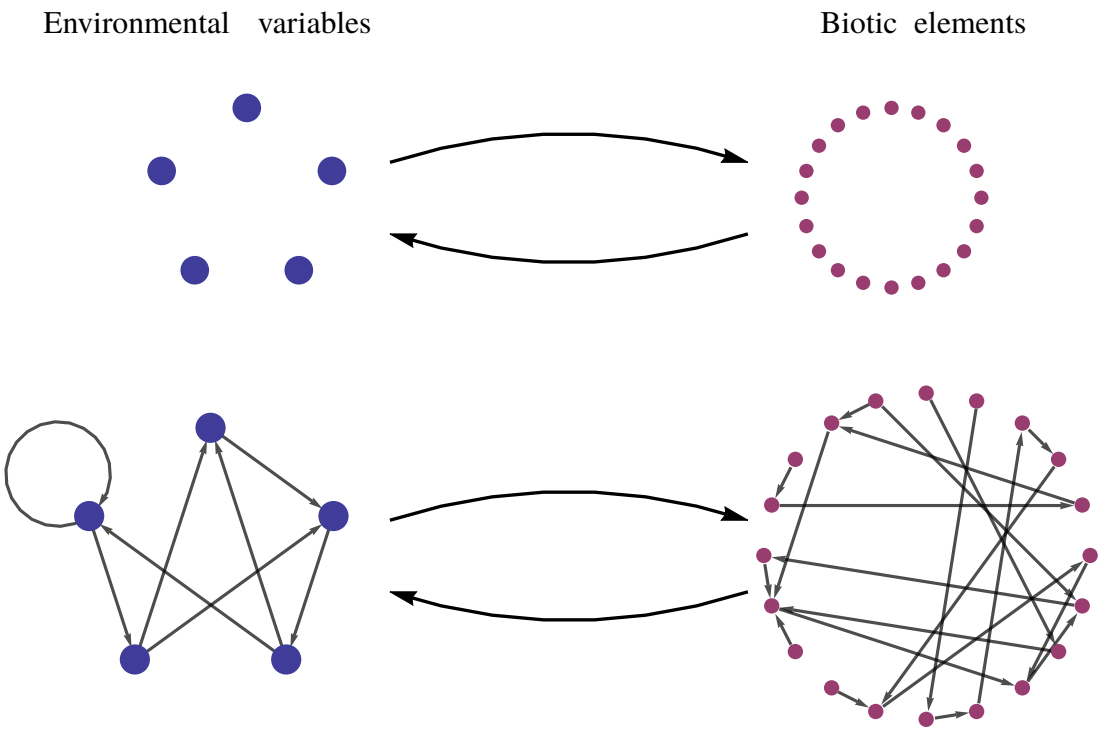

FIGURE 8.3: Our model is simple in the sense that the biota interact only through feedback from their shared environment, and the environmental variables interact only through their shared effect on the composition of the biota. In a more general model, both elements of the model have some internal interaction network.

by Chapter 3, network theory could be used to address perhaps the most obvious criticism of this model; that the model biota and environment contain no self-interactions. The elements of the biota are only able to interact through their shared environment. One species may have an increasing influence on one environmental dimension, which may help or hinder the growth of another species indirectly. This is certainly atypical for biosphere models; even Watson and Lovelock's (1983) Daisyworld allows species to compete for space by factoring both populations into the daisies' logistic growth terms. However, this type of competition has been shown to result in very simple behaviour when considering a larger diversity in the biota; only the two fittest species may coexist in a reign control state. While other species could in principle contribute, this simple type of competition entirely excludes less fit species from the steady state (Lovelock, 1992). Interconnectedness of environmental factors is also absent from the model; abiotic feedback such at between global temperature and chemical weathering processes may have a significant influence on climatic conditions globally. A comparison of the non-interactive model detailed in Chapter 3 to a more general model in illustrated by Fig. 8.3.

8.6 Chapter 7

Up to this point we have not considered any physical laws in establishing collective, or large scale dynamics of our models. In closed systems with a well defined energy, we can exploit the second law of thermodynamics to predict the equilibrium state as that of minimum free energy. Free energy is that which can be converted to perform work such as temperature or chemical gradients. However, there is no corresponding principle known to apply to non-equilibrium systems like atmosphere and ocean circulation within the Earth system. These dynamics are largely driven by a latitudinal temperature gradient, established by energy flux from the sun. This is not to say that there are no candidates for non-equilibrium thermodynamic principles. In this chapter we have compared two in the context of a minimal model of Bénard convection; the principle of Maximum Entropy Production (MEP) and maximum dissipation which postulate that the long time steady state of non-equilibrium systems are those which maximise entropy production, or total dissipation within the physical constraints of the system behaviour (Paltridge, 1975). Bénard convection is chosen as an iconic phenomenon where a fluid subjected to a vertical temperature gradient may exist in one of several potential stable states from a linear vertical temperature gradient to any integer number of convection cells.

We begin with a FHP lattice gas, so named after its authors, Frisch et al. (1986). This is a mass, momentum and energy conserving cellular automaton established on a hexagonal grid (or equivalently a triangular lattice). The model is well known for its fluid-like properties. Indeed the conservation rules stated mean it obeys fundamental transport equations (Succi, 2001, section 2.3). In its original form, thermal anomalies are assumed small and so thermal effects are ignored, although as Frisch (1987) simulate the buoyancy force on relatively hot fluid by modifying the original FHP dynamics. Unfortunately, extracting stationary behaviour from lattice gas models is troublesome, and long runs on a very fine grid are required for a reasonable degree of accuracy (Schaffenberger et al., 2001). Transforming the lattice gas into a continuum model, called a Lattice Boltzmann model, allows for much more efficient and accurate calculation of steady states (Bhatnagar et al., 1954). We add Frisch's (1987) buoyancy modification to the lattice Boltzmann model in order to determine the steady state convective configurations, and assess their stability, entropy production and total dissipation.

We found that the more stable convective configurations appeared to optimise the total dissipation between the hot and cold reservoirs used to establish a temperature gradient rather than the total entropy production although in almost all MEP studies, there remains the question as to which sources of entropy production should be considered as open systems have no clear boundary (Ozawa and Ohmura, 1997; Paltridge, 1975; Lorenz et al., 2001). This type of modelling work is a candidate for experimental validation which may cover some distance towards an understanding of how non-equilibrium systems with multiple stable configurations favour different states.

8.7 Closing remarks

This thesis has made strides towards the marriage of simple, intuitive and most importantly, tractable Earth system models with more complicated, richer models which frequently defy a general analysis. We have seen that separating the time scales of various model processes can be used to greatly simplify its analysis, and that careful model formulation leads to complex models which unite diverse and life-like behaviour with transparent and broad analysis. Furthermore we have developed on existing tools for treating the additional complexities which arise from embedding complex models in space or on networks such as by providing the first ever renormalisation of a cellular automata with an absorbing state in two dimensions, a method which aims to provide a short cut to understanding the large scale behaviour of a breadth of complex models given only their local interaction rules.

This thesis remains open ended. There are a multitude of possibilities to build upon the work of this thesis, which have been detailed in this closing chapter. There are also paths to progress beyond this work, and establish new avenues of research. Particularly the ecosystem model of Chapter 3 which, in its most general formulation, is well posed to find applications beyond those discussed in this thesis. Its large scale behaviour shares features common to many complex systems; socio-economic systems for example may have success in exploiting this type of representation to explore the emergence of stable configurations and an investigation into the important dimensions of control for transitions or regime shifts could have important implications for polity makers.

Bibliography

G J Ackland. Maximization principles and daisyworld. *Journal of theoretical biology*, 227(1):121–128, March 2004. ISSN 0022-5193.

G J Ackland, M A Clark, and T M Lenton. Catastrophic desert formation in Daisyworld. *Journal of Theoretical Biology*, 223(1):39–44, July 2003. ISSN 00225193.

B Adams and J Carr. Spatial pattern formation in a model of vegetation-climate feedback. *Nonlinearity*, 16(4):1339–1357, July 2003. ISSN 0951-7715.

B Adams, J Carr, T M Lenton, and A White. One-dimensional daisyworld: spatial interactions and pattern formation. *Journal of Theoretical Biology*, 223(4):505–513, August 2003. ISSN 00225193.

A Alder and V Strassen. On the algorithmic complexity of associative algebras. *Theoretical Computer Science*, 15(2):201–211, January 1981. ISSN 03043975.

D L Anderson. Theory of the Earth. 1989.

W R Ashby. *Design for a brain*. Chapman and Hall, London, 2nd edition, 1960.

D D Baldocchi, T Krebs, and M Y Leclerc. "Wet/dry Daisyworld": a conceptual tool for quantifying the spatial scaling of heterogeneous landscapes and its impact on the subgrid variability of energy fluxes. *Tellus B*, 57(3):175–188, July 2005. ISSN 0280-6509.

H Balzter, P W Braun, and W Köhler. Cellular automata models for vegetation dynamics. *Ecological Modelling*, 107(2-3):113–125, April 1998. ISSN 03043800.

A-L Barabási. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, October 1999. ISSN 00368075.

D S Battisti and A C Hirst. Interannual Variability in a Tropical AtmosphereOcean Model: Influence of the Basic State, Ocean Geometry and Nonlinearity. *Journal of the Atmospheric Sciences*, 46(12):1687–1712, June 1989. ISSN 0022-4928.

R A Berner. A model for atmospheric CO 2 over Phanerozoic time. *American Journal of Science*, 291(4):339–376, April 1991. ISSN 0002-9599.

P Bhatnagar, E Gross, and M Krook. A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems. *Physical Review*, 94(3):511–525, May 1954. ISSN 0031-899X.

S Brush. History of the Lenz-Ising Model. *Reviews of Modern Physics*, 39(4):883–893, October 1967. ISSN 0034-6861.

D Burraston and E Edmonds. Cellular automata in generative electronic music and sonic art: a historical and technical review. *Digital Creativity*, 16(3):165–185, January 2005. ISSN 1462-6268.

D S Callaway, J E Hopcroft, J M Kleinberg, M E J Newman, and S H Strogatz. Are randomly grown graphs really random? *Physical Review E*, 64(4):41902, September 2001. ISSN 1063-651X.

S R Carpenter and W A Brock. Rising variance: a leading indicator of ecological transition. *Ecology letters*, 9(3):311–8, March 2006. ISSN 1461-0248.

R J Cicerone and R S Oremland. Biogeochemical aspects of atmospheric methane. *Global Biogeochemical Cycles*, 2(4):299–327, December 1988. ISSN 08866236.

M Clynes. Cybernetic implications of rein control in perceptual and conceptual organisation. *Annals of the New York Academy of Sciences*, 156(2 Rein Control,):629–664, April 1969. ISSN 0077-8923.

A Colorni, M Dorigo, V Maniezzo, and Others. Distributed optimization by ant colonies. In *Proceedings of the first European conference on artificial life*, volume 142, pages 134–142. Paris, France, 1991.

P M Cox, R A Betts, C D Jones, S A Spall, and I J Totterdell. Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model. *Nature*, 408(6809): 184–7, November 2000. ISSN 0028-0836.

C C Craig. On the Frequency Function of xy. *The Annals of Mathematical Statistics*, 7(1):1–15, March 1936. ISSN 0003-4851.

V Dakos, M Scheffer, E H van Nes, V Brovkin, V Petoukhov, and H Held. Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences of the United States of America*, 105(38):14308–14312, September 2008. ISSN 1091-6490.

R Dawkins. *The Extended Phenotype*. Oxford University Press, London, 1983.

M J De Oliveira and J E Satulovsky. Renormalization group of probabilistic cellular automata with one absorbing state. *Physical Review E*, 55(6):6377, June 1997. ISSN 1063-651X.

R C Dewar. Maximum entropy production and the fluctuation theorem. *Journal of Physics A: Mathematical and General*, 38(21):L371–L381, May 2005. ISSN 0305-4470.

R C Dewar. Maximum Entropy Production as an Inference Algorithm that Translates Physical Assumptions into Macroscopic Predictions: Dont Shoot the Messenger. *Entropy*, 11(4):931–944, November 2009. ISSN 1099-4300.

E Domany and W Kinzel. Equivalence of cellular automata to Ising models and directed percolation. *Physical review letters*, 53(4):311–314, July 1984. ISSN 0031-9007.

W F Doolittle. Is nature really motherly. *CoEvolution Quarterly*, 29:58–63, 1981.

J G Dyke. Entropy production in an energy balance Daisyworld model. *Artificial Life*, 2008.

J G Dyke. The Daisystat: A model to explore multidimensional homeostasis. In *Artificial Life XI, Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 349–359. MIT Press, Cambridge MA, 2010.

J G Dyke and I R Harvey. Hysteresis and the Limits of Homeostasis: from Daisyworld to Phototaxis. In M Capcarrere, A Freitas, J Bentley, C Johnson, and J Timmis, editors, *Proceedings of VIIIth European Conference on Artificial Life, ECAL 2005*, pages 241–246. Berlin: Springer, 2005.

J G Dyke, J McDonald-Gibson, E A Di Paolo, and I R Harvey. Increasing complexity can increase stability in a self-regulating ecosystem. In F e Costa, L M Rocha, E Costa, I R Harvey, and A Coutinho, editors, *Advances in Artificial Life*, pages 133–142. Springer, 2007.

J G Dyke and I S Weaver. The emergence of environmental homeostasis in complex ecosystems. *PLoS computational biology*, 9(5):e1003050, May 2013. ISSN 1553-7358.

E Edlund and M Nilsson Jacobi. Renormalization of Cellular Automata and Self-Similarity. *Journal of Statistical Physics*, 139(6):972–984, April 2010. ISSN 0022-4715.

C S Elton. The ecology of invasions by animals and plants. 1958.

J W Essam. Directed compact percolation: cluster size and hyperscaling. *Journal of Physics A: Mathematical and General*, 22(22):4927–4937, November 1999. ISSN 0305-4470.

C Essex. Radiation and the Irreversible Thermodynamics of Climate. *Journal of the Atmospheric Sciences*, 41(12):1985–1991, June 1984. ISSN 0022-4928.

A Free and N H Barton. Do evolution and ecology need the Gaia hypothesis? *Trends in ecology & evolution*, 22(11):611–9, November 2007. ISSN 0169-5347.

U Frisch. Lattice Gas Hydrodynamics in Two and Three Dimensions. *Complex Systems*, 1:649–707, 1987.

U Frisch, B Hasslacher, and Y Pomeau. Lattice-Gas Automata for the Navier-Stokes Equation. *Physical Review Letters*, 56(14):1505–1508, April 1986. ISSN 0031-9007.

M R Gardner and W R Ashby. Connectance of Large Dynamic (Cybernetic) Systems: Critical Values for Stability. *Nature*, 228(5273):784, November 1970. ISSN 0028-0836.

K J Gaston. Global patterns in biodiversity. *Nature*, 405(6783):220–227, May 2000. ISSN 0028-0836.

A V Getling. *Rayleigh-Bénard convection: structures and dynamics*, volume 11. World Scientific Pub Co Inc, 1998. ISBN 9789810226572.

V L Girko. Circular Law. *Theory of Probability & Its Applications*, 29(4):694–706, January 1985. ISSN 0040-585X.

N Goldenfeld. *Lectures on Phase Transitions and the Renormalization Group*. Addison-Wesley, 1992. ISBN 9780201554090.

B E Goodison, R D Brown, R G Crane, R B Alley, R Bales, D Barber, R Barry, C Bentley, T Carrol, D Cline, and Others. Cryospheric systems. *EOS Science Plan: T he State of Science in the EOS Program*, pages 261–307, 1999.

S J Gould and N Eldredge. Punctuated equilibria: The tempo and mode of evolution reconsidered. *Palaeobiology*, 3:115–151, 1977.

S J Gould and N Eldredge. Punctuated equilibrium comes of age. *Shaking the Tree: Readings from Nature in the History of Life*, 17, 2000.

S Gregorio, R A Pielke, and G A Dalu. Feedback between a simple biosystem and the temperature of the Earth. *Journal of Nonlinear Science*, 2(3):263–292, September 1992. ISSN 0938-8974.

G Grimmett. *Probability: an introduction*. Oxford University Press, 1986.

R Guimerà, S Mossa, A Turtschi, and L A N Amaral. The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *Proceedings of the National Academy of Sciences of the United States of America*, 102 (22):7794–9, May 2005. ISSN 0027-8424.

R K S Hankin and N Mitchell. Resolving the anomaly of bare habitable ground in Daisyworld. *Tellus B*, 63(1):140–143, February 2011. ISSN 02806509.

J Hardy. Time evolution of a two-dimensional model system. I. Invariant states and time correlation functions. *Journal of Mathematical Physics*, 14(12):1746, December 1973. ISSN 00222488.

J Hardy, O de Pazzis, and Y Pomeau. Molecular dynamics of a classical lattice gas: Transport properties and time correlation functions. *Physical Review A*, 13(5):1949–1961, May 1976. ISSN 0556-2791.

I R Harvey. Homeostasis and rein control: From daisyworld to active perception. In *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems, ALIFE*, volume 9, pages 309–314, 2004.

Alan M. Haywood. From Greenhouse to Icehouse: The Marine Eocene-Oligocene Transition. *Antarctic Science*, 16(4):585–585, December 2004. ISSN 0954-1020.

H Hinrichsen. Non-equilibrium phase transitions. *Physica A: Statistical Mechanics and its Applications*, 369(1):1–28, September 2006. ISSN 03784371.

P F Hoffman, A J Kaufman, G P Halverson, and D P Schrag. A Neoproterozoic Snowball Earth. *Science*, 281(5381):1342–1346, August 1998.

P Hogeweg. Cellular automata as a paradigm for ecological modeling. *Applied Mathematics and Computation*, 27(1):81–100, July 1988. ISSN 00963003.

R Horner, S F Ackley, G S Dieckmann, B Gulliksen, T Hoshiai, L Legendre, I A Melnikov, W S Reeburgh, M Spindler, and C W Sullivan. Ecology of sea ice biota. *Polar Biology*, 12(3-4), September 1992. ISSN 0722-4060.

A Hunt. A new conceptual model for forest fires based on percolation theory. *Complexity*, 13(3):12–17, 2008. ISSN 1099-0526.

A Hunt. Relevance of percolation theory to power-law behavior of dynamic processes including transport in disordered media. *Complexity*, 15(2):13–27, 2009. ISSN 1099-0526.

A Ilachinski. *Cellular automata: a discrete universe*. World Scientific, 2001.

J Imbrie, E A Boyle, S C Clemens, A Duffy, W R Howard, G Kukla, J Kutzbach, D G Martinson, A McIntyre, A C Mix, B Molfino, J J Morley, L C Peterson, N G Pisias, W L Prell, M E Raymo, N J Shackleton, and J R Toggweiler. On the Structure and Origin of Major Glaciation Cycles 1. Linear Responses to Milankovitch Forcing. *Paleoceanography*, 7(6):701–738, December 1992. ISSN 08838305.

IPCC. IPCC Fourth Assessment Report: Working Group II Report "Impacts, Adaptation and Vulnerability". *Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, 2007.

T Ito and A Kleidon. Entropy Production of Atmospheric Heat Transport. In A Kleidon, R Lorenz, and R D Lorenz, editors, *Non-equilibrium thermodynamics and the production of entropy: life, earth, and beyond*, pages 93–106. Springer Verlag, 2005. ISBN 978-3-540-22495-2.

E. T Jaynes. Information Theory and Statistical Mechanics. *Physical Review*, 106(4): 620–630, May 1957. ISSN 0031-899X.

E T Jaynes. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*, 70(9):939–952, 1982. ISSN 0018-9219.

A M Jellinek and A Lenardic. Effects of spatially varying roof cooling on thermal convection at high Rayleigh number in a fluid with a strongly temperature-dependent viscosity. *Journal of Fluid Mechanics*, 629(1):109, June 2009. ISSN 0022-1120.

H Jeong, B Tombor, R Albert, Z N Oltvai, and A-L Barabási. The large-scale organization of metabolic networks. *Nature*, 407(6804):651–654, October 2000. ISSN 0028-0836.

V Kalapala, V Sanwalani, A Clauset, and C Moore. Scale invariance in road networks. *Physical Review E*, 73(2):026130, February 2006. ISSN 1539-3755.

M Kanakidou, J H Seinfeld, S N Pandis, I Barnes, F J Dentener, M C Facchini, R Van Dingenen, B Ervens, A Nenes, C J Nielsen, E Swietlicki, J P Putaud, Y Balkanski, S Fuzzi, J Horth, G K Moortgat, R Winterhalter, C E L Myhre, K Tsigaridis, E Vignati, E G Stephanou, and J Wilson. Organic aerosol and global climate modelling: a review. *Atmospheric Chemistry and Physics*, 5(4):1053–1123, March 2005. ISSN 1680-7324.

J W Kirchner. The Gaia hypothesis: Can it be tested? *Reviews of Geophysics*, 27(2): 223, 1989. ISSN 8755-1209.

J W Kirchner. The Gaia Hypothesis: Fact, Theory, and Wishful Thinking. *Climatic Change*, 52(4):391–408, 2002.

J W Kirchner. The Gaia hypothesis: conjectures and refutations. *Climatic Change*, 58 (1):21–45, 2003.

T Kleinen, H Held, and G 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, June 2003. ISSN 1616-7341.

P Krapivsky, G Rodgers, and S Redner. Degree Distributions of Growing Networks. *Physical Review Letters*, 86(23):5401–5404, June 2001. ISSN 0031-9007.

L D Landau and E M Lifshitz. *Fluid mechanics*. Course of theoretical physics. Pergamon Press, 1987. ISBN 9780080339337.

J S Lansing, J N Kremer, and B B Smuts. System-dependent Selection, Ecological Feedback and the Emergence of Functional Structure in Ecosystems. *Journal of theoretical biology*, 192(3):377–391, June 1998. ISSN 1095-8541.

T M Lenton. Gaia and natural selection. *Nature*, 394(6692):439–447, July 1998. ISSN 0028-0836.

T M Lenton. Early warning of climate tipping points. *Nature Climate Change*, 1(4): 201–209, June 2011. ISSN 1758-678X.

T M Lenton, M Crouch, M Johnson, N Pires, and L Dolan. First plants cooled the Ordovician. *Nature Geoscience*, 5(2):86–89, February 2012. ISSN 1752-0894.

T M Lenton and J E Lovelock. Daisyworld is Darwinian: constraints on adaptation are important for planetary self-regulation. *Journal of Theoretical Biology*, 206:109–114, 2000.

T M Lenton and M van Oijen. Gaia as a complex adaptive system. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 357(1421): 683–695, May 2002. ISSN 0962-8436.

T M Lenton and D M Wilkinson. Developing the Gaia theory - A response to the criticisms of Kirchner and Volk. *Climatic Change*, 58(1-2):1–12, May 2003. ISSN 0165-0009.

I Licata. Almost-anywhere theories: Reductionism and universality of emergence. *Complexity*, 15(6):11–19, 2010. ISSN 1099-0526.

T M Liggett. *Interacting Particle Systems*, volume 276 of *Grundlehren der mathematischen Wissenschaften*. Springer New York, New York, NY, 1985. ISBN 978-1-4613-8544-8.

E N Lorenz. Predictability a problem partly solved. In Tim Palmer and Renate Hagedorn, editors, *Predictability of Weather and Climate*. Cambridge University Press, 2006.

R D Lorenz. Planets, life and the production of entropy. *International Journal of Astrobiology*, 1(1):3–13, May 2002. ISSN 1473-5504.

R D Lorenz, J I Lunine, P G Withers, and C P McKay. Titan, Mars and Earth : Entropy production by latitudinal heat transport. *Geophysical Research Letters*, 28 (3):415–418, February 2001. ISSN 00948276.

J Lovelock. *The ages of Gaia: A biography of our living earth*. Oxford University Press, 1995.

J E Lovelock. A physical basis for life detection experiment. *Nature*, 207:568–570, 1965.

J E Lovelock. *A new look at life on earth*. Oxford University Press, 1979. ISBN 0192862189.

J E Lovelock. A Numerical Model for Biodiversity. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 338(1286):383–391, December 1992. ISSN 0962-8436.

J E Lovelock and L Margulis. Atmospheric homeostasis by and for the biosphere: the gaia hypothesis. *Tellus*, 26(1-2):2–10, February 1974. ISSN 00402826.

R H MacArthur and E O Wilson. *The theory of island biogeography*. Princeton Univ Press, 1967.

F T Mackenzie and J A Mackenzie. *Our changing planet: an introduction to earth system science and global environmental change*. Prentice Hall New Jersey, 1998.

L Margulis and D Sagan. *What Is Life?* University of California Press, 2000. ISBN 9780520220218.

L M Martyushev and V D Seleznev. Maximum entropy production principle in physics, chemistry and biology. *Physics Reports*, 426(1):1–45, April 2006. ISSN 03701573.

R M May. Will a Large Complex System be Stable? *Nature*, 238(5364):413–414, August 1972. ISSN 0028-0836.

K S McCann. The diversity-stability debate. *Nature*, 405(6783):228–233, May 2000. ISSN 0028-0836.

J McDonald-Gibson, J G Dyke, E A Di Paolo, and I R Harvey. Environmental regulation can arise under minimal assumptions. *Journal of theoretical biology*, 251(4):653–66, April 2008. ISSN 1095-8541.

M Mitzenmacher. A Brief History of Generative Models for Power Law and Lognormal Distributions. *Internet Mathematics*, 1(2):226–251, January 2004. ISSN 1542-7951.

K Nagel and M Schreckenberg. A cellular automaton model for freeway traffic. *Journal de Physique I*, 2(12):2221–2229, December 1992. ISSN 1155-4304.

R P Neilson. A Model for Predicting Continental-Scale Vegetation Distribution and Water Balance. *Ecological Applications*, 5(2):362, May 1995. ISSN 10510761.

C Nevison, V Gupta, and L Klinger. Self-sustained temperature oscillations on Daisyworld. *Tellus B*, 51(4):806–814, September 1999. ISSN 0280-6509.

M E J Newman. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98(2):404–409, January 2001. ISSN 0027-8424.

M E J Newman. Assortative Mixing in Networks. *Phys. Rev. Lett.*, 89(20):208701, October 2002a.

M E J Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(1): 16128, July 2002b. ISSN 1063-651X.

M E J Newman. Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, 46(5):323–351, September 2005. ISSN 0010-7514.

E P Odum. *Fundamentals of ecology*. WB Saunders Co, 1971. ISBN 0-7216-6941-7.

L Onsager. Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition. *Physical Review*, 65(3-4):117–149, February 1944. ISSN 0031-899X.

H Ozawa. The second law of thermodynamics and the global climate system: A review of the maximum entropy production principle. *Reviews of Geophysics*, 41(4):1018, 2003. ISSN 8755-1209.

H Ozawa and A Ohmura. Thermodynamics of a Global-Mean State of the AtmosphereA State of Maximum Entropy Increase. *Journal of Climate*, 10(3):441–445, March 1997. ISSN 0894-8755.

G W Paltridge. Global dynamics and climate - a system of minimum entropy exchange. *Quarterly Journal of the Royal Meteorological Society*, 101(429):475–484, July 1975. ISSN 00359009.

R Peierls and M Born. On Ising's model of ferromagnetism. *Mathematical Proceedings of the Cambridge Philosophical Society*, 32(03):477, October 1936. ISSN 0305-0041.

J-R Petit, J Jouzel, D Raynaud, N I Barkov, J-M Barnola, I Basile, M Bender, J Chappellaz, M Davis, G Delaygue, and Others. Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature*, 399(6735):429–436, 1999.

I C Prentice, S P Harrison, and P J Bartlein. Global vegetation and terrestrial carbon cycle changes after the last ice age. *The New phytologist*, 189(4):988–998, March 2011. ISSN 1469-8137.

T Pujol. The Consequence of Maximum Thermodynamic Efficiency in Daisyworld. *Journal of Theoretical Biology*, 217(1):53–60, July 2002. ISSN 00225193.

D Punithan, D-K Kim, and R I B McKay. Daisyworld in two dimensional small-world networks. In *Database Theory and Application, Bio-Science and Bio-Technology*, pages 167–178. Springer, 2011.

D Robertson and J Robinson. Darwinian Daisyworld. *Journal of theoretical biology*, 195 (1):129–34, November 1998. ISSN 0022-5193.

M T Rosing. 13C depleted carbon microparticles in ¿3700-Ma sea-floor sedimentary rocks from West Greenland. *Science*, 283:674–676, 1999.

P T Saunders. Evolution without natural selection: further implications of the daisyworld parable. *Journal of theoretical biology*, 166(4):365–73, February 1994. ISSN 0022-5193.

W Schaffenberger, A Hanslimeier, and M Messerotti. Lattice Gas Model for Two-dimensional Boussinesq Convection. *Hvar Observatory Bulletin*, 25(1):49–60, 2001.

M Scheffer. Multiplicity of stable states in freshwater systems. In Ramesh D. Gulati, Eddy H. R. R. Lammens, Marie-Louise Meijer, and Ellen van Donk, editors, *Biomanipulation Tool for Water Management*, pages 475–486. Springer Netherlands, Dordrecht, 1990. ISBN 978-90-481-4074-9.

M Scheffer, J Bascompte, W A Brock, V Brovkin, S R Carpenter, V Dakos, H Held, E H van Nes, M Rietkerk, and G Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, September 2009. ISSN 1476-4687.

M Scheffer, S R Carpenter, J A Foley, C Folke, and B Walker. Catastrophic shifts in ecosystems. *Nature*, 413(6856):591–6, October 2001. ISSN 0028-0836.

M Scheffer and E Jeppesen. Alternative Stable States. In Erik Jeppesen, Martin Sø ndergaard, Morten Sø ndergaard, and Kirsten Christoffersen, editors, *The Structuring Role of Submerged Macrophytes in Lakes*, volume 131 of *Ecological Studies*, pages 397–406. Springer New York, New York, NY, 1998. ISBN 978-1-4612-6871-0.

D W Schwartzman and T Volk. Biotic enhancement of weathering and the habitability of Earth. *Nature*, 340(6233):457–460, August 1989. ISSN 0028-0836.

J R Scott, J Marotzke, and P H Stone. Interhemispheric Thermohaline Circulation in a Coupled Box Model. *Journal of Physical Oceanography*, 29(3):351–365, March 1999. ISSN 0022-3670.

R Shields. Cultural Topology: The Seven Bridges of Konigsburg, 1736. *Theory, Culture & Society*, 29(4-5):43–57, October 2012. ISSN 0263-2764.

S Shimokawa and H Ozawa. On the thermodynamics of the oceanic general circulation: Irreversible transition to a state with higher rate of entropy production. *Quarterly Journal of the Royal Meteorological Society*, 128(584):2115–2128, July 2002. ISSN 00000000.

D Sornette. Predictability of catastrophic events: material rupture, earthquakes, turbulence, financial crashes, and human birth. *Proceedings of the National Academy of Sciences of the United States of America*, 99 Suppl 1:2522–9, March 2002. ISSN 0027-8424.

H E Stanley. Introduction to phase transitions and critical phenomena. *Introduction to Phase Transitions and Critical Phenomena, by H Eugene Stanley, pp. 336. Foreword by H Eugene Stanley. Oxford University Press, Jul 1987. ISBN-10: 0195053168. ISBN-13: 9780195053166*, 1, 1987.

S Stocker. Regarding mutations in Daisyworld models. *Journal of Theoretical Biology*, 175:495–501, 1995.

H Stommel. Thermohaline Convection with Two Stable Regimes of Flow. *Tellus*, 13(2): 224–230, May 1961. ISSN 00402826.

S Succi. *The Lattice-Boltzmann Equation*. Oxford university press, Oxford, 2001.

C D Thomas, A Cameron, R E Green, M Bakkenes, L J Beaumont, Y C Collingham, B F N Erasmus, M F De Siqueira, A Grainger, L Hannah, L Hughes, B Huntley, A S Van Jaarsveld, G F Midgley, L Miles, M A Ortega-Huerta, A T Peterson, O L Phillips, and S E Williams. Extinction risk from climate change. *Nature*, 427(6970): 145–148, January 2004. ISSN 1476-4687.

T Tomé and M J de Oliveira. Renormalization group of the Domany-Kinzel cellular automaton. *Physical Review E*, 55(4):4000–4004, April 1997. ISSN 1063-651X.

T Tyrrell. *On Gaia: A Critical Investigation of the Relationship between Life and Earth*. Princeton University Press, 2013. ISBN 9781400847914.

G J Van Geest, H Coops, M Scheffer, and E H van Nes. Long Transients Near the Ghost of a Stable State in Eutrophic Shallow Lakes with Fluctuating Water Levels. *Ecosystems*, 10(1):37–47, April 2007. ISSN 1432-9840.

E H van Nes and M Scheffer. Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *The American naturalist*, 169(6):738–47, June 2007. ISSN 1537-5323.

V I Vernadsky. *Biosfera [The biosphere]*. Nauka, 1926.

A Vespignani, S Zapperi, and V Loreto. Renormalization of Nonequilibrium Systems with Critical Stationary States. *Physical Review Letters*, 77(22):4560–4563, November 1996. ISSN 0031-9007.

G Y Vichniac. Simulating physics with cellular automata. *Physica D: Nonlinear Phenomena*, 10(1-2):96–116, January 1984. ISSN 01672789.

N Virgo. From Maximum Entropy to Maximum Entropy Production: A New Approach. *Entropy*, 12(1):107–126, January 2010. ISSN 1099-4300.

P M Vitousek, J D Aber, R W Howarth, G E Likens, P A Matson, D W Schindler, W H Schlesinger, and D G Tilman. Human alteration of the global nitrogen cycle: sources and consequences. *Ecological Applications*, 7(3):737–750, August 1997. ISSN 1051-0761.

V Volterra. Fluctuations in the Abundance of a Species considered Mathematically. *Nature*, 118(2972):558–560, October 1926. ISSN 0028-0836.

W von Bloh, A Block, M Parade, and H J Schellnhuber. Tutorial Modelling of geosphere-biosphere interactions: the effect of percolation-type habitat fragmentation. *Physica A: Statistical Mechanics and its Applications*, 266(1-4):186–196, April 1999. ISSN 03784371.

W von Bloh, A Block, and H J Schellnhuber. Self-stabilization of the biosphere under global change: a tutorial geophysiological approach. *Tellus B*, 49(3):249–262, July 1997. ISSN 0280-6509.

G Wachtershauser. Evolution of the first metabolic cycles. *Proceedings of the National Academy of Sciences*, 87(1):200–204, January 1990. ISSN 0027-8424.

J C G Walker, P B Hays, and J F Kasting. A negative feedback mechanism for the long-term stabilization of Earth's surface temperature. *Journal of Geophysical Research*, 86(C10):9776, 1981. ISSN 0148-0227.

A J Watson and J E Lovelock. Biological homeostasis of the global environment: the parable of Daisyworld. *Tellus B*, 35B(4):284–289, September 1983. ISSN 02806509.

I S Weaver. Preferential attachment in randomly grown networks. *Journal of Statistical Physics (submitted)*, 2014.

I S Weaver and J G Dyke. A novel approach to analysing fixed points in complex systems. In Thomas Gilbert, Markus Kirkilionis, and Gregoire Nicolis, editors, *Proceedings of the European Conference on Complex Systems 2012*, pages 523–533. Springer, 2012a.

I S Weaver and J G Dyke. The importance of timescales for the emergence of environmental self-regulation. *Journal of Theoretical Biology*, 313(0):172–180, 2012b. ISSN 0022-5193.

I S Weaver and J G Dyke. Tipping points in Complex Coupled Life-Environment Systems. In *Advances in Artificial Life, ECAL 2013*, pages 387–394. MIT Press, September 2013. ISBN 9780262317092.

I S Weaver, J G Dyke, and K Oliver. Can the principle of Maximum Entropy Production be used to predict the steady states of a Rayleigh-Bérnard convective system? In Roderick C Dewar, Charley H Lineweaver, Robert K Niven, and Klaus Regenauer-Lieb, editors, *Beyond The Second Law: Entropy Production and Non-Equilibrium Systems*. Springer, 2013.

I S Weaver and A Prügel-Bennett. Renormalisation of two-dimensional cellular automata with an absorbing state. *Journal of Statistical Mechanics (in press)*, 2013.

I S Weaver and A Prügel-Bennett. Renormalization group approach to 1D cellular automata with large updating neighborhoods. *Complexity*, (1), May 2014. ISSN 10762787.

H T P Williams and T M Lenton. The Flask model: emergence of nutrient-recycling microbial ecosystems and their disruption by environment-altering rebel organisms. *Oikos*, 116(7):1087–1105, July 2007. ISSN 00301299.

H T P Williams and T M Lenton. Evolutionary regime shifts in simulated ecosystems. *Oikos*, 119(12):1887–1899, December 2010. ISSN 00301299.

D A Wolf-Gladrow. *Lattice Gas Cellular Automata and Lattice Boltzmann Models*. Springer, 2005. ISBN 978-3-540-66973-9.

A J Wood, G J Ackland, J G Dyke, H T P Williams, and T M Lenton. Daisyworld: a review. *Reviews of Geophysics*, 46(1):RG1001, January 2008. ISSN 8755-1209.

A J Wood, G J Ackland, and T M Lenton. Mutation of albedo and growth response produces oscillations in a spatial Daisyworld. *Journal of Theoretical Biology*, 242(1): 188–198, 2006. ISSN 0022-5193.

A J Wood and J B Coe. A fitness based analysis of Daisyworld. *Journal of Theoretical Biology*, 249(2):190–197, 2007. ISSN 0022-5193 (Print).

P Yodzis. The stability of real ecosystems. *Nature*, 289(5799):674–676, February 1981. ISSN 0028-0836.

X Zeng, R A Pielke, and R Eykholt. Chaos in daisyworld. *Tellus B*, 42(4):309–318, September 1990. ISSN 0280-6509.

G A Zielinski, P A Mayewski, L D Meeker, S Whitlow, M S Twickler, M Morrison, D A Meese, A J Gow, and R B Alley. Record of Volcanism Since 7000 B.C. from the GISP2 Greenland Ice Core and Implications for the Volcano-Climate System. *Science (New York, N.Y.)*, 264(5161):948–52, May 1994. ISSN 0036-8075.