THE DAISYWORLD CONTROL SYSTEM

James Dyke

Submitted for the degree of DPhil

University of Sussex

October 2009

Summary

The original Gaia Hypothesis proposed that life on Earth, along with the oceans, atmosphere and crust, forms a homeostatic system which reduces the effects of external perturbations, so that conditions are maintained to within the range that allows widespread life. Daisyworld is a simple mathematical model intended to demonstrate certain aspects of this planetary homeostasis. There have been a considerable number of extensions and developments to the original Daisyworld model. Some of this work has been produced in response to criticism of the Gaia Hypothesis and Daisyworld specifically and some has been produced by using Daisyworld as a testbed to explore a range of issues.

This thesis examines the Daisyworld control system and in doing so explains how Daisyworld performs homeostasis. The control system is classified as a rein control system which is potentially applicable to a wide range of scenarios from physiological and environmental homeostasis to robotic control. A series of simple Daisyworld models are produced and aspects of the original Daisyworld are explained, in particular the inverse response to forcing: why temperature goes down on Daisyworld when the brightness of the star increases.

The Daisyworld control system is evaluated within an evolutionary context. A key result is that environmental regulation emerges not despite of Darwinian evolution but because of it. Within an ecological context, it is found that increasing the complexity of a self-regulating ecosystem can increase its stability. An energy balance climate model is developed to assess the effects of non-equilibrium thermodynamic processes on the Daisyworld control system. Results are presented that support the hypothesis that when the system is in a state of maximum entropy production, homeostasis is maximised.

Declaration

I hereby declare that this thesis has not been submitted, either in the same or different form, to this or any other University for a degree.

Signature:

Acknowledgments

First, I would like to recognise the contributions of my doctoral supervisor Inman Harvey to this thesis. The time spent with tea and tiffin cake while working on the white board were some of the most productive and happiest moments of my research. Tim Lenton, who acted as an unofficial supervisor was a constant source of support and encouragement and introduced me to the wonders of the Bengal Spice. I would like to thank Axel Kleidon, for his understanding and support during the very final stages of this thesis. I would like to thank Jim and Sandy Lovelock for their generosity, hospitality and inspiration.

Many thanks are due to those that have been kind enough to suffer my contributions to collaborative projects: Graeme Ackland, Jamie McDonald-Gibson, Ezequiel di Paolo and Hywell Williams. Of particular note is Jamie Wood. I am very appreciative of his efforts to get me a job.

This thesis would not exist without Mandi Dyke's constant and at times heroic support, encouragement and ruthless proof reading and copy editing. Thank you. You continue to believe in me. You continue to amaze me.

I would not exist without my parents. So thanks Mum and Dad. Further thanks are due for giving me the opportunities that reason alone would have most probably concluded would have been wasted.

I was funded while doing the research for this thesis by the Engineering and Physical Sciences Research Council. I was very fortunate to have been paid to study.

Finally I dedicate this thesis to Leo and Alexander. It is a natural desire to bring children into a world better than the one you were born into. I hope that the following pages may contain an idea or result or insight that may, at least, not make things any worse.

Publications

Parts of this thesis have been previously published in journal articles and conference proceedings:

Dyke J. G. and Harvey I. R., (2005) Hysteresis and the Limits of Homeostasis: from Daisyworld to Phototaxis, *Proceedings of VIIIth European Conference on Artificial Life*, ECAL 2005, eds Capcarrere, M. and Freitas, A. and Bentley, J. and Johnson, C. and Timmis, J, pp241-246, Springer, Berlin.

Dyke J. G. and I. R. Harvey, (2006) Pushing up the daisies, *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, ALife X, eds L.M. Rocha and L.S. Yager and M.A. Bedau and D. Floreano and R.L. Goldstone and A. Vespignani, pp426-431, MIT Press, Cambridge MA.

Dyke, J. G. and McDonald-Gibson, J. and Di Paolo, E. and Harvey, I. R., (2007) Increasing complexity can increase stability in a self-regulating ecosystem, *Proceedings of IXth European Conference on Artificial Life*, ECAL 2007, pp133-142, Springer, Berlin.

Wood, A. J. and Ackland, G. J. and Dyke, J. G. and Williams, H. T. P. and Lenton, T. M. (2008) Daisyworld: a review, *Reviews of Geophysics* 46 doi:10.1029/2006RG000217.

McDonald-Gibson, J. and Dyke, J. G. and Di Paolo, E. and Harvey, I. R. (2008), Environmental Regulation Can Arise Under Minimal Assumptions, *Journal of Theoretical Biology* 254(4), pp653-666.

Dyke, J. G. (2008), Entropy production in an energy balance Daisyworld model, *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, ALIFE'11, MIT Press, Cambridge MA.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Why Did I Write This Thesis?

This thesis is about Daisyworld [Lovelock, 1983]. It attempts to understand why Daisyworld 'works', why the life forms on Daisyworld (simple black and white daisies) regulate planetary temperature such that it is maintained to within the range of temperatures that the daisies are able to grow over.

There have been a number of critiques that have effectively argued that the Daisyworld model and Gaia Theory more generally, while interesting, require specially crafted conditions or involve not particularly plausible sets of physical assumptions. Consequently they will be of only marginal interest to scientists who are motivated to understand the relationship between biotic and abiotic elements of the Earth's biosphere.

It is important to clearly identify what I will and will not attempt to address with respect to such critiques. I will not assess the biogeochemical plausibility of Daisyworld, Gaia Theory, Geophysiology or associated theories and hypotheses. One common response to criticisms of these theories has been to attempt to increase the plausibility of Daisyworld by incorporating more real-world processes into the model. Increasing the complexity of the model by increasing the numbers of parameters, functions and variables demonstrates that Daisyworld homeostasis can arise under different levels of complexity and physical modelling assumptions. I will adopt a different approach. I will attempt to simplify Daisyworld in order that the mechanism that is responsible for its homeostatic behaviour has no place left to hide. This can be understood as a process of 'opaque thought experiments' [Di Paolo et al., 2000] in which the simple modelling paradigm allows exploration of the mechanisms of interest. Having firmly established

why Daisyworld does what it does, I then analyse a number of more complex models with the knowledge gained from this process.

I will assess the Daisyworld control mechanism in three contexts that are arguably fundamental to all biological processes: Darwinian evolution[1], ecological dynamics and thermodynamics. I will consider how plausible the Daisyworld homeostatic mechanism is when applied to evolving populations. Biological organisms do not live in a vacuum (I ignore here extremophiles that may venture into space due to meteor or asteroid impacts with the Earth). Organisms live within an ecosystem and there is a constant flux of energy and matter between organisms and their environment. I will consider the effects on ecosystem dynamics that Daisyworld homeostasis may produce. Any real-world system must adhere to the laws of thermodynamics. I will consider the effects that certain thermodynamic processes have on the Daisyworld homeostatic mechanism. It has been proposed that these processes confer a notion of directionality on the long-term evolution of the Earth's biosphere. I will assess Daisyworld homeostasis in the light of this aspect of thermodynamics.

A continual if at times implicit theme of this thesis is that the type of homeostasis observed in Daisyworld is a potentially widespread phenomenon and in particular a phenomenon that does not require purposeful design or a 'lucky' coincidence of particular feedback mechanisms or circumstances[2]. I will claim that biological populations and their environments can reduce the impact of external perturbations in ways analogous to the dampening of environmental stress by individual organisms and that this homeostasis is the emergent effect of the interactions between organisms, organismic effects and environmental variables. This claim is not novel. What is new is the analysis of this type of homeostasis and its application to new Daisyworld-type models. Associated with this theme is an evaluation of the role of natural selection in the construction, maintenance and destruction of homeostatic states. I will show that, contrary to initial criticisms of the original Gaia Hypothesis, biologically-mediated homeostasis is compatible with natural selection. Indeed, I will show that Daisyworld homeostasis may arise not despite of but *because* of natural selection.

[1] I use the term 'Darwinian evolution' to be synonymous with 'neo-Darwinism' or the 'modern synthesis'.

[2] There is a sense in which Daisyworld homeostasis requires luck, in that such biologically-mediated homeostasis can only occur on planets with life and there are arguments that life in the universe is extremely improbable.

1.2 Target Audience

1.2.1 Earth system scientists

Daisyworld can be understood as a mathematical proof of concept for the Gaia Hypothesis, [Lovelock and Margulis, 1974], that the Earth's crust, oceans, atmosphere and biota form a homeostatic system that regulates conditions on the surface of the planet that allows widespread life. However, it would be a mistake to assume that Gaia Theory is somehow dependent on Daisyworld. 'Disproving' Daisyworld would not 'disprove' Gaia. Neither would 'proving' Daisyworld necessarily increase the plausibility of Gaia. As I and others have previously argued, Daisyworld has matured to the extent that it is possible to consider it separately from Gaia Theory [Wood et al., 2008]. That said, as the thesis is centred around Daisyworld any reader with an interest in Daisyworld or Gaia will gain deeper insight into Daisyworld and how its homeostatic mechanism can be seen as an instance of a control system that may operate under very different conditions.

The models and analysis of this thesis may be relevant to a systems-level approach to Earth system science. While the models could only be considered as extremely simple representations of extremely complex real-world processes operating within real worlds, the analysis of coupled differential equations and their steady state properties could, in principle, be applied to a wide range of biogeochemical processes on Earth and other planets. In adopting simple or conceptual modelling techniques, what one loses in specificity one may gain in generality.

Earth system scientists will be well versed in analysing models in order to determine their equilibrium and non-equilibrium steady states. What they are typically less inclined to do is posit target values, or goal functions for self-regulating models. This thesis continually attempts to assess simple models in these terms. There may be direct or indirect relevance to other Earth system models. The discussion of non-equilibrium thermodynamics in chapter 8 attempts to form a connection between entropy production in non-equilibrium systems and homeostasis and so may generate new insights into how life interacts with its environment.

1.2.2 Artificial Life practitioners

I will show how the original Daisyworld and a number of its variants are examples of a particular class of homeostatic system. This system may be instantiated in a

population of environmentally altering organisms, physiological processes or robotic controllers. Therefore those researchers motivated to explore the notion of homeostasis, self-regulation and evolution, will, I hope, find the results interesting. This set of research interests is often found within the broad and perhaps poorly defined domain of artificial life [Langton, 1997]. While Daisyworld was not initially conceived as an artificial life model, it can be considered as consonant with artificial life methodologies as Daisyworld explores biospheres as they are and as they could be. In that respect, Daisyworld and associated models can be seen as extending the artificial life paradigm by explicitly modelling the environment and so provide a more comprehensive account of life as it is and life as it could be. Once one fully appreciates the effects that life can have on its environment then an entire vista of behaviour opens up for investigation. For example, previously static fitness landscapes are transformed into plastic surfaces over which life and landscape co-evolve. Fitness peaks can be transformed into valleys and vice versa. Capturing the evolution of both life and environment by explicitly modelling the effects organisms can have on their environments challenges the paradigm which casts the process of evolution as a series of challenge-response interactions between environment and life. Rather than changing environmental conditions posing challenges or questions that evolving populations of organisms respond to or solve, organisms are free to pose challenges to the environment.

Evolutionary computational approaches such as genetic algorithms assume that the environment is fixed and that it is the changing population that must explore it in order to find the highest peaks. Consequently only half of the evolutionary process may be captured. Allowing a population to alter the nature of the problem that they evolving in order to solve, initially sounds absurd. However, the adage of 'there is more than one way to skin a cat' (there is often more than one way or method to achieve a goal) is applicable here. Imposing a fixed landscape limits the population to optimising what has been *assumed* (for whatever reasons) to be the property that requires optimisation.

1.2.3 Ecological and evolutionary theorists

I will claim that a deeper understanding of the Daisyworld control system will have a potentially profound impact on evolutionary and ecological theory. Intimations of this claim can be found in my discussion of particular theories within population genetics and theoretical ecology in chapters 6 and 7. The models I propose and examine are based on very general biological assumptions. I will show that a critical aspect of these models

is the feedback that the biological organisms have on their environment. Explicitly modelling such feedback introduced non-trivial complexities to otherwise simple models. However, the rewards for analysing such models prove to be a deeper appreciation of how organisms and their environments can be seen as co-evolving systems with emergent homeostatic properties. This thesis may be of particular interest to theorists examining frequency dependent selection, niche construction and ecological engineering.

1.3 Layout of Thesis

This is a two-part thesis: In chapters 2 to 5, I will consider the theoretical and analytical basis of what I term the Daisyworld control system and in doing so I will describe the homeostatic mechanism at the core of the Daisyworld model; in chapters 6 to 8 I will apply this analysis to models within the domains of artificial life, ecology, evolution and non-equilibrium thermodynamics.

1.3.1 Chapter 2

In chapter 2 I provide the required background information on Daisyworld and Gaia Theory more generally. I present the original Daisyworld equations and briefly overview their analytical solutions. I review the developments of Daisyworld and how the original model has been extended and used to address a number of different scientific questions.

1.3.2 Chapter 3

In chapter 3 I begin the process of simplification that will highlight the core homeostatic mechanism at work in Daisyworld. A number of the original Daisyworld equations are replaced with simplified, and at times, linear versions. Insights into the steady state properties of the original model are used to produce a simplified algebraic solution for the model. A central objective of chapter 3 is to explore Daisyworld's inverse response to external driving; as the star that warms Daisyworld increases in brightness, there is a range of conditions under which the temperature on the surface of the planet *decreases*. Understanding the inverse response allows deeper insights into the original Daisyworld model and how the homeostatic mechanism operates. In particular, I show that the inverse response is due to a particular modelling assumption with regard to how the daisies' temperature is differentiated from the global temperature. Rather than this

being of limited interest, this feature of the original Daisyworld model illuminates a number of fundamental aspects of the control system that will be explored further within chapter 5.

1.3.3 Chapter 4

In chapter 4 I introduce the Two Box Daisyworld model. This continues the process of simplification begun in chapter 3 and further explains the origins of the inverse response to perturbations. A secondary objective of chapter 4 is to investigate hysteresis in Daisyworld. Hysteresis is observed in Daisyworld with different steady state solutions being found when the star increases or decreases in brightness. A number of counter-intuitive results related to hysteresis and homeostasis in Two Box Daisyworld models are presented. Understanding these results highlights elements of the Daisyworld control system. Finally, I present a simple Daisyworld-type robotic controller that performs phototaxis (light following behaviour) in order to demonstrate that the homeostatic mechanism under consideration is, in principle, applicable to a range of different domains.

1.3.4 Chapter 5

In chapter 5 I further distil the Daisyworld model into a minimal model of homeostasis. I will argue that this much simplified homeostatic system is a rein control system. My analysis in this chapter is based on the previous work of Peter Saunders and his formulation of integral rein control. I introduce the notion of antagonistic rein control (mathematically equivalent to a subset of Saunder's integral rein control formalism) and contrast it with benign rein control. I show how the original Daisyworld is an example of an antagonistic rein control system. A key finding of chapter 5 is that under certain conditions, antagonistic rein control systems will typically exhibit zero steady state error. However, unlike classical and modern control theory techniques that produce equivalent behaviour, pre-defined set points or target values in antagonistic rein control systems can be entirely absent.

1.3.5 Chapter 6

A criticism of the original Daisyworld model was that the choice of physical modelling assumptions were in some way specially crafted in order to ensure that homeostasis

emerged. In chapter 6 I present a simplified version of a Daisyworld model first proposed by Jamie McDonald-Gibson that is largely immune to such criticisms. Numerical results from individual-based simulations are used to identify the core homeostatic mechanism of this version of Daisyworld. I apply the previous chapter's analysis in order to show that this mechanism is a rein control system and emerges not despite natural selection, but *because* of it. I make a number of connections between this model and population genetics models used to explore the proposed mechanism of niche construction.

1.3.6 Chapter 7

In chapter 7 I continue the analysis of the previous chapter's model, with particular emphasis on investigating the relationship between complexity and stability. I show how the McDonald-Gibson model can produce multiple steady states and that a number of processes are responsible for the transitions between homeostatic states. Over suitably long time-scales, the behaviour of the model is similar to that of a random walk. I examine the role of natural selection in the creation, maintenance and destruction of these homeostatic states. I evaluate the relationship between mutation rates and homeostasis and explain why increasing mutation rates can increase homeostasis. I argue that these findings inform the complexity - stability debate within theoretical ecology as a more complex ecosystem is more stable when stability is defined in terms of ability to regulate environmental conditions and when complexity is defined in terms of species diversity. I make a number of connections between the model and the notion of ecosystem engineering.

1.3.7 Chapter 8

In chapter 8 I return to the Two Box Daisyworld in order to assess the Daisyworld control system within the context of thermodynamics. The principle of maximum entropy production, proposes that the Earth's atmosphere is in a state of maximum entropy production with respect to the flow of heat from the hot equators to the cold poles. I develop an energy-balance Daisyworld model and impose maximum entropy production on it in order to assess the effects of the principle on Daisyworld homeostasis. I show that when this Daisyworld model is in a state of maximum entropy production, it is maximally homeostatic.

1.3.8 Chapter 9

The thesis concludes with chapter 9 where I summarise the main results and indicate future work that could develop from this thesis.

1.4 Main results of thesis

The following is a list of the main results of this thesis:

- The origins of the inverse response to driving in the original Daisyworld model are explained and in doing so the question of why temperature in Daisyworld goes down whilst the luminosity of the star goes up is answered. (Chapters 3 & 4)
- A range of Daisyworld-type models are classified in terms of *rein control* and in particular how the original Daisyworld is an example of an *antagonistic rein control* system. It is shown that in certain rein control systems, increasing the robustness of individual elements to perturbations can lead to a decrease in the system's robustness to perturbations. (Chapters 4 & 5)
- Within the domain of evolutionary theory and in particular the population genetics theory of *niche construction* it is shown how evolving populations can alter their environments in ways that lead to environmental homeostasis. This homeostasis emerges via the operation of natural selection on diverse populations. (Chapter 6)
- Within the domain of theoretical ecology and in particular the theory of *ecosystems engineering* it is shown how a more diverse and complex ecosystem can be more robust to external perturbations as such an ecosystem possesses a more effective rein control mechanism. (Chapter 7)
- Within the domain of non-equilibrium thermodynamics the relationship between homeostasis and *entropy production* in a Daisyworld model is explored. It is shown that a Daisyworld that maximises the rate of entropy production is maximally homeostatic. (Chapter 8)

1.5 Summary

In this chapter I have stated that the central objective of the thesis is to answer the question: How does Daisyworld work? I have claimed that answering this question will provide insights not only to Daisyworld and Gaia Theory, but to the domains of theoretical ecology, evolutionary theory and artificial life. I have stated that the central homeostatic mechanism operating within Daisyworld is a rein control system and that this system is not limited to biospheres or living systems and could, in principle, be quite general. I have listed the main findings and layout of the thesis.

The following chapter will begin by providing the required background information on Daisyworld and Gaia Theory and provide definitions for central concepts.

Chapter 2

Background

2.1 Overview

In this chapter I have assumed that some readers will be more familiar with certain concepts and theories than others. For example, an artificial life practitioner may be well versed in classical control theory but be largely unaware of the developments in Gaia Theory. Consequently the chapter is largely structured as a checklist of concepts, definitions and background theory that is required for later chapters.

I begin by defining 'homeostasis' as it is employed in this thesis. I also introduce the notion of 'homeorhesis'. Homeostasis and homeorhesis will be discussed within the context of classical control theory. I then introduce Gaia Theory which proposes that the biota of the Earth is a component in a planetary system that exhibits homeostasis in the face of external perturbations. I present the original Daisyworld model which can be understood as a mathematical proof of concept for Gaia Theory. I review the analytical solutions of the original model and then a number of developments and extensions that have allowed a range of issues to be explored.

2.2 Definitions and Concepts

2.2.1 Homeostasis

I do not propose to offer any overarching definition of homeostasis. Indeed, in this thesis I will employ more than one method to assess a system's homeostatic properties. That said, I would begin any definition of homeostasis with the claim that it would have to include an appreciation of dynamical processes of *feedback*. The provisional definition of

a homeostatic system is a system that responds to perturbations via feedback processes such that the effects of the perturbations on particular aspects of the system are reduced.

Under this definition, a lump of granite is not a homeostatic system with respect to thermal driving[1]. Perturbations in the form of heating will, in time, increase the temperature of the rock. Endothermic animals such as mammals are homeostatic in the respect that changes in external temperature are regulated against 'in order to' maintain core body temperature to within a certain range[2]. Mechanisms to maintain a minimum core body temperature could be shivering and reduction of blood circulation to the extremities. If core body temperature increases to the upper limits of this essential range, then sweating and dilation of capillaries will lower core body temperature. Exothermic animals such as reptiles are not homeostatic in that respect. They do not regulate internal body temperature via physiological processes. However, they can regulate internal body temperature via behavioural processes. They also possess a range of physiological processes that regulate other metabolic variables. Any single system may possess a range of different homeostatic processes.

It is important to be clear about over what time scales the homeostatic properties of systems will manifest and be assessed. Over suitably short time scales a lump of granite can be evaluated as a homeostatic system with respect to external heating as the heat capacity of granite will produce a much smaller increase in the temperature of the rock than a material with a lower heat capacity. Over suitably long time scales no biological organism is a homeostatic system as no organism lives forever and the matter and energy that comprise it will return to lower energy states via decomposition or by being consumed by another organism. As this thesis is not focussed on a particular real-world system, there is an ever present danger that the time scales selected, along with other parameter values and modelling assumptions, could significantly bias the assessment of homeostasis. Wherever possible I will make unambiguous declarations of the time scales under consideration.

In this thesis I will discuss dynamical systems that have a single variable that is subject to regulation. The formalism developed and discussed is in principle applicable to

[1]However, a lump of granite may be a component of a simple homeostatic system. For example, a rock at the bottom of a valley will, if perturbed by being moved up the sides of the valley, move due to the force of gravity back down to the lower energy state at the bottom of the valley.

[2]The expression 'in order to' was contained within scare quotes as it is important to be clear that homeostatic systems need not be teleological. Human designed and built control systems clearly have goal functions and purposes. One should use teleological language with care when discussing non-human created homeostatic systems.

multiple variable systems which feature multiple dimensions of perturbations and feedback. Discussions of dynamical systems with a single regulated variable allow simpler and consistent exposition of the essential ideas of the thesis. Homeostasis is observed in these systems when a perturbing force or input attempts to change the value of the regulated variable and internal processes respond in such a way as to reduce the effect of this perturbing input. This variable is critical to the survival of the system. If the variable decreases or increases below certain thresholds, the survivability of the system is endangered. Following Ashby [Ashby, 1960] I define such internal variables as *essential variables*. The range of values that the essential variable must be maintained within is the *essential range*. For example the essential variable of core body temperature in Homo sapiens must be maintained within the essential range of approximately 35-41 degrees Celsius.

2.3 Homeorhesis

In this thesis I will examine systems in which the essential range may change within the lifetime of a system. Arguably the concept of *homeorhesis* is more relevant here and it is this concept rather than homeostasis that Margulis, co-founder of the Gaia Hypothesis uses to describe Gaian self-regulation [Margulis and Sagan, 1997]. The term was coined by Waddington within a context of developmental biology in order to describe the tendency of an organism to continue development towards a particular given state [Waddington, 1957]. Associated with homeorhesis is the concept of *chreod* or necessary path which can be understood as the trajectory which the system returns to after a perturbation. Such notions can be visualised as the path rain water may take through a series of mountains, valleys and watersheds. Rain falling onto a mountain will run down its sides and end up in a particular watershed such as a valley between two mountains. Streams coalesce into rivers that make their way to the sea. The watershed, the path down valleys between mountains, is a particular trajectory of the system. Large boulders and even dams can impede the progress of a river, but the overall trajectory can still be maintained. However, there may be points along that trajectory where relatively small perturbations could move an amount of water into a different trajectory. For example, a small change in initial position on the side of the mountain could lead to the water running into a different watershed and so follow a different trajectory. An interesting mathematical treatment of homoerhesis is given by [Mamontov, 2007]

who develops an ordinary differential equation model of homeorhesis. He argues that homeostatic systems are in fact a particular subset of homeorhetic systems, with homeostatic systems being necessarily constrained to a single chreod. Within this formalism the original Daisyworld model is a single chreod constrained homeorhetic system. In later chapters I will discuss multiple chreod homeoretic systems that are able to occupy more than one trajectory.

2.4 Control Theory

Physiological control mechanisms such as mammalian temperature regulation are often described in terms of control theory, that branch of mathematics and engineering that analyses dynamical systems in order that they produce a certain output given certain inputs. Classical control theory models regulation in terms of *open loop* and *closed-loop* controllers. A simple open-loop controller would be a brick placed on a car's accelerator pedal. An input (the amount the pedal is depressed) is applied irrespective of the output (the car's velocity). Closed loop control feeds back the effect of the output into the input. An example of a closed-loop controller is a car's cruise control that feeds back the velocity into the accelerator input so as to maintain a particular velocity. An example of a closed-loop controller is shown in figure 2.1.

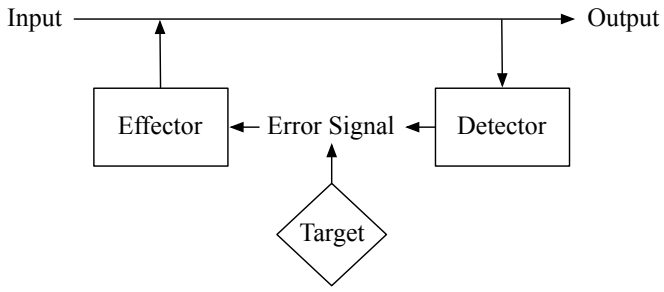

Figure 2.1: A Closed Loop Controller

A detector monitors the output of the system. The output of the detector is compared with a target value. Differences between the detector output and the target value are fed to an effector that modulates the input so that the output is moved towards the target value.

A number of different strategies may be employed in the calculation of the effector force that is applied to the input in order that the output moves towards the target value. A common method is the proportional integral derivative (PID) controller. The

PID employs three effector forces: an amount that is *proportional* to the error signal (if the error is large a large corrective force is applied); an *integral* amount that is the integral of a time series of error signals (this can be understood as a 'momentum' amount in which previous errors will feature in the instantaneous corrective actions of the effector); a *derivative* term that modulates the effector force as the rate of change of the error signal varies and thus reduces overshooting (the output passing the target output) as the effector force decreases the nearer the output is to the target. In the following chapters I will present results from non-PID systems that exhibit the same behaviour. Understanding how these systems produce similar behaviour using different methodologies will provide deeper insights into the Daisyworld control system.

2.4.1 Fixed points and steady-state 'error'

There is a rather subtle but important difference between Ashbian homeostasis and classical control theory homeostasis and I will return to this difference at a number of points within this thesis. Figure 2.2 represents the differences in paradigm between the two. Classical control theory closed-loop controllers will attempt to minimize the error signal produced as the variable subject to control deviates from the prescribed target value. An Ashbian controller may exhibit the same behaviour, but this is produced by attempting to ensure that the controlling variable is held away from the limits of the essential range. The point equidistant from these limits would be the centre, which could be equivalent to the target value for the classical controller. The difference between these two paradigms can be captured within the following analogy that I term 'Harvey's drunken walk'[3]. Imagine a hard working academic returning from a serious lunchtime debating session at a local hostelry. In order to reach his office he must negotiate a connecting corridor. Due to transient physiological imbalances, he experiences a certain amount of difficulty maintaining a straight path. By putting out both arms the academic is able to maintain an approximately centre line trajectory. If he begins to move towards the left, his left arm will push against the left wall with a force that is inversely proportional to the distance between his body and the wall. Similarly the right arm produces a force that maintains his body away from the right wall and so the academic proceeds down the corridor in an approximately straight line. This analogy encapsulates many of the important aspects of the Daisyworld control

[3]I thank Inman Harvey for suggesting a similar analogy while noting that any similarity between Harvey in the Harvey Drunken Walk and any other real or fictional Harvey is entirely coincidental.

system.

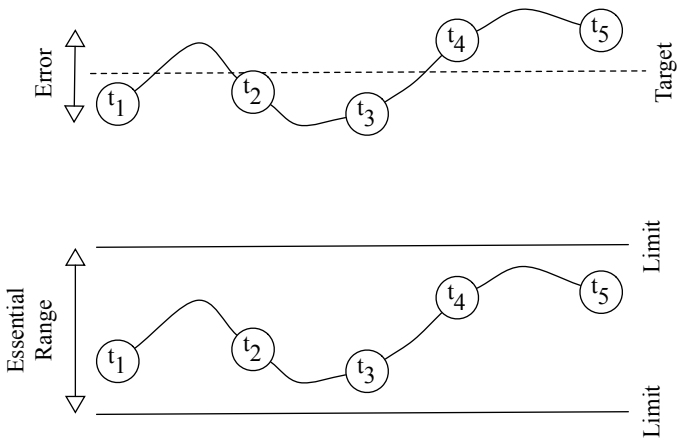

Figure 2.2: Optimal and Within Limits Regulation

The top schematic represents the operation of a closed-loop controller that seeks to minimise the error signal that is produced as the circular variable deviates from the target over time. The bottom schematic represents the operation of an 'Ashbian controller' that seeks to maintain the circular variable to within limits defined by the variable's 'essential range'. The behaviour of both systems may be equivalent.

2.5 Planetary Homeostasis and Gaia Theory

The original Gaia Hypothesis [Lovelock and Margulis, 1974],
[Margulis and Lovelock, 1974] proposed that the system consisting of life, the oceans, atmosphere and crust of the Earth is homeostatic, maintaining certain environmental variables within the limits necessary to sustain widespread life. A precursor to the original hypothesis can be found in Vernadsky's use of the term 'biosphere' and his stress on how life on Earth has a significant impact on abiotic processes [Vernadsky, 1926]. The Gaia Hypothesis goes further in stressing that these biosphere interactions will lead to a homeostatic system. For example, on the surface of the Earth, concentrations of atmospheric gases are sustained within limits that eukaryotic multicellular organisms can tolerate [Lenton, 1998]. Gaian systems are characterised by resistance to both internal and external perturbations [Lenton and van Oijen M., 2002]. There is evidence that the Earth system has demonstrated these properties. For example, biologically-mediated cloud formation over oceans: [Andreae and Crutzen, 1997], [Charlson et al., 1987], [Liss et al., 1997], [Lovelock, 1997], the temperature - biotic

rock weathering feedback loop: [Schwartman and Volk, 1991],
[Schwartman and Volk, 1989] and the regulation of marine salinity and phosphorous to within a narrow range: [Lenton and Watson, 2000], [Redfield, 1958], [Volk, 1998]. Events such as volcanic outbursts and planetesimal impacts have caused mass extinctions and climate change. In spite of these perturbations, widespread and diverse life, together with a tolerable climate, has always returned. Additionally, in the face of continuous solar forcing provided by a sun that has warmed by approximately 25% since the origin of the earth, global temperatures have remained remarkably stable. While one must be aware of observer effects[4] the notion that the Earth-biota system possesses homeostatic properties is a coherent and increasingly defensible one.

Nevertheless the theory has remained somewhat controversial and has been the subject of a pointed critique. Some thought the theory teleological, implying foresight and planning on the part of the biota in order to maintain regulation. The debate is centred, however, on why organisms that are the product of Darwinian evolution, a process that focuses on immediate individual advantage, should act in ways that are beneficial to other organisms and the biota as a whole [Dawkins, 1986]. See [Lenton, 1998] for a review of responses to such claims. A form of global regulation which all organisms contribute towards would be susceptible to cheats. If in failing to 'do my part' for global regulation I instead devote my energies to producing more offspring, then such regulation would very rapidly collapse. In order for any formulation of Gaia homeostasis to be plausible, it must be compatible with Darwinian evolution. Spurred on in part by such criticisms, there have been a series of developments of Lovelock's and Margulis' original hypothesis into a number of related Gaia Theories. Lovelock's books chart this development: [Lovelock, 1979], [Lovelock, 1988], [Lovelock, 1991], [Lovelock, 2006]. In a series of articles — [Kirchner, 1989], [Kirchner, 2002], [Kirchner, 2003] — Kirchner performs a taxonomic critical evaluation of the following Gaia Theories:

- Co-evolutionary Gaia: The biota influences its abiotic environment and the environment in turn influences the evolution of the biota via Darwinian processes.

- Homeostatic Gaia: The biota influences the abiotic world in a way that is stabilizing. The dominant linkages between the biota and environment are negative feedback loops.

[4]Given that there is only one planet with which to test Gaian hypotheses there is always the possibility that the Earth's apparent homeostatic behaviour is nothing more than lucky coincidence. If there had been a significant perturbation that had made extinct all life on Earth then we would not be here to puzzle over how life on Earth has managed to survive such perturbations.

- Geophysical Gaia: The biosphere can be compared with a single immense organism which, like other organisms, may exhibit both homeostatic and unstable behaviour.
- Optimizing Gaia: The biota manipulates its physical environment in ways that create biologically favourable conditions.

Kirchner did not consider these theories to be mutually exclusive. For example a Gaia theory may be considered as an example of Co-evolutionary Gaia and Homeostatic Gaia.

2.6 Daisyworld

Lovelock first presented Daisyworld as 'A Cybernetic Proof of the Gaia Hypothesis', [Lovelock, 1983]. A joint paper with Watson [Watson and Lovelock, 1983] developed Daisyworld into a simple mathematical model describing an imaginary grey planet orbiting a star similar to the Sun. It is home to two daisy types that are identical in all respects apart from their albedo. Albedo is a measure of the reflectivity of an object. An object with an albedo of 0 would be a perfect black body and would absorb all radiation that was received onto its surface. An object with an albedo of 1 would be a perfect mirror and would reflect all radiation that was received onto its surface. In Daisyworld the black daisy types have a low albedo (0.25), the grey bare earth intermediate albedo (0.5) and the white daisy types a high albedo (0.75). The white daisies, having the highest albedo in the model, reflect more of the energy from the star and so have a lower temperature than either the grey planet or the black daisies. The same applies to the black daisies but in reverse. The black and white daisies share an essential range of temperature. They are only able to grow when the local global temperature is within $5 - 40°C$. Within this range growth rates of the daisies vary, with optimum growth being achieved when the temperature is 22.5°C. Simulations begin when the star is dim and the temperature of the planet is below 5°C. As the star increases in brightness the planetary temperature reaches 5°C and black daisies begin to grow. The increase in the number of black daisies is further increased via a feedback loop: more black daisies leads to a lower planetary albedo and so more energy is absorbed which warms the ground and leads to an increase of the growth rate.

This feedback loop is regulated by the parabolic growth rate of the daisies. As the temperature increases past 22.5°C, the daisy growth rate decreases. At steady state, the

global temperature, growth rate and death rate are at equilibrium. As the luminosity (the brightness) of the star continues to increase, coverage of the black daisies decreases and white daisies begin to grow. This initiates a feedback loop that is the inverse of the effect of the black daisies and so decreases their temperature. Again, this feedback loop is regulated by the parabolic growth rate of the daisies. Increasing the luminosity results in a progressive increase in white daisies (and decrease in black daisies) until the maximum coverage of white daisies is reached. Any further increase in luminosity takes the global temperature past the point where growth rates balance death rates and so the coverage of white daisies decreases. This leads to a rapid collapse of white daisies similar in nature to the population explosion of the black daisies. The differential coverage of white and black daisies results in a system that effectively regulates global planetary temperature to within the essential range. Whereas the temperature of a bare lifeless planet would increase in an approximately linear fashion with increases in luminosity, when black and white daisies are present, global temperature remains within the essential range over a wide range of luminosity.

2.6.1 Daisyworld Homeostasis

Given the previous chapter's definition of homeostasis, I define Daisyworld homeostasis as the maintenance of the essential variable of global temperature to within the essential range of $5 - 40°\text{C}$ in the presence of external perturbations. These perturbations are changes in the luminosity of the star that Daisyworld orbits. This homeostasis is subtly but importantly different from the planetary temperature being maintained at a particular value, for example it being maintained at 22.5C, the temperature that produces maximum growth.

2.7 Original Daisyworld

I will now detail the Daisyworld model first published in [Watson and Lovelock, 1983]. Hereafter I will refer to this original Daisyworld as ODW. The model begins with a pair of differential equations that determine the rate of change of the area of the planet that is covered by the black and white daisies, α_b and α_w respectively:

$$\frac{d\alpha_b}{dt} = \alpha_b(x\beta_b - \gamma) \tag{2.1}$$

$$\frac{d\alpha_w}{dt} = \alpha_w(x\beta_w - \gamma) \tag{2.2}$$

where γ is the death rate parameter, taking a value of 0.3, x is the area of bare land, $x = 1 - \alpha_w - \alpha_b$ and $\beta_b = \beta_w$ which are the birth rate functions defined with:

$$\beta_i = max\left\{\ 1 - (\tfrac{5-T_i}{17.5})^2, 0\ \right\} \tag{2.3}$$

where T_i is the temperature of the daisies. This birth rate function is an inverted parabola, in which maximum growth is achieved when the daisy temperature is 22.5°C with growth decreasing to zero at 5°C and 40°C. The temperature of the daisies, T_i, is arrived at by first determining the albedo of the planet:

$$A_p = xA_x + \alpha_w A_w + \alpha_b A_b \tag{2.4}$$

where A_x, A_b and A_w are parameters that determine the albedo of bare ground (0.5), black (0.25) and white (0.75) daisies respectively. The temperature of the planet is determined by the amount of luminosity, how much energy radiated from the star is received on the surface, and the albedo of the surface:

$$\sigma T_p^4 = I(1 - A_p) \tag{2.5}$$

where I is insolation given in units of Wm^{-2}. Insolation is the amount of energy emitted by the star that is received on the surface of Daisyworld. The original study on Daisyworld modulated insolation via a luminosity term, $L \in [0, 2]$ so that $I = LS$, where S is the amount of radiation emitted from the star that is received on the surface of the planet and is fixed at 952.56 Wm^{-2}. When $L = 1$ the amount of energy received on the surface of Daisyworld is approximately that which is received on Earth. I will follow this convention of keeping S fixed and altering I by varying L. σ is the Stefan-Boltzmann constant taking a value of 5.67×10^{-8} $\mathrm{Wm}^{-2}\mathrm{K}^{-4}$. The temperature of the daisies is found through the difference between planetary and daisy albedo parameterised by an amount of heat transference q:

$$T_b^4 = q(A_p - A_b) + T_p \tag{2.6}$$

$$T_w^4 = q(A_p - A_w) + T_p \tag{2.7}$$

The units of T_i and T_p are degrees Kelvin, therefore it is necessary to subtract 273 from both to convert to degrees Celsius. The difference between daisy and planetary albedo and the heat flux parameter, q, determines the daisy to planet temperature differential. For the model to have sensible behaviour (white regions colder than black regions), $q < I/\sigma$ (16.17×10^9 K^4). Watson and Lovelock set q to the value 2.06×10^9 K^4 which is comfortably below this. This produces a temperature differential between the daisies and planet of approximately 5 degrees even with effectively no coverage of daisies - the black daisies are 5 degrees warmer, the white daisies are 5 degrees cooler than the ambient, planetary temperature.

Watson and Lovelock used the forward Euler method to numerically integrate the above equations for a range of luminosity in order to examine the behaviour of Daisyworld to increased radiative forcing from the star. Luminosity was initially fixed at 0.5 and steady state coverage of the black and white daisies was then used to initialise daisy coverage for a small increase in luminosity. It was assumed that the rate of change of the star was sufficiently slow for luminosity to be fixed while the population converged to steady state. Figures 2.3 to 2.6 reproduce these numerical solutions using the same techniques.

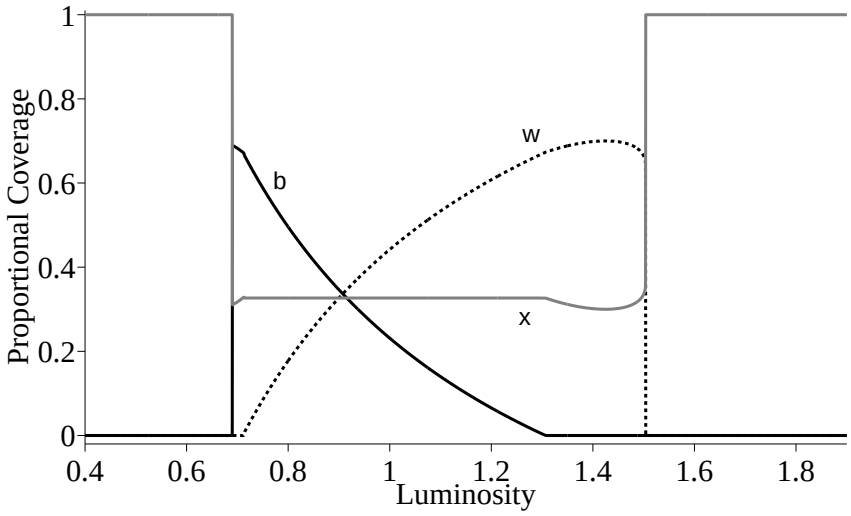

Figure 2.3: ODW Daisy Coverage

Proportional coverage of black daisies is plotted with a solid black line. Proportional coverage of the white daisies is plotted with a dashed black line. The total amount of bare ground, $x = 1 - b - w$, is plotted with a solid grey line. Within the both-daisy region (the region where both black and white daisies grow), the amount of bare ground is constant. The explosion of the black daisies is matched by the sudden increase in planetary temperature as shown in figure 2.5. As luminosity increases, planetary temperature decreases until only white daisies exist, at which point, it increases in response to increased luminosity. There is a discontinuous increase in temperature at the point of the white daisy collapse.

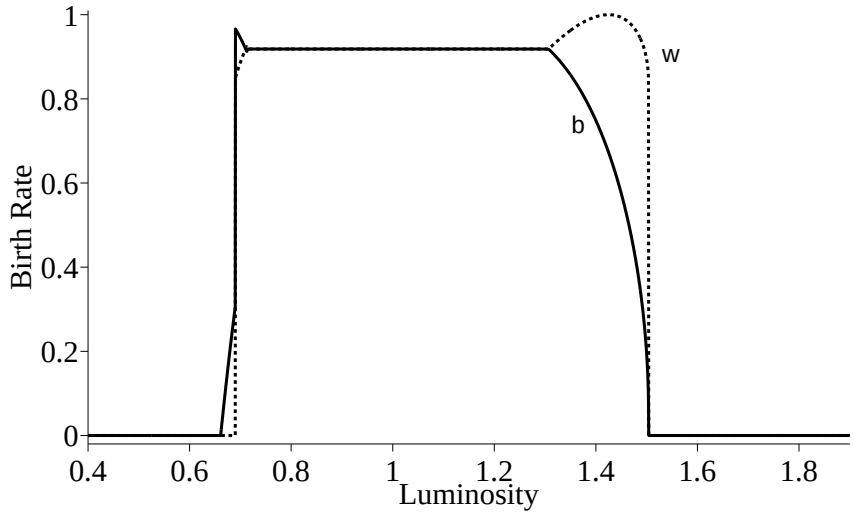

Figure 2.4: ODW Daisy Birth Rates

The birth rate of the black daisies is plotted with a solid black line. The birth rate of the white daisies is plotted with a dashed black line. During the both-daisy range (the region where both black and white daisies grow) the daisy birth rates are fixed and of equal value.

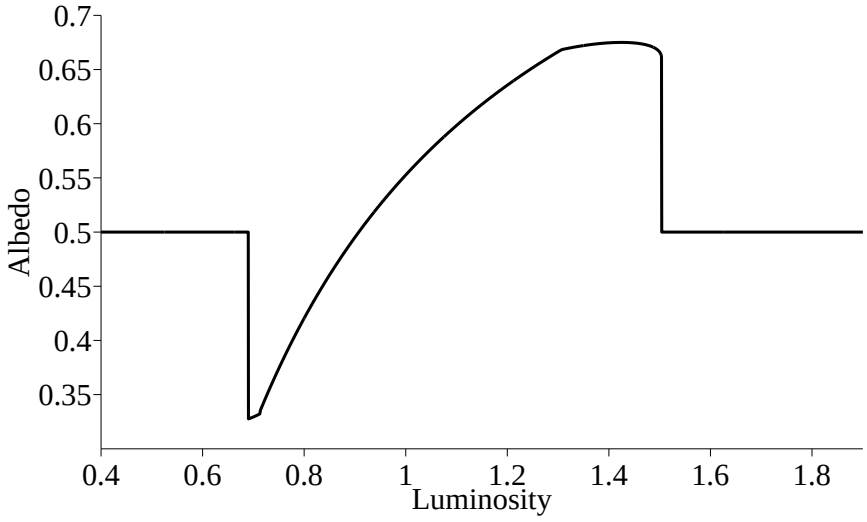

Figure 2.6: ODW Planetary Albedo

Outside of the any-daisy range (the region where either black or white daisies grow) the albedo of the planet is that of bare ground, 0.5. The sudden increase in black daisies and decrease in white daisies is accompanied with sudden changes in planetary albedo.

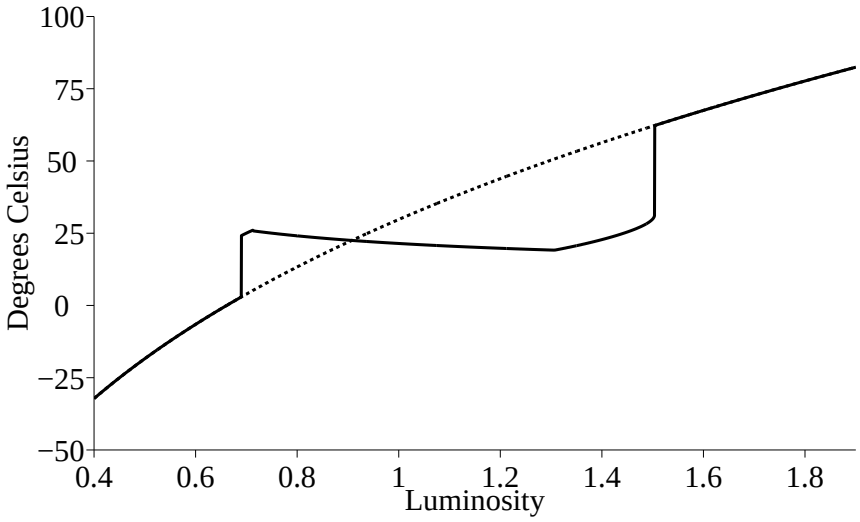

Figure 2.5: ODW Planetary Temperature
Planetary temperature with black and white daisies present is plotted with a solid black line. Planetary temperature on a lifeless planet is plotted with a dashed black line. On a lifeless planet, the temperature increases approximately linearly with increasing luminosity. With black and white daisies present, planetary temperature is regulated to within the essential range of the daisies over a wide range of luminosity.

2.7.1 Hysteresis

A system that exhibits hysteresis is one in which the behaviour or output of the system is path dependent. The history of the system as well as the current input determines the output. In ODW hysteresis is manifest in the bistable steady state solutions towards the limits of daisy coverage. White daisies can only grow at higher luminosity if there is an established white daisy population that cools the planet. If luminosity increases such that the white daisy population crashes, the white daisies do not re-grow (assuming white daisy seeds are continually present on the surface of the planet) if luminosity is decreased a little below this value. Rather, a large decrease in luminosity must occur before the white daisies begin to grow. Similarly there is a hysteresis loop of black daisy coverage at lower luminosity.

2.8 Analytical Solutions

Solutions for ODW are now available, these being largely based on two independent studies: an examination of the ODW dynamical system [DeGregorio et al., 1992] and

exact solutions [Saunders, 1994]. The latter study is of more relevance to this thesis as I shall base my analysis of a simplified Daisyworld model in the following chapter on Saunders' approach but with a number of simplifications that makes this analysis more general. Here I will give a brief synopsis of Saunders' approach. For a range of luminosities, Saunders examines the system's steady state in terms of four possible outcomes: no daisies, black only, white only and black and white daisies. Considering single-daisy solutions, Saunders notes that at equilibrium the growth rate must equal the death rate which implies that:

$$\beta_b = \frac{\gamma}{x} = \frac{\gamma}{1 - \alpha_b} \tag{2.8}$$

and that the birth rate function equation (2.3) leads to:

$$T_b = 22.5 \pm 17.5 \left[\frac{\alpha_b + \gamma - 1}{\alpha_b - 1} \right]^{1/2} \tag{2.9}$$

This can then be used to find T_b and α_b for luminosity. The same approach is used to solve T_w and α_w for luminosity. Considering the both-daisies stable state, Saunders makes the observation that providing the death rate is constant, the growth rates of the black and white daisies are constant with the black and white daisies being an equal distance from the optimal growth temperature of 22.5°C (black daisies being a fixed amount warmer, white daisies a fixed amount cooler):

$$T_b - 22.5 = 22.5 - T_w \tag{2.10}$$

The temperature of the black (and white daisies) during the co-existence range is not determined by the luminosity but on the value of q. As the temperatures of the daisies are constant, so too are the birth rates and the amount of bare ground. What alters is the ratio of black to white daisies. Finding the real roots for the equations gives the following values: $T_b = 27.5$°C, $T_w = 17.5$°C, $\beta_b = \beta_w = \beta = 0.918$ and $\alpha_b + \alpha_w = 0.673$. It is then straightforward to find black and white daisy coverage and planetary temperature for luminosity.

2.9 Developments of Daisyworld

In the 25 years since its inception there have been numerous developments of the ODW, each seeking to explore a particular issue or perhaps prove a particular point. These can be classified as models that altered the original model in the following ways:

- Spatial extensions.
- Ecosystem dynamics including higher trophic levels.
- Evolution.
- Non-radiative models.
- Thermodynamic models.
- Conceptual models.

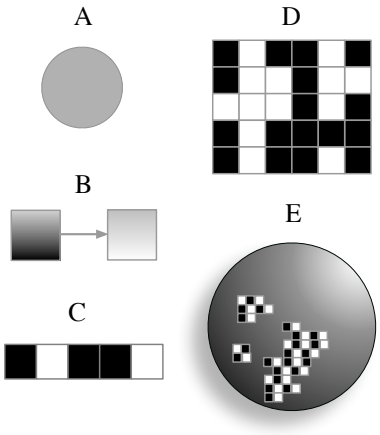

Figure 2.7: Spatial Extensions

Five main spatial categories are represented schematically: A - original 'zero dimensional' Daisyworld [Watson and Lovelock, 1983]; B - Two-Box Daisyworld that features heat diffusion between boxes [Harvey, 2004]; C - one-dimensional Daisyworld that may (or may not) feature luminosity varying over 'latitude' [Adams et al., 2003]; D - two-dimensional cellular automaton Daisyworld [von Bloh et al., 1997]; E - three-dimensional cellular automaton Daisyworld with luminosity varying over latitude [Ackland et al., 2003].

2.9.1 Spatial Extensions

Figure 2.4 represents the five main spatial categories of Daisyworld models. The original Daisyworld can be considered as a zero-dimensional model because where the black and white daisies grow on the surface of the planet is not explicitly modelled. However, there is an implicit notion of spatial structure in the model as black and white daisies and daisy seeds have different temperatures to the planetary global temperature. It is of course unrealistic to suppose that a daisy seed would be some five degrees Celsius cooler or warmer than the ground in which it is embedded. This can make sense if black daisies are seeded in areas that are warmer and white daisies seeded in areas that are cooler. This introduces a notion of spatial heterogeneity into the model and proves to be a critical element of the homeostatic mechanism as I will explain in the following chapters.

The two-box model initially proposed by Harvey, [Harvey, 2004] and then developed in a series of papers: [Dyke and Harvey, 2005], [Dyke and Harvey, 2006], [Dyke, 2008] can be regarded as the simplest example of a one-dimensional model.

One-dimensional Daisyworlds, [Adams et al., 2003], [Adams and Carr, 2003] feature varying luminosity as the modelling assumption is that the 'strip' of daisy cells runs latitudinally along a curved planet. Implicit diffusion using the q parameter is retained. Desert formation and counter-intuitive striping of black and white daisy cells can emerge.

The simplest two-dimensional models employ cellular automaton update rules to kill and seed black and white daisies [Lenton and van Oijen M., 2002].
[von Bloh et al., 1997] models a continuous shade of daisies. These models feature explicit heat diffusion equations.

Three dimensional models are two-dimensional Daisyworlds that model the effects of varying luminosity with latitude over a sphere as in one-dimensional models, [Ackland et al., 2003].

2.9.2 Ecosystem Dynamics

The original Daisyworld model can be understood as an example of a simple ecological model. Ecological dynamics are captured in the competition for the shared 'resource' of bare ground with which to seed new daisies. A number of extensions have allowed richer dynamics to be explored. These studies have typically been focussed

on what effects certain ecological mechanisms may have for environmental homeostasis, although attempts have been made to feed results and analysis into mainstream ecology [Wilkinson, 2003b], [Wilkinson, 2003a]. Higher trophic levels in the form of 'rabbits' that ate daisies was introduced by [Lovelock, 1992] with further herbivore effects investigated by [Harding and Lovelock, 1996] and [Harding, 1999] with a core result being that ecosystems that feature more species are more stable. Two-dimensional Daisyworld models were used to explore the actions of herbivores on environmental homeostasis [von Bloh et al., 1997]. Studies on Daisyworld that consider alternative ecological resources such as concentrations of particular chemicals rather than bare ground are discussed in section 2.9.5 below.

2.9.3 Adaptation of Albedo

Darwinian evolution has been introduced into Daisyworld by allowing variation in albedo and/or optimal growth temperature. [Lovelock, 1992] added an intermediate grey daisy type to the ODW in order to counter criticism that regulation would be destroyed by intermediate albedo 'cheaters'. This initial study was extended in [Lansing and Smuts, 1998] with the original result being confirmed. A series of papers: [Lenton, 1998], [Lenton and Lovelock, 2001] and [Ackland et al., 2003] have further investigated the effects of increasing the number of intermediate daisy types. A common result of these studies is that despite the number of intermediate albedo daisies, stable states are typically those which feature only two daisy types that straddle the optimal growth temperature: one slightly darker and one slight lighter than the albedo that would be perfectly adapted to the current amount of luminosity. In chapter 5 I will explain these results and show why Daisyworld homeostasis will typically require pairs of opposing daisies that increase and decrease planetary temperature. Daisyworlds that feature continuous rather than discrete albedo types were originally investigated in [Stocker, 1995] and later extended in [Wood and Coe, 2007].

2.9.4 Adaptation of Optimal Growth Temperature

A number of studies have investigated the effects of altering the fixed birth rate functions of the ODW. [Keeling, 1991] introduced a type of daisy that had the same albedo as the black daisies but which produced greatest growth at 33°C rather than 22.5°C. This resulted in a significant reduction in the homeostatic abilities of the model.

Progressive adaptation of the daisies' optimal growth temperature was first included in [Robertson and Robinson, 1998] which found that rather than regulate temperature against external driving, the black daisies underwent a sequence of adaptations that tracked the increase in luminosity which led to a progressively lighter population of daisies and a decrease in the range of luminosity over which the daisies grew. A number of studies have in effect replied to [Robertson and Robinson, 1998], with [Sugimoto, 2002] questioning certain central assumptions of the study and [Lenton and Lovelock, 2000] pointing out that [Robertson and Robinson, 1998] assume that daisies are able to respond to an effectively unlimited range of environmental changes such that daisies have optimal growth below freezing or above the boiling point of water. Later studies have shown that homeostasis can be maintained when daisies are allowed continuous mutations of optimal growth temperatures [Williams and Noble, 2005], [Dyke et al., 2007] and [McDonald-Gibson et al., 2008].

2.9.5 Non-radiative Models

There has been a number of studies that have examined Daisyworld homeostasis where the environmental variable being regulated is not temperature but some other feature of the abiotic environment. Bacterial regulation that employed Daisyworld-type modelling assumptions was investigated in the Guild model of [Downing and Zvirinsky, 1999]. Chemo-thermodynamic constraints were introduced to the Guild model with the METAMIC (Metabolically Abstract Microorganism) model, [Downing, 2002] and [Downing, 2003]. A Daisyworld-type model was used to assess productivity of rice farming in [Baldocchi et al., 2005]. [Dyke et al., 2007] and [McDonald-Gibson et al., 2008] present abstract ecosystems which regulate a resource variable that could be seen as a token for a range of abiotic elements such as pH, salinity or concentration of a particular compound.

2.9.6 Conceptual Models

There has been a series of studies that have attempted to distil the essential elements of Daisyworld: The 'cut down' Daisyworld of [Harvey, 2004] featured simplified ecological interactions and physics. This was extended in [Dyke and Harvey, 2005] which implemented a robotic control system based on Daisyworld and [Dyke and Harvey, 2006] which examined the limits of homeostasis of Daisyworld.

2.10 Summary

In this chapter I have defined homeostasis and homeorhesis, situated Daisyworld within the context of Gaia, presented the equations and numerical and analytical results of the original Daisyworld model and reviewed how Daisyworld has been developed in order to counter initial criticisms as well as a platform for more general enquiries.

In the next chapter I will present a simplified ODW model that allows a more general analytical solution which will be used to explain the ODW's inverse response to external driving (why planetary temperature on Daisyworld decreases while luminosity increases) as well as providing the basis for the following chapter's distillation of the core homeostatic mechanism.

Chapter 3

Simplified Daisyworld

3.1 Overview

My first objective in this chapter is to begin the process of simplification that will allow me to highlight the core homeostatic mechanism of Daisyworld. I have assumed that the reader is not familiar with the exact analytical solutions of the original Daisyworld model. Solving a simplified version will allow a more direct appreciation of the essential elements of Daisyworld and how the original model regulates planetary temperature to within the essential range in the face of external perturbations. My second objective is to explain a particular aspect of the original Daisyworld's behaviour - its inverse response to perturbations. Planetary temperature *decreases* when luminosity *increases*. I will show that the inverse response is not caused by the original choice of birth rate function or non-linear relationship between albedo and temperature. Understanding the inverse response highlights the importance of how the daisies are coupled to each other via their effects on temperature and competition for available space. These interactions are important elements of the Daisyworld control system.

3.2 Simplifying Daisyworld

Although Daisyworld is a simple model of a very complex system, it has a number of non-linear functions and modelling subtleties that can impede the analysis of its homeostatic mechanism. In this chapter I will proceed with a number of simplifications to the original Daisyworld model (hereafter referred to as ODW). The intention is to simplify ODW in order to understand what is relevant in a broader context rather than

what is physically plausible with regard to planetary climates. To that end I will replace parabolic birth rate and quartic temperature functions with linear functions and so formulate a simpler ODW model (hereafter referred to as SDW). I will present steady state solutions for SDW found using numerical techniques as well as exact analytic solutions.

The simplifications performed in the process of constructing SDW will render the model physically implausible with respect to the behaviour of objects that absorb and emit radiation. However if for only didactic purposes I will continue to phrase the model in terms of daisies, planets, luminosity, albedo and temperature. In doing so scare quotes should, arguably, be applied to such terms but to do so would risk labouring the point. In chapter 5 I develop a more abstract model that dispenses with any terms that relate to real world processes.

3.3 Model

The equations governing the rate of change of the daisies in SDW are identical to ODW. They are reproduced in this chapter for convenience:

$$\frac{d\alpha_b}{dt} = \alpha_b(x\beta_b - \gamma) \tag{3.1}$$

$$\frac{d\alpha_w}{dt} = \alpha_w(x\beta_w - \gamma) \tag{3.2}$$

where α_b and $\alpha_w \in [0, 1]$ and are the proportional coverages of the black and white daisies respectively. γ is the death rate parameter, fixed at 0.3, x is the area of bare land, $x = 1 - \alpha_w - \alpha_b$ and $\beta_b = \beta_w$ are the birth rate functions. ODW employs a parabolic type birth rate function. I will replace this parabolic function with a linear 'hat function' birth rate function, first used in Daisyworld by [Harvey, 2004]:

$$\beta_i = max\left\{\ 1 - 2\frac{|T_o - T_i|}{R}, 0\ \right\} \tag{3.3}$$

where T_o is the optimum temperature, measured in dimensionless units, and is fixed at 22.5. R is the width of the essential range, the range of temperatures over which the daisies have non-zero growth. As in ODW this is fixed at 35 as the highest and lowest temperatures that support daisies are 5 and 40 respectively. The ODW and SDW birth rate functions are plotted in figure 3.1. The temperature of the black and white daisies

is found with linear versions of the ODW temperature equations:

$$T_b = q(A_p - A_b) + T_p \tag{3.4}$$

$$T_w = q(A_p - A_w) + T_p \tag{3.5}$$

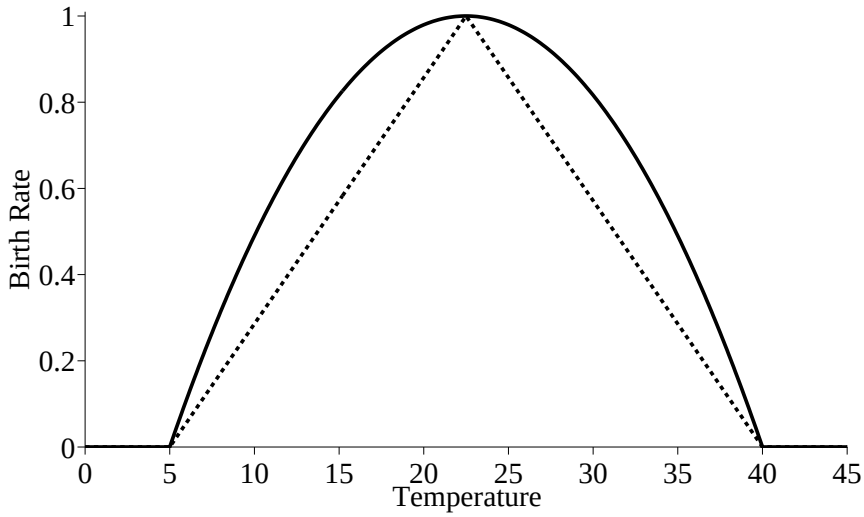

Figure 3.1: Original and Linear Birth Rate Functions

Two birth rate functions are plotted: ODW (solid black line) and SDW (dashed black line). Both birth rate functions have maximum values when temperature is at 22.5 and zero values when temperature is 5 or 40.

ODW employs quartic temperature and energy balance functions. Linearised non-quartic Stefan-Boltzmann functions in Daisyworld have been discussed previously: [Saunders 1994], [Pujol et al 2005], [Staley, 2002], [Weber and Robinson, 2004] and [Wood et. al. 2008]. It has been shown that linearising the temperature function does not significantly change results produced by the ODW as the range of temperatures of interest are sufficiently far from the repeated fourth root as to be essentially linear. Following [Harvey 2004] I will take this process further and remove the scaling parameter and Stefan-Boltzmann constant that is used in calculating the temperature of a black body subject to incident radiation. The temperature of the planet will become the amount of radiation received, multiplied by one minus the albedo of the planet:

$$T_p = I(1 - A_p) \tag{3.6}$$

Insolation, I, is now a dimensionless scalar that represents the amount of radiation emitted from the star and received on the surface of the planet. As in ODW, insolation is the product of luminosity, L and stellar radiation, S. In the following results S is fixed at the ODW value of 952.56 and $L \in [0, 2.5]$ The increase in the range of luminosity from $[0, 2]$ to $[0, 2.5]$ is required by the reduction of the white daisy albedo which means that they are able to survive at higher luminosities. A_p is the albedo of the planet and is found with:

$$A_p = xA_x + \alpha_w A_w + \alpha_b A_b \tag{3.7}$$

where A_b and A_w are the albedo of the black and white daisies. The ODW values for these parameters are 0.25 and 0.75. These will be changed to 0 and 1 which represent a perfect black-body object and a perfect mirror respectively. Steady state values for daisy coverage and daisy and planet temperatures were found with Euler's method as for ODW. Results are plotted in figures 3.2 - 3.5. The results are qualitatively the same as the ODW but with larger regions of hysteresis due to the increase in the difference between black and white daisy albedo values. This will be discussed further in chapter 4. Of particular interest to this chapter is the region of luminosity over which the temperature of the planet goes down. The model exhibits the same inverse response to forcing. This response will be explained in the following sections together with analytical solutions to the model.

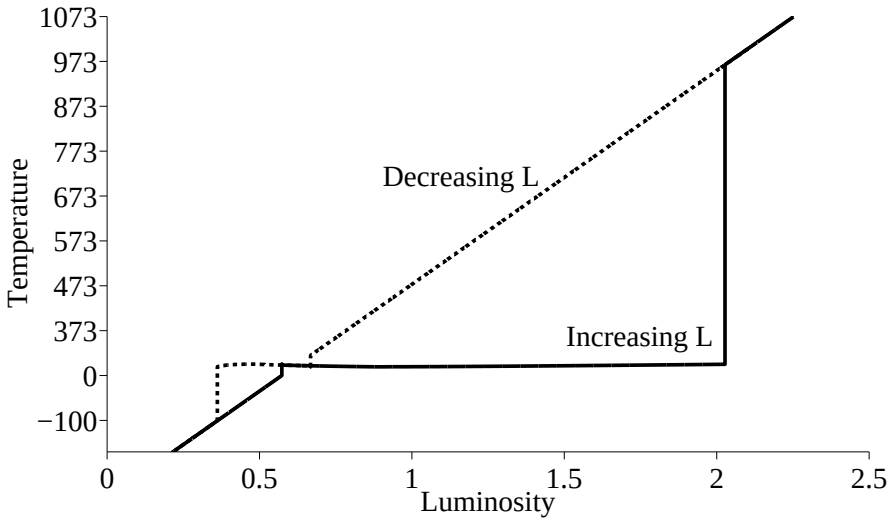

Figure 3.2: Numerical Solutions for Planetary Temperature

Computed numerical steady state planetary temperature is plotted over a range of luminosity values. Results for increasing luminosity are plotted with a solid black line and results for decreasing luminosity are plotted with a dashed black. The increase in the area of hysteresis is primarily due to the change in luminosity of the black daisies from 0.25 to 0 and the white daisies from 0.75 to 1.

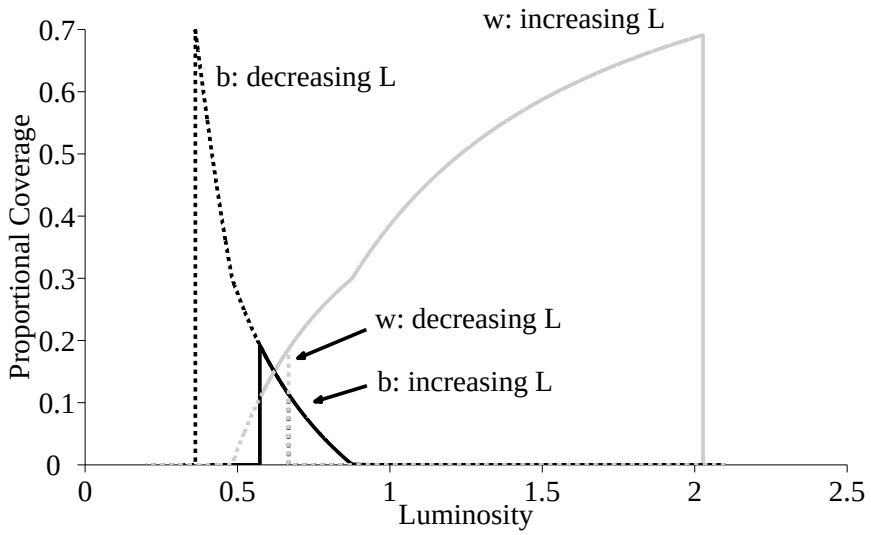

Figure 3.3: Numerical Solutions for Daisy Coverage

Computed numerical steady state black and white daisy coverage is plotted over a range of luminosity values. Results for increasing luminosity are plotted with a solid lines and results for decreasing luminosity are plotted with a dashed lines. Results for black daisies are plotted with black lines and results for white daisies are plotted with grey lines.

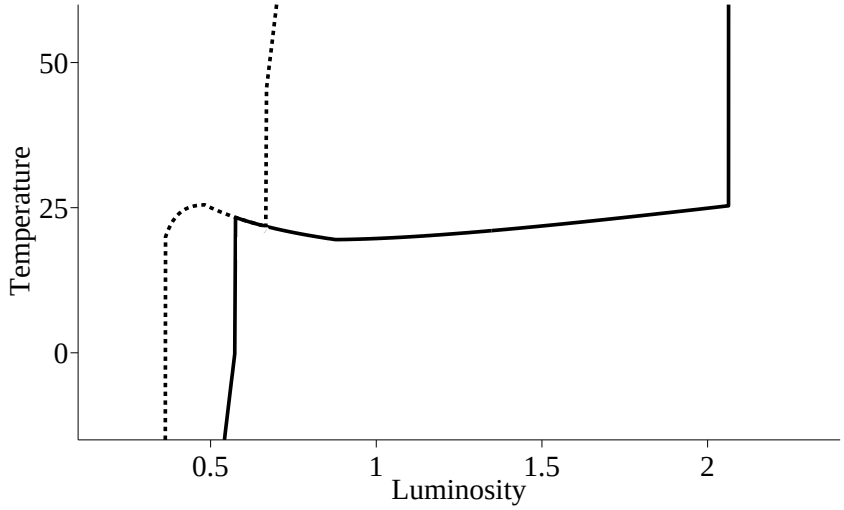

Figure 3.4: Numerical Solutions for Planetary Temperature

This figure can be understood as a closer view of figure 3.2. This clearly shows that planetary temperature decreases over a range of increasing luminosity. The inverse response to forcing is observed. Results for increasing luminosity are plotted with a solid black line and results for decreasing luminosity are plotted with a dashed black.

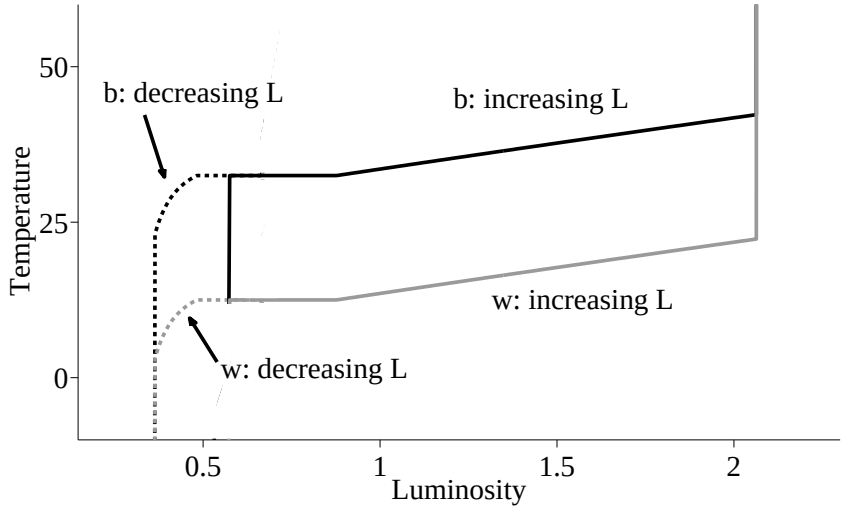

Figure 3.5: Numerical Solutions for Daisy Temperature

This figure accompanies 3.4 in that it is a closer view of daisy temperatures. Computed numerical steady state black daisy temperatures are plotted in black lines, white daisy temperatures in grey lines. Results for increasing luminosity are plotted with a solid lines and results for decreasing luminosity are plotted with a dashed lines. The inverse response to forcing is not observed as black and white daisy temperatures do not decrease with increasing luminosity.

3.4 Analysis

In this section I will present analytical solutions to SDW and their derivations. These derivations are important as understanding how this simple model is solved will go a considerable way to increasing our understanding of how the Daisyworld control system operates. The general approach I adopt is similar to that employed in [Saunders 1994]. However, by using linear functions and examining only stable steady states I am able to produce simpler solutions which allow more intuitive insights into Daisyworld. This also facilitates an examination of the inverse response. All numerical and exact analytic solutions will be presented to four decimal places.

3.4.1 White Daisies Only

There are four possible stable steady states for SDW: no daisies, white daisies only, black daisies only and both black and white daisies. Solving for no daisies is trivial and not particularly illuminating as the temperature of the planet is a simple, linear function of luminosity. In the following sections I will solve SDW for the other three possible stable steady states. Considering the white-only and black-only states is necessary in order to understand the entire range of homeostasis observed in the model, as at the limits of homeostasis only one daisy type will grow. Also, considering solutions to these states is the precursor to the analysis I conduct in chapter 5 where I produce a minimal Daisyworld control system that can be comprised of one or two elements. In this and following sections, temperature will be given in Kelvin units and I will formulate the temperature of the daisies and planet as functions of insolation, I, the amount of energy received on the surface of the planet. When plotting data, temperature will be given in dimensionless Celsius units (add 273 for Kelvin units) over luminosity.

To solve for white-only states we begin by realizing that the linear birth rate function, equation (3.3), can be simplified further. As was noted by [Saunders, 1994], stable steady state birth rates for white daisies are always within the range $[278, 295.5]$. While there are steady state values for white daisy birth rates within the higher range $[295.5, 313]$ these will in practice never be found as any perturbation will either lead to a runaway increase in temperature and so zero birth rate, or a decrease in temperature towards the stable steady state values within the lower range $[278, 295.5]$. This is shown geometrically in figures 3.6 and 3.7.

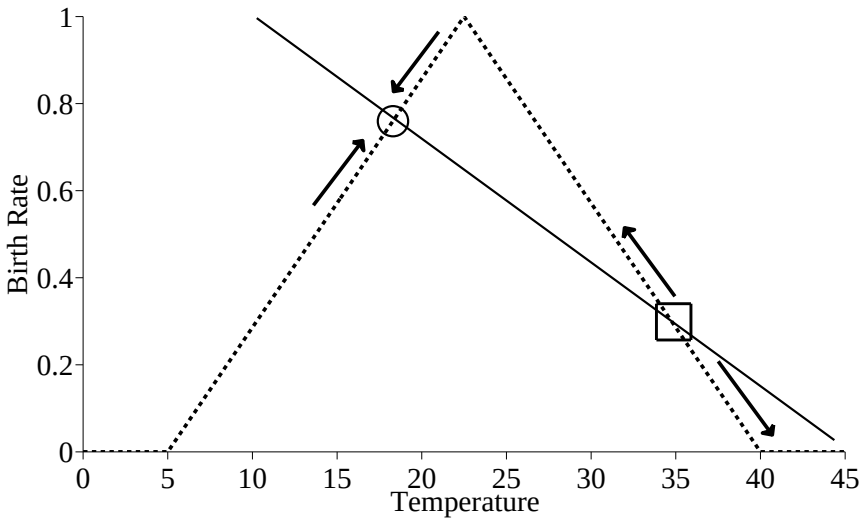

Figure 3.6: Stable and Unstable Steady States - White Daisies

The linear birth rate function shared by both black and white daisies is represented by a triangle. The solid sloping line represents the negative feedback the white daisies exert on temperature for a fixed value of insolation. The stable steady state is denoted by a circle, the unstable steady state is denoted by a square. The arrows denote the change in the birth rate towards and away from the stable and unstable stead states. Any increase in the birth rate from the square would lead to a positive feedback effect that would rapidly increase birth rate to the maximum and then back down the other side of the birth rate function finally coming to rest at the circle. Any decrease in the birth rate function from the square would again lead to positive feedback, but this time a stable steady state of zero birth rate would be found.

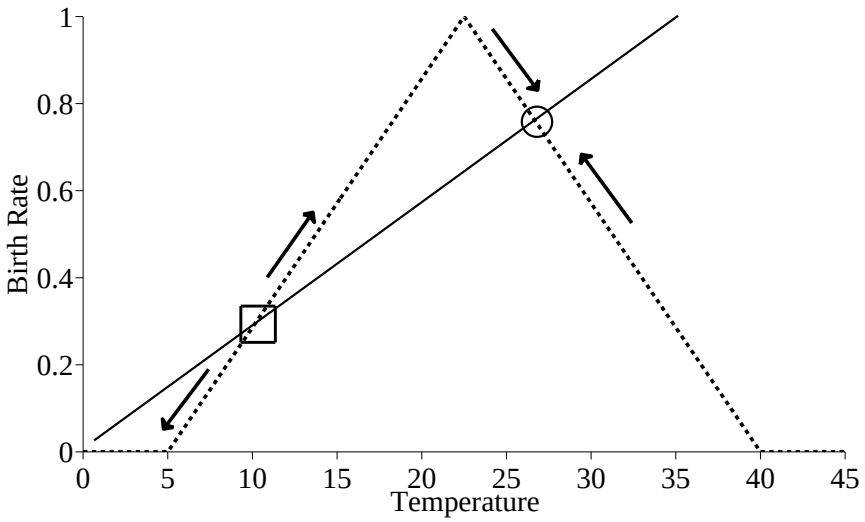

Figure 3.7: Stable and Unstable Steady States

This figure accompanies 3.6. The linear birth rate function shared by both black and white daisies is represented by a triangle. The solid sloping line represents the positive feedback the black daisies exert on temperature for a fixed value of insolation. The stable steady state is denoted by a circle, the unstable steady state is denoted by a square.

As luminosity increases T_w from 278 towards 295.5 there is an increase in daisy growth until maximum coverage is obtained. Any further increase in insolation and white daisy temperature leads to a collapse in the white daisy population. If we ignore unstable steady state solutions, equation (3.3) for white daisies can be formulated as:

$$\beta_w = 1 - \frac{295.5 - T_w}{17.5} \tag{3.8}$$

This can be expressed as a function of how far the daisy temperature is from the start of the essential range, the lowest temperature that gives non-zero growth:

$$\beta_w = \frac{T_w - 278}{17.5} \tag{3.9}$$

Next, we recall that steady state white daisy coverage is found when the increase of daisies balances the decrease of daisies:

$$0 = \beta_w(1 - \alpha_w) - \gamma \tag{3.10}$$

When the death rate, γ equals 0.3, substituting equation (3.9) into (3.10) and rearranging to find α_w gives:

$$\alpha_w = 1 - \frac{5.25}{T_w - 278} \tag{3.11}$$

This can be rearranged so the temperature of the white daisies can be found as a function of their coverage:

$$T_w = \frac{5.25}{1 - \alpha_w} + 278 \tag{3.12}$$

In order to find the temperature of the white daisies as a function of insolation we first note that the albedo of bare ground is fixed at 0.5 and the white daisies at 1. Consequently the albedo of the planet is:

$$A_p = \alpha_w + (1 - \alpha_w)0.5 = 0.5 + 0.5\alpha_w \tag{3.13}$$

Next we reproduce (3.5):

$$T_w = q(A_p - A_w) + T_p$$

where q is the insulation parameter and determines how insulated the black and white daisies are from the ambient, planetary temperature. We can replace the A_p term by substituting in (3.13) and express T_p in terms of A_p:

$$T_w = 0.5q(\alpha_w - 1) + 0.5I(1 - \alpha_w) \tag{3.14}$$

Next, we substitute (3.11) for α_w in (3.14) in order to find the temperature of the white daisies as a function of insolation:

$$T_w = 139 + 0.5\sqrt{278^2 - 10.5(q - I)} \tag{3.15}$$

We can also find α_w for I by substituting (3.12) for T_w in (3.14):

$$\frac{5.25}{1 - \alpha_w} + 278 = 0.5q(\alpha_w - 1) + 0.5I(1 - \alpha_w) \tag{3.16}$$

which is a quadratic:

$$\alpha_w = \frac{I - q - 278 - 0.5\sqrt{40(I - q) + 556^2}}{I - q} \tag{3.17}$$

3.4.2 Limits of White Daisy Growth

Equation (3.12) shows us that with increasing insolation the lowest white daisy temperature that produces non-zero coverage is 283.25 as temperatures less than 283.25 produce a birth rate less than 0.3 and so less than the fixed death rate. Therefore with increasing luminosity, white daisies begin to grow when their temperature reaches 283.25. Using (3.5) and (3.6) it is trivial to find the value of luminosity that produces a white daisy temperature of 283.25 because in the absence of any daisies the albedo of the planet is 0.5. This value of luminosity is 0.6157. The birth rate of the white daisies is greatest when $T_w = 295.5$. The maximum possible coverage of white daisies is 0.7. The albedo of the planet with the maximum possible coverage of white daisies is 0.85. This gives a planetary temperature of 298.5 and so the maximum value of luminosity that can support the white daisies is 2.0891.

When luminosity is initialised at a high value and then *decreases*, hysteresis is observed in ways qualitatively similar to ODW. White daisies will begin to grow when the birth rate is equal to or greater than the fixed death rate. This occurs when the temperature of the white daisies is lower than 307.75 which is produced when luminosity decreases to 0.6671. Analytical solutions for white-only daisy coverage are shown in figures 3.8, planetary temperature in figure 10 and white-only daisy temperatures in figure 3.11.

3.4.3 Black daisies only

The route to finding black-only stable steady solutions is similar to that employed for white-only solutions. As with white-only daisy scenarios, we can simplify black-only daisy scenarios by simplifying the black daisy birth rate so as to consider only stable steady states. Stable black birth rates are within the range $[295.5, 313]$. This allows the birth rate for black daisies to be formulated as:

$$\beta_b = 1 - \frac{T_b - 295.5}{17.5} \tag{3.18}$$

Black daisy temperatures less than 295.5 or greater than 313 will produce a zero birth rate. This can be simplified further by expressing the birth rate as a function of how far the daisy temperature is from the end of the viability range (with increasing luminosity):

$$\beta_b = \frac{313 - T_b}{17.5} \tag{3.19}$$

Next, we recall that steady state black daisy coverage is found when the increase of daisies balances the decrease of daisies:

$$0 = \beta_b(1 - \alpha_b) - \gamma \tag{3.20}$$

When the death rate, γ equals $= 0.3$, substituting equation (3.19) for β_b in (3.20) and rearranging to find α_b gives:

$$\alpha_b = 1 - \frac{5.25}{313 - T_b} \tag{3.21}$$

This can be rearranged so the temperature of the black daisies can be found as a function of their coverage:

$$T_b = \frac{5.25}{\alpha_b - 1} + 313 \tag{3.22}$$

The albedo of the black daisies was fixed at 0 and the bare ground 0.5. Therefore the albedo of the planet is:

$$A_p = 0.5 - 0.5\alpha_b \tag{3.23}$$

Recall (3.4):

$$T_b = q(A_b - A_p) + T_p$$

The temperature and albedo of the planet can be expressed in terms of α_b by using (3.23):

$$T_b = 0.5q(1 - \alpha_b) + 0.5I(1 + \alpha_b) \tag{3.24}$$

We can substitute (3.21) for α_b in (3.24) to find the black daisy temperature as a function of insolation:

$$T_b = 156.5 + 0.5I + 0.5\sqrt{(-313 - I)^2 - 1241.5I + 10.5q} \tag{3.25}$$

Using (3.22) and (3.24) allows us to find the coverage of the black daisies as a quadratic function of insolation:

$$\alpha_b = \frac{313 - q - 0.5\sqrt{626^2 - I(2462 - 4I) + 42q}}{I - q} \tag{3.26}$$

3.4.4 Limits of Black Daisy Growth

With low and then *increasing* luminosity, black daisies will grow when their temperature reaches 283.25. This will be 10 temperature units higher than the planetary temperature. Therefore black daisies will begin to grow when luminosity reaches 0.5737. The maximum temperature that gives a non-zero steady state coverage is 307.75 and not 313 as temperatures greater than 307.75 produce a birth rate less than 0.3 and so less than the fixed death rate. At the limit of black daisy growth, the albedo of the planet is again 0.5 as the coverage of black daisies is infinitesimally small and so planetary temperature will be 10 temperature units lower than that of the daisy temperature. Therefore black daisies will be back to zero coverage when luminosity increases to 0.6252.

When luminosity is initialised at a high value and then *decreases*, hysteresis is observed in ways qualitatively similar to ODW. Black daisy growth will extend below luminosity values of 0.5737 until the temperature of the black daisies is 295.5 and coverage is at the maximum 0.7. The albedo of the planet at that point will be 0.15 and planetary temperature 292.5. Therefore the lowest value of luminosity that produces non-zero black daisy coverage is 0.3613. Analytical solutions for black-only daisy coverage are shown in figures 3.8, planetary temperature in figure 10 and black-only daisy temperatures in figure 3.11.

3.4.5 Black and White Daisies

As in [Saunders, 1994], analytical solutions to black and white daisy scenarios begin with the appreciation that at steady state the rate of increase in daisy coverage must equal the rate of decrease. Therefore, $x\beta = \gamma$. As the rate of change of the black and white daisies are functions of the amount of bare ground, x, when both daisies have non-zero growth, at steady state they are identical. $\beta_w = \beta_b$. This symmetry is used as the basis for providing stable steady state solutions for black and white daisies. The difference between the two daisy beds is found by subtracting the white daisy from the black daisy temperature:

$$T_b - T_w = q(A_w - A_b) \tag{3.27}$$

The difference in temperature between the black and white daisies is a function of their fixed albedo values, A_b, A_w and the insulating parameter, q. As the black daisy birth rate must equal the white daisy birth rate, the temperatures of each daisy type are the same distance from the optimal temperature. This can be visualised as the

temperatures of the daisies coming to rest on opposite sides of the birth rate function, the distance from the peak being determined by the difference between the respective daisy albedo and the fixed albedo of bare ground and q. Section 3.5.1 will discuss this further. If we assume $A_w = 1, A_b = 0$ the temperatures are maximally different:

$$T_b = T_o + 0.5q \tag{3.28}$$

$$T_w = T_o - 0.5q \tag{3.29}$$

Where $T_o = 295.5$ which is the temperature that produces the maximum birth rate. The temperatures of the black and white daisies are 305.5 and 285.5 respectively with both having a constant birth rate of 0.4285. While the birth rates are fixed, the coverage of the daisies is not. In order to find steady state daisy coverage and temperature values when both daisy types are present we recall the differential equation that governs daisy growth:

$$\frac{d\alpha_i}{dt} = \alpha_i(x\beta - \gamma) \tag{3.30}$$

At steady state, growth equals death and so:

$$x\beta = \gamma \tag{3.31}$$

We can now substitute in the value for the birth rate which is the same for both daisies and rearrange in order to find the amount of bare ground which is fixed during the period of both-daisy growth:

$$x = \frac{\gamma}{\beta} = \frac{0.3}{0.4285} = 0.7001 \tag{3.32}$$

Next, we note from (3.7) that with both daisies present the albedo of the planet is:

$$A_p = 0.5 - 0.5\alpha_b + 0.5\alpha_w \tag{3.33}$$

and using (3.4) and (3.6) that the temperature of the black daisies is:

$$T_b = q(A_p) + I(1 - A_p) \tag{3.34}$$

We know that the temperature of the black daisies is fixed at 305.5, therefore:

$$305.5 = q(A_p) + I(1 - A_p) \tag{3.35}$$

which rearranges to:

$$A_p = \frac{305.5 - I}{q - I} \tag{3.36}$$

The coverage of the daisies is found by recalling that the amount of bare ground, x, where $x = 1-\alpha_b-\alpha_w$, is constant when both daisy types are present. Planetary albedo in (3.33) can be related to x with:

$$\alpha_w + \alpha_b = 1 - x = 2A_p + 2\alpha_b - 1 \tag{3.37}$$

Substituting (3.36) for A_p in $1-x = 2A_p+2\alpha_b-1$ allows black and white daisy coverage to be expressed as functions of insolation:

$$\alpha_b = 0.64994 - \frac{305.5 - I}{q - I} \tag{3.38}$$

$$\alpha_w = 0.35006 + \frac{305.5 - I}{q - I} \tag{3.39}$$

Finally the temperature of the planet can be found by substituting equation (3.36) into (3.6) $T_p = I(1 - A_p)$:

$$T_p = \frac{I(q - 305.5)}{q - I} \tag{3.40}$$

Black and white daisy coverage when both daisy types are present is shown in figure 3.9. Figure 3.10 shows planetary temperature when both daisy types are present.

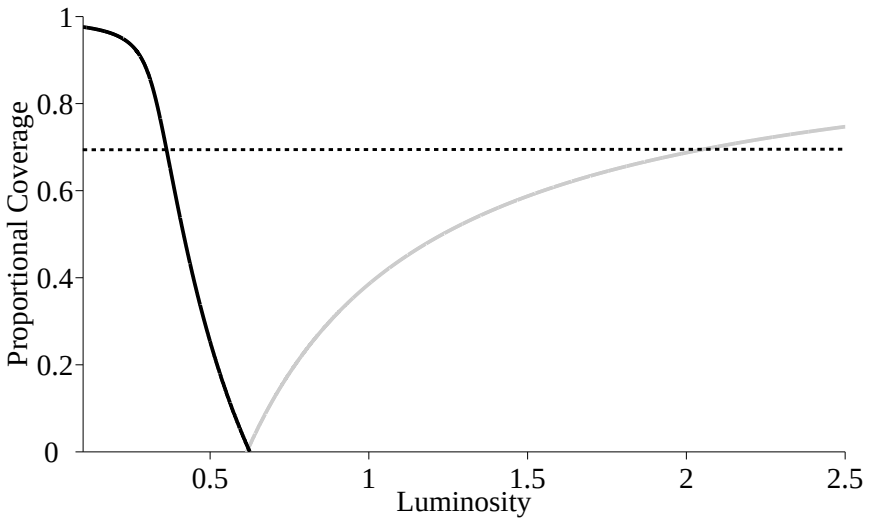

Figure 3.8: Analytical Coverage Solutions - Single Daisy Type

Analytical solutions for black daisies are plotted with a black solid line, white daisies a solid grey line when only black or white daisies are present. Hysteresis is not observed in the analytical solutions. The dashed line at 0.7 represents the maximum coverage possible for either daisy type as $\alpha_{max} = 1 - \gamma$. The stable steady state solutions for black and white daisy coverage is zero where the analytical solutions intersect the dashed black line.

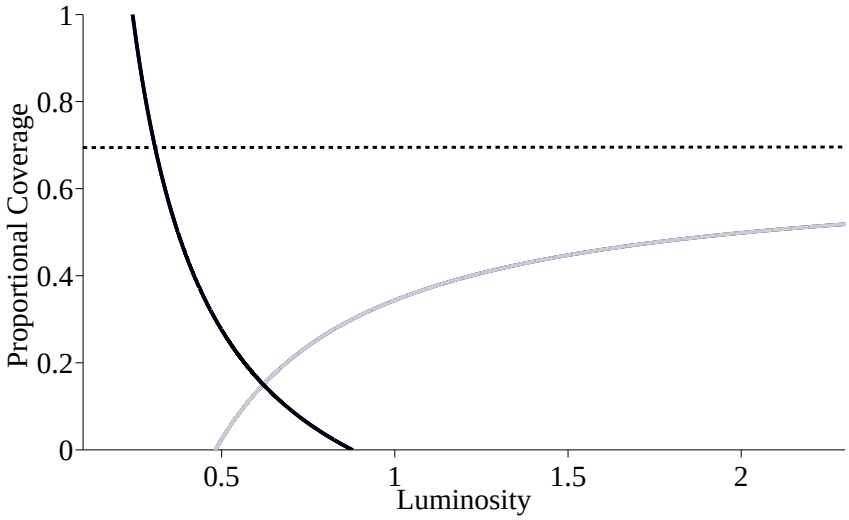

Figure 3.9: Analytical Coverage Solutions- Both Daisy Types

Analytical solutions for black daisies are plotted with a black solid line, white daisies a solid grey line when both black and white daisies are present. Hysteresis is not observed in the analytical solutions. The dashed line at 0.7 represents the maximum coverage possible for either daisy type as $\alpha_{max} = 1 - \gamma$. The stable steady state solutions for black and white daisy coverage is zero where the analytical solutions intersect the dashed black line.

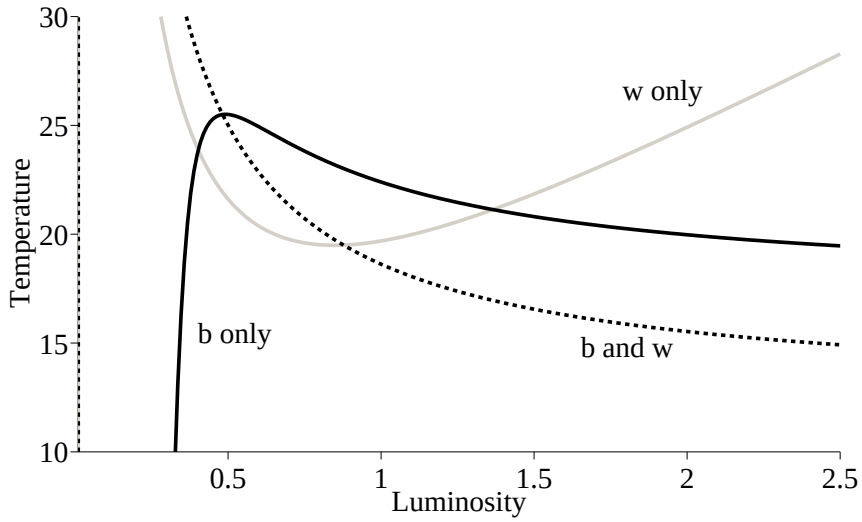

Figure 3.10: Analytical Planetary Temperature Solutions

Analytical solutions for planetary temperature with only black daisies present is plotted with a solid black line, with only white daisies present is plotted with a solid grey line and with both daisy types present is plotted with a dashed black line. The range of temperatures plotted is comparatively narrow. This allows the inverse response to be observed. All planetary temperatures have a point of inflection which corresponds to the inverse response to external perturbations observed in ODW. This shows that planetary temperature will decrease when there are black-only, white-only or both black and white daisies present.

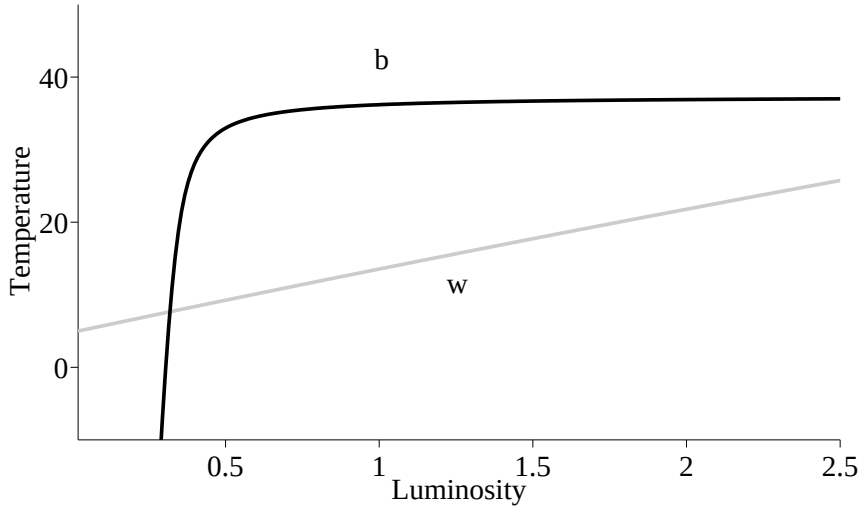

Figure 3.11: Analytical Daisy Temperature Solutions - Single Daisy Type
Analytical solutions for black daisy temperature with only black daisies present is plotted with a solid black line and with only white daisies present is plotted with a solid grey line. Both black and white daisy temperatures increase monotonically with increasing luminosity.

3.5 Investigating the Inverse Response

In this section I will highlight the causes of the inverse response to perturbations in ODW and SDW which will allow deeper insights into the homeostatic mechanism operating in Daisyworld. Figure 3.11 shows analytical solutions for planetary temperature with black only, white only and black and white daisy scenarios. It is important to note that all these scenarios feature a range of luminosity over which planetary temperature decreases with increasing luminosity. I will focus on the rate of change of planetary temperature when both daisies are present as the inverse response is most pronounced then.

It is straightforward to formulate a differential equation for the rate of change of the planetary temperature. Planetary temperature during the both-daisy range, equation 3.40, can be differentiated with respect to insolation:

$$\frac{dT_p}{dI} = \frac{q^2 - 305.5q}{(q - I)^2}. \tag{3.41}$$

Figure 3.12 demonstrates the effect different values for q has on planetary temperature and the rate of change of planetary temperature with respect to luminosity. The larger the value of q the greater the inverse response. The value of 305.5 in the $305.5q$ term is in part determined by the difference in the black and white daisy albedos. The values used in this chapter represent the greatest possible difference. If this is decreased by increasing the albedo of the black daisies and/or decreasing the albedo of the white daisies, then the range of insolation over which the inverse response is observed, as well as magnitude of this rate of change, is decreased. As $I \to q$, the absolute value of the rate of change of planetary temperature, $\dot{T}_p \to \infty$. Lower values of insolation will give larger rates of change in planetary temperature. $\dot{T}_p$ will be negative when $0 < q < 305.5$. As I increases it will dominate the term and so $\dot{T}_p \to 0$. With low insolation, increasing q produces greater positive values for T_p and greater negative values for $\dot{T}_p$.

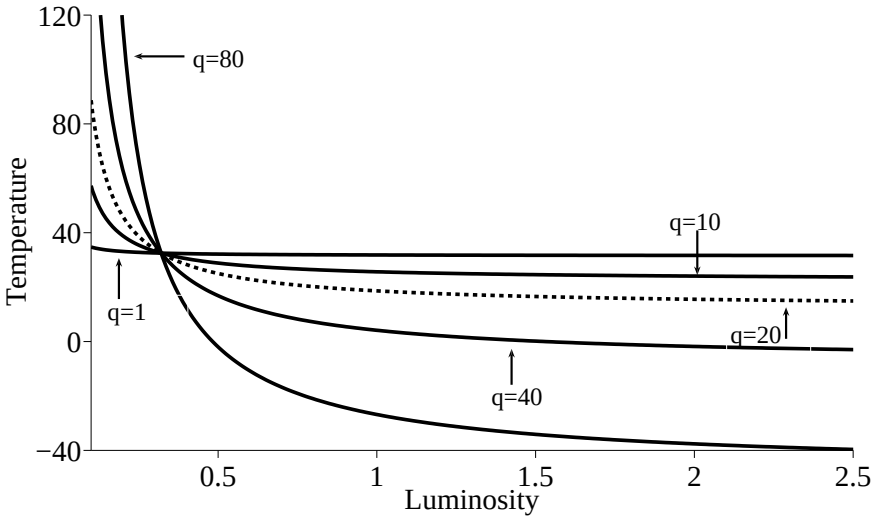

Figure 3.12: Planetary Temperature With Different Values for q

Analytical solutions for planetary temperature with both daisies present for five different values of q: 1, 10, 20, 40, 80. $q = 20$ is plotted with dashed line as this was the value used in SDW. Increasing the value of q increases the amount of insulation between the daisies and the bare ground and increases the rate of change of planetary temperature. Using the parameters described in this chapter, the range of luminosity over which both daisies would be present is approximately $[0.5, 0.9]$. Altering parameters such as albedo of the daisies or bare ground would alter this range.

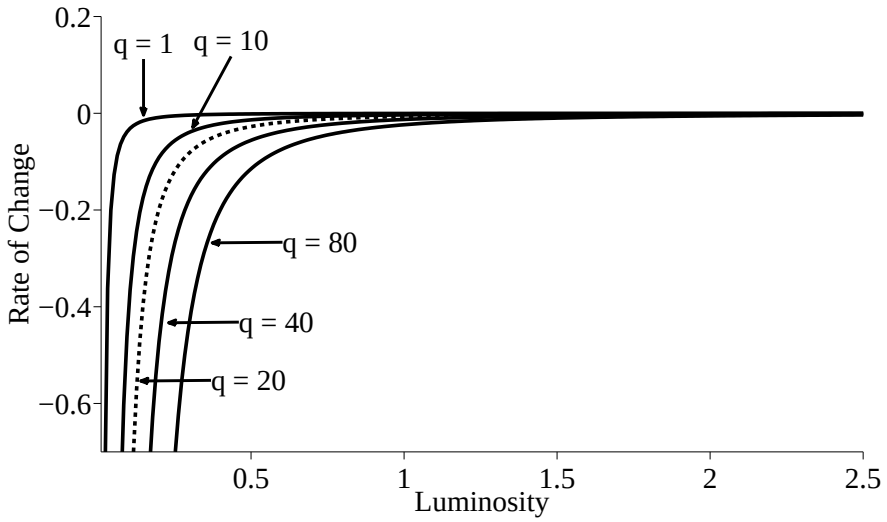

Figure 3.13: Rate of Change of Planetary Temperature With Different Values for q Analytical solutions for the rate of change of planetary temperature for five different values of q: 1, 10, 20, 40, 80. $q = 20$ is plotted with dashed line as this was the value used in SDW. With lower luminosity, increasing the value of q increases the magnitude of the negative rate of change of planetary temperature. As luminosity increases, all functions plotted converge to a value that is slightly less than zero.

3.5.1 What does q mean?

The method of finding daisy temperatures in both ODW and SDW leads to the greatest temperature difference between the daisies and their surroundings being found with the least coverage of daisies. This is arguably implausible as a single seed would, presumably, be the temperature of the ground that surrounds it. This modelling assumption has been defended by alluding to spatial factors that would mean black daisy seeds were more likely to be located in areas that would be warmer than the planetary average and white daisy seeds tending to be in areas cooler than the planetary average [Saunders, 1994]. This then introduces spatial assumptions into what was purported to be a zero-dimensional model.

We can put the plausibility of particular physical assumptions to one side for the moment and instead consider the effects of the particular method used to differentiate the daisy temperatures from ambient temperatures. I have shown that the q parameter is a critical component of the inverse response exhibited by the model. The q parameter

determines the degree of insulation between daisies and their surroundings, the ease by which heat will flow from hot to cold areas. As the value of q *increases*, insulation *increases*, the flow of heat *decreases* and so the greater the temperature difference between daisy and planetary temperature. When $q = 0$ there is no insulation and the system is isothermal with $T_b = T_w = T_p$. As the black and white daisies have the same temperature, any change in the proportional coverage of black daisies is matched with an equal change in white daisies. This produces a composite 'grey' daisy population with an albedo of 0.5 that is equal to that of the bare ground. Consequently the daisies no longer affect planetary albedo and so planetary temperature is simply a function of luminosity. Homeostasis is no longer observed. As q increases, the insulating effect increases and the temperatures of the black and white daisies are shifted further apart. This can be visualised as the black and white daisy temperatures moving further to the left and right on the birth rate function as in figure 3.14. There are effective limits to the increase in insulation if one is concerned with maintaining physical plausibility. For ODW, in order for the model to feature realistic heat flux, q must be set at a value less than $I\sigma$, where σ is Stefan-Boltzmann's constant, as values greater than this would result in heat flowing from cold to hot areas. The purpose of this discussion is not to assess the real-world plausibility of any particular value of q but rather the effects different values for this parameter will have on the model's behaviour.

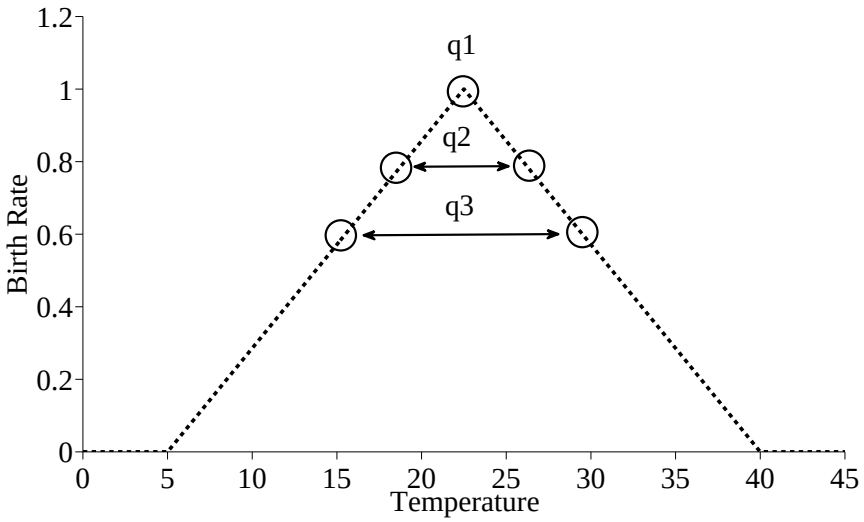

Figure 3.14: Black and White Daisy Temperatures With Different q Values

Steady state temperatures for the black and white daisies (denoted in circles) are shown for three values of the insulating parameter, q, where $q1 < q2 < q3$, and $q1 = 0$. White daisy temperatures are on the left hand side of the triangle, black daisy temperatures on the right hand side. For a fixed amount of luminosity, the steady state temperatures of the black and white daisies will be affected by the value of the q parameter. With no insulation the black and white temperature are the same. As the insulation effect increases, the steady state temperatures of the black and white daisies move further apart.

3.5.2 The inverse response and q

The inverse response to temperature can now be viewed as the reduction in the initially large temperature difference between the daisies and the planet. The range of luminosities under consideration will produce black-only, black and white then white-only steady states. The decrease in planetary temperature can be understood as a change from the hot black-only to cold white-only steady states. The inverse response can also be affected by the choice of albedo for the daisies. Increasing the difference between the black and white daisies, making the black daisies darker and the white daisies lighter, will increase the temperature difference between the daisies and the planet and so again increase the inverse response.

Varying the value for the q parameter varies the rate of the inverse response. To

fully appreciate the model's behaviour we must consider the role that bare ground plays. With both daisies present we were able to identify a symmetry in the black and white daisy birth rates. This symmetry was produced by both birth rates being a function of not only the daisy temperatures but the shared variable of bare ground. This led to the birth rates, and so temperatures, of the daisies being fixed within the both-daisy range of luminosity. The value of q determined how near or far the fixed birth rates were and consequently how near or far the black and white daisy temperatures were. This led to the birth rates, total coverage and daisy temperatures being fixed within the both-daisy range. It is planetary albedo and temperature that vary with the latter exhibiting the inverse response to perturbations.

In the following chapters I will show that the inverse response is not a central or necessary element of the Daisyworld control system. I will show that alternative mechanisms that do not use a 'q-type' parameter are available to couple the daisies and these do not show an inverse response to forcing. While the q parameter and its value are central to the homeostatic properties of the Original Daisyworld, the Daisyworld control system can be formulated in alternative ways which maintain homeostasis but remove the inverse response. I will also show that the fixed birth rates and daisy temperatures can be understood in terms of a general control system that will typically exhibit zero steady state error.

3.6 Summary

In this chapter I have presented a simplified version of the Original Daisyworld (ODW) model - Simplified Daisyworld (SDW). The first objective behind the formulation of SDW was to begin the process of simplification by which the core homeostatic mechanism operating within Daisyworld would be examined. The simplification process involved replacing the non-linear temperature and birth rate functions with linear functions. The second objective was to investigate the inverse response to perturbations observed in ODW. Numerical and analytical results showed that SDW also exhibited an inverse response to perturbations in qualitatively the same way as ODW: there was a range of luminosity over which the temperature of the planet *decreased* when luminosity was *increased*. This inverse response was not a result of any particular non-linear functions. When these were replaced by simplier, linear functions, the inverse response was still observed. I have shown that the method of differentiating the black and white

daisy temperatures from the ambient planetary temperature is the root cause of the inverse response. The greatest difference between daisy and planetary temperature is found with the least daisy coverage. Increasing q increases this temperature difference. Analytic solutions showed that the temperature of the planet is strongly non-linear at low luminosity, with increasing values of q increasing this non-linear response to luminosity. It is this non-linear response that is responsible for the decrease of planetary temperature with increasing luminosity. Figure 3.10 showed that the phenomenon of decreasing planetary temperature with increasing luminosity was also observed in black-only and white-only daisy scenarios although the rate of decrease was less. Within a more general context, the q parameter can be regarded as modulating the coupling strength between the two elements of the homeostatic system with higher values producing looser couplings.

In the next chapter I will continue my process of simplifying Daisyworld by presenting a Two Box Daisyworld. This model not only features simplified functions, but also removes some interactions present in ODW and SDW. Of particular interest, is that the method of differentiating the daisy temperature from the planetary temperature is altered. This removes the inverse response to perturbations.

Chapter 4

Two Box Daisyworld

4.1 Previous Publications

The results of section 4.5 are reproduced from [Dyke and Harvey, 2006]. The cable car model description and results are taken from [Dyke and Harvey, 2005].

4.2 Overview

In the previous chapter I presented a simplified version of the original Daisyworld. The simplifications allowed more insight into the workings of the model as well as identifying the mechanism behind the original model's inverse response to external perturbations. In this chapter I continue to simplify Daisyworld by developing a number of Two Box Daisyworld models (hereafter referred to as TBDW). The first TBDW model will use a different method to differentiate the daisy temperatures. I will present results from this model that confirm the finding that the inverse response was caused by the method of differentiating the daisy temperatures. I will also introduce a second TBDW which removes competition for a shared resource will be presented. This highlights an important aspect of the Daisyworld control system as the competition for a shared resource determines whether the model exhibits zero steady state errors with respect to planetary temperature.

A third version of TBDW is used to explore hysteresis in Daisyworld and the relationship between hysteresis and homeostasis. When increasing the essential range of the daisies in this model - increasing the range of temperatures that the daisies are able to grow over - the range of luminosity that the daisies are able to grow over is *decreased*.

Finally, in order to demonstrate the general applicability of the Daisyworld control system, I develop a simple simulated robotic agent. The control system for this agent is very nearly equivalent to the TBDW model. However instead of regulating temperature, the agent perform phototaxis (light following behaviour). I will show that a number of model dynamics are identical to TBDW and that there are close analogues between the phototactic agent and Daisyworld.

4.3 Two Box Daisyworld

The TBDW model was first proposed by Harvey [Harvey, 2004] who sought to radically simplify the original Daisyworld. Realistic modelling assumptions were replaced by 'toy physics' and the method of finding steady state solutions for temperatures and daisy coverage was simplified. TBDW also changed the way the two daisy types interact by removing competition for bare ground. I will show that this has significant consequences. To begin with I will present a model of intermediate complexity that shares the same physical modelling assumptions of Harvey's TBDW but continues to feature resource competition between the two daisy types as in O/SDW. Resource competition will be removed later.

TBDW can be conceptualised as a one-dimmensional model that features two grey boxes or cells that are connected energetically but not materially; black and white daisies only grow within their respective boxes and there is no migration between the boxes. TBDW is represented schematically in figure 4.1. When the daisy seeds are dormant both boxes have zero coverage of daisies and the system is isothermal. As the seeds germinate and daisies begin to cover each box, a temperature gradient is established. The two boxes are coupled via a heat conducting medium that allows heat to flow between the two boxes.

4.4 Model

TBDW begins with a pair of differential equations that govern the rate of change of the black and white daisies. These equations are identical to ODW and SDW and are reproduced here for convenience:

$$\frac{d\alpha_b}{dt} = x\alpha_b(\beta_b - \gamma) \tag{4.1}$$

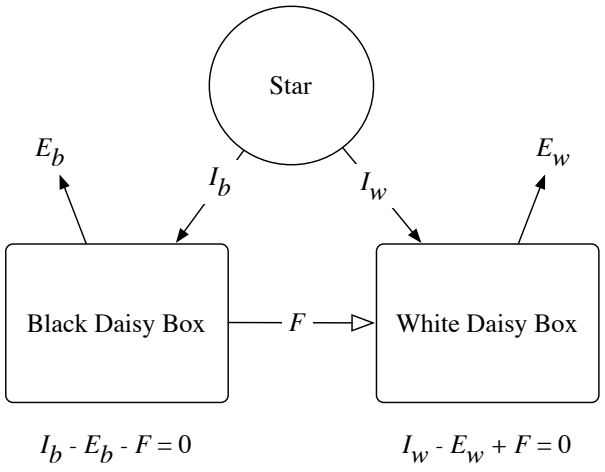

Figure 4.1: Schematic of the Two-Box Daisyworld

Insolation, I_i, warms each box and is radiated back into space as longwave emissions, E_i, with an amount of heat flux, F, proportional to the temperature gradient between the boxes. Only black daisies are seeded in the black daisy box and only white daisies in the white daisy box. White daisies, being lighter than the black daisies will reflect more energy from the star. Black daisies, being darker, absorb more energy. Hence the black daisies are warmer and the white daisies are cooler than the grey bare earth.

$$\frac{d\alpha_w}{dt} = x\alpha_w(\beta_w - \gamma) \tag{4.2}$$

where β_i are linear birth rates functions similar to those used in SDW in chapter 3. In O/SDW x represented bare ground that was available for new daisies to be seeded into. Given that there is no migration between the two boxes, x can instead be understood as a resource that determines the carrying capacity for the entire two-box system. This resource could be a globally distributed food or chemical. As in the previous chapter, it is important to note that I am not attempting to plausibly model a particular real world system and that terms such as 'temperature' and 'insolation' are used figuratively. In later chapters Daisyworld models will be used in a more broader biological context. The temperature of each daisy box is found with:

$$T_b = I(1 - A_b) - F \tag{4.3}$$

$$T_w = I(1 - A_w) + F \tag{4.4}$$

where A_i is the albedo of the boxes and I is insolation. As in O/SDW the amount of insolation will be parameterized by a luminosity variable, L, so that $I = LS$ where $L \in [0, 2]$ and S is the amount of energy emitted from the star that is received on the

surface of the planet and is fixed at 952.56. F is the heat flux between the two boxes and is found with:

$$F = D(T_b - T_w) \tag{4.5}$$

The flux of heat is proportional to the temperature difference between the two boxes and a diffusion parameter, D, which is fixed at a value within the range of [0,1]. When D equals 0 the two boxes are perfectly insulated and no heat flows. When D equals 1 there is maximum heat flow and the two boxes are isothermal. The temperature of the planet is found with:

$$T_p = 0.5I(A_b + A_w) \tag{4.6}$$

When the albedo of the bare ground, black daisies and white daisies are 0.5, 0.25 and 0.75 respectively, the albedos of each daisy box, A_i, are found with:

$$A_b = 0.5(1 - \alpha_b) + 0.25\alpha_b \tag{4.7}$$

$$A_w = 0.5(1 - \alpha_w) + 0.75\alpha_w \tag{4.8}$$

Figure 4.2 and 4.3 shows numerically computed steady state values for daisy coverage and planetary temperature when the diffusion parameter, D equals 0.5.

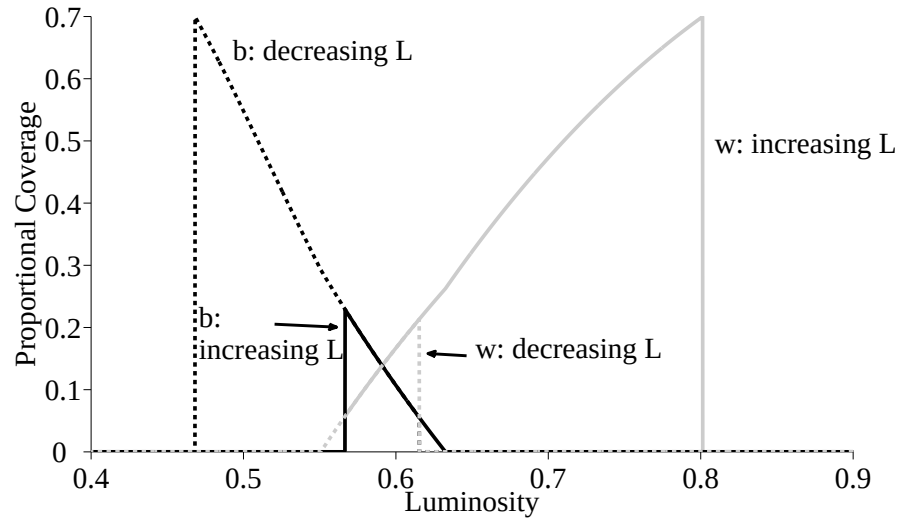

Figure 4.2: TBDW Numerical Daisy Coverage Solutions

Numerically computed results for daisy coverage for increasing luminosity are plotted in solid lines, for decreasing luminosity in dashed lines, for black daisies in black lines and white daisies for grey lines.

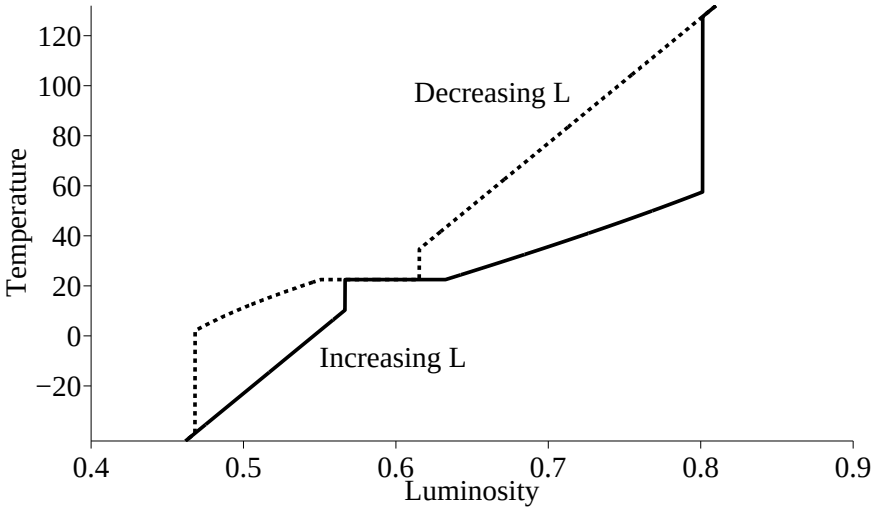

Figure 4.3: TBDW Numerical Planetary Temperature Solutions

Numerically computed results for planetary temperature with increasing luminosity is plotted with a solid line, for decreasing luminosity is plotted with a dashed line.

4.5 Resource Competition and Zero Steady State Error

Figures 4.4 and 4.5 gives a closer view of the both-daisy range and shows that the system does not exhibit the inverse response to perturbations observed in ODW and SDW. As luminosity increases, the planetary temperature no longer decreases during the both-daisy range, but remains *fixed* at the optimal growth temperature, T_o, of 22.5. The system is exhibiting a zero steady state error with respect to planetary temperature and the optimal birth rate temperature. However, it is important to note that while the temperature of the planet is at the optimum for growth, the daisy boxes are not. Black daisies are warmer and white daisies are cooler than T_o.

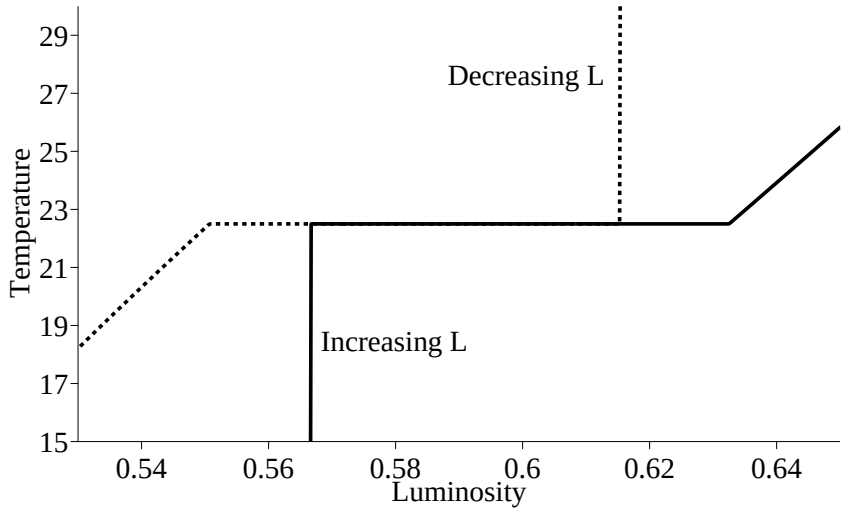

Figure 4.4: Planetary Temperature for TBDW With Resource Competition
Numerically computed results for planetary temperature with increasing luminosity is plotted with a solid line, for decreasing luminosity is plotted with a dashed line. Planetary temperature remains fixed at the optimal growth temperature of 22.5.

One can speculate that this effective regulation of planetary temperature at the optimal value over a range of luminosity values was the intention of Watson and Lovelock, but that the method of modelling heat flux resulted in the inverse response observed in O/SDW. It may be argued that the simpler implementation in TBDW provides a commensurately simpler result and more direct insight into the sort of biologically-mediated homeostasis that Watson and Lovelock were investigating. However it is important to

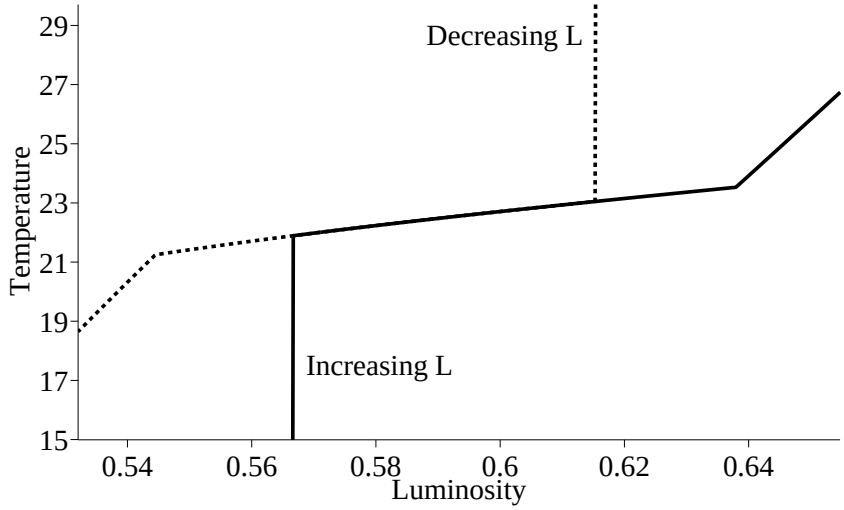

Figure 4.5: Planetary Temperature for TBDW Without Resource Competition Numerically computed results for planetary temperature with increasing luminosity is plotted with a solid line, for decreasing luminosity is plotted with a dashed line. Planetary temperature no longer remains fixed at the optimal growth temperature of 22.5. The rate of change of planetary temperature within the both-daisy range (when both black and white daisies are present) is less than the any-daiy range (when black only or white only daisies are present) which is less than when no daisies are present.

appreciate the context in which the Original Daisyworld model was developed and that the model is based on established zero dimensional climate models. As in O/SDW, resource competition results in the birth rates of the black and white daisies being equal as $x\beta = \gamma$. However, unlike O/SDW the birth rates are not fixed during the both-daisy range. The birth rates are not fixed because the temperatures of the daisies are not fixed. It was the q parameter that pegged the daisy temperatures a fixed amount from the optimal growth temperatures. In the previous chapter this was visualised as increasing q values moving the black and white daisy temperatures further apart on opposite sides of the birth rate function. Daisy temperatures are no longer fixed in TBDW as heat flux is modelled more explicitly by it being a function of the temperature gradient between the two boxes and a fixed diffusivity term. Heat flux varies with varying luminosity as do daisy temperatures. While these variables, including the birth rates of the black and white daisies, vary with varying luminosity, the birth rates are equal. The competition for a shared resource leads to a 'zero sum game' in which an increase in the proportional coverage of black daisies is matched by an equal decrease in the proportional coverage of the white daisies and vice versa. An exact increase in black daisy coverage is matched by an exact decrease in white daisy coverage that leads to a zero sum change in the albedo of the planet and so planetary temperature remains fixed at a point equidistant from T_b and T_w.

Removing resource competition from TBDW is straightforward and demonstrates that in the absence of zero sum game dynamics between the daisies, planetary temperature no longer remains fixed during the both-daisy range. The replicator equations now feature separate x_i terms:

$$\frac{d\alpha_b}{dt} = x_b\alpha_b(\beta_b - \gamma) \tag{4.9}$$

$$\frac{d\alpha_w}{dt} = x_w\alpha_w(\beta_w - \gamma) \tag{4.10}$$

where $x_b = 1 - \alpha_b$ and $x_w = 1 - \alpha_w$. The effects of this change can be seen in figure 4.5. Planetary temperature is no longer fixed at T_o within the both daisy range but increases slightly.

4.6 Hysteresis and the Limits of Homeostasis

TBDW exhibits qualitatively the same hysteresis observed in O/SDW. In this section I will examine how hysteresis in TBDW affects homeostasis. This will be progressed by altering the viability range of the daisies. Intuitions may suggest that the width of the essential range determines the range of luminosity over which the daisies grow. If black and white daisies are able to grow at colder and hotter temperatures then increasing the width of the birth rate function should increase the any-daisy range of non-zero growth. Daisies should be present at lower and higher values of luminosity.

I will show that is not necessarily the case. There are circumstances where *increasing* the essential range leads to a *decrease* of the any-daisy and both-daisy range. Moreover there are circumstances where decreasing the essential range can have no effect on the any-daisy and both-daisy range. This can lead to the situation where a TBDW that features daisies with an essential range that is approaching zero has the same any-daisy range as a TBDW with an original essential range or even wider essential range. What typically changes with a varying essential range is the area of hysteresis.

4.7 Harvey TBDW

In order to explore these issues, I will use what I will call the 'Harvey TBDW'. The difference between TBDW and Harvey TBDW is the way the daisy coverage steady state values are found. TBDW uses the same methodology as O/SDW in that replicator equations that feature a variable birth rate with a fixed death rate determine the rate of change of the daisy coverage. Harvey TBDW dispenses with these replicator equations and instead employs a single 'hat' function that has the same form of the linear birth rate function used in S/TBDW, but now determines the steady state coverage of the daisies:

$$\alpha_i = H(T_i) = max\left\{\ 1 - 2^{\frac{|T_o - T_i|}{R}}, 0\ \right\} \tag{4.11}$$

It is now no longer necessary to define a variable such as the amount of bare ground that in TBDW produces density dependent growth, as the hat function itself limits the maximum coverage of the daisies. A pair of differential equations are implicit in the model and it is these that are numerically integrated to steady state:

$$\frac{d\alpha_b}{dt} = H(T_b) - \alpha_b \tag{4.12}$$

$$\frac{d\alpha_w}{dt} = H(T_w) - \alpha_w \tag{4.13}$$

The method of integration can be understood as a variant of the Euler feed forward method. The algorithm is detailed below.

1. Calculate box albedos from black and white daisy proportional coverage.
2. Calculate box temperatures.
3. Calculate heat flux F.
4. Adjust each box temperature by $\pm F$.
5. Calculate new daisy proportional coverage, α_i'.
6. Adjust proportional daisy coverage: $\alpha_i = \alpha_i(1-\delta) + \delta\alpha_i'$.
7. Go to 1.

When $\delta = 0.001$, 200,000 iterations produce changes in coverage, albedo, heat flux and temperature that are no greater than 10^{-22}. It is assumed, as with other studies on Daisyworld, that the rate of change of daisy coverage is sufficiently faster than that of luminosity so as to allow the above equations to be numerically integrated to steady state whilst luminosity is fixed.

The original Harvey TBDW initialised daisy coverage to the maximum amount for each value of luminosity and then integrated the equations to steady state. This removed any hysteresis from the system. As in TBDW and O/SDW I will instead assume that the coverage of the daisies evolve over time. Daisy coverage will not be reset to the maximum value at the start of the integration algorithm but will be set to the previous luminosity steady state values. This recovers hysteresis and allows examination of this aspect of the model's dynamics.

Although conceptually simpler than O/SDW, Harvey TBDW contains a number of assumptions that make it significantly more difficult to find exact analytical solutions. Of particular importance is the removal of any resource competition and a single hat function rather than pair of differential equations. This removes the symmetry in temperature and birth rates that was the route into solutions for O/SDW while the linear hat function complicates the algebraic solution. Exact analytical solutions to the

model were found and published in [Wood et al]. This solution was the work of Jamie Wood. They are reproduced below. I made minor contributions:

$$\alpha_b = \frac{k}{k + \frac{4T_I B}{(1+2D)}} - \frac{2(T_I - T_o)}{k + 4T_I B} \tag{4.14}$$

$$\alpha_w = \frac{k}{k + \frac{4T_I B}{(1+2D)}} + \frac{2(T_I - T_o)}{k + 4T_I B} \tag{4.15}$$

$$T = T_I \left[\frac{T_o + k/4B}{T_I + k/4B} \right] \tag{4.16}$$

where B is based on the assumption that the albedos of the black and white daisies are equally distant from the 0.5 albedo of bare ground. $A_b = 0.5 - B$ and $A_w = 0.5 + B$. T_I is the 'imposed' temperature - the temperature of Daisyworld with no daisies. k is a constant that takes the value of 20. Of particular interest to this chapter's discussion is that the rate of change of planetary temperature is now positive for all luminosity.

4.7.1 Results

Two versions of TBDW were numerically integrated to steady state. The first version features a 'normal' essential range. The black and white daisies had non-zero coverage over the range of [90,110] so the essential range was 20. The second version featured a 'wide' essential range. The black and white daisies had non-zero coverage over the range of [85,115] so the essential range was 40. For both versions the diffusivity parameter was set to 0.8 and the albedos of the black and white daisies were 0.25 and 0.75 respectively. Solutions were found for luminosity values that started low and then increased, and when luminosity values started high and then decreased.

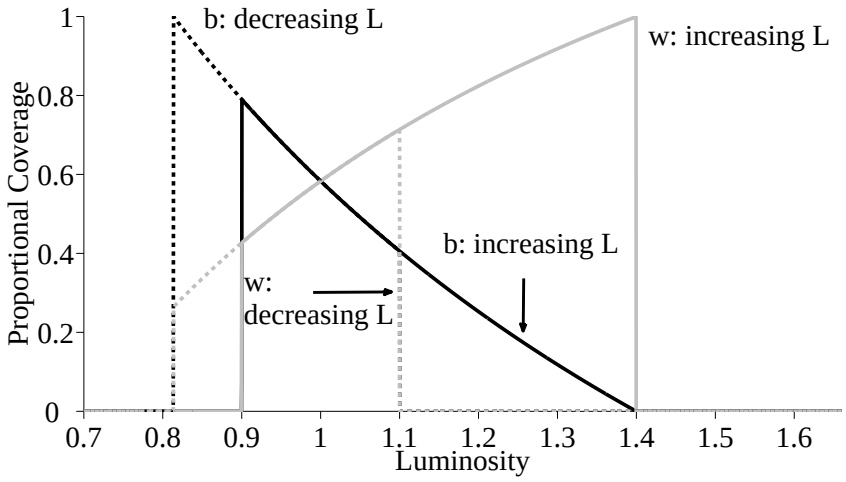

Figure 4.6: Daisy Coverage With Narrow Essential Ranges

Numerically computed results for daisy coverage for increasing luminosity are plotted in solid lines, for decreasing luminosity in dashed lines, for black daisies in black lines and white daisies for grey lines when the essential range of both daisies is set to 20. When the essential range is 20, daisies are present over the luminosity range [0.81,1.4]. When the essential range is 40, daisies are present over the narrower luminosity range of [0.83,1.31].

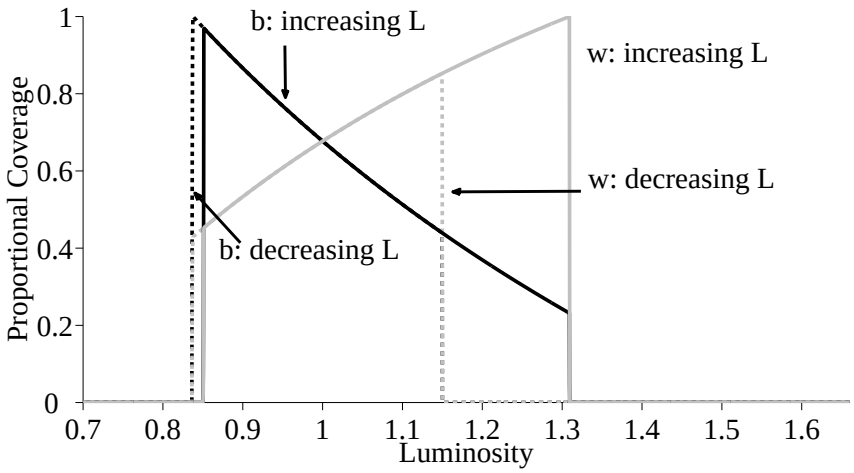

Figure 4.7: Daisy Coverage With Wide Essential Ranges
Numerically computed results for daisy coverage for increasing luminosity are plotted in solid lines, for decreasing luminosity in dashed lines, for black daisies in black lines and white daisies for grey lines when the essential range of both daisies is set to 40. When the essential range is 40, daisies are present over the narrower luminosity range of [0.83,1.31].

Figures 4.6 and 4.7 show that allowing the daisies to grow over a wider range of temperatures has resulted in a system that is less able to regulate at the extremes of luminosity. To understand this result, there are two important points to note. Firstly, the width of the essential range is largely irrelevant in determining the any-daisy range. What will change with a varying essential range is the area of hysteresis. Secondly, at the limits of the any-daisy range, a single daisy type is responsible for the maintenance of homeostasis.

Figure 4.8 geometrically represents the qualitative differences produced with two different essential range values with only white daisies present. When luminosity is low and then increases, the wider essential range daisies will begin to grow at lower luminosity than the narrower essential range daisies. However, both narrow and wider essential range daisies will continue to grow until the same maximum coverage is reached. It is the *height* of the hat function not its slope that determines the limits of the any-daisy range. The limits of white daisy growth are determined by the strength of the feedback the daisies exert on temperature (which is determined by the albedo of the white

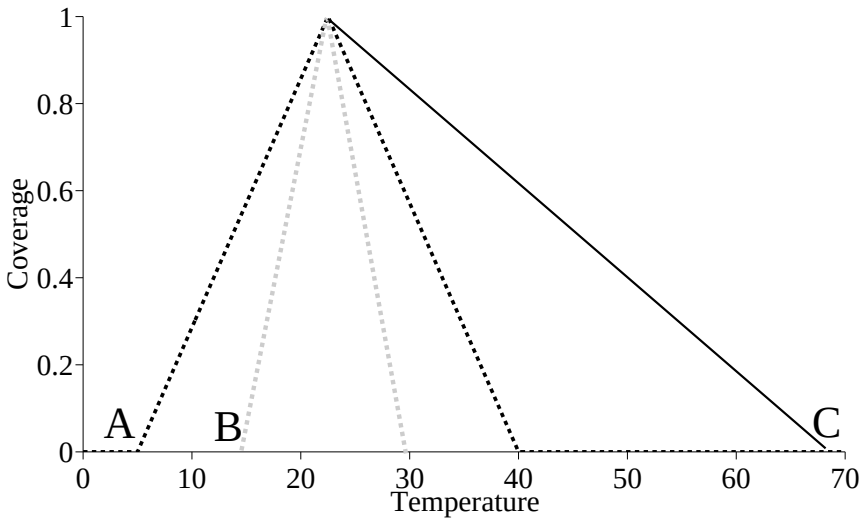

Figure 4.8: Luminosity Range With Different Essential Range Values

The 'normal' coverage function is plotted with a black dashed line; the 'narrow' coverage function is plotted with a grey dashed line. The sloping solid black line represents the negative feedback white daisies exert on temperature. With low then increasing luminosity white-only daisies will begin to grow at point A for the original birth rate and point B for the narrower coverage function. Irrespective of the point at which the white daisies begin to grow, both original and narrower coverage functions will persist up to point C as the original and wide birth rates have the same maximum value. It is this maximum value and the strength of feedback the daisies exert (which is determined by their albedo) that determines the maximum amount of luminosity that the daisies are able to withstand.

daisies) and the maximum value of the birth rate function. With a fixed maximum birth rate, increasing the albedo of the white daisies increases the the maximum luminosity that produces non-zero growth.

The considerations above explain why the any-daisy range does not increase with an increase in the essential range. To understand why the any-daisy range has decreased requires an appreciation that at the limits of regulation, only a single daisy type is actively maintaining homeostasis. With increasing luminosity, it is the increase in the coverage of white daisies that maintains temperature to within the habitable range. For a fixed amount of luminosity, increasing the coverage of the black daisies will increase the temperature and so move the system closer towards the limits of homeostasis. Such an increase in black daisy coverage is produced by increasing the essential range. The black daisies are able to grow at higher temperatures and so increase the temperature

of the white daisies. Therefore the white daisy bed temperature will reach the optimal temperature with lower luminosity and thus the system is less able to maintain homeostasis at high luminosity values.

With decreasing luminosity the same situation applies, but in reverse. At the limits of homeostasis, the black daisies are increasing towards maximum coverage. Increasing the coverage of the white daisies cools the black daisy bed and so decreases the temperature, moving it closer to the optimal coverage temperature. Once the black daisy bed temperature falls below the optimal coverage temperature, the system collapses.

It is the interference at the limits of habitability that is responsible for the impairment of homeostasis. This interference was produced by increasing the essential range. It is also important to note that the diffusivity value, D, can also affect the limits of homeostasis. Increasing the value of D produces a more tightly coupled system. The value of diffusivity used in this section's results was selected in order to highlight the effects of altering the essential range. The values for diffusivity and the essential range that result in this decrease in the any-daisy range will vary with differing model assumptions and parameter values.

As well as showing a difference in the any-daisy range, figure 4.8 shows that altering the essential range alters the area of hysteresis in the system. Increasing the essential range decreases the area of hysteresis as the daisies will begin to grow at lower and higher luminosities. In a biological context, this change in the area of hysteresis may be irrelevant. For example, if the essential range represented the range of temperatures that a mammal must regulate its core body temperature to within, then once the temperature moves outside of this range, the animal will die. Decreasing the area of hysteresis will not mean that the animal will spontaneously spring back to life 'sooner'! There may be other biological and artificial contexts in which the area of hysteresis will be an important factor. With respect to the original Daisyworld, scenarios typically only considered increasing luminosity. Therefore increasing the essential range would result in an increase in the any-daisy range as black daisies would grow at lower luminosities.

4.8 Cable Car Model

In order to demonstrate the general applicability of the Daisyworld models presented in this chapter I will present the Cable Car Model (hereafter referred to a CCM) initially proposed in [Dyke and Harvey, 2005]. The CCM is a simulated simple robotic

controller that performs phototaxis (light following behaviour). I will show that the CCM is essentially a TBDW model, with phototaxis being an analogous behaviour to temperature homeostasis. It displays behaviour similar to Harvey TBDW. I will also use the CCM in order to explore the relationship between hysteresis and homeostasis in Daisyworld.

The CCM is based on the cable cars found in San Francisco. Unlike the systems used in the Alps and other mountainous regions, the San Francisco system consists of cables that are located under the road surface and connect to tram-like cars. In the CCM, a photovoltaic cell (a solar panel) is located on the roof and supplies current to an electric motor which instead of being located in a winding house is carried within the cable car itself. As the motor turns, it winds in a cable that moves the car up the side of a valley. The output of the solar panel and therefore the force that the motor produces, changes linearly with varying inclination from a moveable light source. When the light source is directly overhead, maximum output is produced and so maximum motor output is achieved. Deviations left or right by either the cable car or light source result in decreasing motor output. The range of light source locations that produce current in the solar panel is the activation range and is analogous to the essential range of viable daisy temperatures in Daisyworld. Figure 4.9 shows a geometrical overview of the CCM.

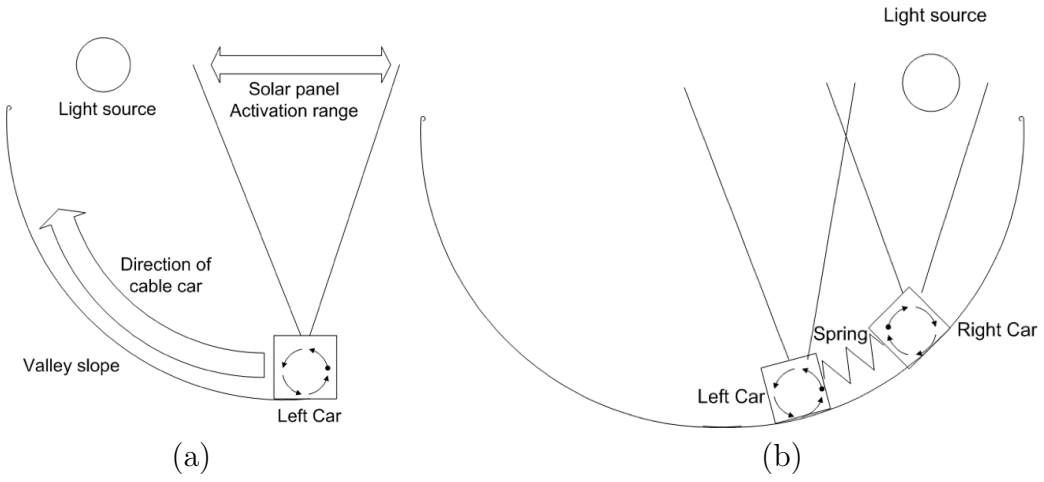

Figure 4.9: The Cable Car Model
One simple simulated agent travels up the left side of a valley (the left car) while another travels up the right side (the right car). The distance up the valley is determined by the output of the motor which is powered by a solar cell. As a light source moves from left to right, the positions of the car vary. The two cars are connected with a spring so that either car can pull the other car up the opposite slope.

4.9 Model

The CCM consists of two cable cars. The left car travels up the left side of the valley while the right car travels up the right side of the valley. Both cable cars have a dimensionless mass of 1 unit. As the cars move further from their starting position, the gradient of the slope increases and so the resisting force due to gravity, γ, pulling the cars back to the bottom of the valley increases:

$$\gamma_{left} = \eta \frac{X_{left}}{l} \tag{4.17}$$

$$\gamma_{right} = \eta \frac{X_{right}}{l} \tag{4.18}$$

Where X_i is the position of the cars in dimensionless x-units, l is the length of the slope and η measures the rate of increase of the resisting force as the car travels higher up the slope. η was set to 1 and so increases linearly from 0 when the car is at the bottom, to 1 when the car is at the top of the valley. The output of the cable car motor, O, is a function of the current position of the cable car and the light source. The form of this function is essentially equivalent to the linear piece-wise hat function used in SDW and TBDW:

$$O_i = max\left\{\ 1 - 2\frac{|X_{light} - X_i|}{R}, 0\ \right\} \tag{4.19}$$

where X_{light} is the location of the light source in dimensionless x-units and R is the activation range of the solar panel. The energy provided by the solar panel turns the motor in the left car anti-clockwise which moves it to the left and the motor in the right car clockwise which moves it to the right. A spring is attached between the cars. As the cars move apart, the spring is stretched and a force, F, is exerted that pulls the cars back together. This force is essentially equivalent to the flux of heat in TBDW:

$$F = E(X_{left} - X_{right}) \tag{4.20}$$

where E is the elasticity of the spring and is parameterized from 0 (infinitely elastic, giving F = 0), to 1 (completely rigid so that both cars move as a single unit). This is essentially equivalent to the diffusion parameter D in TBDW.

Distance is measured in dimensionless x-units. The furthest left side of the valley is at x-position 0. The centre of the valley is at x-position 50. The furthest right side of the valley is at position 100.

Table 4.1 lists the components of the cable car and Daisyworld models and allows a comparison of the two.

Cable Car Model	**Daisyworld**
Motor output	Daisy coverage
Light location	Luminosity of star
Solar panel activation range	Range of viable growth temperatures
Left car	Black daisies
Right car	White daisies
Connecting Spring	Flow of heat from hot to cold

Table 4.1: CCM and TBDW Comparisons

4.10 Results

Simulations consisted of initialising the light on the right hand side of the valley at x-position 72. The light is slowly moved to the left, past the centre of the valley at position 50 to position 28 on the left hand side of the valley. It then reverses direction and travels back to the starting x-position at which point is reverses and repeats the movement. Over time the movement of the light source is sinusoidal. Figures 4.10 and 4.11 shows numerical solutions for the output of the left and right cable cars and their position as the light source moves back and forth. The motor outputs are analogous to the proportional coverage of black and white daisies with the left motor then right motor output changing as the light moves back and forth. The differential motor outputs lead to the light being maintained between each cable car. The system is exhibiting phototaxis in a way similar to a Braitenberg Vehicle [Braitenberg, 1984]. This is analogous to the the planetary temperature being maintained within the essential range in Daisyworld as the light source is maintained to within the activation range of the cable cars. The cable cars are able to 'survive' over a wider range of light positions.

4.11 Hysteresis and the Limits of Phototaxis

Hysteresis is observed when simulations are performed with the light source outside of the activation range of either cable car as shown in figure 4.10. Once the cable cars begin to move, they are able to track the light over the entire range of light source positions. If the light source begins within the activation range of either car, or if the right/left

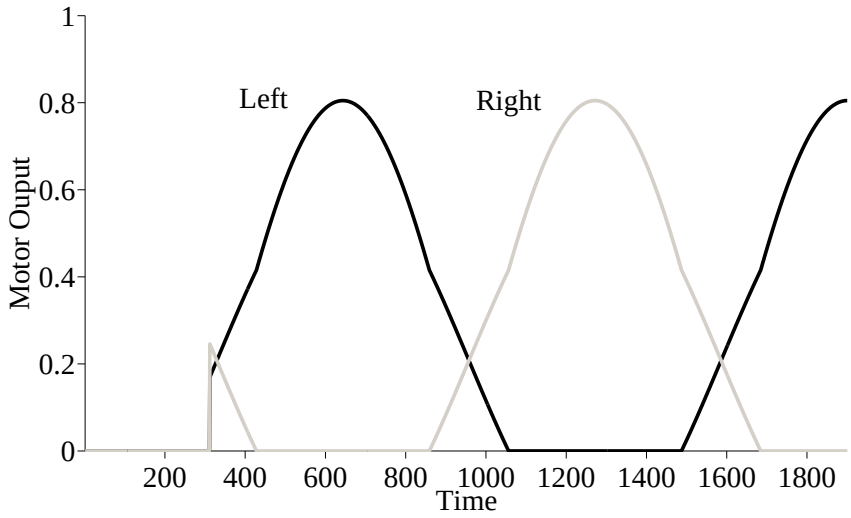

Figure 4.10: Cable Car Motor Ouput
The motor output of the left cable car is plotted with a black line, the motor of output of the right cable car is plotted with a grey line. As the light moves back and forth across the centre of the valley there is a differential activation of the left and right motors.

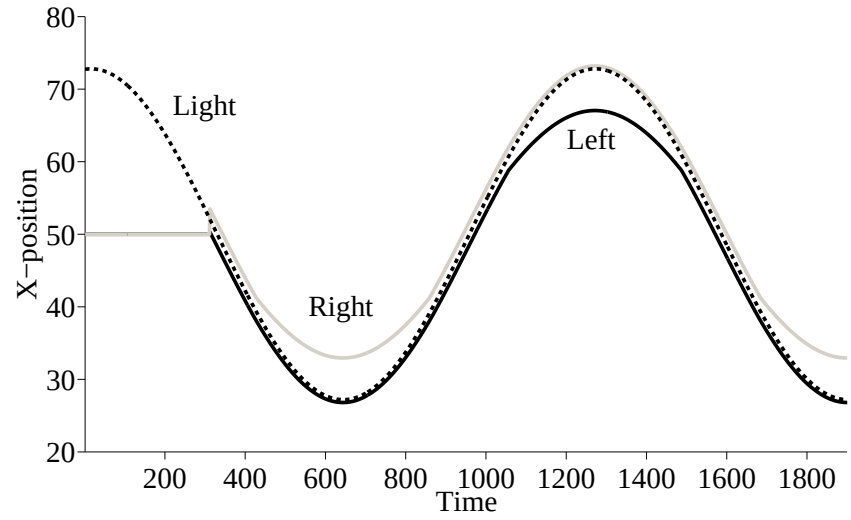

Figure 4.11: Cable Car Car Positions
The position of the left cable car is plotted with a solid black line, the position of the right cable car is plotted with a solid grey line. The position of the light source is plotted with a dashed black line. The differential response of the motors shown in figure 4.10 leads to the light source being maintained between the two cars that track the light source.

hand car is held at its maximum distance from the bottom of the valley and the light is introduced from the right/left, then there is no period when the cars are inactive. This means that decreasing the activation range may not decrease the range of light positions over which phototaxis can be achieved. For example, the hat function could be transformed into a function which is arbitrarily sharp with the result that the light source must be directly overhead in order to produce solar panel output and so motor force. This is equivalent to the the finding in TBDW that the width of the essential range did not determine the any-daisy range, but rather the area of hysteresis. For a single cable car, the limits of activation are not determined by the form of the activation function but rather the maximum output of the motor. This means that the function that determines the limits of phototaxis, X_{max}, does not contain any information about the form of the activation function:

$$X_{max} = O_{max}\frac{l}{\eta} \tag{4.21}$$

where O_{max} is the maximum possible motor output. If the left hand car starts at the bottom of a slope of length 50 and has a maximum motor output of 1 with η set to 1, the furthest left it can travel is 50 units. It is not necessary to specify the width or any other aspect of the solar panel activation range in order to determine X_{max}. The presence of both cars complicates the situation. For example, it is possible that *increasing* the activation range of both cable cars lead to a *decrease* in the limits of phototaxis. The mechanism behind this behaviour is analogous to that described in section 4.7.1 in this chapter.

4.12 Analysis

The strength of each feedback channel, in other words how hard the motors pull, in part determines the range of external forcing over which the cable car model is able to perform phototaxis. What the channels are pulling against is just as important. Indeed it is the resistance that the cable cars experience as they travel up the valley slope that allows phototaxis to be performed. To explain this will require a moment's anthropomorphising. The roots of hysteresis are found in the different cable car behaviours in response to the light source that are dependent on the direction from which it enters the activation range. For example, if the light source approaches the *left* car from the *right* hand side, the car attempts to run away up the slope. It is light-phobic. This is

referred to as behaviour A. If however the light source approaches from the *left* hand side, it is a light-phile. The car runs up the slope to meet the light and then follows it back down the slope. This is referred to as behaviour B. The response of the car is the same under both situations. As the cable car motor is only able to move the car up the slope, the moment the light is detected, the car attempts to escape to the left. During behaviour B, this moves the light across the centre of the activation range, past the mid-point, coming to rest on the right hand side where the output of the motor is balanced by the resisting forces pulling the car back to the bottom of the valley. This is equivalent to the balance of daisy birth rates and death rates in Daisyworld. As the light continues to move to the right, the car follows as the motor output steadily decreases. Due to the coupling spring, the left hand car will continue to move to the right and travel up the opposite slope as the right hand car performs its equivalent of behaviour A. Behaviour B is produced by the expenditure of the potential energy that the car stored when the light source first entered the activation range. It is the storing then release of this energy, the balance between γ and O as the cars move up and down the energy gradient that produces both behaviours.

Increasing the essential range can have a dramatic impact on the limits of phototaxis. As the essential range of the cars increases, the sum of resisting forces on the active car increases as the passive car pulls more strongly away from the light source. The active car not only hauls the mass of the passive car up the valley slope, but has to drive against the passive cars motor which attempts to pull the cars in the opposite direction. This can lead to a reduction in the ability of the cars to follow the light just as an increase in the range of temperatures that gave non-zero growth in TBDW led to a reduction in the range of luminosity over which the daisies grew.

4.13 Summary

In this chapter I have presented a number of Daisyworld models of decreasing complexity. The Two Box Daisyworld (TBDW) was presented. This differed from SDW in that the insulating q parameter was replaced with a diffusion parameter, D, that was responsible for modulating the flux of heat across a temperature gradient produced by black and white daisy coverage. The inverse response to external driving was no longer observed and planetary temperature did not decrease during the both-daisy range but instead remained fixed at the optimum growth temperature. It was shown that this

zero steady state error was the result of a zero sum game being played between the black and white daisies. When the birth rate of each daisy type is a function of a shared resource variable, the black daisies can only increase in coverage at the expense of the white daisies and vice versa. Removing resource competition resulted in planetary temperature no longer remaining fixed, but slightly increasing with increasing luminosity.

The Harvey TBDW was presented and used to explore the effects of altering the essential range of the daisies, the range of temperatures that the daisies are able to grow over. It was shown that increasing the essential range may not increase the range of luminosity over which the daisies grow. Indeed for a range of parameter values, *increasing* the essential range *decreases* the any-daisy and both-daisy ranges.

The Cable Car Model (CCM) was presented and used to demonstrate that the Daisyworld control system is in principle applicable to non-biogeochemical domains. Rather than regulate temperature, the CCM performed phototaxis via the differential response of two cable cars in ways analogous to the differential proportional coverage of black and white daisies in response to varying luminosity. A number of behaviours of the CCM were shown to be essentially equivalent to TBDW. In particular, the limits of phototaxis with a single cable car were shown to be independent of the width of the solar panel activation range.

In the following chapter the simplification process reaches its conclusion with the formulation of a minimal Daisyworld control system. I will show that this control system is an example of a *rein control* system. A number of results in this and previous chapters will be explained in terms of rein control. This analysis will then be used in later chapters to explain the homeostatic properties of more complex models.

Chapter 5

Rein Control

5.1 Overview

The previous chapter developed a series of simpler Daisyworld models. This chapter concludes the simplification process with a minimal model of the Daisyworld control system. I will show that this system is an example of a *rein control* system. After a brief review of the origins of rein control I will construct a simple dynamical system that retains the bare essentials of the Daisyworld control system. I will present analytical solutions for this system that demonstrate the significance of the different types of interactions between rein controllers. In particular I will demonstrate the difference between rein control and what I call *antagonistic rein control*. This will highlight the importance of the nature of the coupling between elements of a rein control system and will explain the result of the previous chapter that competition for a shared resource led to a zero steady state error with respect to planetary temperature and the optimal growth temperature. The chapter concludes by raising the possibility that Daisyworld represents a selective use of feedback elements in a rein control system in order to ensure negative feedback predominates and homeostasis is obtained. The charge put to Daisyworld is that if these assumptions are relaxed in order to allow both positive and negative feedback elements, homeostasis will no longer be observed. This chapter allows us to quantitatively formulate such charges while the following chapter will assess them.

5.2 Unidirectional Rein Control

Manfred Clynes postulated that many physiological homeostatic processes operate on the basis of separate channels which can be likened to the pair of reins used in controlling a horse [Clynes, 1969]. He argued that the origins of unidirectional rein control can be found in the non-polarity of molecular signaling. Molecules are only ever transmitted in positive numbers. One does not observe a negative number of calcium ions or potassium molecules being sent or received over a particular channel. 'Negative-molecules' are not transmitted in order to cancel out the effects of 'positive-molecules'. Rather there are separate channels that work in opposing directions in order to affect a system variable in two directions. These channels are called reins, as reins can only pull not pull and push.

Clyne's notion of rein control was developed in two papers by Saunders, Koeslag and Wessels. [Saunders et al., 1998] formulate a rein control model in order to identify pathological features of type I and type II diabetes while [Saunders et al., 2000] extends this analysis into a more general model that can conceptually model an arbitrary number of variables. Saunders et al model the rate of change of blood glucose, V over time with:

$$\frac{dV}{dt} = I + \alpha(A, V) - \eta(B, V) - \lambda(k, V) \tag{5.1}$$

where, I is glucose entering the blood stream from food and the λ function is the uptake of glucose by muscles. A and B represent the levels of the hormones glycogen and insulin respectively. Functions α and η model the effects of the two hormones. The rate of change of A and B are governed by the following equations:

$$\frac{dA}{dt} = A(\phi(V)x_1(A, B) - \gamma_1(A)) \tag{5.2}$$

$$\frac{dB}{dt} = B(\psi(V)x_2(A, B) - \gamma_2(B)) \tag{5.3}$$

where x is an 'inhibitory' function defined as $x_1 = x_2 = 1 - A - B$. This means the production of one hormone suppresses the other. γ_i determines the rate of decay of A and B. Functions ϕ and ψ are increasing and decreasing functions of V so that as glucose concentrations increase, insulin increases and as glucose decreases, glycogen increases. By clamping the output of either A or B in the model, Saunders et al were able to explain features of the disease as the loss of one or two rein controlling elements. For example, if the production of insulin is impaired, blood glucose will be

maintained at a higher threshold and perturbations in input may be sufficient for type I hypoglycaemia.

In the following sections I will develop a series of simple rein control models. These are effectively a sub-set of Saunders et al's formalism. The simplifications are made in order to explore several aspects of rein control and the Daisyworld control system that could otherwise get lost in some of the details. I will show how rein control systems that feature a single rein controller will lead to a system variable being maintained at a higher or lower boundary, just as in type I diabetes. Rein control systems that feature two rein controllers will maintain a system variable closer to a 'set point', but this set point is an emergent property of the system. I will show that ensuring that the rein controllers exert inhibitory effects on each other produces zero steady state errors very much like those observed in the TBDW in chapter 4 that featured competition for the shared resource of space.

5.3 Daisyworld Effect Relationships

I start by distilling the fundamental elements of Daisyworld. While Daisyworld is a very simple model of a very complex system it still contains a number of 'ancillary' elements with regard to the central homeostatic mechanism. A schematic of the Daisyworld system is shown in figure 5.1. The effect the daisies have on planetary temperature and so their own growth rate is mediated via their effect on the planetary albedo which determines the overall energy budget with respect to received and emitted radiation. This can be simplified into a dynamical system in which a rein control element is affected by, and in turn, affects a system variable that is also affected by an external perturbing input.

In the simplified Daisyworld system, the environmental variable, V, and the population of organisms, O, are distinct and separate from the perturbing input of insolation, I, as the planet can be seen as a materially closed system. However, it is energetically open and to a certain extent the dashed circle is an arbitrary distinction. I am concerned with this system's ability or propensity to reduce the effect of the perturbing input, I, on V. This arises via the effects the organisms have on their environment and the effects the environment has on the organisms. There are four possibilities for the overall feedback of this system. These are shown schematically in figure 5.2. Considering the case where an increase in V either increases or decreases O and an increase in O either

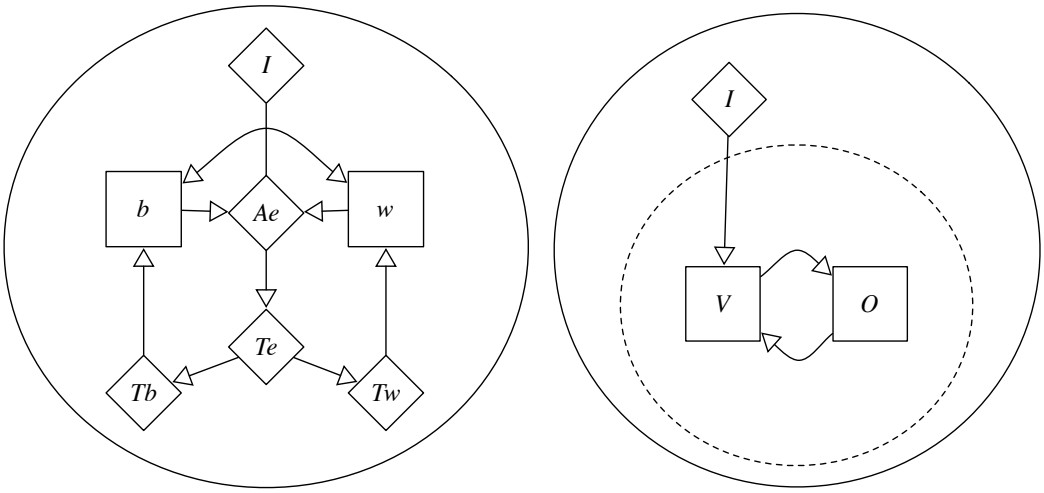

Figure 5.1: Daisyworld Effect Relationships

Plot (a) shows how insolation, I, affects planetary temperature, Te, via planetary albedo, Ae. Planetary temperature determines the temperature of the black and white daisies, Tb and Tw. This determines the coverage of the black and white daisies, b and w which in turn affects planetary albedo. The coverage of black and white daisies affect each other via their competition for available bare ground in which to seed new daisies. Plot (b) significantly simplifies this system. Now, an external perturbing input, I, affects a system variable, V, that in turns affects a population of organisms, O, that in turn affect the system variable.

increases or decreases V, the four possibilities are:

i. $+V \rightarrow +O$ and $+O \rightarrow +V$: positive feedback
ii. $+V \rightarrow +O$ and $+O \rightarrow -V$: negative feedback
iii. $+V \rightarrow -O$ and $+O \rightarrow +V$: negative feedback
iv. $+V \rightarrow -O$ and $+O \rightarrow -V$: positive feedback

It is interesting to highlight the clear parallels with this level of simple analysis and the notion of *feedback circuits* developed in a series of studies by Thomas:
[Thomas and D'Ari, 1990], [Kaufman and Thomas, 2003], [D'Ari and Thomas, 2003], [Thomas et al., 2004]. Thomas' analysis is conducted within the context of gene interactions, such as synergistic epistasis, and analysing the possible steady states for systems that consist of multiple feedback paths. This approach is in principle applicable to any system that can be characterised as an ensemble of individual elements with the overall system behaviour being largely determined by the nature of the interactions between these elements. Thomas calls such systems 'feedback circuits' and he and his collaborators develop mathematical and logical formalism that is used to predict the nature of steady states of different feedback circuits. Of particular relevance here is Thomas' observation that feedback circuits that feature an odd number of positive sign or negative sign interactions will be those circuits that exhibit negative feedback and are self-regulating. This is consistent with the analysis of the simplified Daisyworld effect relationships I have defined as negative feedback is produced by one, not two positive or negative interactions.

5.4 A Simple Rein Control Dynamical System

In this section I will develop a simple dynamical system based on the feedback circuits classified in the previous section. Four different systems will be shown. The first will feature a single rein controller which, as it increases in output, increases the system variable, which in turn determines the output. The second will feature a single rein controller which, as it increases in output, decreases the system variable, which in turn determines the output. The third will feature both increasing and decreasing effect rein controllers that interact benignly via their effects on the system variable. The fourth will feature both increasing and decreasing effect rein controllers that interact

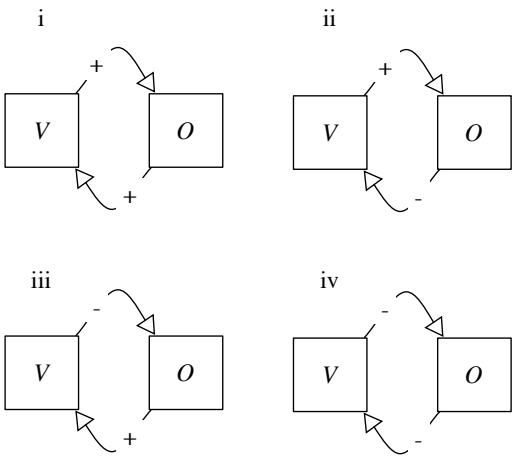

Figure 5.2: Feedback Circuits

Four feedback circuits are shown. Negative feedback is produced with an odd number of positive or negative feedback effects. Positive feedback is produced with an even number of positive or negative effects.

directly. I will define direct and benign interactions within the relevant sections. By this process I will build rein control systems that share central aspects of Daisyworld and demonstrate, in the simplest terms, how Daisyworld 'works'.

The simple Daisyworld system shown in figure 5.1(b) can be formulated as a dynamical system. I begin by defining the system variable, V, as a simple function of I and O:

$$V = I + O \tag{5.4}$$

where I is the external perturbing input and O is the rein controller effect function and represents the effects the rein controllers have on the system variable V. There could be an arbitrary number of arguments for the rein controller effect function: $O = f(\alpha_1, \alpha_2, ...\alpha_n)$. In the Daisyworld models that have been discussed so far in this thesis there are two arguments: black and white daisies, $O = f(\alpha_b, \alpha_w)$. In chapters 6 and 7 I will consider models that have an effectively infinite number of arguments for function O. As in O/S/TBDW, the rate of change of the ith rein controller is determined by the following differential equation:

$$\frac{d\alpha_i}{dt} = \alpha_i(\beta_i x - \gamma) \tag{5.5}$$

where β_i is a response function analogous to the Daisyworld birth rate function but of a simpler, monotonic form. x is a resource variable and γ is a decay parameter. γ will

be fixed at 0.1 for all the models presented in the following sections.

5.5 Single Rein Controller

I will consider two single rein control systems. Both systems feature an odd number of negative effects and are homeostatic. These can be thought of as Daisyworlds in which only white or only black daisies grow.

5.5.1 Type ii System

Using the classification from the previous section, a type ii system will feature a rein controller that decreases V and has a form of response to V that leads to the output of this rein controller increasing with increasing values of V. I call this rein controller B. A simple form for the output function is $O = f(B) = -\eta B$ where η is a scalar that modulates the strength of the rein controller effect on O. In the following I set η to unity and so $O = -B$. The rate of change of the rein controller is given with the following differential equation:

$$\frac{dB}{dt} = B(\beta_B x - \gamma) \tag{5.6}$$

where $x = 1 - B$. β_B increases with increasing V:

$$\beta_B = \begin{cases} 0, & V < V_1 \\ 1 - (V_1 - V)/\phi, & V_1 > V < V_2 \\ 1, & V > V_2 \end{cases} \tag{5.7}$$

where $\phi = V_2 - V_1$ and is the essential range for the rein controller - the range of variables that give non-zero output. If we set $V_1 = 0$ and $V_2 = 1$, then $\beta_B = V$ when $V \in [0, 1]$. This is shown in figure 5.3. Steady state values for B will be where the output of the controller is balanced by the decay rate:

$$B = 1 - \frac{\gamma}{V} \tag{5.8}$$

The output of B with varying I is shown in figure 5.4. With only the B rein controller in operation, the variable is $V = I - B$. Substituting in equation (5.8) for V in $V = I - B$

allows V to be expressed as a quadratic function in terms of the perturbing input I:

$$V^2 - V(I-1) + \gamma = 0 \tag{5.9}$$

Steady state solutions for V over a range of I are shown in figure 5.5.

5.5.2 Type iii System

Using the classification of the previous section, a type iii system will feature a rein controller that increases V and has a form of response function that leads to the output of this rein controller decreasing with increasing values of V. I call this rein controller A. A simple form for the output function is $O = f(A) = \eta A$ where η is a scalar that modulates the strength of the rein controller effect on O. We set η to unity and so $O = A$. As before, the rate of change of the rein controller is given with the following differential equation:

$$\frac{dA}{dt} = A(\beta_A x - \gamma) \tag{5.10}$$

where $x = 1 - A$ represents a fixed resource, and β_A being defined as:

$$\beta_A = \begin{cases} 1, & V < V_1 \\ 1 - (V - V_1)/\phi, & V_1 > V < V_2 \\ 0, & V > V_2 \end{cases} \tag{5.11}$$

As V increases, A decreases. If we set $\eta = 1$, $V_1 = 0$ and $V_2 = 1$, then $\beta_A = 1 - V$ when $V \in [0, 1]$. With γ fixed at 0.1, as V decreases towards zero, A increases with the value of A saturating when $V = 0$. As V increases towards one, A decreases with the value of A being zero when $V = 0.9$. This is shown in figure 5.3. Steady state values for A will be where the output of the controller is balanced by the decay rate:

$$A = 1 - \frac{\gamma}{1 - V}. \tag{5.12}$$

The output of A with varying I is shown in figure 5.4. We can solve V for I and so show how the system variable changes over varying perturbing input. Substituting equation (5.12) for V in $V = I + A$ allows V to be expressed as a quadratic function of I:

$$V^2 - V(2 + I) + I + \gamma = 0. \tag{5.13}$$

Steady state solutions for V over a range of I are shown in figure 5.5.

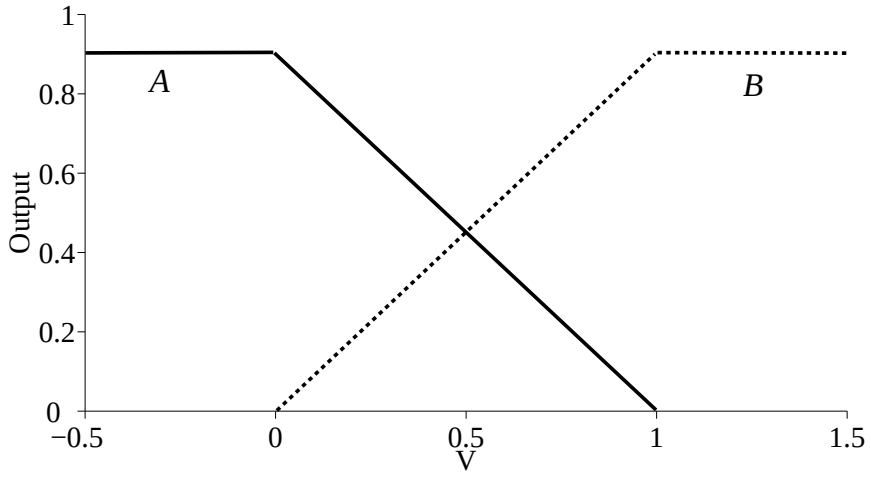

Figure 5.3: Rein Controller Output - Varying V

The output of two rein controllers for different values of the system variable, V are shown. Rein controller A output is plotted with a solid black line, rein controller B is plotted with a dashed black line. A decreases in output as V increases. B increases in output as V increases. A and B saturate when $V = 0$ and $V = 1$ respectively.

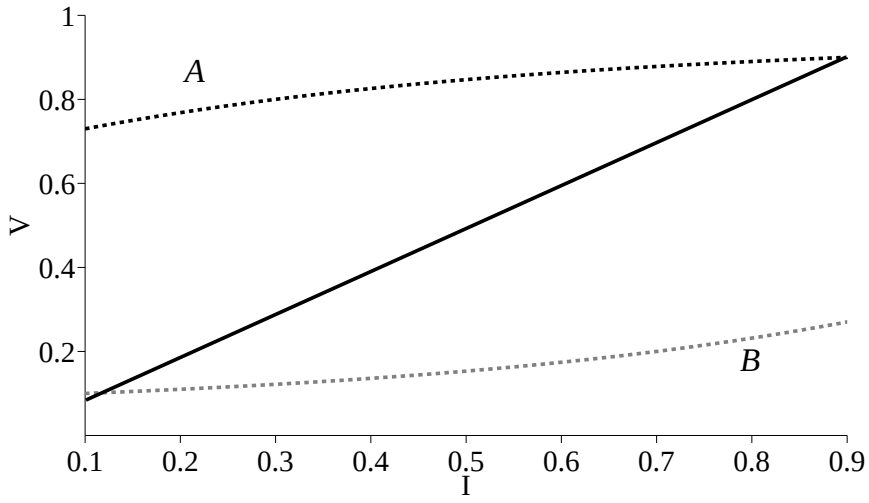

Figure 5.4: Steady State Values for V for Type ii and Type iii Systems
Steady state solutions for values of the system variable, V, with either A only (dashed black line) or B only (dashed grey line) for a range of I. The solid black line shows values for V in the absence of A or B. With a single rein controller operating, V is maintained at upper or lower thresholds.

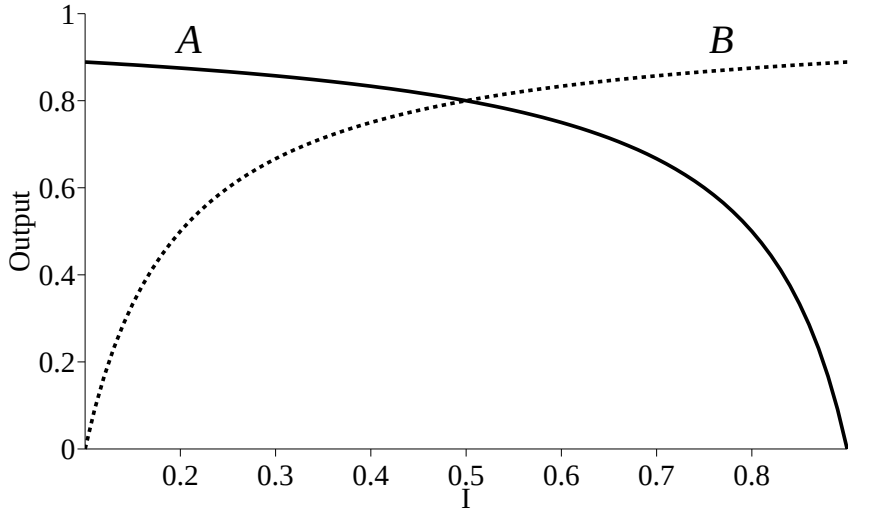

Figure 5.5: Rein Controller Output - Varying I
The output of two rein controllers for different values of the perturbing input, I are shown. Rein controller A output is plotted with a solid black line, rein controller B is plotted with a dashed black line.

The effect of A-only and B-only systems was a reduction in the rate of change of the system variable V as it was perturbed by the external input I. A-only steady states feature V being held at an upper threshold with B-only steady states featuring V being held at a lower threshold. If I were to decrease below 0.1 or increase beyond 0.9 then the A and B rein controllers would saturate and the rate of change of V would equal that of I. This assumes that η remains fixed at unity. If η increases then the range of I over which V is maintained to within 0.1 or 0.9 increases. As η modulates the strength of feedback between the rein controllers and the system variable it can be understood as being analogous to daisy albedo in Daisyworld.

5.6 Two Rein Control Systems

This section can be understood as a more abstract analysis of the role of resource competition in Daisyworld. In chapter 4 I showed the competition for a shared resource led to planetary temperature remaining fixed at the optimal birth rate temperature. By producing even simpler models for these types of interactions I am able to provide exact analytical solutions that clearly identify the mechanisms at work and increase their applicability to other domains.

I will detail rein control systems that have both rein controllers operating at the same time. Two types of interaction between the two reins can be identified. I call these 'benign' and 'antagonistic'. Benign interaction arises from the effects each rein has on the system variable that in turn affects each rein via the response function. As A increases V, it will affect B and vice versa. Antagonistic interaction corresponds to the value of A being a function of B and vice versa. This will arise via the rate of change of A and B being modulated by a shared resource variable. Systems that feature this sort of antagonistic interaction will lead to the system variable, V, remaining fixed while I varies.

5.6.1 Benign Interactions

I will first consider a system that features A and B interacting benignly. $O = f(A, B) = A - B$. Equation (5.4) now becomes:

$$V = I + A - B \tag{5.14}$$

The rate of change of both rein controllers is governed by differential equations that each feature individual resource variables, x_A and x_B:

$$\frac{dA}{dt} = A(\beta_A x_A - \gamma), \tag{5.15}$$

$$\frac{dB}{dt} = B(\beta_B x_B - \gamma) \tag{5.16}$$

where $x_A = 1 - A$ and $x_B = 1 - B$. We can solve V for I and so observe how the system variable changes as the perturbing input changes by substituting steady state values for A and B into equation 5.14. This gives:

$$V = I - \frac{\gamma}{1 - V} + \frac{\gamma}{V} \tag{5.17}$$

which is a cubic equation:

$$V^3 - V^2(1 + I) + V(I - 2\gamma) + \gamma = 0 \tag{5.18}$$

Steady state solutions for V over a range of I are shown in figure 5.6.

5.6.2 Antagonistic Interactions

I will now change the previous benign interaction system to one in which the rein controllers A and B interact antagonistically via a shared resource variable. As before, the response of both reins is governed by differential equations of the form:

$$\frac{dA}{dt} = A(\beta_A x - \gamma) \tag{5.19}$$

$$\frac{dB}{dt} = B(\beta_B x - \gamma) \tag{5.20}$$

but now $x = 1 - A - B$. Steady states will be found when the output of each rein controller balances the decay rate:

$$x\beta_A = \gamma, x\beta_B = \gamma \tag{5.21}$$

and so with a shared resource:

$$\beta_A = \beta_B = \beta \tag{5.22}$$

This identity between the two rein controller output functions is similar to the identity between daisy birth rates and, just as in Daisyworld, it results in the variable that is subject to regulation remaining fixed while both rein controllers are operating. Finding the value of β can be understood geometrically as the point where the A and B output functions intersect. V will be held at this value with the output of A and B varying with varying I. This can be understood in terms of classical closed loop control. A and B represent 'effector inputs' that seek to reduce an 'error signal', ϵ, defined as:

$$\epsilon = I - T \tag{5.23}$$

where T is the set-point 'target' which is where the two output functions intersect. The output of the rein controllers will be that which leads to the error signal returning to zero. Using the simple formulations for β_A and β_B, the steady state outputs of A and B are:

$$A = 0.5(0.8 + \epsilon) \tag{5.24}$$

$$B = 0.5(0.8 - \epsilon) \tag{5.25}$$

Values for A and B for a range of I are plotted in figure 5.7. With antagonistic interactions, the rate of change of the rein controllers is linear with $\dot{A} = -\dot{B}$.

An interesting consequence of antagonistic rein control is that the particular form of the output functions for the rein controllers is irrelevant in determining the rate of change of the system variable. Equations (5.24) and (5.25) contain no information about how the output of A and B varies with varying V. All that is required is where the two output functions intersect. These functions may be non-linear. If they are unimodal there will be a single intersection point and V will be fixed at that point. Higher order polynomial output functions may lead to multiple solutions to the pair of simultaneous equations. Consequently there could be n fixed point values for V where n is the degree of the polynomial output functions. There are similarities here to Thomas' observation that feedback circuits with non-linear output functions such as $\sin V$ will have a potentially infinite number of steady states within the feedback circuit schema and are classified as 'ambiguous' circuits.

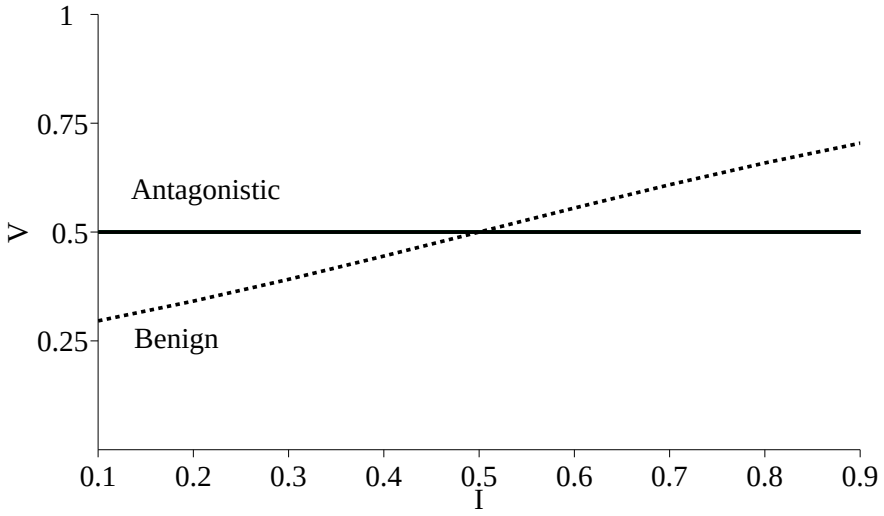

Figure 5.6: System Variable Values With Benign and Antagonistic Interaction
Steady state solutions are shown for the system variable, V, when both controllers are operating and interacting benignly (dashed black line) and antagonistically (solid black line). Antagonistic interactions lead to V being fixed at 0.5, which is where the two rein controller output functions intersect.

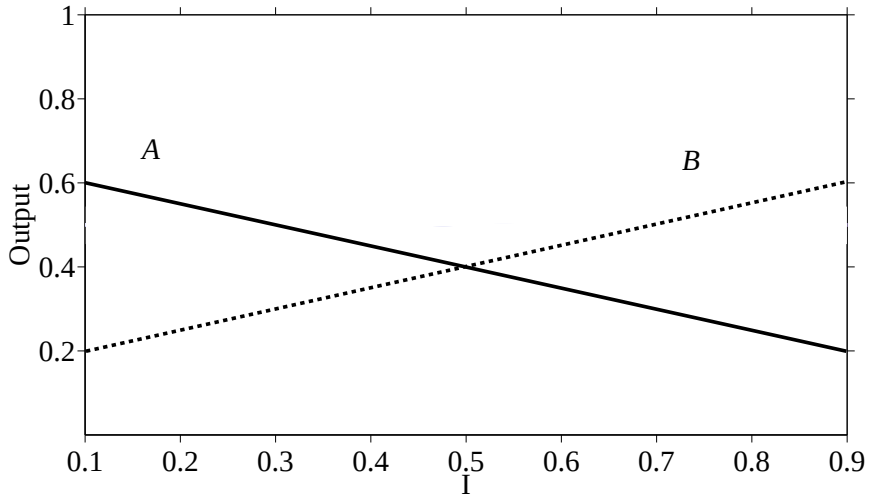

Figure 5.7: Output of A and B With Antagonistic Interaction
Steady state solutions are shown for the output of A (solid black line) and B (dashed black line) rein controllers when both controllers are operating and interacting antagonistically via a shared resource.

5.7 Daisyworld and Rein Control

While the simple rein control models developed in this chapter share a number of important properties of O/S/TBDW there is an important difference that would appear to limit the rein control analysis. All the rein control models featured controllers that had differing responses to a shared system variable. As V increased, A decreased and B increased. In Daisyworld, both black and white daisies share the same birth rate function. However as was shown in chapter 3 with SDW, black daisy and white daisy steady states are always on the opposite sides of this shared birth rate function. If one decomposes the Daisyworld birth rate function into two sloping, opposing functions, two opposing functions very similar to the output functions for A and B arise. This highlights a critical aspect of the Daisyworld control system. Homeostasis requires a *differential* response of the rein controllers. If the output function of A and B were identical, homeostasis would not be observed as any change in A would be matched with a change in B. This is equivalent to the black and white daisies sharing the same albedo *or* the temperature of the black and white daisies being equal.

In O/SDW this difference in temperature is achieved via the daisies being insulated from the ambient, planetary temperature. I showed how this was responsible for the inverse response to external perturbations - why planetary temperature decreased when luminosity increased. In TBDW the temperature difference between black and white daisies was the result of the the black and white daisies inhabiting separate boxes that were thermally connected. With non-infinite rates of heat flux between the two beds, a temperature gradient was established. In higher dimensional Daisyworlds such as cellular automatons, daisies experience local patch temperatures that are determined by the albedo of the daisies and there is a non-infinite flux of heat across neighbouring tiles thus ensuring temperature differences are maintained. Therefore, the black and white daisies do not in fact share a single system variable, but rather respond to local variants of it with the black daisies always being warmer than the white daisies.

5.8 Limits of Regulation

Chapter 4 showed how the width of the birth rate function was largely irrelevant in determining the range of luminosity over which daisies grow. In Daisyworld it was the maximum value of the birth rate and the albedo that determined the any-daisy range. In the simple rein control model it was the maximum value of the output function and

the value set for η that determined the range of perturbations that the system variable was regulated over. In all the results presented, η was fixed at unity. Increasing η would increase the range of I over which A and B produced outputs within the range $[0.1, 0.9]$. The η parameter modulates the strength of the feedback between the rein controllers and the system variable, similar to albedo in Daisyworld modulating the strength of the temperature effect.

5.9 'Lucky Gaia'

There has been a noticeable omission from the models discussed in this chapter. Of the four types of feedback circuit I identified in section 5.3 I have examined only the two that resulted in negative feedback and homeostasis. The positive feedback and so non-homeostatic type i and type iv systems have not been discussed. A common criticism of the original Daisyworld and its variants is that a disproportionate amount of importance is given to negative feedback systems with the possibility of positive feedback systems being ignored. The power of simple, conceptual models is that they can be used to explore a very wide range of phenomena. The danger is that this flexibility leads to exclusion of those assumptions that produce phenomena that do not support preconceptions. Simply put, the charge put to Daisyworld is that it is the result of selective assumptions that ensure homeostasis emerges. Just as in this chapter I did not examine positive feedback systems, Daisyworld models that feature black daisies which decrease planetary temperature or white daisies which increase planetary temperature do not really exist[1]. Therefore, if we are interested in considering to what extent negative feedback systems will emerge via natural processes, there would initially appear to be a 0.5 probability that a rein control system will be homeostatic. If we are unable to demonstrate that negative feedback systems are more likely to emerge than positive feedback systems, then the proposed negative systems operating within the biosphere are potentially improbable. This echoes Kirchner's 'Lucky Gaia' claim [Kirchner, 2002]. Kirchner goes further and argues that if negative and positive feedback systems are equally likely then we should typically never observe homeostatic states because positive feedback would tend to eject the system away from any stable,

[1]That is not entirely correct as there have been previous studies that allow adaptation of the birth rate function of the daisies and so, in principle, allow the emergence of cooling/black and warming/white daisies. However, as I will explain in the following chapter, these models either no longer regulate temperature or are in fact biased against such positive feedback cicuits with negative feedback rein controllers being selected for.

homeostatic state. In the following chapter I will show how homeostasis can emerge when negative as well as positive feedback are possible. Therefore it is not necessary to invoke a version of Lucky Gaia in order to explain biologically-mediated homeostasis within the biosphere.

5.10 Summary

In this chapter I have classified a range of rein control systems on the basis of the feedback between a rein controller and a system variable. Four possible configurations were identified with two featuring negative feedback and homeostasis, and two featuring positive feedback and non-homeostatic behaviour. With a single negative feedback rein controller operating, the rate of change of the system variable was less than the external perturbing input that operated on the system variable. With two rein controllers operating I highlighted the difference between benign and antagonistic interactions. Benign interactions result from the two rein controllers only interacting via their effects on the system variable. Antagonistic interactions result from the two rein controllers interacting via their effects on the system variable and via a shared resource variable that modulated the rate of change of the output of both rein controllers. Antagonistic rein control systems were characterised in terms of zero steady state error with respect to the system variable and a target value that was the result of the intersection of the two rein control output functions.

I showed how a number of apparently important differences between the rein control models and Daisyworld were, under closer examination, different formulations of the same underlying mechanisms. In particular it was shown that in order for homeostasis to emerge the rein control elements must respond differently to a shared system variable. In Daisyworld this differential response is the result of the difference between black and white daisy temperatures being maintained via insulating forces. In the rein control model the differential response was explicitly modelled with different output functions.

In the following chapter I will formulate a rein control model that will have equal potential to form negative and positive feedback circuits. This will directly address Kirchner's arguments. If Kirchner is right, such a system will not be homeostatic. I will show why Kirchner is wrong and how negative feedback homeostatic systems can emerge and persist under such assumptions.

Chapter 6

Evolving Rein Control

6.1 Previous Publications

The model developed in this chapter is a simplification of the model first proposed in [McDonald-Gibson, 2006] and further developed, with my contributions, in [McDonald-Gibson et al., 2008]. Parts of section 6.5 have been taken from my contributions to [Dyke et al., 2007].

6.2 Overview

The previous chapter concluded by highlighting a significant omission from the discussion on rein control systems. I had only considered negative feedback and so homeostatic circuits. Given there was an equal number of positive feedback circuits identified, then ignoring those that do not produce homeostasis may appear to support an argument formulated by Kirchner [Kirchner, 2002]. He argued that if there is equal probability for negative and positive feedback to arise in natural systems, then observing negative feedback operating in the Earth's biosphere must in certain respects be 'lucky'. An element of luck must be involved as positive feedback systems will tend to eject any system away from a homeostatic steady state.

In this chapter I will show that it is possible for homeostatic states to emerge in systems that feature the potential for both positive and negative feedback and that negative feedback is 'selected' for, leading to homeostasis being observed. The selection mechanism is natural selection that operates on a population of diverse individuals. Rein control homeostasis arises not in spite of, but *because* of natural selection. I will

present the results from stability analysis that shows the homeostasis is robust to a range of parameter values. This will involve simulations that feature different mutation rates. However, such mutation rate changes will be limited as the following chapter will explore the effects that different mutation rates have on the homeostatic properties of the model.

6.3 Reduced Assumption Daisyworld

The model I present is a version of a Daisyworld model proposed by [McDonald-Gibson, 2006]. The primary motivation behind McDonald-Gibson Daisyworld (hereafter referred to as MGDW) was to produce a Daisyworld model that did not feature a number of what were previously considered as necessary assumptions. In particular, McDonald-Gibson sought to address the following three assumptions.

First, black daisies are always warmer and the white daisies are always cooler than their environments. As I argued in the previous chapter each individual daisy in O/S/TBDW is affected by a local temperature that is dependent on the albedo of the daisy. This allows a differential response to planetary temperature and that differential response is necessary for rein control systems.

Second, black daisies only ever increase planetary temperature and white daisies only ever decrease planetary temperature.

Third, the optimal environmental conditions are predefined and inflexible. In other words the growth response of the daisies is fixed. This assumption was addressed by [Lenton and Lovelock, 2000] which featured daisies with a range of different growth responses to temperature. Natural selection was able to operate over such variation so that those daisies that had optimal growth temperatures closest to the prevailing conditions would have faster growth rates. This would allow the emergence of black daisies that had faster growth when it was hot and white daisies that had faster growth when it was cold. Both would produce positive feedback, non-homeostatic states. However, [Lenton and Lovelock, 2000] contained the additional assumption that as daisies adapted to higher or lower temperatures, the maximum possible fitness decreased as it was argued that there would be a particular temperature that facilitated the most efficient metabolism.

A more general examination of the mutation of growth response to temperature was produced by [Wood et al., 2006] where it was found that allowing mutations in albedo

and the growth response produced oscillations in the population and environment. The population and environment can be viewed as two objects orbiting each other as selection first operates on the albedo of daisies (the model was a spatial Daisyworld in which each individual daisy had a local temperature that was different from the ambient temperature dependent on the albedo of the daisy) and then the growth response of the daisies after which selection operates on albedo and so on.

These three assumptions can be understood as ensuring that only type ii and type iii negative feedback can arise. In chapter 5 I classified four rein control circuits in terms of their overall negative or positive feedback. I reproduce it below for convenience. Black daisies that had higher growth at higher temperatures and white daisies that had higher growth at lower temperatures would produce type i and type iv systems respectively.

i. $+V \rightarrow +O$ and $+O \rightarrow +V$: positive feedback
ii. $+V \rightarrow +O$ and $+O \rightarrow -V$: negative feedback
iii. $+V \rightarrow -O$ and $+O \rightarrow +V$: negative feedback
iv. $+V \rightarrow -O$ and $+O \rightarrow -V$: positive feedback

MGDW did not feature any of these three assumptions. First, the model featured a single, global temperature variable. All daisies experienced the same temperature. Second, the model featured daisies that had a warming effect and birth rate functions that would produce higher growth at higher temperatures, and daisies that had a cooling effect and birth rate functions that would produce higher growth at lower temperatures. Third, adaptation of the response to environmental conditions was allowed so that the optimal temperature that would produce maximum growth could take any value over a range of possible values.

As in Daisyworld, the environment was perturbed by an increasing or decreasing force. Natural selection operated on the population as those organisms who were closely matched to the current conditions would have higher fitness. Given such assumptions it was very surprising to observe Daisyworld-type homeostasis emerging and persisting. It was not immediately apparent *how* such homeostasis arose. My main contribution to [McDonald-Gibson et al., 2008] was to explain the emergence and persistence of the observed homeostatic states and to cast this explanation in terms of rein control. This involved the development of a series of simplifications of the original MGDW and it is

one of these models that will be presented and analysed in this chapter.

6.4 Model

MGDW is in principle applicable to a wide range of real-world scenarios. For didactic purposes I will develop and discuss the model in the context of simple bacterial organisms that inhabit a simple, homogeneous environment. The ecosystem consists of a fixed population of two-locus haploid individuals. Each individual has a simple genome that consists of a single copy of two loci or 'slots' for particular values for genes. These individuals interact with their environment by increasing or decreasing an environmental variable that in turn determines the fitness of the organisms. We can think of the environmental variable as the concentration of a particular chemical.

A proportion of the population reproduce by cloning offspring that are subject to mutation with a fixed probability. The same proportion of individuals die thus maintaining a constant population. Reproduction and death occur simultaneously at every dimensionless unit of time. This may represent an annual, daily or even hourly cycle if we are considering bacterial populations. The ecosystem is perturbed over time by an external force. We can think of the perturbing force as an imposed chemical input that alters the environment. The volume of the ecosystem does not change, so we also assume that an equal amount of material is removed. Consequently we will be concerned with the concentration of the fitness-changing environmental variable. This is similar to real world artificial ecosystems such as 'chemostats' [Novick and Sziliard, 1950]. Figure 6.1 shows an idealised chemostat.

The environment is represented by a single variable, $V \in [-50, 150]$. The fitness, F, of individuals is a function of V. Each individual has a genotype that specifies its effect on V and the value of V that produces maximum fitness. The A locus specifies the individual's effect on V. This is analogous to albedo in Daisyworld. The B locus specifies the value of V to which the individual is best adapted. This is similar to the optimal birth rate temperature in Daisyworld.

The A locus has two possible alleles: a decreasing effect e which reduces V and an increasing effect E which increases V. The discrete A locus alleles are modelled with double floating point numbers in the range [-1,1]. An individual will have the e allele if $-1 < A < 0$ and the E allele if $0 < A < 1$. The B locus is represented with a double floating point number in the range [15,85] and specifies the point that the individual

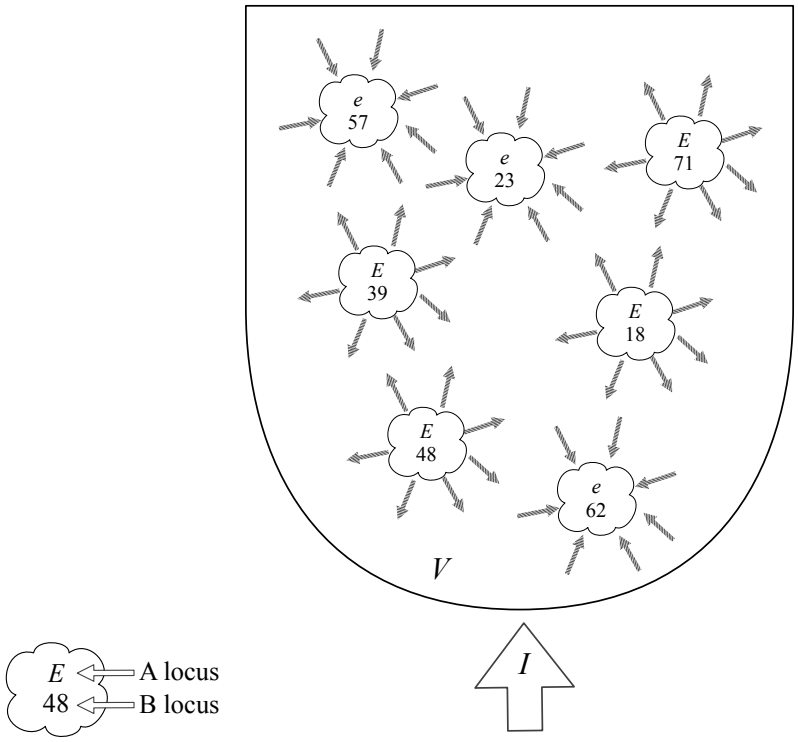

Figure 6.1: MGDW: An Idealised Chemostat

MGDW consists of an idealised ecosystem that features asexually reproducing, haploid organisms or 'bugs'. A single bug is shown on the left. It has a simple genome that consists of two loci. The A locus determines the effect the bug will have on the environmental variable, V. The B locus determines how V will affect the bug. A population of bugs live within a simple chemostat or bioreactor. Arrows away from particular bugs denote an increasing effect on V. Arrows pointing towards particular bugs denote a decreasing effect on V. An external perturbing input I affects V.

is best adapted to within the range of V values. Equation (6.1) is used to calculate an individual's fitness which is a parabolic function of the variable. This function peaks at the B value and declines sharply on either side towards zero. The fitness of the ith individual is a function of V and λ which determines the width of the essential range:

$$F_i = \begin{cases} 1 - \lambda(B_i - V)^2, & |B_i - V| < \lambda^{-\frac{1}{2}} \\ 0 & \text{otherwise} \end{cases} \tag{6.1}$$

Natural selection is modelled using a genetic algorithm with a proportion of the population subject to death, selection and reproduction at each unit of time. A constant death rate, γ, is applied. Therefore, on average, the lifetime of any individual is $1/\gamma$. For each death, a tournament is held between two randomly selected members of the population and the victor is the individual with higher fitness. The winner of a tournament replaces the loser in the population. If both individuals fitness are the same, no replacement occurs. Mutation occurs with a probability of μ at each locus. This is performed by adding a number to the allele drawn from a Gaussian distribution of mean 0 and standard deviation 0.05 to the allele. External perturbations are modelled with a dimensionless scalar input variable, I. The effect the individual organisms have on the environment, O, is the sum of the decreasing and increasing effects:

$$O = \sum e + \sum E \tag{6.2}$$

As in chapter 5, the environmental variable is a simple function of the organisms' effect and perturbing input:

$$V = O + I \tag{6.3}$$

The rate of change of the environmental variable can be expressed in terms of organisms' effects and perturbing input:

$$\frac{dV}{dt} = O + (I - V) \tag{6.4}$$

6.4.1 Algorithm

The operation of the algorithm is given below.

1. Initialise external perturbing input I.
2. Create random population.
3. Calculate phenotype values from genotype values.

4. Calculate population effects on environmental variable, V.
5. Calculate each individual's fitness, F.
6. Select n individuals for tournament reproduction.
7. Replace n/2 loosers with n/2 winners within population
8. Increment I.
9. Goto 3.

6.5 Results

In the following results, unless otherwise stated, the maximum population size, K, is 2×10^3 with the population being initialised with 2×10^3 randomly generated individuals. λ is fixed at 0.04 for all individuals which gives an essential range of 10. Mutation rates, μ, for alleles at both loci is fixed at 0.1. The death rate, γ, is fixed at 0.01, and consequently 40 individuals are selected for tournament reproduction at every unit of time with the winners replacing the losers within the population. I is initialised at 0 and increases to 150 over 10^5 discrete time steps. Two simulation results will be presented. The first will feature individuals that have no effect on their environment: $e = -E = 0$. The second simulation will feature individuals that have small effects on their environment: $e = -E = 0.05$.

6.5.1 Simulation 1: Without Organism Effects

Figures 6.2 and 6.3 shows results where the individuals respond to the environmental variable and do not affect it. As the fitness of each organism is a simple function of the environmental variable there is a clear selection pressure operating on the B allele. If the perturbing input I is set to a value within the range [15,85] then a selective sweep will lead to the B allele closest to I increasing in frequency in the population until the frequency of this B allele approaches 100% (dependent on the mutation rate). As I changes it would be expected that the population would respond to this change in selection pressure with varying B allele frequencies. If I were to increase slightly, individuals with slightly higher value B alleles would have higher fitness and so these slightly higher B allele values would increase in fitness, and so on as I varies over the

range [15,85]. Such a selective sweep is observed in that the mean B allele value for the population, $\bar{B}$, moves rapidly from approximately 50 to closely match V, which in the absence of any organism effects equals I. Consequently the V, I and $\bar{B}$ data lines are superimposed in figure 6.2(a) while V is within the essential range.

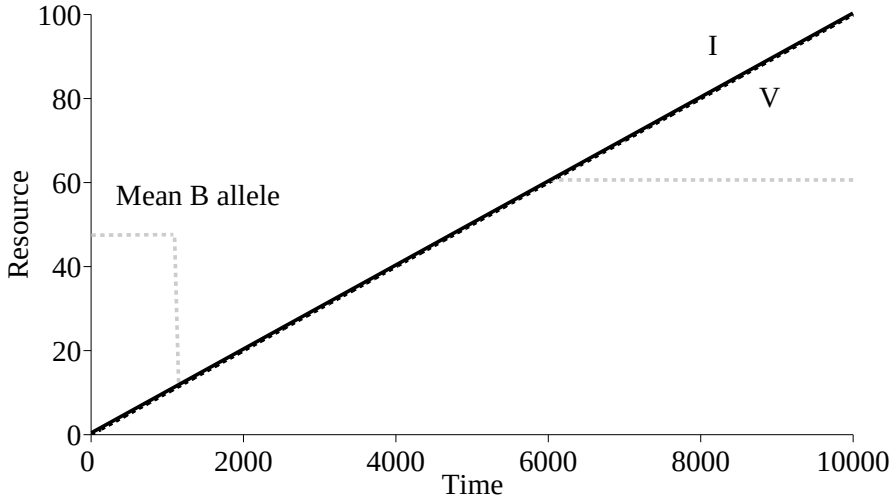

Figure 6.2: Without Organism Effects

The perturbing input, I, is plotted with a dashed black line, the mean B allele is plotted with a dashed grey line and the environmental variable, V is plotted with a solid black line. V initially increases with I until I drives V into the range [15,85] at which point the population rapidly converges to a single B trait so that $\bar{B} \approx 15$. As I continues to increase, $V = \bar{B} = I$ until $I \approx 95$ at which point $\bar{B}$ no longer increases.

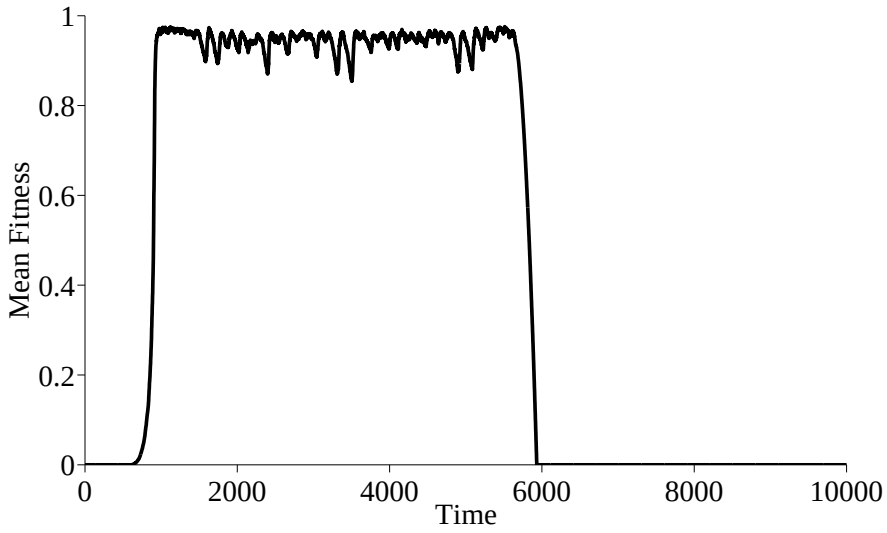

Figure 6.3: Without Organism Effects - Mean Fitness

The mean fitness of the population is held near the maximum value while V is within the essential range.

6.5.2 Simulation 2: With Organism Effects

The second simulation features the same assumptions as the first, but with the difference that now each individual has a small effect on V so that $e = -E = 0.05$. Recall that the individual effects on V are global and that if individual x increases V by 0.05 then all other individuals experience the same value of V that x experiences. Consequently there is no selection pressure for the A allele. Therefore, although the population has the capacity to produce a total force of ± 100 there appears to be no mechanism that will lead to the emergence of population effects that counteract changes in I and it would be expected that O will randomly wander around a mean of zero (as with a randomly initialised population, the number of e alleles is approximately the same as the number of E alleles).

Figures 6.4 - 6.6 shows that as in simulation 1, V initially increases with I. However when V enters the range [15,85] it no longer increases with increasing I, but stabilises at ≈ 15. As I increases, V remains at the same value whereas the sum of population effects decreases. The supposedly non-adaptive A allele appears to be responding to changes in selection pressure while the adaptive B allele is 'ignoring' the change in the environmental variable and remaining fixed. This appears to be fundamentally

non-Darwinian.

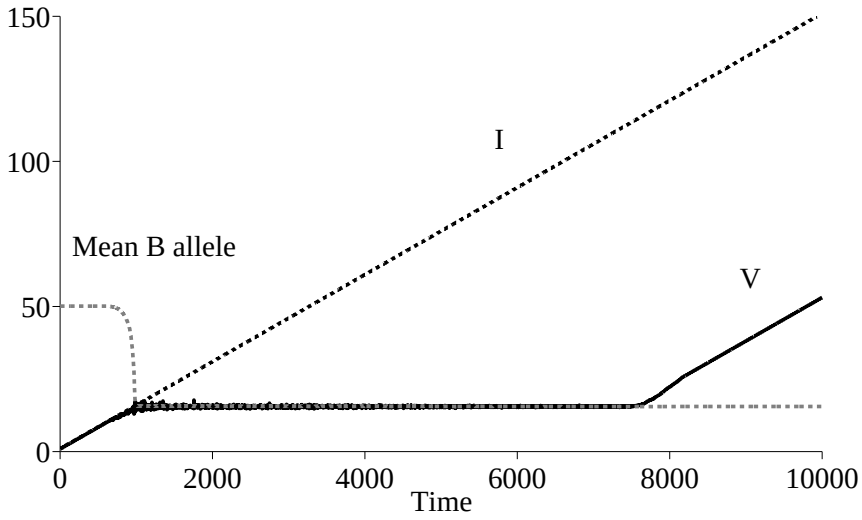

Figure 6.4: With Organisms Effects

The perturbing input, I, is plotted with a dashed black line, the mean B allele is plotted with a dashed grey line and the environmental variable, V is plotted with a solid black line. V initially increases with I until I drives V into the range [15,85] at which point the population rapidly converges to a single B trait so that $\bar{B} \approx 15$. As I continues to increase, V no longer increases but remains fixed until $I \approx 110$. As I increases, mean fitness is held near the maximum value for the period of homeostasis.

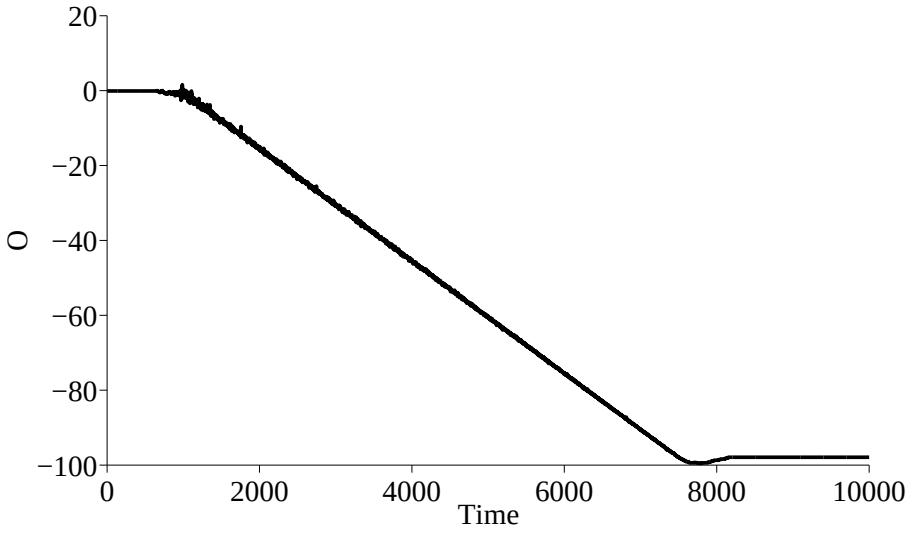

Figure 6.5: With Organism Effects - Organism Effects

Organism effects, O, is shown for the same simulation as shown in figure 6.4.

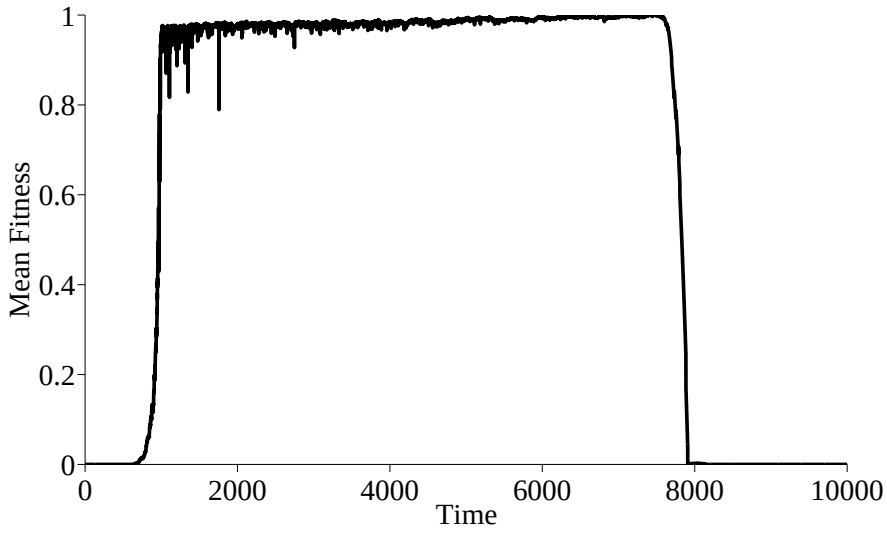

Figure 6.6: With Orgamism Effects - Mean Fitness

The mean fitness of the population is shown for the same simulation as shown in figure 6.4. As I increases, mean fitness is held near the maximum value for the period of homeostasis.

6.6 Stability Analysis

As MGDW is a stochastic model, it is necessary to run ensembles of computer simulation experiments in order to ensure that the observed homeostasis is neither improbable, nor the result of a particular set of parameter values. Table 6.1 increases our confidence that the observed homeostasis is robust to a range of parameter values. The effects of different mutation rates will be considered further in the following chapter. The averaged results of 24 parameter sets are shown. For each parameter set, 100 simulations were performed with the average results being computed. In order to scale the individual organism effects so that the same range of I could be regulated against, the values of e and E varied as the population size, k varied. $e = -E = 100/k$. Therefore as population size increases, the strength of each individual's effect on the environmental variable decreases. Each parameter set was evaluated in terms of 'average homeostasis'. The method of evaluating each simulation is analogous to regression of the environmental variable, V, time series data. For the purposes of this chapter, I define homeostasis as being the situation where V remains around a particular value. A simulation will exhibit homeostasis when $dV/dt < x$ where x is approximately a magnitude less than dI/dt. As V will change in response to the stochastic tournament selection method, there will be a certain amount of 'noise'. This noise can be greater than dI/dt. Consequently it is necessary to average V over a number of discrete time steps in order to smooth out the data. The procedure used to compute the percentage homeostasis is given below.

1. Store current value of V as V_{start}.
2. After n iterations of main model algorithm store current value of V as V_{end}.
3. $dV/dI = (V_{end} - V_{start})/(I_{end} - I_{start})$.
4. If $dV/dI \times x < dI$ then model has been homeostatic for n iterations of main algorithm.
5. Goto 1

where $x = 10$ and $dI = I_{end} - I_{start}$. All simulations for all parameter sets were initialised with $I = 15$. I then increased to 100. Maximal homeostasis would be with V being regulated around a fixed point for this entire period. Parameter set scores are expressed as percentages of this maximum.

It is necessary to explain at this point why this method of evaluating parameter sets was chosen when alternative methods are available. For example, it would be possible to evaluate the model in terms of the range of I values that the system is able to maintain V to within the habitable range. The particular values of V would not be important, only that it is within the range $[15, 85]$. In the following sections I will show why regulation of V to within the habitable range by the population is equivalent to the regulation of V at a fixed point. What is of particular interest is that the *particular* value that V is regulated around is an emergent property of the model. There are no explicit or implicit fixed points woven into the model assumptions, unlike O/S/TBDW.

6.6.1 Rates of Change

A number of results from the sensitivity analysis are of immediate interest. It can be seen that the number of discrete time steps can have a significant impact on homeostasis. In particular when the number of time steps, τ, equals 10^4, homeostasis can be significantly reduced. The ability of the system to exhibit homeostasis is determined not only by the rate of change of the perturbing input, but the rate of change of the alleles in the population and the strength of the individual organism effects. With low τ the rate of change of I is high. Consequently the population needs to respond quickly and the strength of that response needs to be high. Large populations subject to rapid changes in the perturbation will be less homeostatic than smaller populations. This is simply due to more time being required for a large population to respond to external perturbations. If the population size is increased by 10, in order to respond to a change in I of 10, 10 times more individuals need to be changed in the population. This results in the population requiring approximately 10 times longer for the same change in the frequency of alleles to be observed. This can be observed in the significant increase in homeostasis in large populations with an increase in the number of discrete time steps.

6.6.2 Essential Range Values

All the results presented and discussed in this chapter feature fixed values for the essential range of the individuals, λ. It is important to note that varying the essential range by decreasing the range of V values that produce non-zero fitness do not make substantive changes to the behaviour of the model. Parameter λ can be reduced by a magnitude and regulation will still be observed over the same range of external perturbations. This supports the findings of chapter 4 where it was found that the essential

Mutation Rate	Time Steps	Population Size	Avg. Homeostasis
$\mu = 0$	$\tau = 10^4$	k = 10^3	91%
$\mu = 0$	$\tau = 10^4$	k = 2×10^3	55%
$\mu = 0$	$\tau = 10^4$	k = 10^4	0%
$\mu = 0$	$\tau = 10^5$	k = 10^3	97%
$\mu = 0$	$\tau = 10^5$	k = 2×10^3	96%
$\mu = 0$	$\tau = 10^5$	k = 10^4	90%
$\mu = 0$	$\tau = 10^6$	k = 10^3	99%
$\mu = 0$	$\tau = 10^6$	k = 2×10^3	99%
$\mu = 0$	$\tau = 10^6$	k = 10^4	99%
$\mu = 0.1$	$\tau = 10^4$	k = 10^3	75%
$\mu = 0.1$	$\tau = 10^4$	k = 2×10^3	47%
$\mu = 0.1$	$\tau = 10^4$	k = 10^4	0%
$\mu = 0.1$	$\tau = 10^5$	k = 10^3	85%
$\mu = 0.1$	$\tau = 10^5$	k = 2×10^3	86%
$\mu = 0.1$	$\tau = 10^5$	k = 10^4	84%
$\mu = 0.1$	$\tau = 10^6$	k = 10^3	58%
$\mu = 0.1$	$\tau = 10^6$	k = 2×10^3	70%
$\mu = 0.1$	$\tau = 10^6$	k = 10^4	89%
$\mu = 0.2$	$\tau = 10^4$	k = 10^3	76%
$\mu = 0.2$	$\tau = 10^4$	k = 2×10^3	46%
$\mu = 0.2$	$\tau = 10^4$	k = 10^4	0%
$\mu = 0.2$	$\tau = 10^5$	k = 10^3	88%
$\mu = 0.2$	$\tau = 10^5$	k = 2×10^3	90%
$\mu = 0.2$	$\tau = 10^5$	k = 10^4	81%
$\mu = 0.2$	$\tau = 10^6$	k = 10^3	66%
$\mu = 0.2$	$\tau = 10^6$	k = 2×10^3	71%
$\mu = 0.2$	$\tau = 10^6$	k = 10^4	85%

Table 6.1: MGDW Sensitivity Analysis Results

24 parameter sets were evaluated in terms of average homeostasis. The number of time steps describes the number of iterations of the main algorithm over which the perturbing input is incremented. The effects that varying mutation rates have on homeostasis will be discussed within the following chapter. The 0% values are the product of insufficient time for larger populations to respond to the increasing perturbing input.

range of the rein controller did not determine the range of external perturbations that the system was able to regulate against.

6.7 Explaining the Homeostatic Mechanism

To introduce a personal note to this chapter, I would like to state my initial reactions to the MGDW results: I thought they were possibly the consequences of specially selected parameter values that would not prove to be robust to alternative values. I was wrong. In my defence, the original MGDW model was rather more complex than the model presented in this chapter. There were, perhaps, more places for the mechanism that produced the counter-intuitive results to hide. What prompted my further interest in the model was a series of email discussions with Jamie McDonald-Gibson, Ezequiel Di Paolo (Jamie's MSc thesis supervisor), Inman Harvey, Tim Lenton and Hywel Williams. I think it fair to say that we were collectively stumped as to why the model was exhibiting homeostasis. I then decided to implement the model myself. I produced the same homeostatic results. Curious. I then produced a number of simpler versions of the model. I produced the same homeostatic results. Curiouser and curiouser. This, in conjunction with discussions with Inman Harvey slowly led me to understand the model in terms of rein control. There was a particularly illuminating exchange held with Inman as we drove to meet the creator of Daisyworld, Jim Lovelock. While filling up with petrol, I was describing the emergence of homeostasis in terms of a crinkled landscape of individual effects that could produce negative feedback regions. Inman characterised this as 'zero crossing points' that could lead to rein control. The population's response to changes to the environment can be understood in terms of rein control. This explanation will feature negative and positive feedback systems. It is important to note that the following analysis will ignore the effects of mutation on either loci. The establishment of a homeostatic state is sufficiently fast for the effects of mutation to be ignored. As the following chapter will show, mutation can significantly affect to what extent homeostatic states are maintained and the probability of whether a homeostatic states is established. For the time being, such details are not relevant to understanding the core, rein control mechanism.

6.7.1 Type ii Rein Control

In the simulations presented thus far, I is initialised at a low value and is slowly increased. Eventually V will reach the lower limits of the viability range and so there will be an individual that has non-zero fitness. This individual will have the lowest value B allele. I call this individual B_{best}. Given a random initial population, it is equally likely to have a decreasing effect e allele or an increasing effect E allele. Let us initially assume that it has an e allele. Consequently as it increases in frequency, the effect the population has on the variable, O, becomes negative and so dV/dI decreases. If V were to be driven back below the viability range all individuals would have zero fitness and there would be no further change in the population until I increased V back into the viability range. In this way V is maintained around the lower limit of the viability range and the mean fitness of the population is close to zero. As I increases, B_{best} increases in frequency and $\bar{A}$ progressively moves down towards B_{best}. The limits of homeostasis are reached when B_{best} reaches approximately 100% of the population (dependent on the mutation rate) and $\bar{A} = B_{best}$. Further increases in I lead to an increase in V at the same rate as the maximum population effect has been reached. This produces a hat shaped response in mean fitness as V traverses the essential range of B_{best}. This is a type ii rein control system in that as the rein controller output increases, it has a decreasing effect on the system variable. Consequently the environmental variable is maintained at the lower limits of the essential range of the population. Figures 6.7 - 6.9 shows a typical type ii simulation. This single rein control homeostasis proves to be uncommon as it requires a single type to maintain the environment at the very limits of regulation. If I is initialised within the essential range, or if slightly different circumstances occur at the very start of the homeostatic state, then an antagonistic form of rein control will emerge.

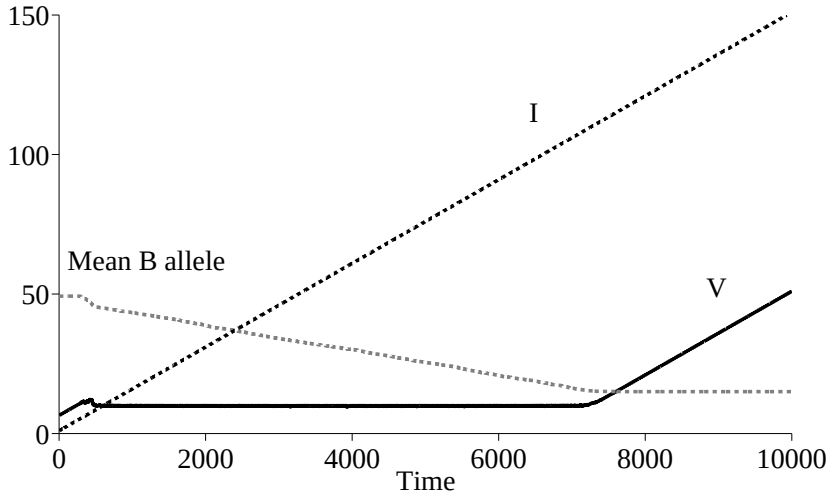

Figure 6.7: Type ii Rein Control

The perturbing input, I, is plotted with a dashed black line, the mean B allele is plotted with a dashed grey line and the environmental variable, V is plotted with a solid black line. V initially increases with I until I drives V into the range [15,85] at which point the population begins to slowly converge to a single B allele. As I continues to increase, V no longer increases but remains fixed around 15 until $I \approx 115$.

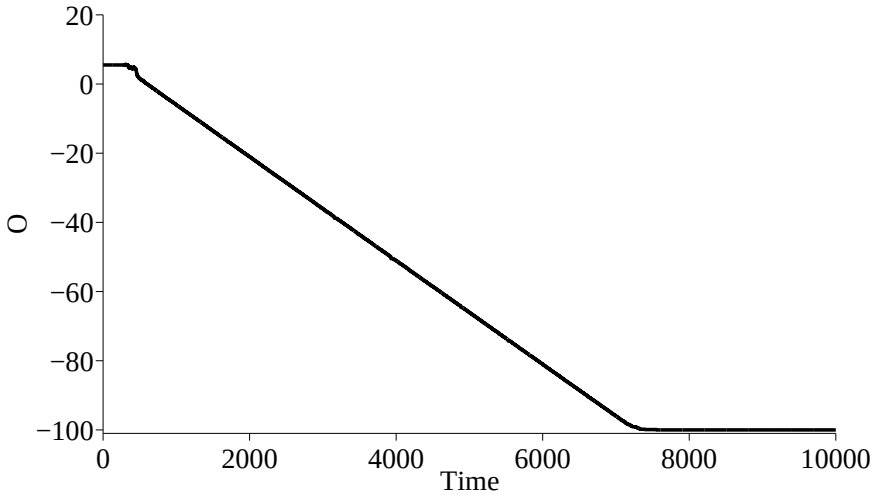

Figure 6.8: Type ii Rein Control - Organism Effects

Organism effects, O, is shown for the same simulation as shown in figure 6.7. As the perturbing input, I, increases, O decreases such that the environmental variable, V if held at the lower limits of the essential range.

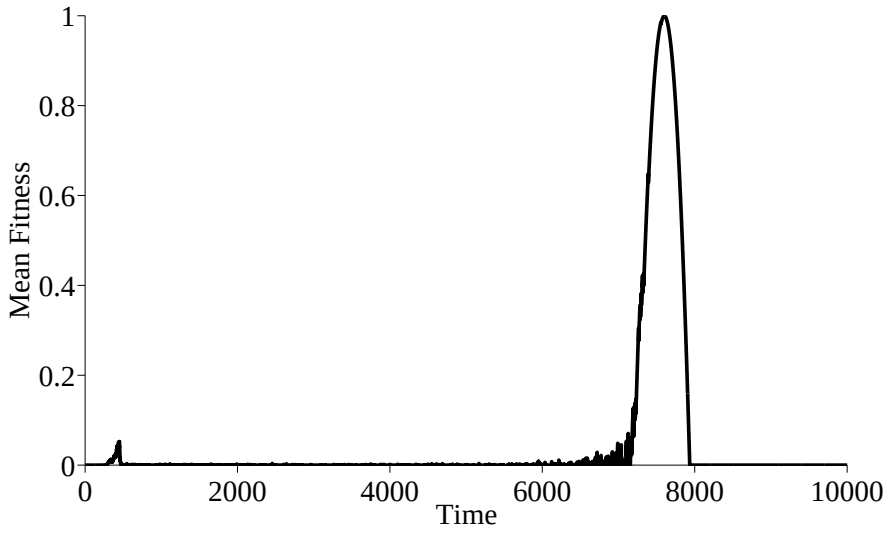

Figure 6.9: Type ii Rein Control - Mean Fitness

The mean fitness of the population is shown for the same simulation as shown in figure 6.7. Mean fitness is held near the minimum value until the end of the homeostatic period at which point it increases to a maximum and then minimum as I drives V past the essential range of B_{best}.

6.7.2 Type ii & Type iii Antagonistic Rein Control

Given that there is an equal probability of B_{best} having the increasing E allele or the decreasing e allele, there would initially appear to be a probability of 0.5 for the establishment of a homeostatic state because if B_{best} had an E allele, it would increase the rate of change of V and so produce a positive feedback type 1 circuit.

However, as V transits the optimal value of B_{best} it will move closer to the optimal value of the genotype that has the next lowest B allele value in the population and so a new B_{best} is established. Once again this individual will have an e or E allele with equal probability. If it too has the increasing E allele then V will continue to increase, transit this B_{best} and move towards the individual with the next highest B allele. See figure 6.10 for a geometrical explanation of this process. Given the equal chance of moving towards an individual with an e allele, homeostasis will be rapidly established as an individual with a reducing e allele will soon be encountered. As this type increases in frequency, a stable environmental variable value, $V*$, will be established between the [E, B_{low}] type and the [e, B_{high}] type. This produces two dominant sub-populations

that straddle V and force it in opposing directions. See figure 6.11 for a geometrical explanation of this second part of the process. Varying the strength of these reins (changing the numbers of individuals within each sub-population) leads to the variable being regulated to within a narrow range over varying external perturbations. If I were to increase, V would increase and move closer to the optimum of the [e, B_{high}] type. This would increase the fitness of [e, B_{high}] which would increase in frequency and so lead to a decrease of O and so reduce the increase in V. If I were to decrease, V would decrease and move closer to the optimum of the [E, B_{low}] type. This would increase the fitness of [E, B_{low}] which would increase in frequency and so lead to an increase of O and so reduce the decrease in V.

This is a type ii and type iii antagonistic rein control system that is comprised of two rein controllers that move the system variable *away* from their optimal output values and *towards* the optimal value of the opposing rein controller. It is an antagonistic rather than benign rein control system as rein controller 1 can only increase in output with a commensurate decrease in rein controller 2 output and vice versa. This arises due to the fixed population size and the fact that within a homeostatic state, the entire population resides in either one or the other of the rein control sub-populations. This is analogous to the birth rate of the daisies in Daisyworld being modulated by the amount of bare ground. A zero sum game is played between the [E, B_{low}] (type iii rein controller) and [e, B_{high}] (type ii rein controller). Consequently any change in the frequency of the e allele is matched by an opposite change in frequency of the E allele. A typical type ii antagonistic rein control system result was shown in figures 6.4 - 6.6.

6.8 Niche Construction

An evolutionary theorist may regard MGDW as a Moran-type model that uses Monte Carlo methods to change the frequency of alleles over time. What is perhaps less likely, is for the connection between MGDW and the theory of niche construction to be made, as niche construction is something of a 'niche' area of study within population genetics. I will argue that MGDW shares a number of fundamental assumptions with certain niche construction models and so shows how the evolutionary mechanism of niche construction can lead to environmental homeostasis. If we extend the discussion to the more general notion of frequency dependent selection, then we may begin to discern the outlines of a theory that would incorporate Gaia Theory into certain fundamental aspects of

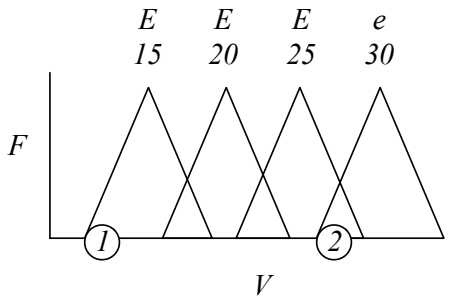

Figure 6.10: Emergence of Type ii & Type iii Antagonistic Rein Control: Part 1
This is the first part of a geometrical explanation for the emergence of antagonistic rein control. See figure 6.11 for the second part of this explanation. The main idea here is that if an individual has an increasing effect on the environmental variable, V, then it will increase V past the value that gives maximum fitness and so lead to a decrease in its fitness. As it drives V away from its optimal value, it will at the same time drive it closer to another individual. If this too has an increasing effect, then the rapid increase of V will continue. Four individuals within the population are ranked in order of increasing B allele value. The individual on the left has the lowest B allele value at 15, the individual on the right the highest at 30. The change in each individual's fitness as the environmental variable, V, changes is represented with a linear hat function. The peak of each hat function shows what V value will produce maximum fitness. Fitness, F, is plotted on the y-axis. The first three individuals have increasing effect E alleles. The fourth individual has the decreasing effect e allele. The circled number 1 denotes the value of V at time = 1. As V is within the essential range of the leftmost individual, this will increase in frequency in the population. Consequently, the frequency of E alleles would increase which would impart a further increase in V. Positive feedback results in V rapidly increasing and transiting the optimal value of this individual and moving towards the $B = 20$ individual. As this also has the E allele it will further increase V. This process of positive feedback would continue until V reached the position denoted with the circled number 2 at time = 2.

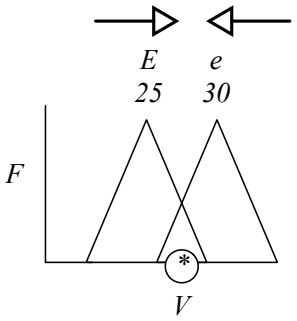

Figure 6.11: Emergence of Type ii & Type iii Antagonistic Rein Control: Part 2
This is the second part of a geometrical explanation for the emergence of antagonistic rein control. See figure 6.10 for the first part of this explanation. The main idea here is that given a randomly created population, an individual with a decreasing effect on the environmental variable, V, should be quickly encountered. As V is driven into the essential range of this individual, it will increase in fitness and so arrest the increase of V. A steady state is found where V is maintained at a value $V*$ by two sub-populations that have opposing effects on V. Plot (b) shows that the [e, $B = 30$] type produces a resisting force to increases in V. If V were to increase past $V2$ the frequency of [e, $B = 30$] would increase and so the frequency of e alleles would increase and the sum of the population effect operating on V would become negative and V would decrease. If V were to decrease below $V2$, the fitness of the [E, $B = 25$] type would increase and so V would increase. This results in V remaining fixed at $V*$ for varying perturbing input, I, and thus is equivalent in this respect to the proportional integral controller elements of detector, regulator and effector which seek to maintain a zero steady state error for the variable. The arrows represent the action of the effectors that drive the environmental variable, V, in opposite directions.

evolution.

6.8.1 The Evolution of Environmental Effects

A common assumption in evolutionary studies is that adaptation occurs in response to changes in the environment. The effect that organisms may have on the environment, while appreciated for some time [Darwin, 1881], introduces a significant feedback loop, the consequences of which have only recently been explored within evolutionary theory, [Lewontin, 1982], [Lewontin, 2001]. Organisms through their activity and metabolism will create, modify and at times destroy their environmental niches. The scale of these niches may differ significantly in space and time. For example burrowing earth worms change the composition of local soil whereas photosynthetic cyanobacteria have an effect on the global climate via their emission of molecular oxygen. The unifying concept is that such changes to the environment will in turn have an effect on the selection pressures operating on organisms, which in turn have an effect on the environment through their niche constructing activities and so on. The theory of niche construction attempts to incorporate such feedback into the established theories of population genetics and evolutionary theory [Odling-Smee et al., 2003]. Notable results are that niche construction can maintain or destroy stable polymorphisms, allow otherwise deleterious alleles to reach fixation, introduce an effect of evolutionary momentum and influence dis-equilibrium [Laland et al., 1999], [Laland et al., 2000], [Laland and Sterelny, 2006], [Silver and Di Paolo, 2006]. An important concept in niche construction is the acknowledgement that environmental conditions can, in conjunction with genetic information, form a dual system of inheritance. Just as an organism will inherit its parent's genetic information, so it will inherit its parent's environment. This environment may have been altered by its parent (and other organisms) and may in turn be altered by the offspring. Figure 8.7 provides an overview of this dual inheritance niche construction system.

Following [Odling-Smee et al., 2003] I identify two types of niche construction effects that change environmental conditions: inceptive - organisms initiate a change in their environment; and counteractive - organisms counteract a prior change in the environment. Both inceptive and counteractive niche construction are visible in MGDW. Inceptive niche construction can lead to positive feedback effects whereby a genotype increasing in frequency in the population leads to an increasing effect on the environmental variable that further increases the fitness of this genotype and so on. This was

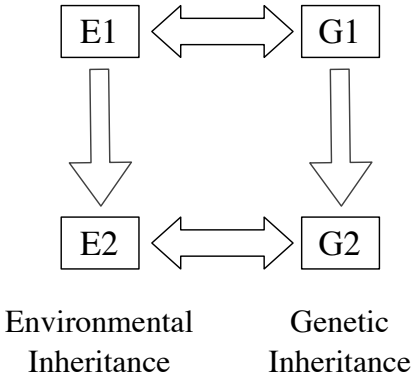

Figure 6.12: Niche Construction Dual Inheritance

The box labelled G1 represents the genetic information of a population of organisms at time 1. This population is affected by its environment, denoted by box E1, which in turn is affected by the population. This dual effect is represented by the double-headed arrow. Assuming discrete, non-overlapping generations, the box labelled G2 denotes the offspring of G1. This new population would have inherited G1's genetic information. It would also have inherited the modified E1 environment, denoted by box E2, which has an effect on G2 and is in turn affected by it.

observed in the establishment of type ii antagonistic rein control where the increase in frequency of the E allele led to a positive feeback effect that rapidly increased the environmental variable. Counteractive niche construction can lead to negative feedback whereby the change in the frequency of a genotype or genotypes leads to a counteractive force being applied to the environmental variable. This was observed in type ii rein control where the increase in frequency of the e allele counteracted the increase in V produced by I and in type ii antagonistic rein control where the two rein controller sub-populations counteracted any changes in V produced by I increasing or decreasing.

The similarities to niche construction go further in that MGDW shares a number of assumptions with a landmark niche construction study [Laland et al., 1999]. This population genetics model featured genotypes that determine a phenotypic effect on a global resource that in turn determined the fitness of phenotypes. Laland et al found that when introducing this feedback loop from organism to environment, a number of surprising results were found. In particular, stable polymorphisms could emerge. Rather than a single genotype reaching 100% frequency in the population, other 'sub-optimal' types could persist. MGDW exhibits such stable polymorphisms. When in a type ii antagonistic rein control state, neither sub-population is able to reach 100% in the population as this would drive the environmental variable beyond its essential range

and so lead to zero fitness. Each sub-population drives the environment away from it and towards the other sub-population and so a stable polymorphism is produced. While Laland et al capture the population dynamics of this mechanism what is not appreciated is that this stable polymorphism may be achieved in conjunction with environmental homeostasis. This is not surprising as a population genetics study will typically be primarily interested in the dynamics of populations. However, once we appreciate the consequences that niche construction effects can have on shared environments, then it should come as no real surprise to see that environmental homeostasis can accompany population stabilizing mechanisms.

6.9 Discussion

In this chapter I have presented a version of the McDonald-Gibson Daisyworld (MGDW) that had the potential for both positive and negative feedback circuits to emerge (types i - iv rein control using the classification from the previous chapter). When the organisms had no effect on their environment, a selective sweep led to the population converging to the current conditions. As the perturbing input further increased the environmental variable, the population tracked this change. The situation was markedly different when the individuals exerted small effects on their environment. The environment then remained fixed as the net effect of the population responded to changes to the perturbing input.

This homeostasis was explained in terms of type ii and type ii & type iiii antagonistic rein control. Negative feedback circuits emerge that stabilise the environmental variable about a particular value. Although positive feedback effects are equally probable as negative feedback effects, positive feedback does not dominate and the environmental variable is neither driven beyond the essential range, nor progressively increased with increases in the perturbing input. Some readers may still think that some trick has been performed here. They would be correct in so far as the trick involves appreciating that the rein controllers in MGDW have *unimodal* responses to the environmental variable. The rein controllers discussed in chapter 5 had output functions that saturated at particular system variable values. While an established rein control system in MGDW is effectively equivalent to a type ii or type ii & type iii rein control system, there is an important difference in how such rein control is established in MGDW. This may require unimodal output functions or indeed any output function that can have at least

one point of inflection so that it can be both an increasing and decreasing function of the environmental variable. And therein lies the trick because the MGDW unimodal fitness function will transform a positive feedback effect into a negative feedback effect. If V were to increase the fitness of an E allele type, then this type will further increase V until it transits the value that gives maximum fitness. Further increases to V will decrease the fitness of this type. If V were to now *decrease*, the fitness of this type would *increase*. A type i, positive feedback circuit rein controller would be transformed into a type iii, negative feedback circuit rein controller. This is essentially equivalent to the anthropomorphic explanation of the phototaxis behaviour of the simple cable car model in chapter 4. Depending on the direction the light source enters the activation range of the solar panel, either cable car can be either a light-phile or light-phobe. In MGDW, any individual can attempt to move the current environmental conditions towards its B allele value, or away from it. Given the unimodal nature of the fitness function, any transit of the B allele value will result in a negative feedback effect rein controller being established. Within the context of Thomas' feedback circuits, output functions that are both increasing and decreasing functions of a system variable are 'ambiguous' in that they may produce negative or positive feedback effects.

If V were initialised within the essential range and between two opposing rein control individuals, then homeostasis could emerge within non-unimodal output functions. Unimodal and higher order functions allow homeostasis to emerge under a much wider range of conditions. It is arguably a universal characteristic of biological organisms to have such bounded responses to environmental conditions. Organisms can only tolerate certain ranges of certain environmental variables. It is also arguable that all organisms affect their environments.

6.10 Summary

In this chapter I have made a number of connections between MGDW and the theory of niche construction. Whilst incorporating the effects that organisms have on their evolution may produce a number of non-trivial complexities into currently tractable evolutionary models, it may highlight the potential for such co-evolving systems to produce homeostatic states. Biologically-mediated environmental homeostasis may emerge not despite of evolutionary mechanisms, but because of them.

This is of relevance to studies that argue that environmental homeostasis could only

arise as a by-product of organismic function. For example [Volk, 2003], [Volk, 2002] argues that a form of Gaian homeostasis could emerge from the incidental effects that organisms have on their environments. While MGDW is consistent with Volk's thesis as there is no selection pressure for the environmentally altering traits in MGDW, I have shown how natural selection is an important mechanism for the establishment of homeostasis.

In the following chapter I will continue to explore MGDW by evaluating the effects of mutation on homeostasis. I will show how mutation can destroy particular homeostatic states, and also the mechanisms whereby homeostasis can be recovered. This will shed light on the complexity-stability debate within theoretical ecology.

Chapter 7

Complexity and Stability

7.1 Previous Publications

The results in section 7.4 are reproduced from [Dyke et al., 2007]. Figure 7.1 is reproduced from [McDonald-Gibson et al., 2008]. Parts of the discussion on the effects of mutation are taken from my contributions to [McDonald-Gibson et al., 2008].

7.2 Overview

In the previous chapter I introduced the McDonald-Gibson Daisyworld (MGDW). I argued that this is an important model as it demonstrates that homeostasis can arise under less constrictive assumptions than other versions of Daisyworld and in particular that homeostasis can emerge in the presence of both positive and negative feedback effects. I explained the homeostatic behaviour of MGDW in terms of rein control and the emergence of counteracting sub-populations that seek to drive the environmental variable in opposing directions.

In this chapter I will continue to explore MGDW with a focus on identifying the limits of homeostasis and the effects of mutation rates. I will detail two mechanisms whereby mutation can lead to the collapse of any homeostatic state. Understanding these two mechanisms could lead to the suggestion that increasing the mutation rate would lead to a reduction in the model's ability to regulate the environmental variable. However, such intuitions are misleading, as while the model may be simple it has the capacity to surprise us. This is because there is a range of mutation rates over which increasing mutation increases the homeostatic properties of the model. Understanding

this result will shed light on the relationship between the complexity and stability of ecosystems. A key component of the result is that MGDW can inhabit multiple homeostatic states where the environment is regulated around a series of particular values. This further demonstrates the absence of prior defined optimal conditions in MGDW.

As I argued in the previous chapter, once one incorporates the effects that organisms can have on their environment, new results may be returned. The concept of *ecosystem engineering* attempts to incorporate such feedback into theoretical ecology. Just as organismic feedback operating over evolutionary timescales may be significant to the long-term evolution of life and its environment, I will argue that ecosystem dynamics that occur over shorter timescales may also be significantly affected by the effect organisms can have on their environment.

7.3 The Effects of Mutation in MGDW

Mutation operates independently on A locus alleles (which determines the phenotype's effect on the environmental variable - this can be understood as daisy albedo in Daisyworld) and B locus alleles (which determines the phenotype's optimal environmental conditions - this can be understood as the temperature at which the daisies will have the fastest growth rate in Daisyworld). Therefore an organisms may differ from its parent in how it responds to its environment and how it affects its environment.

7.3.1 Mutation of A Alleles

The following discussion will assume binary A alleles. This maintains my adopted convention of formulating MGDW in terms of decreasing e and increasing E alleles and also allows a simpler explanation for the effects of mutation. The core mechanism identified also applies to formulations of MGDW with continuous value A alleles.

Increasing the mutation rate operating on the A locus alleles increases the probability that an offspring from an e allele parent will have an E allele and vice versa and so have the opposite effect on the environment from its parent. Recall that there is no selection pressure for A alleles in that every individual is affected by the shared environmental variable in exactly the same way. So if an individual increases the environmental variable, V, then the change in V it experiences because of its effects are exactly the same as all other individuals within the population. Therefore with non-zero mutation

rates we should expect a homeostatic population with its initial distribution of e and E alleles to smooth out over time. This is because as there is no selection pressure for A alleles, random drift would be likely to mix up the biomodal population. This drift can be observed if we view the distribution of e and E alleles within the population. It is possible to construct a figure that provides a three dimensional view of the population in both A and B allele 'space'. Figure 7.1 is an example of such a view that shows the origins of a rein control state.

Figure 7.1 was produced by 'binning' the population on the basis of the B allele value. Each individual was placed within a discrete B allele bin that was closest to its B allele value. For example, if using integer value bins and individual x had a B allele value of 56 they would be placed in bin 56 and if individual y had a B allele of 57 they would be placed in bin 57. Each bin was then assessed in terms of the sum of e and E alleles. If 13 individuals were in bin 56, then the sum of their e and E alleles would be assigned to that bin. Figure 7.2 shows the results of a simulation in which the perturbing input, I, remained fixed at 40 and mutation is only operating on A alleles. The rapid establishment of a rein control system is followed by the progressive decrease in the sum of e and E alleles over time so that the net effect of the organisms, O, decays to approximately zero.

Figure 7.3 helps to illustrate this mechanism. Recall that homeostasis is maintained by two sub-populations that drive the environmental variable in opposite directions. As the perturbing input changes, it is the relative sizes of the sub-populations that change while the environmental variable remains fixed. We can imagine that each sub-population is playing a modified version of the game of tug-of-war. The two sub-populations are called team X and team Y and their team members are equal in strength. Team X consists of four Strong Men. Team Y consists of four Amazonian Wonder Women. Rather than pull in opposing directions on a rope, each team member holds a rigid bar and pushes *towards* the opposing team. Behind each team lays a drop off into a ditch or body of water or suitably motivational geographical feature of your choice. A drift event in the A allele is analogous to a member of team X turning around and rather than pushing towards team Y begins pushing *away* from Y and towards said ditch, water etc. This event is represented in figure 7.3(b) with a Strong Man turning into a Wonder Woman while at the same time a Wonder Woman turns into a Strong Man. If both teams have the same number of members, then we should expect approximately the same rate of strong men turning into Wonder Women and Wonder Women

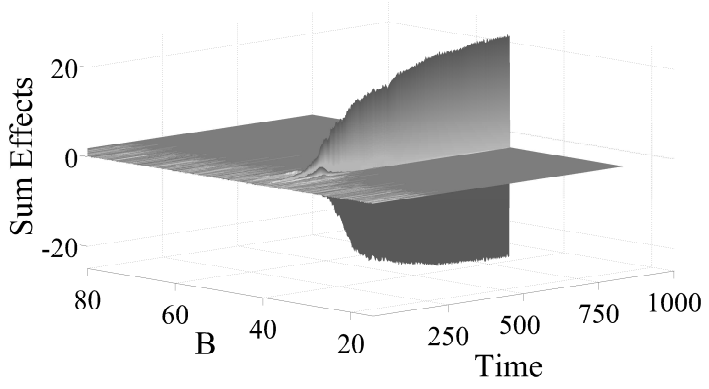

(a)

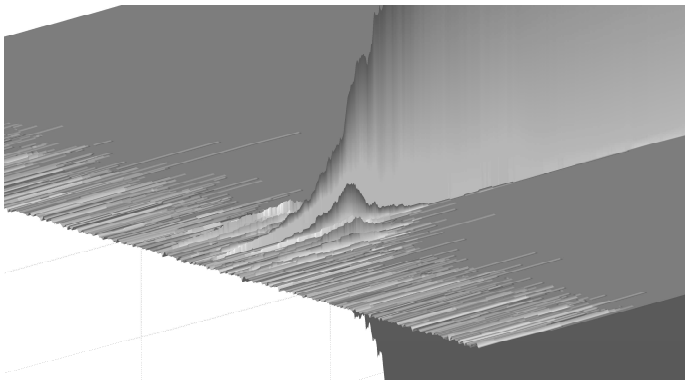

(b)

Figure 7.1: 3 Dimensional Plot Showing the Origins of Rein Control

The main idea here is to show the emergence of a type ii and type iii rein control system. An initially diverse population rapidly converges to either an increasing effect or a decreasing effect sub-population. Plots (a) and (b) show results from the same simulation with the following parameter set: population size, $k = 2,000$, $e = E = -0.05$, mutation rate $= 0.1$ and the perturbing input, I, is initialised at 40 and increases to 55 over 1045 time steps. The bottom plot (b) is a closer view of the central portion of the top plot (a). In order to create the plots, individuals were placed into B allele value bins at 10^{-3} intervals. The sum of e and E alleles was calculated for each bin and plotted on the vertical y axis. Time is plotted on the x-axis (right horizontal axis) and the distribution of B alleles on the z-axis (left horizontal axis). See section 7.3.1 for details on how the data was generated. A initial population of diverse B alleles converges to approximately 40.

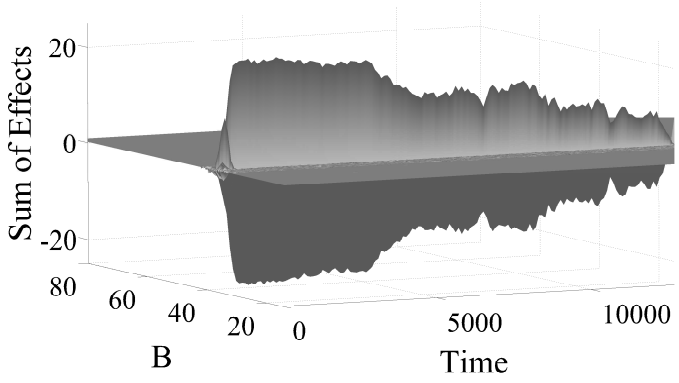

Figure 7.2: Drift of A Alleles

The main idea here is to show the rapid establishment of an antagonistic rein control system, and then the progressive decrease in the strength of the increasing effect and decreasing effect sub-populations due to mutation operating on the A allele. The data was produced with the following parameter set: population size $k = 2,000$, $e = E = -0.05$, mutation rate $= 0.1$ and perturbing input, I was fixed at 40 for 10^5 time steps.

turning into strong men. This will maintain the rigid bar in the same position. Figure 7.3(c) shows one possible equilibrium state for this game which is an equal number of strong men and Wonder Women on both teams both pulling and pushing. In rein control terms, this represents the end of the homeostatic state. Recall that it is the differential change in the frequency of e and E alleles that regulates the environmental variable. Once both sub-populations have the same proportion of e and E alleles, a change in the environmental variable will no longer produce a counteracting change in the proportion of e and E alleles.

This can be characterised as an absorbing state in that once established it is impossible to escape. Natural selection operating on such a population could not recreate the bimodal distribution of e and E alleles that is required for rein control. This bimodal distribution was produced by the original type ii and type iii rein controllers rapidly increasing in frequency in the population. while the drift of A alleles is a much slower process, ultimately any rein control state is doomed to decay towards the absorbing state.

7.3.2 Mutation of B Alleles

In the previous chapter I showed how the effects organisms have on their environment can produce stable polymorphisms that buffer their environment against perturbing

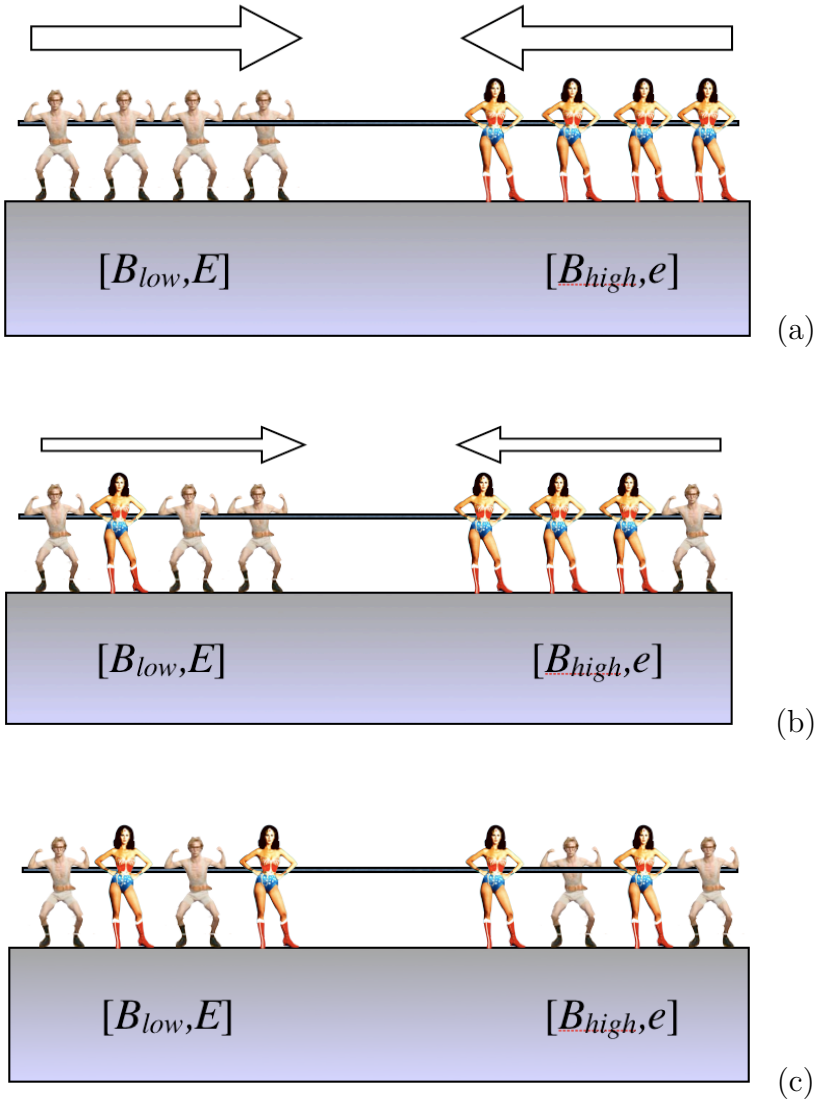

Figure 7.3: Drift in A Allele Leading to Reduction in Organism Effects
Plot (a) shows an established rein control population. The team of strong men represent the increasing $[B_{low}, E]$ sub-population while the team of Wonder Women represent the decreasing $[B_{high}, e]$ sub-population. The sum of the effects of both teams is equal and is denoted by the large equally sized and opposing arrows. Plot (b) shows two mutation events at the A allele locus occurring in two different individuals. One right-pushing Strong Man has mutated into a left-pushing Wonder Woman while on the other team one left-pushing Wonder Woman has mutated into a right-pushing Strong Man. Plot (c) shows a possible equilibrium state where random drift in the A allele has produced an equal number of Strong Men and Wonder Women in both teams. This is an absorbing state as once reached it is not possible to reproduce the bimodal distribution of decreasing e and increasing E alleles in the population.

input. Mutation operating on the B locus also maintains polymorphisms through the simple mechanism that higher mutation rates will increase the probability that a B allele offspring will have a higher or lower B allele value than its parent. However, this increased diversity can be operated on by selection and in doing so can lead to a B allele monomorphic population that is no longer homeostatic.

Figure 7.4 shows that intermediate B allele mutants will have higher fitness than either $[B_{low}E]$ or $[B_{high}e]$ types and will, over time, increase in frequency. This leads to the two sub-populations moving closer together in 'B space'. Given sufficient time, this would lead to the two sub-populations merging into a single population where $[B_{high}] = [B_{low}]$. This spells the end of the antagonistic rein control system as there will be no change in the frequency of the e and E alleles in response to changes in V produced by I. It was the change in frequency of the $[B_{low}, E]$ or $[B_{high}, e]$ types that was responsible for the net change in the sum of the population's effects on their environment. Once all individuals share the same B allele value, the population can offer no differential response to changes in V and therefore homeostasis is no longer observed. This is an absorbing state as once all individuals share the same B allele, natural selection would maintain the B allele monomorphism.

However, this is typically not observed in MGDW because as the two populations converge towards a single optimal, intermediate B allele, the range of fitter intermediate mutants decreases. With a continuous range of B allele values, the distance between the B_{high} and B_{low} genotypes can become infinitesimally narrow and so the probability of producing a mutant offspring with an intermediate B allele can become infinitesimally small. In conjunction with this is the fact that changes to V may be sufficiently large to keep V outside of the fitter intermediate mutant range. A parameter set will demonstrate this mechanism. Let the population size, $k = 1000$, death rate $\gamma = 1/k$ so that two individuals are selected for tournament reproduction at every time step and $e = -E = 0.1$. If a single individual is replaced in the population via tournament selection at every time step, then there will be a maximum change of ± 0.2 and a minimum change of 0 in V at every time step (± 0.2 produced by an e replacing an E or vice versa). Figure 7.5 shows an example of such variation in V. The change in the frequency of e and E alleles at every time step is sufficiently large for V to 'leap frog' over the intermediate range of values that are between the two rein controlling sub-populations. while intermediate mutants would have higher fitness as their B allele would be closer to the time-averaged value of V, the probability of such an intermediate mutant being

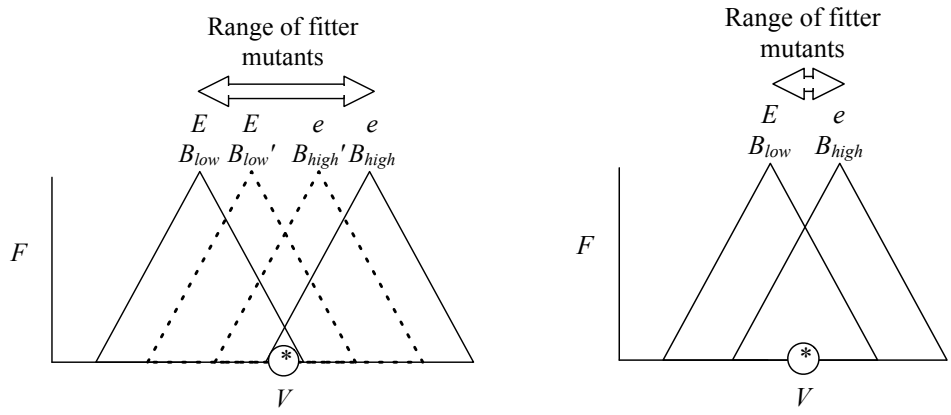

Figure 7.4: Selection of Intermediate B Allele Mutants and Progression Towards Absorbing State

This figure is similar to figure 6.6 in the previous chapter that explained the emergence of type ii and type iii antagonistic rein control. The main idea here is that for any antagonistic rein control system, there will be a range of intermediate mutants that will have higher fitness than the established sub-populations and so will, in time, increase in frequency in the population and ultimately replace the previous sub-populations. The established rein control sub-populations are represented by the solid triangles. The environmental variable, V, is being regulated at $V*$ - the circled asterisk. The double-headed arrow represents the range where mutants would have higher fitness than either of the established sub-populations. The dashed black triangles represent two mutants that have B locus alleles within this range and so are closer to $V*$. The two mutants are labelled $[B'_{low}, E]$ and $[B'_{high}, e]$. Plot (b) shows the loss of the previous dominant sub-populations and the establishment of the previously intermediate mutants at time = 2. The previously dashed line triangles are now solid. The range of higher fitness intermediate mutants has narrowed. Through this process, the two rein control populations would move closer and closer together.

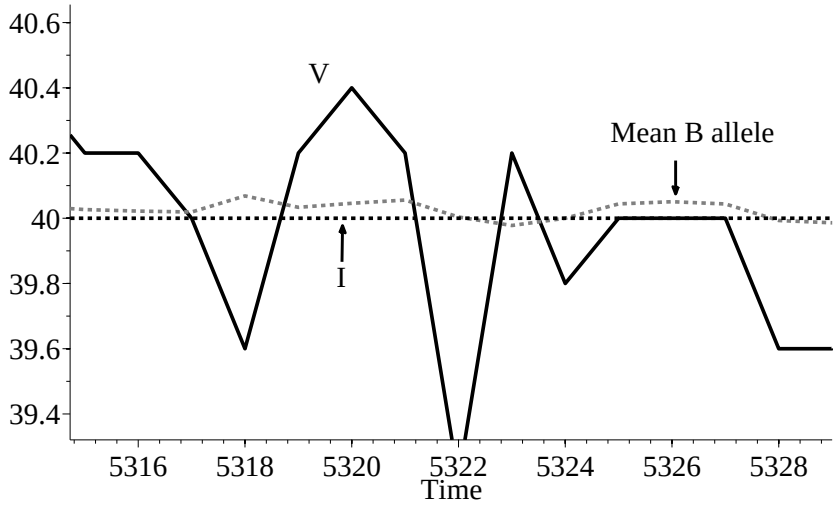

Figure 7.5: Minimum Changes in V

Results are shown from a simulation with the following parameter set: population size, $k = 1,000$, $e = E = -0.1$, mutation rate $= 0.1$, death rate $= 0.01$ and perturbing input, I is fixed at 40. The environmental variable, V, the perturbing input, I and the mean B allele of the population are plotted in dimensionless units on the y-axis. Changes in the environmental variable, V, at each time step are sufficiently large enough to move it above and below the mean B allele value of the population.

selected for tournament selection at a time when V is within the intermediate range is small.

While such changes in the environmental variable can mean that the increase in frequency of intermediate mutants becomes increasingly less probable, it is important to consider more generally the progression of the population to absorbing states. For example, very large populations could produce very small changes in their environment and the essentially discrete genetic code could produce discrete phenotypic changes in the fitness functions such that the convergence to a monomorphic B allele population becomes increasingly probable. The next section will discuss a mechanism whereby absorbing states could be avoided under a much wider range of assumptions.

7.4 Multiple homeostatic states

Selection operating on the A and B locus alleles could give rise to absorbing states and the collapse of homeostasis. In this section I will show that mutation can also lead to the collapse of a homeostatic state through the emergence of positive feedback effects.

Selection can produce type i and type iv positive feedback rein controllers that drive the environmental variable beyond the range of the controlling sub-populations. In chapter 5 I classified a range of possible feedback effects into positive and negative feedback effect rein controllers. Type i rein controllers have an increasing effect on a system variable and as that variable increases, the output of the controller increases. Type iv rein controllers have a decreasing effect on a system variable and as that variable decreases, the output of the controller increases. Both types will lead to positive feedback loops. I will show that such positive feedback does not necessarily lead to the irreversible collapse of homeostasis as MGDW may produce *multiple* homeostatic states with the environmental variable being regulated around *multiple* fixed values.

Figures 7.6 - 7.8 shows the results of a simulation that features multiple homeostatic states. Between each period of homeostasis there is a sharp decrease in fitness, an increase in V and $\bar{B}$ and O. It is important to note that the sum of the population's effect, O, is less than the maximum 100. There is 'surplus' capacity on the part of the population to counteract increasing I. A particular homeostatic state collapses at less than half of the maximum possible O value. Mean fitness is near the maximum during the periods of homeostasis so it is reasonable to assume that very nearly the entire population is held within the rein controller sub-populations which will be close to the current environmental conditions.

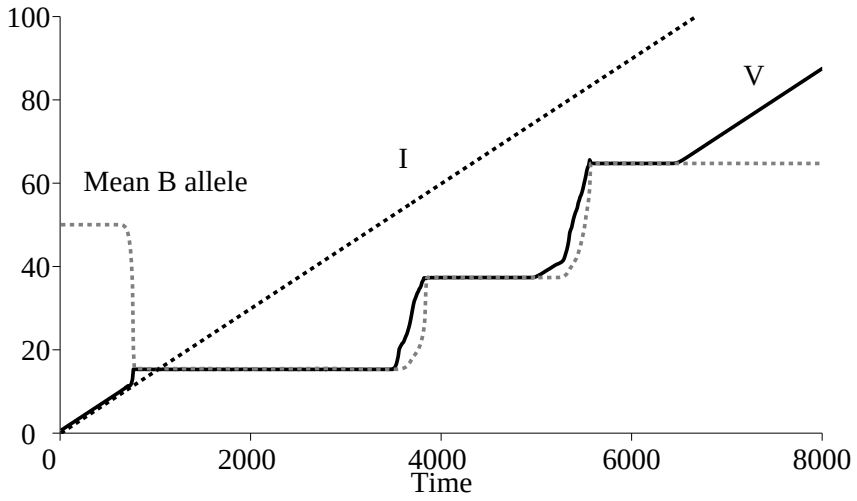

Figure 7.6: Multiple Homeostatic States

Results were produced from the following parameter set: population size, $k = 10,000$, $e = E = -0.01$, mutation rate = 0.1 and perturbing input, I is initialised at 0 and increases to 150 over 8×10^5 time steps. Perturbing input, I, is plotted with a black dashed line. Environmental variable, V is plotted with a solid black line. Mean B allele is plotted with a dashed grey line. Three homeostatic states regulated V at three different values. There are sharp changes in V and the mean B allele between each homeostatic state. There is a similarity between the behaviour of MGDW and the model reported in [Williams and Noble, 2005] which also exhibited sharp transitions between multiple homeostatic states.

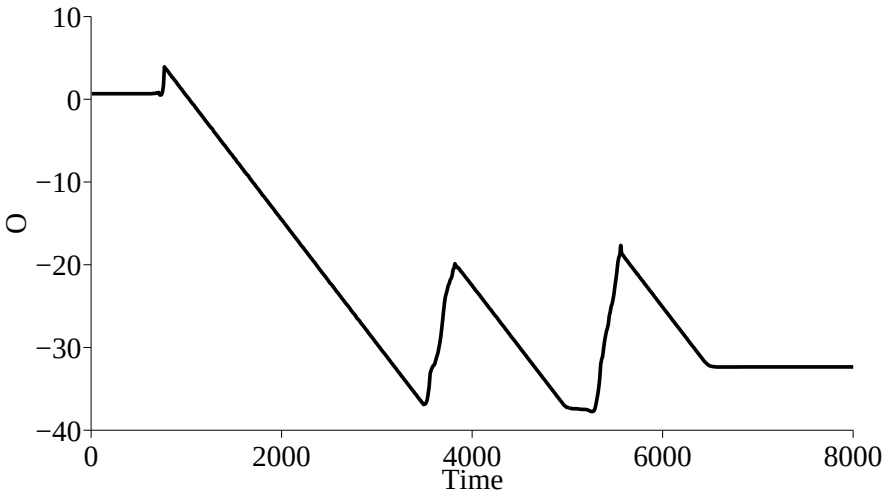

Figure 7.7: Multiple Homeostatic States - Organism Effects
The organism effects, O are shown for the same simulation as shown in figure 7.6. There are sharp changes in O between each homeostatic state.

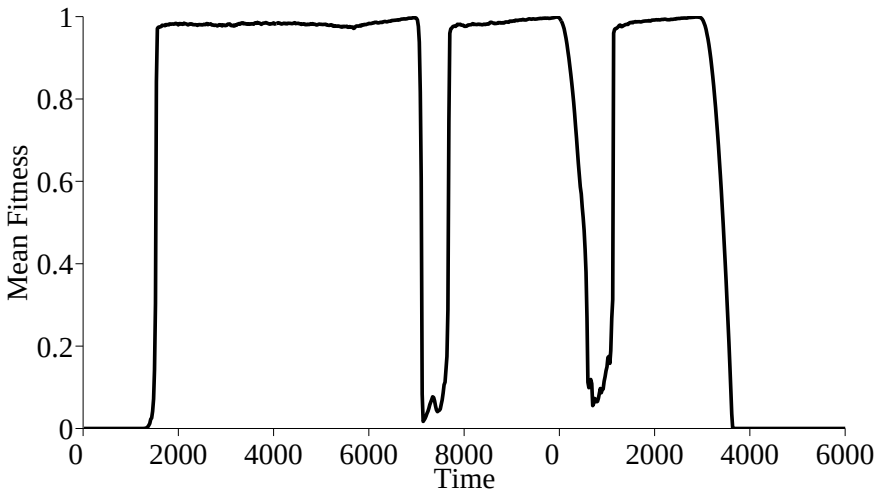

Figure 7.8: Multiple Homeostatic States - Mean Fitness
The mean fitness of the population is shown for the same simulation as shown in figure 7.6. There are sharp changes in mean fitness between each homeostatic state.

The sharp transitions can be explained in terms of positive feedback effects. Figure 7.9 shows a possible scenario in which a type i individual could emerge in the population

via the generation of a $[B_{high}, E]$ mutant which was produced by a $[B_{high}, e]$ parent. Given a small distance between the rein controller sub-populations and an amount of random variance in V, it is possible that this mutant could increase in frequency in the population. Such an increase would result in an increase in the total number of E alleles and so lead to an increase in V that would further increase the fitness of this mutant. A scenario similar to the origins of an antagonistic rein control state is produced in which positive feedback effects rapidly eject V away from the absorbing state and past the essential range of the mutant. This will lead to a collapse in the mean fitness of the population, reduction in the absolute value of O and an increase in V and $\bar{B}$. In this way, mutations operating on the A allele lead to the creation of a positive feedback effect individual which allows the system to escape from an absorbing state.

We also need to consider B allele mutants that are created outside of the intermediate range. For example a $[B_{low}E]$ parent that produces a $[B_{high}E]$ offspring will produce the same type i individual that arose due to mutation operating on the A locus.

An appreciation of the other effects that B allele mutants may have is necessary in order to understand how homeostasis is recovered. The establishment of rein control requires individuals that, as they increase in frequency, produce a counteracting force on V. Within the context of the current discussion, this would be an individual that had a B allele higher than the type i mutant and had the e decreasing effect allele. As the type i mutant drove V towards it, it would increase in frequency and so decrease the rate of increase in V until a new stable value for V was established. The probability of this occurring increases with increasing mutation rate as more mutants are produced. Higher mutation rates will lead to a more diffuse population and counteract selection which will reduce variance. Consequently, while higher mutation rates will typically lead to the faster collapse of any particular homeostatic state, it can increase the chance that a new homeostatic state at different environmental conditions will be achieved. Mutation is the engine of both destruction and creation of homeostasis in MGDW.

7.5 Mutation Rates and Homeostasis

In order to further explore the relationship between mutation rates and homeostasis, 100 simulations were performed for the 11 mutation rates over the range [0,1]. For each simulation, the total amount of time that V was regulated was recorded, along with the ecosystem's biodiversity. Results are plotted in figures 7.10 and 7.11. Biodiversity

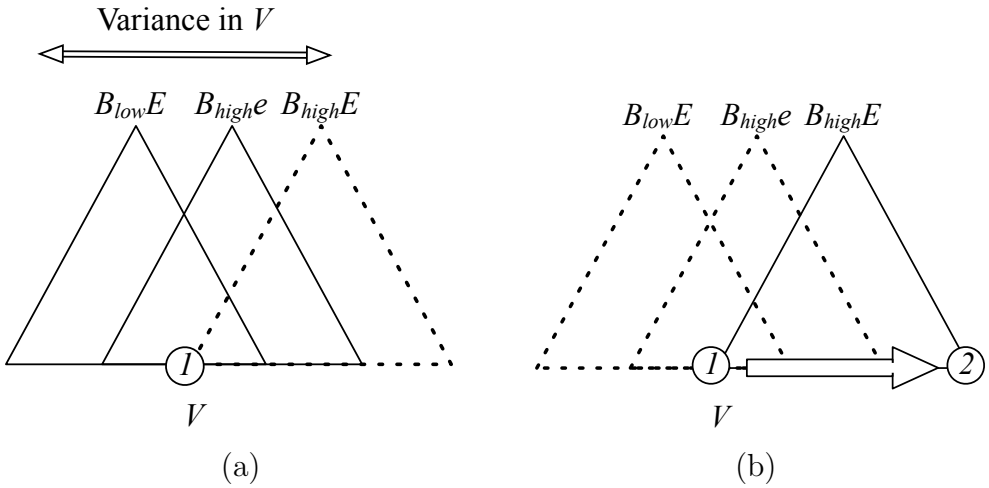

Figure 7.9: Escape from Absorbing State

The main idea here is to detail a mechanism whereby a population would avoid an absorbing state. Such an escape would be facilitated by the emergence of an individual that had an increasing effect on the environmental variable, V, *and* had higher fitness with higher values of V. This would produce a rapid increase in V and eject the population away from the absorbing state. Selection operating on B locus alleles would produce mutants that would have B alleles lower than the $[B_{low}, E]$ type and higher than the $[B_{high}, e]$ type. Plot (a) shows the production of a $[B_{high}, E]$ mutant offspring (dashed black line) at time = 1. Random fluctuations in the level of the environmental variable, V (denoted by the double headed arrow), may be sufficient for this mutant to increase in frequency in the population. Such an increase may produce an increase in the frequency of increasing effect E alleles and so increase V, which would bring it closer to the mutant's optimal value and so increase its fitness and increase its frequency. Plot (b) shows the effects of this positive feedback at time = 2. The environmental variable has been increased from the previous value at $V1$ to a new, higher value at $V2$. The previously dominant rein controller sub-populations (denoted with dashed black lines) are in the process of being lost as their fitness is zero.

was measured in terms of the number of unique B alleles in the population using a similar binning technique detailed in section 7.2.1. The B allele for each individual was rounded to the nearest integer and a sum of present B alleles recorded at every time step (the maximum being 71). This total was then divided by the number of time steps over a simulation. This provided a measure of the diversity of the population over the length of a simulation. Homeostasis was assessed using the same technique as detailed in the previous chapter's section 6.5. Figure 7.11 shows that as mutation rates increase, the diversity of the population increases. Homeostasis increases with increasing mutation rate, μ, until μ equals 0.3. Increasing mutation rates beyond 0.3 leads to a significant decrease in homeostasis.

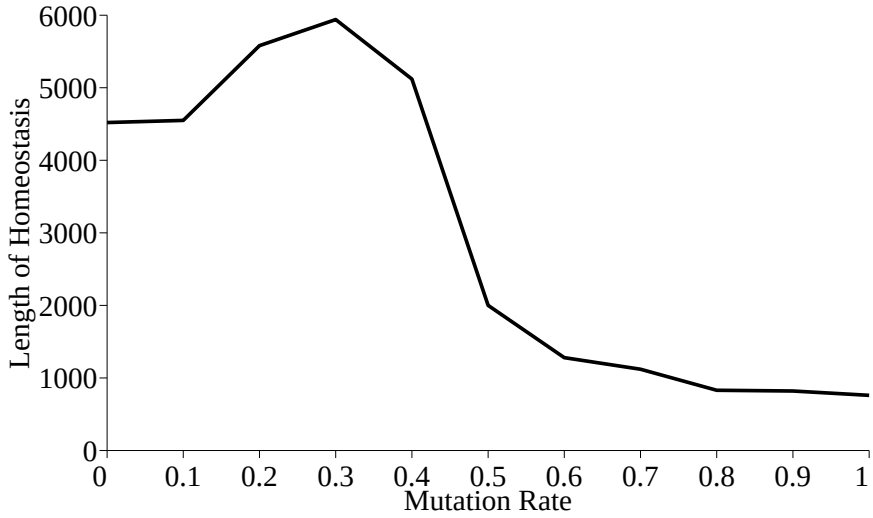

Figure 7.10: Mutation Rate Effects on Homeostasis

100 simulations were performed for the 11 mutation rates over the range [0,1]. All simulations used the same parameter set (apart from mutation rate) of: population size, $k = 2000$, $e = -E = 0.05$, perturbing input, I initialised at 0 and then increasing to 100 over 10^5 time steps. For each simulation, the length of time that the environmental variable, V was regulated was recorded (see chapter 6 section 6.6 for details). The mean number of turns that V was regulated for each mutation rate is plotted. Homeostasis increases as mutation rates increase from 0 to 0.3. Mutation rates higher than 0.3 lead to sharply reduced homeostasis.

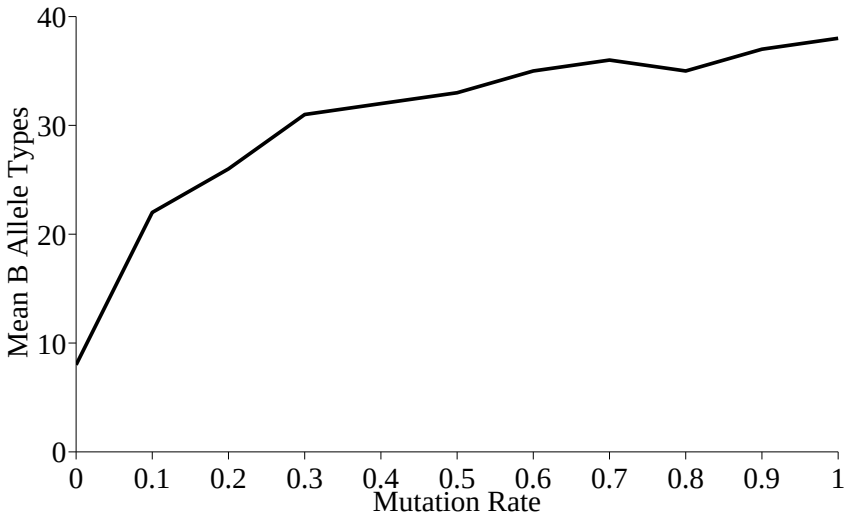

Figure 7.11: Mutation Rate Effects on Population Diversity
The mean amount of diversity for each mutation rate for the same 100 simulations as shown in figure 7.10 is plotted. Diversity was calculated in terms of unique B alleles (see section 7.5 for details). Increasing mutation increases diversity with the greatest rate of increase over the range [0,0.3]

Figures 7.12 and 7.13 show that increasing the mutation rate can lead to more rapid collapse of a particular homeostatic state. However, another homeostatic state rapidly emerges. Consequently, the population is able to regulate the environment over a wider range of perturbing input. With no mutation, the population will rapidly converge to only 2 distinct B allele types. Once this homeostatic state ends, there is very little chance of a new rein control state being found. Mutation opposes selection and ensures that individuals with lower and higher B alleles are continually produced and so increases the probability of new homeostatic states being established.

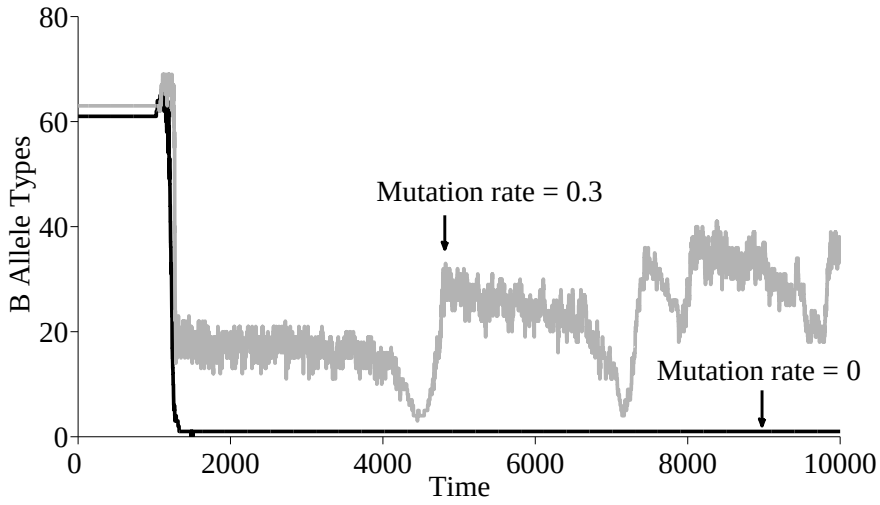

Figure 7.12: Population Diversity For Two Mutation Rates

The change in the number of B allele types in the population over time for two different simulations with different mutation rates, μ, are shown. These results are typical of simulations when the mutation rate is fixed at 0 and 0.3 and are presented in order to illustrate how increasing the mutation rate can increase the total amount of homeostasis observed in the model. In both simulations I was initialised at 0 and increased to 100 over 10^5 time steps. With no mutation, the number of B allele types rapidly decreases to 2 with the establishment of rein control homeostasis. When $\mu = 0.3$ the population rapidly decreases to approximately 20 and then undergoes a series of decreases and sharp increases as new homeostatic states are established.

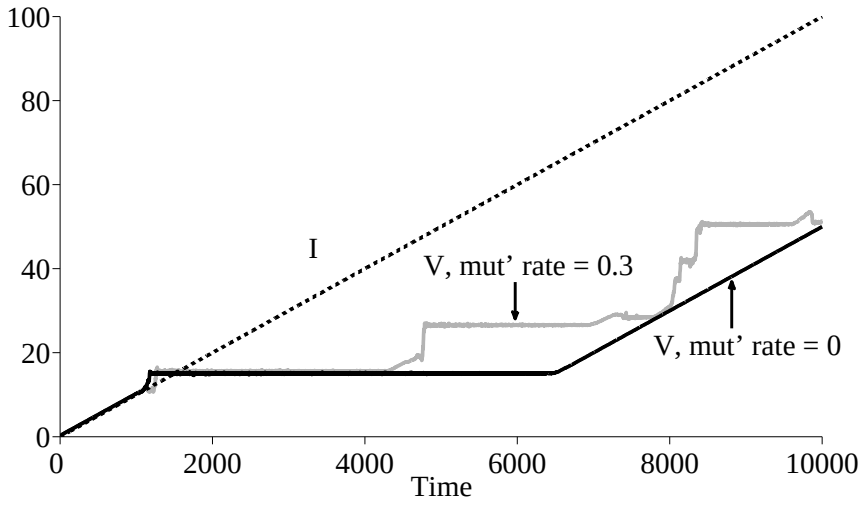

Figure 7.13: Homeostasis For Two Different Mutation Rates

This figure compliments figure 7.12 in that is shows results from the same two simulations. The perturbing input, I, is plotted with a dashed black line and the environmental variable, V, is plotted with a solid black line. The initial homeostatic state with no mutation lasts longer than with a mutation rate of 0.3. However, total homeostasis is greater with a mutation rate of 0.3, as subsequent homeostatic states are established. The establishment of these states is reflected in the sharp spikes in diversity in plot figure 7.12.

The process whereby particular homeostatic states are established, collapse and are then established again, but at under different environmental conditions, can occur in the absence of any change in the perturbing input. This highlights the essential non-directional nature of the long-term evolution of the system. Figures 7.14 - 7.16 show the continual evolution of steady states over 5 million time steps in MGDW when the perturbing input is initialised and fixed at 40. The environmental variable, organism effects and mean B allele undergo a series of sharp changes that are characterised by sudden collapses in mean fitness.

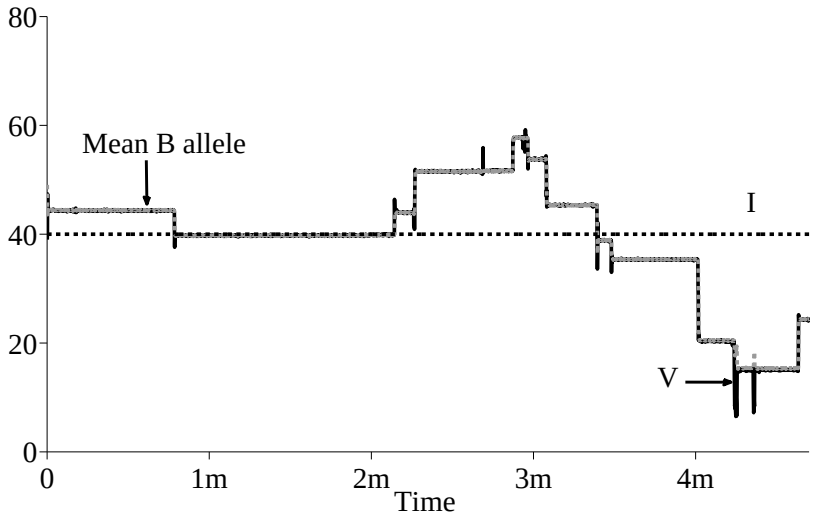

Figure 7.14: Multiple Homeostatic States With Fixed I

Results are shown for the following parameter set: population size, $k = 10,000$, $e = E = -0.01$, mutation rate = 0.1, perturbing input, I, fixed at 40 for 5 million time steps. The perturbing input, I, is plotted with a dashed black line. The environmental variable, V, is plotted with a solid black line. The mean B allele value of the population is plotted with a dashed grey line. This is very close in value to V and so the two lines are effectively superimposed and form a single line. Even when there is no change in the perturbing input, the system can move from different steady states. Over longer timescales the behavour of MGDW begins to approximate a random walk that is bounded by the lower and upper limits of the essential range.

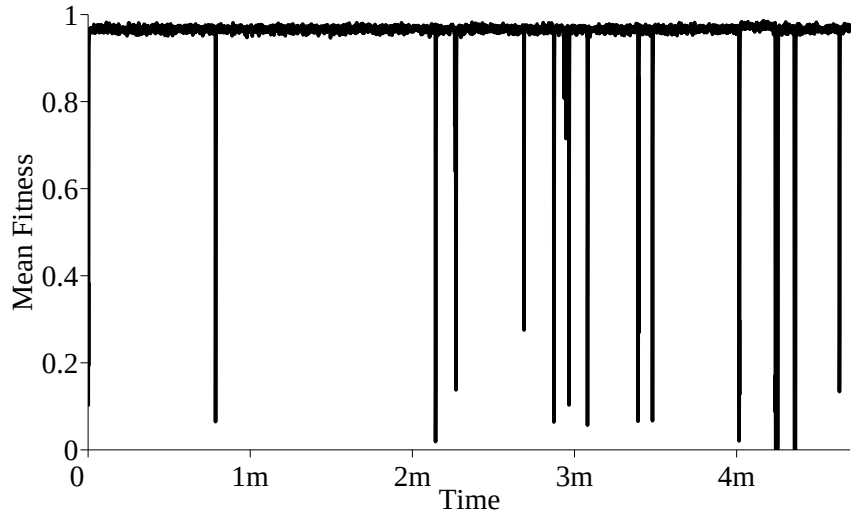

Figure 7.15: Multiple Homeostatic States With Fixed I - Mean Fitness
This figure compliments figure 7.14 in that the mean fitness of the population is shown for the same simulation. The transition to new stable states is marked by rapid decreases then increases in mean fitness.

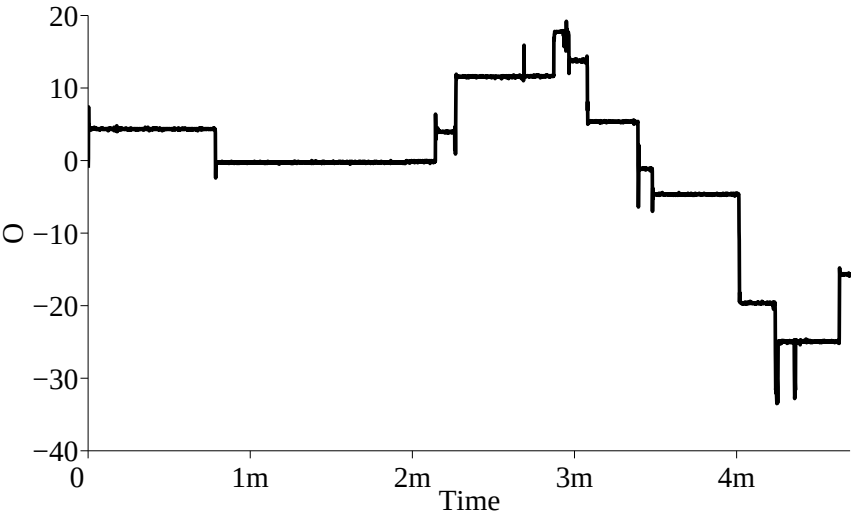

Figure 7.16: Multiple Homeostatic States With Fixed I - Organism Effects
This figure compliments figure 7.14 in that the organism effects are shown for the same simulation.

7.6 Complexity and Stability in Ecosystems

The relationship between mutation rates and homeostasis in MGDW may shed light on the relationship between complexity and stability within theoretical ecology. However, care must be taken when casting the MGDW results in such terms because diversity was increased by mutation and decreased by selection. In ecological studies, diversity may be altered by migration or invasion rates or other factors. These can be considered as quite distinct from evolutionary processes which will typically occur over longer time scales. However, if we interpret the mutation rate as a measure of migration rate in MGDW, equate the number of distinct B allele types with complexity and equate homeostasis with ecosystem stability, then the results of this chapter lend support to the hypothesis that more complex ecosystems (in terms of higher species diversity) will be more stable.

Prior to the 1970s, ecologists generally held the view that the more complex an ecosystem, the more stable it would tend to be [Odum, 1953], [Elton, 1958]. These conclusions were based on the observation that terrestrial ecosystems that had undergone significant simplification (e.g. alien species invasion or human monoculture), were more prone to pronounced fluctuations in population density. This reasoning was challenged by the theoretical study of May [May, 1972] who, building on earlier work by Gardner and Ashby [Gardner and Ashby, 1970], showed that the greater the number and strength of connections between species in a model ecosystem, the less stable it is likely to be. It is important to note that these two studies assumed linear interactions between species. The MGDW model is not compatible with this type of analysis.

Notwithstanding the details of these particular theoretical studies, the relationship between complexity and stability has proved to be more subtle and elusive with the debate being conducted in two broad strands: a search for a general relationship between diversity and stability, and an investigation into the relationship between food web structure and stability. See [McCann, 2000] for a review. The stability of ecosystems can be assessed in terms of constancy, persistence, resistance, inertia and resilience. The results presented in section 7.5 support the hypothesis that increasing ecosystem complexity reduces resistance (the ability to resist external pertubations), but increases resilience (the ability to recover after perturbations). These changes can result in an overall increase in the stability of the ecosystem as measured by its ability to regulate an environmental variable. Complexity in an ecosystems context can be defined and measured in a number of ways. For example, increasing complexity may equate

to increasing the number and type of ecological interactions. Given that increasing diversity in MGDW increases the number of interactions between organisms and their environment, this is arguably a reasonable interpretation of ecosystems complexity.

The results of this chapter support the conclusions of [Lovelock, 1992] and [Harding, 1999] that in order for an ecosystem (a multi-species Daisyworld) to be able to counteract external perturbations, there must be sufficient biodiversity. Lovelock showed that the limits of stability are characterised with decreases in species diversity. Both Lovelock's study and MGDW feature the critical component of feedback between organisms and their environment. Such feedback is captured within the notion of *ecosystem engineering* [Jones et al., 1994], [Jones et al., 1997], [Gutierrez and Jones, 2006]. Jones et al classify two types of ecosystem engineering. Allogenic engineering involves the modification of the environment by the actions of ecosystem engineers. The classic allogenic engineer is the beaver which significantly alters its environment with dam building and in doing so has profound effects on species composition and ecosystem functioning [Wright et al., 2002]. Autogenic engineers alter their environments as a consequence of their development. For example, trees provide a range of habitats and ecological niches within their growing roots, trunk branches and bark. The concept of ecosystem engineering is similar to that of niche construction in that they both share an appreciation of the effects that organisms can have on other organisms via their effects on shared environments. Rather than only ever responding to environmental change, organisms can initiate changes that effectively counteract perturbing inputs. This opens the possibility for new techniques that will predict ecosystem's responses to external perturbations. A finding of chapter 5 was that when dealing with antagonistic rein control systems, the particular details of how rein controller elements respond to changes in the system variable may be irrelevant to determining the steady state solutions. Similarly with ecosystem functioning, if it were possible to identify the broad direction and strength of particular species' effects on their environment, then it may be possible to assemble a coarse-grained feedback circuit that could be analysed in terms of rein control and so provide predictions for how the system will respond to changes in perturbing input.

7.7 Discussion

This chapter continued the examination of McDonald-Gibson Daisyworld (MGDW) with particular focus on the effects that mutation produced on population diversity and homeostasis. It was shown that increasing the mutation rate can decrease the length of any particular homeostatic state, but increase the probability that subsequent states are generated. This can lead to an increase in the total amount of time that the environmental variable is regulated. The sharp transitions between steady states were produced by positive feedback type i and type iv rein controllers. This positive feedback ejected the population away from potentially absorbing states which would have resulted in the cessation of the population's ability to regulate against perturbing inputs.

The relationship between mutation rate and homeostasis in MGDW was discussed within the context of the complexity - stability debate within theoretical ecology. Increasing mutation rates in MGDW increases the system's complexity in terms of population diversity, while reducing the system's resistance (the ability to resist external pertubations). However, this also increases the system's resilience (the ability to recover after perturbations) such that the total values for stability are increased for a particular range of mutation rates. Increasing the mutation rate increased the diversity of the population by seeding individuals outside of the regulating sub-populations. It would be possible to vary the diversity of the population by means other than mutation. For example, the influx of alien species could be modelled with the random replacement of individuals. However, given the very simple individuals modelled in MGDW this, with certain limitations, could be considered as equivalent to the mutation of existing individuals. The previous chapter made a number of connections between MGDW and the notion of niche construction. In this chapter the notion of ecosystem engineering was discussed. Ecosystems engineering seeks to capture the effects that organisms can have on their environment as such effects can be important for species composition and stability of ecosystems.

It was shown that MGDW exhibited multiple homeostatic states. The co-evolving system comprised of life and its environment produced a series of steady states. The periods between these steady states was characterised by large fluctuations in the frequency of alleles in the population and sharp changes in the environment. Over longer timescales the behaviour of MGDW with fixed perturbing input was similar to that of a bounded random walk. This demonstrated the assumption that the organisms were

free to evolve responses to the environmental variable and that there was no notion of a predefined optimal environment. In that respect, the concept of homeorhesis may be more appropriate to describe the behaviour of the model, as the population is never at rest but travelling down a particular trajectory. The system can be ejected from any particular trajectory, or in the terminology of homeorhesis, a particular chreod. Positive feedback leads to the rapid transition to new trajectories, new chreods, and the system proceeds again. The perturbing input can now be understood as an additional 'push' to the system which can move the system up and out of one attractor and into another, and so can accelerate the process of destruction and creation of trajectories. However, such behaviour occurs in the absence of any perturbing input and is a fundamental aspect of the model when mutation is operating.

The sharp transitions and multiple homeostatic states observed in MGDW are reminiscent of interpretations of the fossil record that have been used to support the theory of punctuated equilibrium [Eldredge and Gould, 1972]. Punctuated equilibrium may be defined in a number of ways or be of varying strengths. Its arguably most neutral form is the observation that there are relatively long periods of stasis in the fossil record that are punctuated with relatively rapid periods of evolutionary change. A stronger notion of punctuated equilibrium is that this represents a significant aspect of natural evolution which proceeds in relatively short bursts of change and much longer periods of stasis. The theory of punctuated equilibrium has proved to be controversial. For example, Dawkins effectively dismisses the theory in [Dawkins, 1986]. while I will not comment on how punctuated equilibrium may or may not challenge the gradualist notion of evolution, I will point out that the analysis of this and the previous chapter may provide a mechanism that could explain aspects of punctuated equilibrium. Rapid periods of evolution may be produced by rapid biologically-mediated changes to the environment. Such environmental changes will produce a range of changes in the direction and strength of selection operating not only on those species that produced the initial change, but a potentially wide range of species that may be distantly connected to the initial species, both in space and time. It is now well established that the effects life has had on the oceans, atmosphere and crust of the Earth have had profound effects for the evolution of life. [Lenton et al., 2004] sought to correlate the series of major evolutionary transitions of the biota with transitions in abiotic components of the biosphere. An emphasis was placed on regarding the biosphere as a co-evolving system. Once it is appreciated that life can have profound effects on its environment, then life becomes

an active partner rather than only ever responding to abiotic changes.

Contemporary research into the possible implications of anthropogenic carbon emissions contains the central assumption that species X (in this case Homo sapiens) may produce significant changes to not only future generations of species X, but a much larger set of of species. Recent research on 'tipping elements' attempts to quantify the probability that certain ecosystems may undergo effectively irreversible collapse and in doing so produce a change in the strength and even sign of their feedback effects on other ecosystems and the biosphere as a whole [Lenton et al., 2008]. This can be understood as previously homeostatic type control circuits 'flipping' to positive feedback states as external perturbing inputs drive the environmental variable outside the essential range of those organisms that alter the variable. In chapter 6 I explained the 'trick' of MGDW by explaining the change in the sign of feedback produced by unimodal fitness functions. Changes to the environmental variable can lead to increases and then decreases in the fitness of any particular organism. This same trick can lead to the establishment of positive feedback and the rapid collapse of a particular homeostatic state. In the simple MGDW model, new homeostatic states could be established if there was sufficient diversity of e and E alleles in the population. For real-world ecosystems that have a global impact it is not clear at all how such significant changes in forcing will affect other ecosystems and the rest of the biosphere.

This chapter concludes the examination of MGDW and evaluating the Daisyworld control system in a evolutionary and ecological context. In the following chapter I will change the domain of enquiry to thermodynamics and assess the Daisyworld control system in the light of a proposed maximisation principle.

Chapter 8

Entropy Production and Daisyworld

8.1 Previous Publications

The main results from this chapter are reproduced from [Dyke, 2008]. Parts of the review of previous Daisyworld studies on entropy production are taken from my contributions to [Wood et al., 2008].

8.2 Overview

In this penultimate chapter of the thesis, I will explore the effects of the thermodynamic principle of maximum entropy production (MEP) on homeostasis within Daisyworld. I will consider the hypothesis that planetary homeostasis may be the consequence of negative feedback processes that arise from the maximisation of entropy production. This marks a departure from the previous two chapters' consideration of the Daisyworld control system within an evolutionary and ecological context. However, towards the end of this chapter I will continue the discussion from chapter 7 which highlighted the essentially non-directional evolution of the McDonald-Gibson Daisyworld model (MGDW). It was found that over suitably long time scales, the behaviour of MGDW was similar to a random walk. The system co-evolved to maintain environmental conditions at different levels over the entire essential range. In terms of homeorhesis, there appeared to be a practically infinite number of chreods over which the system could evolve. I will discuss this in terms of the evolution of the Earth's biosphere.

It has been proposed that thermodynamics confers a sense of directionality for the the long-term evolution of the Earth's biosphere [Kleidon, 2004]. This can be considered

true in a 'trivial' sense: a biosphere, just like any other physical system must adhere to the laws of thermodynamics which includes the Second Law and associated changes in entropy. However there is a more subtle and potentially profound implication: a biosphere is a co-evolving system comprised of environmental and biological components that self-organises itself to be in states as far from thermodynamic equilibrium as possible. I will use a version of the Two Box Daisyworld (TBDW) to explore the relationship between entropy production and homeostasis. This approach is novel in that it allows natural formulations and quantitative assessments of entropy production and homeostasis.

There are numerous mechanisms and processes that produce entropy in the Earth's climate. In the following sections I will define entropy, entropy production and briefly review the MEP principle. There have been several studies on the effects of entropy production in Daisyworld and I will summarise these while pointing out their respective limitations. I will show that the TBDW model is ideally placed to explore the effects of entropy production in Daisyworld. TBDW is conceptually similar to energy-balance climate models that have been used to assess the impacts of maximising rates of entropy production in the Earth's and other planets' climates. The significant difference between TBDW and other energy-balance models is that the temperature gradient in energy-balance models is produced by different amounts of energy being received at different latitudes on the surface of spherical planets. The poles are cooler than equatorial regions on the Earth as they receive much less energy from the Sun. In TBDW, temperature gradients are produced by the presence of black and white daisies. Consequently, in TBDW we are able to explore the role that life has on the energy budget of an idealised climate system. I will show that when values for heat diffusion in TBDW are selected in order to maximize the rate of entropy production, the any-daisy range (the range of luminosity over which black or white daisies grow) and the both-daisy range (the range of luminosity over which black and white daisies coexist) are maximized. Consequently, planetary temperature is regulated over the widest possible range of solar forcing.

8.3 Entropy

8.3.1 Equilibrium Systems

Entropy is a central concept within classical thermodynamics, statistical mechanics and information theory. The concept arose during the examination of heat engines in the

18th and 19th centuries. A coal-fired steam engine is a heat engine. Heat flows from a hot reservoir (a boiler) to a cold reservoir (the outside world) and this flux of heat is able to do work such as drive a piston. The greater the temperature gradient between the two reservoirs, the greater the amount of work that can be done by the system. However, not all heat can be translated into work. Friction will reduce the output of a steam engine while heat will radiate 'uselessly' into the outside world. Entropy can then be understood as the amount of energy in a system that is unable to do 'external work' or 'useful work'.

The first law of thermodynamics states that energy can be neither created or destroyed. The second law employs the notion of entropy to keep account of the energy within systems. Entropy increases as the amount of energy not available for work increases. In classical thermodynamics, entropy is not explicitly defined, but rather can be expressed in terms of how it changes as the amount of heat and temperature of a system changes:

$$\frac{dS}{dt} = \frac{Q}{T} \tag{8.1}$$

This tells us that the rate of change of entropy, S, over time increases as the amount of heat, Q, increases over temperature T. A large input of heat into a system that does not increase in temperature would produce a large increase in the entropy of the system. Entropy can be seen as a 'bookkeeping' factor as the energy that was added to the system in the form of heat must be somewhere.

It was the statistical formulation of entropy that shed light on where this 'missing' energy was within such systems. Statistical thermodynamics defines entropy as being proportional to the logarithm of the number of microstates that are consistent with a macroscopic description of a system.

$$S = k_B \ln \Omega \tag{8.2}$$

where k_B is Boltzmann's Constant and Ω is the number of the microstates that correspond to the observed macroscopic description. Increasing the entropy of a system, in the words of J. W. Gibbs, increases the amount of 'mixupedness' of the system. While entropy is not necessarily synonymous with disorder or randomness, both disorder and randomness will typically increase with an increase in entropy. The power of statistical mechanics derives from the typically very large values for Ω. For example, an average living room may contain many billions of air molecules. The probability for all these randomly moving molecules to move to one corner of the room, whilst non-zero, is

extremely small. Probabilistic estimations will become increasingly accurate because the variance from the mean decreases as the number of possible microstates increases. With many systems, the deviations become so small as to be effectively zero and one can ignore the configuration of the particular atoms and instead deal with a system-level property such as entropy.

A consequence of both classical and statistical formulations of entropy is that in a closed system, it can only ever increase. When a heat engine performs work, the amount of energy produced is never greater than the amount of energy that was consumed. $dS/dt \geq 0$. If a closed system consists of a box of two equal partitions and one partition contains n gas molecules while the other partition is empty, removing the partition will lead to an increase in the entropy of the system as the number of possible microstates has increased. At equilibrium we would expect the greatest value for entropy, with a variance from this value that is proportional to the number of possible microstates.

8.3.2 Non-Equilibrium Systems

Classical and statistical thermodynamics have proved to be enormously powerful tools for the analysis of physical systems. However, the discussion so far has been limited to closed systems at equilibrium. Many systems will be materially and/or energetically open and could be far from equilibrium. The Earth's climate is an example of such a system. It is energetically open as it receives shortwave radiation from the Sun which is ultimately radiated back out into space as longwave radiation. There is an imbalance in the amount of energy received on the surface of the Earth from the Sun. This imbalance of energy across the boundary of the Earth's climate sets up motions in the atmosphere and oceans. These motions transport heat from the hot equator to the cold poles and in that respect the Earth's hydrosphere can be viewed as a heat engine. Just like any other heat engine, the amount of work that can be done by this flow of heat depends on the temperatures of the reservoirs: the greater the drop in temperature for a given amount of heat flow, the greater the thermodynamic efficiency of the system and so the greater the amount of work produced. This amount of work is synonymous with the rate at which the system produces entropy:

$$\frac{dS}{dt} = \frac{F}{T_p} - \frac{F}{T_e} \tag{8.3}$$

where T_e is the temperature of the equator and T_p is the temperature of the poles. F is the flux of heat from equator to poles and is found with:

$$F = 2D(T_e - T_p) \tag{8.4}$$

where D is a diffusivity parameter that determines how easily the oceans and atmosphere transport heat. This definition for heat flux is essentially the same as that employed in TBDW and is based on [Lorenz et al., 2001].

Work can be produced by other forms of energy gradients. Figure 8.1 shows the work that can be performed by dissipating energy across a gradient produced by gravity. An important difference between systems that have fixed energy gradients and the Earth's climate is that an increase in the rate of energy dissipation in the atmosphere by an increase in the rate of heat flow can decrease the energy gradient and so the rate of entropy production. As more heat is transported, the temperature of the equator will go down and the temperature of the poles will go up. With an infinite rate of heat flow, the temperature of the poles will equal that of the equator and no entropy will be produced. With perfect insulation, no heat flows, the temperature gradient is maximised and so the rate of entropy production is again zero. There will be a rate of energy dissipation that produces that maximum amount of entropy.

It has been proposed that the Earth's climate maintains steady states in which the production of entropy is maximized as there are sufficient degrees of freedom associated with the atmosphere to allow it to 'select' the maximising entropy state. This is the principle of maximum entropy production (MEP). MEP in the Earth's climate was first proposed by Garth Paltridge who applied the MEP principle in order to solve energy-balance models of the Earth's climate [Paltridge, 1975]. Paltridge was faced with the problem of continually having one more unknown term than equation and consequently being unable to produce exact algebraic solutions. He was aware of the work of Prigogine and the extremum principle of minimum entropy production [Prigogine, 1962]. Paltridge assumed that his climate model's solution would, given the boundary conditions such as the amount of energy received on the surface of the planet at different latitudes, maximise the rate of entropy production via latitudinal transport of heat from the equator to the poles. This allowed Paltridge to solve the model and make predictions for conditions for the Earth's climate at different latitudes. These predictions were remarkably close to empirical data [Paltridge, 1975]. However, in the absence of any proposed mechanism or demonstration of other systems that maximized

Figure 8.1: Cheese Rolling: A Demonstration of Energy Gradients
Every spring in Gloucestershire, UK, large wheels of cheese are rolled down a hill with a 1 in 2 gradient. The tumbling crowd that follows the cheeses do so on the promise of winning the cheese if they are first to the bottom of the hill. Catching the cheese before it reaches the bottom of the hill is a rare event as the cheese can reach speeds in excess of 60mph. This is a useful demonstration of work being performed by the dissipation of energy across a gradient. The steeper the hill and the faster the competitors hurl themselves down it, the greater the amount of work that is produced (which is sometimes used to dislocate limbs or break bones).

entropy production in this manner, Paltridge's results failed to gain significant traction within the scientific community. This situation has changed somewhat as MEP has been observed in the atmospheres of Mars and Titan. [Lorenz et al., 2001] accurately predicts equatorial and polar temperatures on Mars and Titan by assuming that rates of latitudinal heat flux would be those that maximise the rate of entropy production. The models used by Lorenz et al are strikingly simple yet they perform much better than more complex and tuned models. Lorenz later applied the simple MEP models to examine entropy production within the convective medium of planetary mantles [Lorenz, 2002]. Why such systems would select maximum entropy states is subject to debate. Roderick Dewar has derived MEP from an information theoretic formulation of statistical mechanics [Dewar, 2003] [Dewar, 2005]. Dewar employs Jayne's technique of MaxEnt [Jaynes, 1957a] [Jaynes, 1957b] and attempts to show that maximum entropy production states represent those states that are more probable when considering the time evolution of the system. Dewar visualises the possible trajectories of MEP systems as a bundle of fibres with MEP states being achievable with the greatest proportion of fibres. This, in conjunction with the empirical studies discussed previously, has increased the 'respectability' of the principle of maximum entropy production. See [Ozawa et al., 2003] and [Martyushev and Seleznev, 2006] for reviews.

In this chapter I will use the MEP principle as a 'black box' mechanism. I will assume that the modelled climate system has sufficient degrees of freedom to configure itself into a state in which the rate of entropy production is maximised. In that respect, I will use the MEP principle as a selection mechanism. I will select certain parameter values on the basis of their effects on the rate of entropy production.

8.4 Previous Daisyworld Studies

The method of parameterizing heat flow in the original Daisyworld model (the mechanism by which heat moves from the black daisies into the bare ground and then into the white daisies) is fundamentally the same as that employed in energy balance climate models e.g. [North et al., 1981]. The temperature difference between the black and white daisies and the bare ground will be determined by the difference in albedo and the diffusivity of the atmosphere. Simple models that reproduce the Earth's current climate operate on similar lines in that diffusivity will determine how much heat is transported from the hot equator to the cold poles. The actual mechanisms for this

heat transport are complex as they comprise oceanic and atmospheric turbulent forces. The particular utility of this modelling approach is that such details are effectively ignored and are wrapped up into a single parameter that determines the overall amount of diffusivity of the atmosphere at steady state.

Previous studies have investigated entropy production in Daisyworld: [Pujol, 2002] within the original version, [Tonaizzo et al., 2004] with a version that allowed an arbitrary number of daisy types and a two-dimensional cellular automata version in [Ackland, 2004]. In general they pose the same question: what are the implications for Daisyworld (and in particular temperature regulation) if the system were to maximise the rate of entropy production? The first two studies faced certain limitations due to the particular implementation of heat flux in the respective models. However they both concluded that when Daisyworld maximizes the rate of entropy production, there is an increase in the range of luminosity over which daisies grow. Rather than assuming Daisyworld is in a maximizing entropy production state, [Ackland, 2004] tests the contrary hypothesis that Daisyworld self-organises to either maximize entropy production or maximize the total amount of life for any given luminosity (a maximum life principle). It is found that a maximising life not maximising entropy principle is selected. However, due to the modelling assumptions of cellular automata Daisyworlds, computing the rate of entropy production via heat flux is not possible. Consequently it is the rate of biodiversity entropy that Ackland measures.

It is instructive to highlight two issues these studies encounter when assessing Daisyworld in terms of MEP. First, previous MEP studies into planetary entropy production, [Lorenz et al., 2001] and [Ozawa et al., 2003] assume a constant flow of energy into and out of the system. Daisyworld, with its varying daisy coverage and so varying albedo does not have this constant flow of energy into and out of the system. While it may be argued that energy flux through the system is at a steady state when the daisy coverage is at steady state, it is necessary to note this potentially important difference with previous MEP studies. Secondly, heat diffusion in the Daisyworld models presented in [Pujol, 2002], [Ackland, 2004] and [Tonaizzo et al., 2004] is fixed. The flow of heat is adjusted by varying the daisy coverage. This is in contrast with other studies into MEP where it is the diffusivity parameter that is adjusted in order to produce conditions in which maximum entropy production is produced. The model I will develop in this chapter allows adjustment of the diffusivity parameter.

8.5 Entropy production in TBDW

I will use a version of the Harvey TBDW model presented in chapter 4 with the difference that, as in Harvey's original model, the coverage of the black and white daisies is reset to 100% for every value of luminosity and the equations are numerically integrated to steady state. This removes hysteresis from the model as I discussed in chapter 4. This version of the model is used in this chapter as it allows clearer presentation of the results. Qualitatively similar results are returned when the coverage of the daisies is not reset for each value of luminosity. The temperature of each daisy box is found with:

$$\sigma T_b^4 = I(1 - A_b) - F \tag{8.5}$$

$$\sigma T_w^4 = I(1 - A_w) + F \tag{8.6}$$

where A_i is the albedo of the boxes and I is insolation. These quartic equations employ the Stefan-Boltzmann constant, σ, and so are physically realistic in the sense that the balance between longwave and shortwave radiation is accurately represented[1]. As in O/SDW the amount of insolation will be parameterized by a luminosity variable, L, so that $I = LR$ where R (in previous chapters this was denoted by S) is the amount of energy emitted from the star and is received on the surface of the planet and is fixed at 952.56; and $L \in [0, 2]$. F is the heat flux between the two boxes and is found with:

$$F = D(T_b - T_w) \tag{8.7}$$

The flux of heat is proportional to the temperature difference between the two boxes and a diffusion parameter, $D \in [0, 1]$. When $D = 0$ the two boxes are perfectly insulated and no heat flows. When $D = 1$ there is maximum heat flow and the two boxes are isothermal. The temperature of the planet is found with:

$$\sigma T_p^4 = 0.5I(A_b + A_w) \tag{8.8}$$

Harvey TBDW dispenses with fixed death rate and birth rate functions and instead employs a single 'hat' function that determines the steady state coverage of the daisies

[1]The use of these equations, rather than the simplified versions used in TBDW previously was initially motivated in order to make the model more consistent with energy-balance climate models. The same relationship between entropy production and homeostasis that is presented in this chapter is observed when using simplified temperature equations.

and has the same form of the birth rate function in TBDW:

$$\alpha_i = H(T_i) = max\left\{\ 1 - 2^{\frac{|T_o - T_i|}{R}}, 0\ \right\} \tag{8.9}$$

It is now no longer necessary to define a variable such as the amount of bare ground that produces density dependent growth in TBDW, as the hat function itself limits the maximum coverage of the daisies. A pair of differential equations are implicit in the model and it is these that are numerically integrated to steady state:

$$\frac{d\alpha_i}{dt} = H(T_i) - \alpha_i \tag{8.10}$$

The method of integration can be understood as a variant of the Euler feed forward method. The algorithm is detailed below.

1. Set luminosity.
2. Initialise coverage of black and white daisy bed to 100%.
3. Calculate box albedos from black and white daisy proportional coverage.
4. Calculate box temperatures.
5. Calculate heat flux F.
6. Adjust each box temperature by $\pm\ F$.
7. Calculate new daisy proportional coverage, α_i'.
8. Adjust proportional daisy coverage: $\alpha_i = \alpha_i(1 - \delta) + \delta\alpha_i'$.
9. Go to step 3 until changes to variables are sufficiently 'small' (defined below).
10. Increase luminosity.
11. Go to step 2.

When $\delta = 0.001$, 200,000 iterations produce changes in coverage, albedo, heat flux and temperature that are no greater than 10^{-22}. It is assumed, as with other studies on Daisyworld, that the rate of change of daisy coverage is sufficiently faster than that of luminosity so as to allow the above equations to be numerically integrated to steady state whilst luminosity is fixed.

In energy-balance climate models, temperature gradients are established at different latitudes. In TBDW the difference in temperature is due to different albedos. However, heat will flow from warm to cool regions irrespective of how such a situation was produced and the method of calculating entropy production budgets in energy balance box models such as [Lorenz et al., 2001] can be naturally employed in order to calculate entropy production in TBDW. The rate of entropy production in TBDW is a function of heat flux over the difference in temperature between the hot and cold boxes:

$$\frac{dS}{dt} = \frac{F}{T_w} - \frac{F}{T_b}. \tag{8.11}$$

The greatest rate of entropy production will be achieved with the greatest temperature difference and the greatest heat flux between the two daisy boxes. Attempting to increase entropy production by increasing heat flux via increasing diffusion may lead to a decrease in the temperature gradient and thus a decrease in the thermal efficiency of the 'heat engine' and so a decrease in the rate of entropy production. Consequently maximizing entropy production is a balancing act with the value of D required to produce maximum rates of entropy production varying with the driving of the system. Figure 8.2 shows planetary temperature and figure 8.3 shows daisy coverage with two fixed values for D.

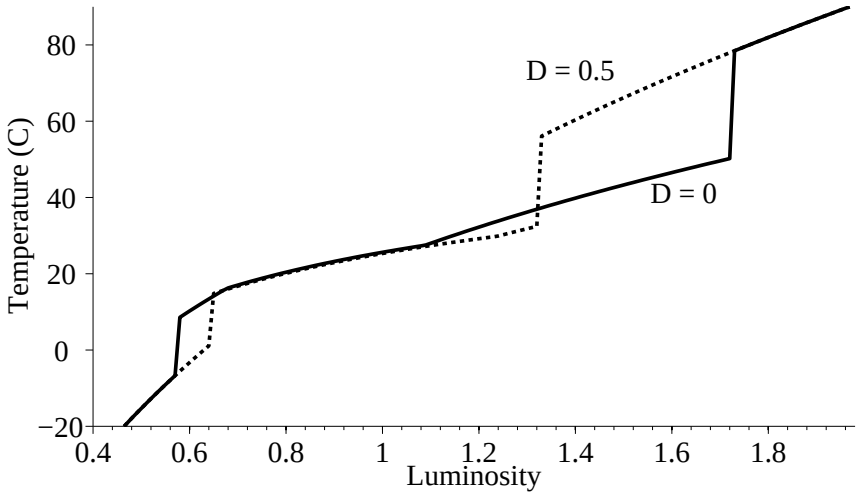

Figure 8.2: Planetary Temperature With Two Values of Fixed Diffusion
Planetary temperature when the diffusion parameter, $D = 0$, is plotted with a solid black line and when $D = 0.5$, is plotted in a dashed black line.

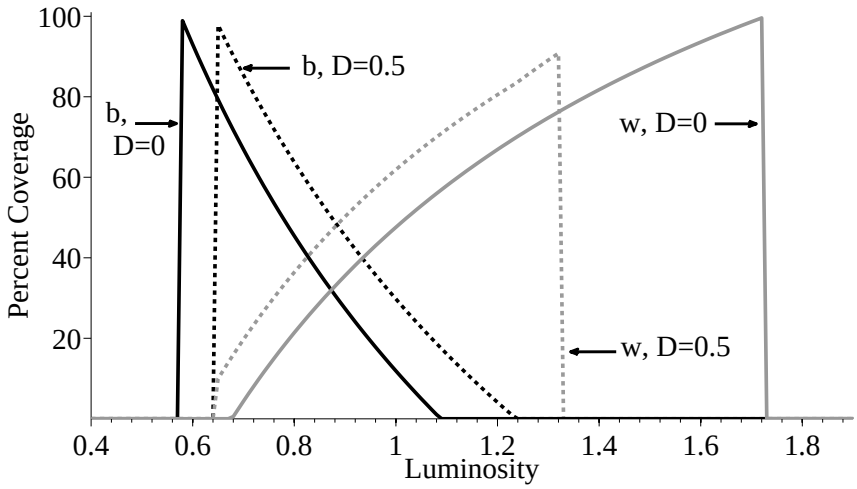

Figure 8.3: Daisy Coverage With Two Values of Fixed Diffusion
This figure accompanies figure 8.2 in that the same simulation results are shown. The coverage of the daisies when $D = 0$ is plotted with solid lines. The coverage of the daisies when $D = 0.5$ is plotted with dashed lines. The coverage of the black daisies is plotted with black lines. The coverage of the white daisies is plotted with grey lines. The any-daisy range (the range of luminosity over which black or white daisies grow) is greater when $D = 0$. The both-daisy range (the range of luminosity over which black and white daisies coexist) is greater when $D = 0.5$.

8.5.1 Maximising Entropy Production

It has been postulated that certain dissipative systems maximize the rate of entropy production. *How* these systems reach such non-equilibrium stable states is a different topic of enquiry. Here I will assume that the imaginary two-box planetary system, like the Earth, possesses an atmosphere with sufficient degrees of freedom to produce heat fluxes that lead to the maximisation of entropy production. In doing so we can construct a thought experiment analogous to Maxwell's Demon which is used to explore certain aspects of the second law of thermodynamics[2]. Figure 8.4 shows that rather than monitoring the velocity of air molecules, our demon monitors the flux of heat and temperature of the daisy boxes. It has a dial at its disposal that modulates the diffusivity of the atmosphere. For any luminosity value, the demon can alter the diffusion and in doing so find the rate of diffusion that maximizes equation 8.11. The behaviour of the demon can be implemented in an algorithm that modulates diffusion for any fixed luminosity in order to find the value of diffusion that produces the highest rate of entropy production.

The effects of the maximising demon are shown in figures 8.5 and 8.6. Luminosity is fixed at 0.7 and diffusivity is adjusted from 0 (no flux of heat) to 1 (maximum flux of heat). Entropy production increases as the flux of heat increases and the black and white daisy box temperatures move closer together. Maximum rates of entropy are produced when diffusivity is approximately 0.7. As luminosity varies, the value for diffusivity that produces maximum entropy production will also vary. There is no significance to the fact that when luminosity equals 0.7, maximum entropy production is achieved when diffusivity is approximately 0.7.

[2]I thank Inman Harvey for suggesting the maximising demon analogy.

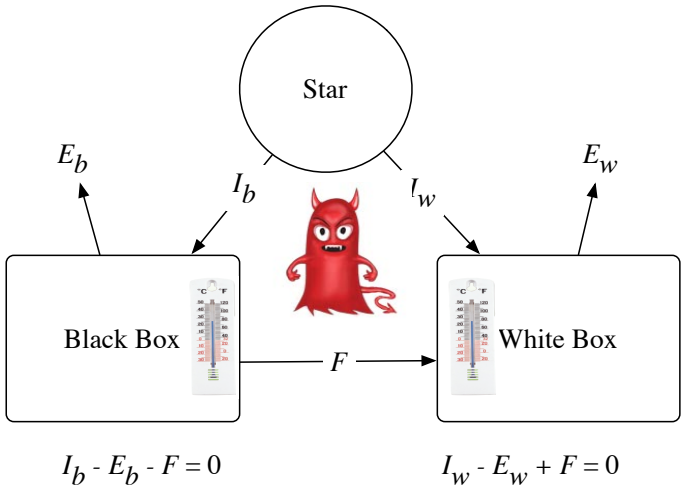

Figure 8.4: The Maximising Demon

A modified TBDW is shown. Between the black and white daisy boxes floats Maxwell's Demon. This demon can control the amount of heat flux, F, between each box by adjusting the diffusivity, D, of the material that connects the two boxes. The demon can see the two thermometers that show the temperatures of each box. The demon also knows the value for F at any given moment of time. The demon adjusts D so that the maximum rate of entropy is produced: $dS/dt = F(1/T_w - 1/T_b)$. Real-world MEP systems such as the Earth do not contain real demons and it has been proposed that they are in states of maximum entropy production due to these states being the most probable non-equilibrium steady states.

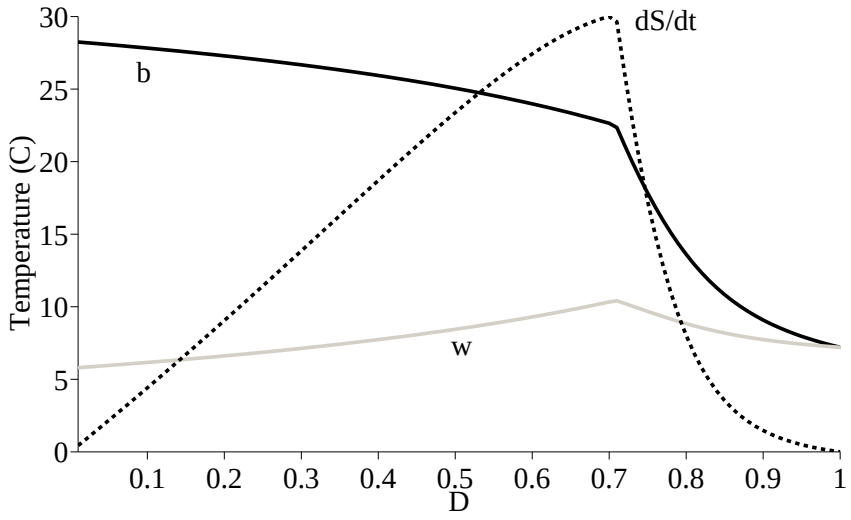

Figure 8.5: Adjusting Diffusivity to Maximise Entropy Production - Temperature The rate of entropy production, dS/dt for different values of diffusivity, D, when luminosity is fixed at 0.7 is plotted with a dashed black line. The temperature of the black daisies is plotted in a solid black line. The temperature of the white daisies is plotted in a solid grey line. The greatest rate of entropy production is produced when $D \approx 0.7$. Increasing diffusion leads to a decrease in the temperature difference between the boxes and entropy production until the daisy boxes are isothermal with no entropy production and steady state coverage being that of a 'grey' daisy type with albedo of 0.5. The value of D required to produce maximum entropy production will vary with L.

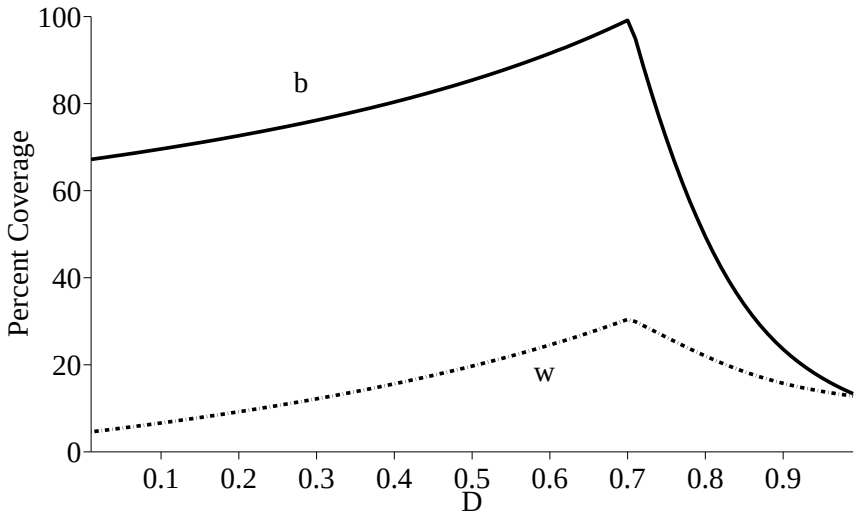

Figure 8.6: Adjusting Diffusivity to Maximise Entropy Production - Coverage
The coverage of the black daisies for various values of heat diffusion where luminosity is fixed at 0.7 is plotted with a solid black line, for white daisies with a dashed black line. Increasing diffusion leads to an increase in the coverage of black and white daisies until the coverage of daisies undergoes a sharp decline as increasing heat flux drives the black and white box temperatures away from the optimal temperature of 22.5C.

8.6 Results

Two experiments were conducted with the maximising demon. The first saw the demon maximise the rate of entropy production for luminosity values over the range [0.3,2]. The second experiment saw the demon maximise the total coverage of the daisies over the same range of luminosity. I will show that this produced identical widths for the any-daisy and both-daisy ranges. On that basis, maximising the rate of entropy production maximises homeostasis because in TBDW planetary temperature is regulated with either daisy type present. The rate of change of planetary temperature is least when both daisies are present. In that respect maximising homeostasis in TBDW is equivalent to maximising the any-daisy and both-daisy ranges.

It is straightforward to show that maximising entropy production will lead to a maximisation of the any-daisy range. In order for there to be non-zero entropy production there must be a temperature gradient between the two daisy boxes. Therefore there

must be non-zero coverage in either or both boxes. This equates to black daisies growing at the lowest possible luminosity and white daisies growing at the highest possible luminosity. In order for black daisies to grow when the planet is cool, the amount of heat flux must be reduced. Increasing heat flux cools the black daisies and so a greater amount of luminosity is required to warm the black daisies. The same applies to white daisies at higher luminosity. Increasing heat flux would increase the temperature of the white daisies and so lead to the maximum birth rate temperature being reached with lower amounts of luminosity. Consequently, at the limits of the any-daisy range diffusivity needs to be as close to zero as possible in order to maximise the any-daisy range, but not actually zero as this would not produce entropy.

The limits of the any-daisy range are also when the daisies reach maximum coverage. The analysis of the Simplified Daisyworld model in chapter 3 and TBDW in chapter 4 showed how luminosity drives the daisy coverage higher until the biota can no longer respond, at which point any further increases lead to a population collapse. The daisy box temperature that produces maximum coverage in this chapter's model is 22.5°C (295°K). Setting $D = 0$ enables us to find the luminosity for the start and end of the any-daisy range.

$$L_{start} = \frac{295.5^4\sigma}{R(1-0.25)} \tag{8.12}$$

$$L_{end} = \frac{295.5^4\sigma}{R(1-0.75)} \tag{8.13}$$

With the parameter and constant values as detailed in section 8.5, the values of luminosity at the start and end of the any-daisy range are 0.5674 and 1.729 respectively (to within 4 significant figures). These are the values returned with numerical results as shown in table 8.1.

Finding the limits of the both-daisy range is not a trivial exercise. At lower luminosities we want to find that combination of luminosity and heat flux that increases the white daisy box to 5°C. This will be achieved at lower luminosities with higher heat flux. But if heat flux is too high, the temperature of the black daisy box can be cooled so far as to lead to a collapse in the black daisy population. Similarly we want to find the greatest heat flux that can maintain the black daisy box to within 40°C without increasing the temperature of the white daisy box beyond 22.5°C and so causing its collapse. Table 8.1 confirms the hypothesis that maximising the both-daisy range is achieved when entropy production is maximized. The both-daisy range for both the maximum entropy and maximum life simulations starts when luminosity reaches 0.6285

	Maximum Entropy	Maximum Life
Any-Daisy Start	0.5674	0.5674
Any-Daisy End	1.729	1.729
Both-Daisy Start	0.6285	0.6285
Both-Daisy End	1.283	1.283

Table 8.1: Maximum Entropy and Maximum Life Results
The results of two simulations are shown: Maximum Entropy - diffusivity was selected in order to maximise the rate of entropy production for each value of luminosity; and Maximum Life - diffusivity was selected in order to maximise the total coverage of black and white daisies. The start and end of the any-daisy range (the range of luminosity over which either black or white daisies grow) and the both-daisy range (the range of luminosity over which both daisies coexist) are the same.

and ends when luminosity reaches 1.283 (to 4 significant figures). Figure 8.7 details the effects of the maximising demon on TBDW.

8.7 Discussion

We have seen that when diffusivity is altered in the two-box Daisyworld model in order to maximise the rate of entropy production via latitudinal heat flux, the range of luminosity over which daisies grow is maximised. I assumed that the rate of change of daisies was sufficiently faster than luminosity to allow luminosity to remain fixed whilst steady state coverage was found. I also assumed that diffusion was fixed whilst the system moved to steady state. A wide range of different model dynamics could be produced by relaxing or altering these assumptions. That said, there are immediate intuitive connections between self-regulation and entropy production. The system can only regulate temperature when daisies are present. Latitudinal heat flux entropy can only be produced when daisies are present. If the two-box Daisyworld atmosphere had sufficient degrees of freedom to maximize entropy production, then it would maximize the range of luminosity over which daisies can grow.

Care must be exercised when interpreting results from such simple models, especially in the absence of empirical data. The model presented in this chapter is based upon established climate models and the observed real world phenomenon of entropy maximisation via latitudinal heat flux. However, the maximisation of entropy production was *imposed* on the model by a notional demon. MEP models can be regarded as 'black box' models in that one need not know the details of the system's dynamics. Within

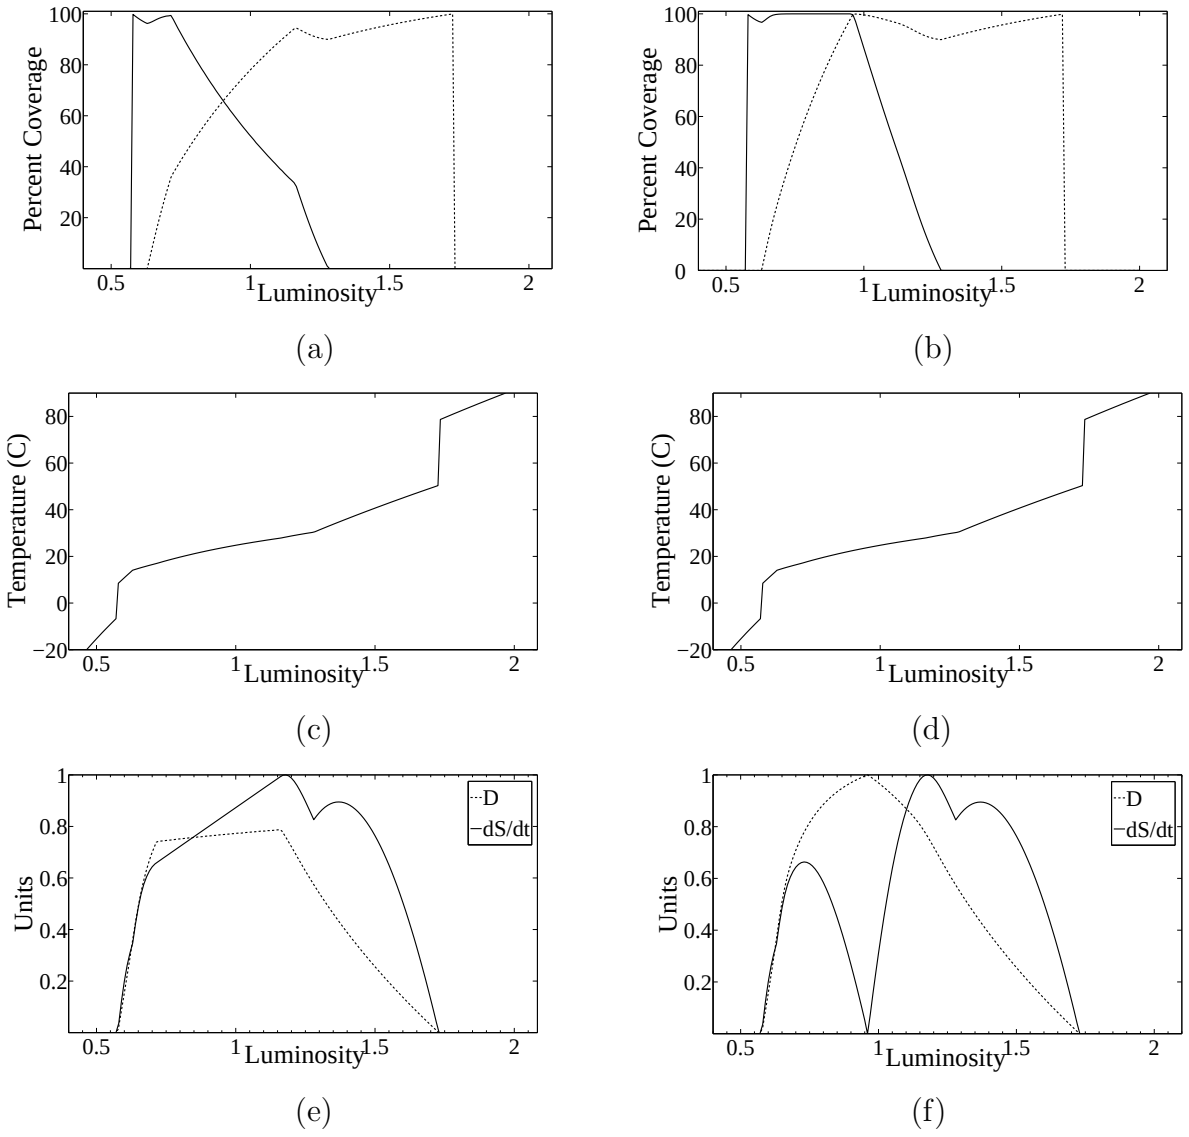

Figure 8.7: TBDW MEP results

Data for plots (a), (c) and (e) were produced by a simulation in which diffusivity, D, was varied in order to maximise the rate of entropy production. Data for plots (b), (d) and (f) were produced by a simulation in which diffusivity, D was varied in order to maximise the total amount of life. The coverage of the black and white daisies is shown in plots (a) and (b). The temperature of the planet is shown in plots (c) and (d). The rate of entropy and value for diffusivity is shown in plots (e) and (f). When adjusting heat flux to maximize the rate of entropy production, the any-daisy and both-daisy ranges are maximsed. Entropy production is greatest when $L \approx 1.2$. Note the sharp decrease in entropy production in plot (f) where maximising the total coverage of daisies leads to a maximum rate of diffusivity and so zero entropy production.

TBDW I have assumed that there are sufficient degrees of freedom within the processes that determine the diffusivity between the two boxes, to afford the system the ability to configure itself into a state that maximizes the rate of entropy production. When TBDW is in a MEP state with respect to latitudinal heat flux, it is not in a MEP state with respect to shortwave to longwave radiative balance. This would be achieved by making Daisyworld as dark as possible thus absorbing as much of the star's energy and so converting the maximum amount of shortwave radiation to longwave radiation. On Earth, entropy is produced by shortwave radiation from the Sun warming the surface and atmosphere of the planet which then radiates this now longwave energy back into space. The Earth does not maximize this rate of this entropy production. This is due to the insufficient degrees of freedom the radiative mechanism possesses, with MEP only to be expected in complex, turbulent, dissipative systems [Ozawa et al., 2003]. Arguably the utility of the results presented here will depend on the plausibility and conceptual coherence of the principle of maximum entropy production. If the MEP is 'real' and applicable to a range of systems then it seems reasonable to investigate its effects on systems such as biologically-mediated changes in climate models.

8.7.1 Entropy production and homeostasis

There is an ongoing research programme that explores the hypothesis that planetary homeostasis is, in large part, a consequence of negative feedback processes that are the result of the Earth's biosphere being in a state of maximum entropy production: [Kleidon, 2004], [Kleidon and Fraedrich, 2005], [Kleidon, 2006],
[Kleidon and Schymanski, 2008]. This is a form of Optimising Gaia - that the biota affects its environment in ways that create biologically favourable conditions. One stumbling block for any form of Optimising Gaia is the appreciation that what is biologically deleterious or favourable at time 1 may prove to be the opposite at time 2. A classic example is the Great Oxidation whereby the percentage of molecular oxygen in the atmosphere was increased from less than 1 % to near current day levels of 21% by photosynthetic cyanobacteria which expired molecular oxygen as a by-product of using sunlight to produce sugars from carbon dioxide[3]. This highly reactive oxygen

[3]A recent study, [Goldblatt et al., 2006], has argued that the relatively sudden increase in the partial pressure of oxygen can be understood as the sudden switch to one of the possible bi-stable states for the atmosphere and that biogenic methane levels were critical for this switch. This rapid, biologically-mediated environmental change is echoed in the results of MGDW where sudden transitions between stable states are produced by the simple life forms driving the environment away from potentially absorbing states.

would have been poisonous to many contemporary life forms, but is now necessary for the aerobic metabolisms that maintain modern eukaryotic life forms such as mammals, fish and reptiles. If Gaia maintains the biosphere to those conditions optimal for life then it would appear necessary to integrate this notion of optimality over sufficiently long time-scales in order to discern this optimisation. There is also the problem of how one quantifies such optimality. Kleidon argues that the form of optimisation is one of primary production, which is the production of chemical energy in organic compounds by living organisms [Leith and Whittaker, 1975]. Such primary production produces entropy and the rate at which this entropy is produced is maximised. This suggests the possibility of a mechanism that could account for the long-term evolution of the biosphere. For example, while the Darwinian algorithm is a directionless process, the principle of maximum entropy production would, given the energy gradient established by low entropy energy from the Sun being received on the surface of the Earth, *predict* the evolution of autotrophic photosynthesis as such organisms would increase the rate of entropy production. Consequently there may be a sense of directionality for the evolution of the Earth's biosphere and so the forms of life that exist today and those that are preserved in the fossil record may in that respect be 'reproducible'. If it were possible to conduct an experiment whereby time was turned back to just after the origins of life on Earth, then over the proceeding billions of years we should expect the same broad patters of biological evolution to emerge. Furthermore, this sense of reproducibility would be homeorhetic. The particular trajectories that the biosphere has evolved along represent the most probable and most stable trajectories.

This hypothesis, like the original Gaia Hypothesis is controversial and the debate is ongoing [Volk, 2007], [Kleidon, 2007]. This can be situated within other attempts to explain the evolution of the biosphere such as Stuart Kauffman's autocatalytic networks and co-constructing biospheres [Kauffman, 1993], [Kauffman, 1998] or Brian Goodwin's complex systems approach to morphogenesis and evolution [Goodwin, 1997]. The acid test for these theories will be formulating hypotheses about natural processes that can be falsified by empirical data. In this chapter I have provided an initial method for quantitatively assessing the relationship between homeostasis and entropy production. This is something that has been considered before within the context of Daisyworld, but as I argued in Section 8.4 all of these studies featured assumptions that limited their ability in this respect.

8.8 Summary

This chapter examined the Daisyworld control system within the context of non-equilibrium thermodynamics and the maximum entropy production principle. I showed that a version of the TBDW initially presented in chapter 4 can be regarded as an example of an energy balance climate model. Energy balance models have been used to explore the hypothesis that planetary atmospheres maximise the rate of entropy production via the transport of heat from the hot tropics to the cold poles. In TBDW, the difference in zonal temperatures was produced by a difference in albedo rather than latitude. When the amount of diffusion was adjusted in order to maximize the rate of entropy production, the range over which any and both daisies grow was maximised and the essential range of life on the planet was maximised. Maximising the rate of entropy production led to a maximisation of the range of luminosity over which homeostasis was observed. Although very simple, the model allowed the introduction of biotically-mediated changes to the imaginary planet's radiative budget. Thermodynamically coupling life and environment in that way extended a simple and powerful climate model and once again highlighted the importance of the effect that organisms can have on their environments.

This chapter concludes the presentation and analysis of results. In the following chapter I will conclude the thesis by reviewing the main findings, assessing the significance of these findings and suggesting possible directions for future research.

Chapter 9

Conclusion

9.1 Overview

In this final chapter I will return to the questions posed at the start of the thesis and argue how these have been addressed and in doing so state the original contribution of my thesis. I then review the findings of the previous chapters and suggest research that could continue a particular aspect of this thesis.

9.2 Main Research Question and Contribution of Thesis

I began this thesis by asking whether the Earth's crust, oceans, atmosphere and biota form a homeostatic system that is robust to perturbations. Gaia Theory may explain why widespread life has always returned after major perturbations such as changes in the brightness of the Sun, meteorite impacts and vulcanism. A possible response to Gaia Theory is to claim that there is nothing more than luck being demonstrated in the persistence of life on Earth. However life started on Earth, and this may have required a certain amount of luck, the fact we are here to wonder over its evolution is a testament to good fortune. Perhaps it also demonstrates the importance of gas giants within a solar system as these massive bodies will tend to 'hoover up' stray objects that could otherwise impact the Earth. Perhaps other aspects of the solar system are more important.

Daisyworld is a simple model in which life and its environment are components of a homeostatic system that resist external perturbations. In chapter 2 I reviewed the quite

numerous developments of Daisyworld and how a common thread of these developments was the attempt to reduce the number of assumptions that Daisyworld homeostasis required. By reducing the number of assumptions, the model was made potentially more supportable by empirical data and increased confidence that temperature regulation was not a quirk or the result of some narrow set of assumptions or modelling parameters. In this thesis I have developed and discussed a series of very simple models that have featured relatively few assumptions whilst still exhibiting homeostasis. The original contribution of this thesis has been to explain how such homeostasis is observed under such relaxed assumptions: how Daisyworld 'works'. I have argued that Daisyworld is a rein control system. This rein control emerges and persists not in spite of natural selection but *because* of it.

This is an important contribution as it directly addresses the concern that natural selection operating on individuals would at best be indifferent and at worst be hostile to any form of biologically-mediated homeostasis. This has led some to consider the possibility that other evolutionary forces may provide a mechanism. The theory of kin selection [Hamilton, 1964a], [Hamilton, 1964b] seeks to explain altruistic behaviour whereby an individual will experience a reproductive cost in helping a relative. Hamilton showed under what general circumstances such altruism could arise. Kin selection could explain how individuals could 'do their bit' for environmental homeostasis, and at some cost to themselves, as maintaining the environment within certain bounds would increase the fitness of their relatives. This would increase the fitness of the genes of the relative, some of which the performer of the altruistic act would share. If some of these genes are responsible for the altruistic behaviour then there is a mechanism whereby they can be maintained.

However, when considering planetary homeostasis the notion of kin selection is of arguably limited utility as the interactions occur over long timescales and often between different species that are only very distantly related. Of perhaps more relevance is the notion of group selection [Wilson and Dugatkin, 1997]. Rather than natural selection only operating on individuals, careful experiments have shown that groups of individuals can respond to selection even when such group selection is opposed by natural selection. In the context of Daisyworld, group selection may provide a mechanism whereby homeostatic groups, by resisting external perturbations, out-compete non-homeostatic groups and so will flourish. A series of studies have sought to situate Gaia Theory within group selection: [Williams and Lenton, 2007a], [Williams and Lenton, 2007b],

[Williams and Lenton, 2008]. In this thesis I have shown why group selection or kin selection are not necessary conditions for Daisyworld homeostasis. Natural selection operating on a diverse population of simple organisms is sufficient for the establishment and maintenance of rein control states that respond to external perturbations in such ways as to reduce the effect of these perturbations on an internal environmental variable.

9.3 Review and Discussion of Chapters

In this section I will review the previous chapters and summarise their main results.

In chapter 2 I reviewed the Daisyworld and Gaia Theory literature. I reproduced Kirchner's taxonomy of Gaia Theories. Of particular relevance to chapters 5, 6 and 7 was the 'Lucky Gaia' theory that claimed that planetary homeostasis, if it were to arise would be considered as improbable. One conclusion of Lucky Gaia is that homeostatic biospheres arising and persisting on other planets with life are not to be expected. In chapter 6 I explored a Daisyworld model that would directly challenge Lucky Gaia and argue that homeostasis may emerge and persist under more relaxed assumptions. I provided an initial Ashbyan definition of homeostasis. I also discussed homeorhesis which was to prove to be an important concept for the explanation of the behaviour of models presented in later chapters. I presented the original Daisyworld model along with a synopsis of its analytical solutions. I classified the wide range of extensions and developments of the original model.

In chapter 3 I developed a simplified version of the original Daisyworld and detailed its analytical solutions. In doing so I shed light on the inverse response observed in Daisyworld. That is, there is a range of luminosity over which, as luminosity increases, the temperature of the planet *decreases*. I showed that this inverse response was not the result of particular non-linear functions in the original model but was caused by the method of ensuring that the black daisies were always warmer than bare ground and that the white daisies were always cooler than bare ground. This temperature difference was established with the use of an insulating parameter that led to the maximum temperature between the daisies and bare ground being found with the least number of daisies. Increasing values for the insulating parameter increased the inverse response.

In chapter 4 I formulated an alternative method for differentiating the daisy temperatures and in doing so produced a model that did not exhibit an inverse response to

forcing. Instead, planetary temperature was fixed when both black and white daisies were present. I explained this zero steady state error as the results of a zero sum game being played between the black and white daisies as they competed for a shared resource. By changing this assumption and removing the competition for a shared resource I showed that planetary temperature no longer remained fixed but increased with increasing luminosity, although at a much smaller rate of change than luminosity. I presented counter-intuitive results where increasing the essential range of the daisies, increasing the range of temperature over which they would grow, *decreased* the range of luminosity over which the daisies grew. I explained this result and highlighted the relationship between homeostasis and hysteresis in Daisyworld. I demonstrated the generality of the Daisyworld control system by formulating a simple robotic controller that performed phototaxis (light following behaviour).

In chapter 5 I completed the process of simplification with the production of a rein control model. I classified a simple set of rein controllers on the basis of their positive or negative feedback effect using the notion of feedback circuits. Of the four circuits identified, two would produce negative feedback and two would produce positive feedback. I developed four simple negative rein control systems and provided analytical solutions for their steady states in response to external perturbations. I highlighted the difference between benign and antagonistic interactions with the latter giving rise to equivalent behaviour in TBDW with competition for a shared resource. When antagonistic interactions were operating in the simple rein control model, it was not necessary to specify how the individual rein controller elements responded to external perturbations. I showed that the system variable that was subject to regulation remained fixed for a range of perturbations and that it was the rein controller output that varied as perturbations varied. This allowed a simpler method for finding steady state values.

Chapter 5 provided the analysis to understand how negative feedback rein control systems could perform homeostasis with the notable omission of any discussion of positive feedback systems. As this thesis is focussed on examining homoestatic systems this is understandable. However, this could have suggested that the Daisyworld control system could only arise in the absence of positive feedback. In chapter 6 I presented the McDonald-Gibson Daisyworld model (MGDW) that featured the capacity to produce both negative and positive feedback. MGDW was an individual-based model where individuals were simple organisms that had a genotype with alleles that specified its response to an environmental variable and its effect on the environmental variable.

By allowing these alleles to mutate during reproduction it was possible for individuals to emerge which would result in positive feedback. Such individuals were similar to black daisies that out-competed white daisies at higher temperatures and white daisies that out-competed black daisies at lower temperatures. The frequency of alleles in the population changed over time due to the operation of natural selection. Despite these assumptions that would allow an equal probability of negative and positive feedback effects emerging, homeostasis was observed and proved to be robust to a range of parameter values. Understanding how homeostasis emerged and persisted in MGDW required the use of the rein control analysis developed in chapter 5. I made a number of connections between rein control and the population genetics theory of niche construction. MGDW can be seen as sharing some important assumptions with a landmark niche construction study that produced a number of surprising results. One of these results was that niche construction can lead to stable polymorphisms whereby rather than a single genotype reaching 100% within the population, other genotypes persist due to the actions of the dominant and sub-dominant genotypes. This is similar to the situation in MGDW where rather than a single optimal genotype that is most closely matched to the current environmental conditions reaching 100% within the population, sub-populations emerge that regulate the environmental variable.

In chapter 7 I continued the discussion of MGDW with a focus on the effects of mutation and how it affected homeostasis. I showed how mutation can lead to an absorbing state where the regulating sub-populations would have no overall impact on the environmental variable. Natural selection is unable to recreate the bimodal distribution of environmentally affecting alleles and so once this state is reached homeostasis will cease. A series of experiments conducted with MGDW showed that these abosrbing states were not reached. I detailed a mechanism whereby the population would escape any particular absorbing state. This would also lead to the collapse of homeostasis. However, homeostasis would be recovered at *different* environmental conditions and with a different population structure. With mutation operating, MGDW inhabited a series of homeostatic steady states. Over suitably long timescales, the behaviour of MGDW was similar to that of a random walk where a series of stable states would be created, degraded, collapsed and then new states created. There are similarities here with the behaviour of Ashby's Homeostat that was created in order to demonstrate and explore the notion of ultrastability. [Ashby 1960]. The homeostat had the potential to reside in any number of stable states and if perturbed past the essential

range would generate a rapid succession of system changes until a new stable state was chanced upon and homeostasis would be established again. I showed that by increasing mutation rates it was possible to increase the overall amount of time that MGDW regulated the environmental variable, as increasing the mutation rate decreased the model's resistance (the ability to resist external perturbations), but increased its resilience (the ability to recover after perturbations). I discussed these findings within the context of the complexity-stability debate within theoretical ecology. I argued that when one incorporates the effects that species can have on other species via direct and environmentally-mediated actions then new results can be returned. The theory of ecosystem engineering incorporates such effects and I drew a number of similarities between rein control and ecosystem engineering.

Chapter 8 saw a return to the TBDW via a discussion on the relationship between homeostasis and entropy production. The previous chapter highlighted the essentially directionless Darwinian algorithm. While Darwinian evolution may be an unguided, purposeless mechanism, the matter that constitute the environments in which organisms evolve and the material that constitute the organisms themselves must adhere to physical laws. In chapter 8 I examined the maximum entropy production principle that may impose boundary conditions on the evolution of open, dissipative, non-equilibrium systems such as planetary atmospheres. I evaluated the effects of maximising the amount of entropy production via latitudinal heat flux in TBDW. I showed that when heat flux was adjusted in order to maximise entropy production, the any-daisy and both-daisy ranges were maximised. This result was explained by appreciating that increasing the coverage of black and white daisies over certain ranges of luminosity increases the temperature gradient between the two daisy boxes. The greater the gradient, the greater the possible rate of energy dissipation and so the greater the rate of entropy production.

9.4 Future Work

The program of research that I have reported in this thesis has been completed. However, this in turn can be seen as opening up new avenues of research. I will conclude my thesis by suggesting future work.

In developing the very simple rein control models of chapter 5, I detected similarities between the formalism of rein control and frequency dependent selection. Whilst [Haldane, 1932], [Fisher, 1930] and [Wright, 1948] can be seen as doing most of the hard

work when it comes to understanding the mathematical basis of frequency dependent selection, what has not been considered is the role of frequency dependent selection in environmental homeostasis. Later I was surprised to learn that there have in fact been no studies that have cast Daisyworld or Daisyworld-type models in terms of some of these well established results of population genetics although some have come close, see for example [Maddock, 1991]. Such an absence could indicate that there is in fact no sensible relationship. However, personal correspondence and further consideration of my initial hunch suggests that there is a productive line of enquiry here.

It will become immediately apparent, that once one begins to incorporate such effects into models of evolutionary processes, then previously tractable models may become hopelessly complex. Previous techniques may no longer work. Previous results may no longer be valid. This raises the question: is it worth it? For example, if niche construction is not a significant evolutionary mechanism, why invest the considerable time and effort into incorporating it into existing, successful models, when it can be effectively ignored? This response to niche construction has been made by a number of evolutionary theorists and Laland has been at some pains to respond to them [Laland and Sterelny, 2006]. I am in no position to decide whether or to what extent niche construction is 'true'. However I am able to report that all biologists would agree that the established theories of population genetics proceed on the basis of a number of simplifying assumptions that render certain problems amenable to mathematical formalism. For example:

"there is nothing artificial about assuming that fitnesses are frequency-dependent. The artificial assumption is that relative fitnesses are constant: its justifcation is mathematical convenience, not truth." [Maynard-Smith, 1989] p70.

Arguably *all* selection is frequency dependent selection as *all* organisms, in some way, great or small, affect their environment. To assume the default stance that such effects are sufficiently small to be ignored is not necessarily justified. Chapters 7 and 8 showed that allowing organisms to affect shared environments led to model ecosystems in which competitive exclusion [Gause, 1934] was not observed. Rather than one genotype increasing in fitness and so excluding other genotypes, stable polymorphic populations emerged.

With specific interest to the Gaia Hypothesis, I believe that providing a mathematical connection between rein control and frequency dependent selection could significantly support the hypothesis that organisms, in interacting with each other and their environments, form homeostatic systems that are robust to perturbations. This would cast the Gaia Hypothesis in terms of of population genetics and imbed biologically-mediated homeostasis within the foundations of the modern evolutionary synthesis. This thesis can be considered as a 'proof of concept' for such ideas.

9.5 Concluding Remarks

The very simplest dynamical systems can surprise us. A seemingly innocent equation, once iterated a modest number of times can soon wander through the labyrinthine halls of chaos being at once entirely deterministic and completely unpredictable. If the Earth truly is uniquely privileged to be the only home to life throughout the cosmos, then we are inhabitants of the most complex system in the universe. A fundamental assumption of this thesis, and arguably any scientific theory, is that major mechanisms can be identified within such complex systems. However, maintaining a sense of humility when considering our abilities to fully understand such systems is, I would argue, vital when considering how we interact with them and how we may wish to guide them. Very complex systems can produce very big surprises. The global financial events of 2007 - 2009 demonstrate that even the Masters of the Universe and their mathematical mages may be as blind to impending crises as other mere mortals.

In conducting my research I have been fortunate to have access to computing power that allows me to conduct thousands of experiments while sipping a coffee or conduct millions of experiments while sleeping. Reviewing and sifting results can guide new experiments. Probing the regions where homeostasis collapses can identify important conditions and aspects of the mechanisms under investigation. However, I am also conducting experiments while powering my computer in order to run simulations or in order to surf for research papers online. I am making a personal contribution to a civilisation-wide experiment when I drive my car to the supermarket and purchase groceries. Global emissions of anthropogenic carbon dioxide is estimated at 6.3 giga tons a year. The release of carbon stores due to deforestation is estimated at 1.6 giga tons a year. From the years 1750 to 2000, the concentration of carbon dioxide increased by $31 \pm 4\%$, and that of the much stronger greenhouse gas, methane, rose by $151 \pm 25\%$.

It is considered 'likely' that the rate of increase of the northern hemisphere surface temperature was greater in the 20th century than during any other in the last 1,000 years and that the 1990s was the warmest decade of the millennium [IPCC, 2007]. Recent research supports the hypothesis that ecological impacts caused by Homo sapiens are currently producing one of the greatest extinction events in the history of life on Earth [Cadotte et al., 2008].

There is a profusion of model and theory that seeks to understand these changes and predict future conditions. It is unfortunate that the Earth's biosphere cannot be reset with different starting conditions if results are found to be incompatible with our and other species' requirements for survival.

...

Bibliography

[Ackland, 2004] Ackland, G. J. (2004). Maximization principles and daisyworld. *Journal of Theoretical Biology*, 227(1):121–128.

[Ackland et al., 2003] Ackland, G. J., Clark, M. A., and Lenton, T. M. (2003). Catastrophic desert formation in daisyworld. *Journal of Theoretical Biology*, 223(1):39–44.

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

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

[Andreae and Crutzen, 1997] Andreae, M. O. and Crutzen, P. J. (1997). Atmospheric aerosols: Biogechemical sources and role in atmospheric chemistry. *Science*, 276:1052–1058.

[Ashby, 1960] Ashby, W. R. (1960). *Design for a brain*. Chaman and Hall, 2nd edition.

[Baldocchi et al., 2005] Baldocchi, D. D., Krebs, T., and Leclerc, M. Y. (2005). 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 Series B-Chemical and Physical Meteorology*, 57(3):175–188.

[Braitenberg, 1984] Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic*. MIT Press.

[Cadotte et al., 2008] Cadotte, M. W., Cardinale, B. J., and Oakley, T. (2008). Evolutionary history and the effect of biodiversity on plant productivity. *Proceedings National Acadamies of Science*, 105(44):17012–17017.

[Charlson et al., 1987] Charlson, R. J., Lovelock, J. E., Andreae, M. O., and Warren, S. G. (1987). Oceanic phytoplankton, atmospheric sulphur, cloud albedo and climate. *Nature*, 326:655–661.

[Clynes, 1969] Clynes, M. (1969). Cybernetic implications of rein control in perceptual and conceptual organization. *Annals of New York Academy of Science*, 156:629–670.

[D'Ari and Thomas, 2003] D'Ari, R. and Thomas, R. (2003). Hardware (dna) circuits. *C. R. Biologies*, 326:215–217.

[Darwin, 1881] Darwin, C. (1881). *The formatin of vegetable mold through the action of worms, with observations on their habits*. John Muray, London.

[Dawkins, 1986] Dawkins, R. (1986). *The blind watchmaker*. Longman, London.

[DeGregorio et al., 1992] DeGregorio, S., Pielke, R. A., and Dalu, G. A. (1992). Feedback between a simple biosystem and the temperature of the earth. *Journal of Nonlinear Science*, 2:263–292.

[Dewar, 2003] Dewar, R. (2003). Information theory explanation of the fluctuation theorem, maximum entropy production and self-organized criticality in non-equilibrium stationary states. *Journal Physics A*, 36(631-641).

[Dewar, 2005] Dewar, R. (2005). Maximum entropy production and the fluctuation theorem. *Journal Physics A*, 38(21):371–381.

[Di Paolo et al., 2000] Di Paolo, E., Noble, J., and Bullock, S. (2000). Simulation models as opaque thought experiments. In *Artificial Life VII, Proceedings of the Seventh International Conference on the Simulation and Synthesis of Living Systems*.

[Downing, 2002] Downing, K. (2002). The simulated emergence of distributed environmental control in evolving microcosms. *Artificial Life*, 6(2):291–318.

[Downing, 2003] Downing, K. (2003). Gaia in the machine: The artificial life approach. In Schneider, S. H., editor, *Scientists Debate Gaia: The Next Century*. MIT Press, Cambridge MA.

[Downing and Zvirinsky, 1999] Downing, K. and Zvirinsky, P. (1999). The simulated evolution of biochemical guilds: Reconciling gaia theory and natural selection. *Artificial Life*, 5(4):291–318.

[Dyke, 2008] Dyke, J. G. (2008). Entropy production in an energy balance daisyworld model. In *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems, ALIFE'11*. MIT Press.

[Dyke and Harvey, 2005] Dyke, J. G. and Harvey, I. R. (2005). Hysteresis and the limits of homeostasis: from daisyworld to phototaxis. In Capcarrere, M., Freitas, A., Bentley, J., Johnson, C., and Timmis, J., editors, *Proceedings of VIIIth European Conference on Artificial Life, ECAL 2005*, pages 241–246. Springer-Verlag.

[Dyke and Harvey, 2006] Dyke, J. G. and Harvey, I. R. (2006). Pushing up the daisies. In Rocha, L. M., Yager, L. S., Bedau, M. A. Floreano, D., Goldstone, R. L., and Vespignani, A., editors, *Artificial Life X, Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 426–431. MIT Press, Cambridge MA.

[Dyke et al., 2007] Dyke, J. G., McDonald-Gibson, J., Di Paolo, E., and Harvey, I. R. (2007). Increasing complexity can increase stability in a self-regulating ecosystem. In *Proceedings of IXth European Conference on Artificial Life, ECAL 2007*, pages 133–142.

[Eldredge and Gould, 1972] Eldredge, N. and Gould, S. J. (1972). Models in paleobiology. chapter Punctuated equilibria: an alternative to phyletic gradualism. Freeman, Cooper and Co, San Francisco.

[Elton, 1958] Elton, R. M. (1958). *Ecology of Invasion by Animals and Plants*. Chapman and Hall, London.

[Fisher, 1930] Fisher, R. A. (1930). *The Genetical Theory of Natural Selection*. Oxford University Press.

[Gardner and Ashby, 1970] Gardner, M. A. and Ashby, W. R. (1970). Connectance of large dynamic (cybernetic) systems: critical values for stability. *Nature*, pages 228–784.

[Gause, 1934] Gause, G. F. (1934). *The struggle for existence*. Williams and Wilkins, Baltimore.

[Goldblatt et al., 2006] Goldblatt, C., Lenton, T. M., and Watson, R. A. (2006). The great oxidation as a bistability in atmospheric oxygen. *Nature*, 443:683–686.

[Goodwin, 1997] Goodwin, B. C. (1997). *How the Leopard Changed its Spots: The Evolution of Complexity*. Cambridge University Press, Cambridge.

[Gutierrez and Jones, 2006] Gutierrez, J. and Jones, C. G. (2006). Physical ecosystem engineers as agents of biogeochemical heterogeneity. *BioScience*, pages 227–236.

[Haldane, 1932] Haldane, J. B. S. (1932). *The Causes of Evolution*. Longmans, Green and Co, London.

[Hamilton, 1964a] Hamilton, W. D. (1964a). The genetical evolution of social behaviour. i. *Journal of Theoretical Biology*, 7(1):1–16.

[Hamilton, 1964b] Hamilton, W. D. (1964b). The genetical evolution of social behaviour. ii. *Journal of Theoretical Biology*, 7(1):17–52.

[Harding, 1999] Harding, S. P. (1999). Food web complexity enhances community stability and climate regulation in a geophysical model. *Tellus Series B-Chemical and Physical Meteorology*, 51:815–829.

[Harding and Lovelock, 1996] Harding, S. P. and Lovelock, J. E. (1996). Exploiter-mediated co-existence and frequency-dependent selection in a numerical model of biodiversity. *Journal of Theoretical Biology*, 182:109–116.

[Harvey, 2004] Harvey, I. R. (2004). Homeostasis and rein control: From daisyworld to active perception. In Pollack, J., Bedau, M., Husbands, P., Ikegami, T., and Watson, R. A., editors, *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems, ALIFE'9*, pages 309–314. MIT Press, Cambridge MA.

[IPCC, 2007] IPCC (2007). *Climate Change 2007: Synthesis Report, Contribution of Working Groups I, II and III to the Fourth Assessment of IPCC*.

[Jaynes, 1957a] Jaynes, E. T. (1957a). Information theory and statistical mechanics. *Physical Review*, 106(4):620–630.

[Jaynes, 1957b] Jaynes, E. T. (1957b). Information theory and statistical mechanics. ii. *Physical Review*, 108(2):171–190.

[Jones et al., 1994] Jones, C. G., Lawton, J. H., and Shachak, M. (1994). Organisms as ecosystem engineers. *Oikos*, 69:373–386.

[Jones et al., 1997] Jones, C. G., Lawton, J. H., and Shachak, M. (1997). Positive and negative effects of organisms as physical ecosystem engineers. *Ecology*, 78:1946–1957.

[Kauffman, 1993] Kauffman, S. A. (1993). *Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, Oxford.

[Kauffman, 1998] Kauffman, S. A. (1998). *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. Oxford University Press, Oxford.

[Kaufman and Thomas, 2003] Kaufman, M. and Thomas, R. (2003). Emergence of complex behaviour from simple circuit structures. *Comptes Rendus Biologies*, 326:205–214.

[Keeling, 1991] Keeling, R. F. (1991). Mechanism for stabilization and destabilization of a simple biosphere: Catastrophe on daisyworld. In Schneider, S. H. and Boston, P., editors, *Scientists on Gaia*, pages 118–120. MIT Press, Cambridge MA.

[Kirchner, 1989] Kirchner, J. W. (1989). The gaia hypothesis: can it be tested? *Reviews of Geophysics*, 27(2):223–235.

[Kirchner, 2002] Kirchner, J. W. (2002). The gaia hypothesis: fact, theory, and wishful thinking. *Climatic Change*, 52:391–408.

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

[Kleidon, 2004] Kleidon, A. (2004). Beyond gaia: Thermodynamics of life and earth system functioning. *Climatic Change*, 66(3):271–319.

[Kleidon, 2006] Kleidon, A. (2006). Maximum entropy production and the strength of boundary layer exchange in an atmospheric general circulation model. *Geophysical Research Letters*, 33(doi: 10.1029/2005GL025373).

[Kleidon, 2007] Kleidon, A. (2007). Thermodynamics and environmental constraints make the biosphere predictable. *Climatic Change*, 85(3):259–266.

[Kleidon and Fraedrich, 2005] Kleidon, A. and Fraedrich, K. (2005). Biotic entropy production and global atmosphere-biosphere interactions. In Kleidon, A. and Lorenz, R. D., editors, *Non-Equilibrium Thermodynamics and the Production of Entropy: Life, Earth, and Beyond*. Springer-Verlag, Berlin.

[Kleidon and Schymanski, 2008] Kleidon, A. and Schymanski, S. J. (2008). Thermodynamics and optimality of the water budget on land: A review. *Geophysical Research Letters*, 35(doi:10.1029/2008GL035393).

[Laland et al., 1999] Laland, K. N., Odling-Smee, F. J., and Feldman, M. W. (1999). Evolutionary consequences of niche construction and their implications for ecology. *Proceedings National Acadamies of Science*, 96(18):10242–10247.

[Laland et al., 2000] Laland, K. N., Odling-Smee, J., and Feldman, M. W. (2000). Niche construction, biological evolution, and cultural change. *Behavioural and Brain Sciences*, 23(1):131–146.

[Laland and Sterelny, 2006] Laland, K. N. and Sterelny, K. (2006). Perspective: seven reasons (not) to neglect niche construction. *Evolution: International Journal of Organic Evolution*, 60(9):1751–1762.

[Langton, 1997] Langton, C., editor (1997). *Artificial Life: An overview*. MIT Press, Cambridge MA.

[Lansing and Smuts, 1998] Lansing, J. and Smuts, J. K. B. (1998). System-dependent selection, ecological feedback and the emergence of functional structure in ecosystems. *Journal of Theoretical Biology*, 192(3):337–391.

[Leith and Whittaker, 1975] Leith, H. and Whittaker, R. H. (1975). *Primary Productivity on the Biosphere*. Springer-Verlag, Berlin.

[Lenton, 1998] Lenton, T. (1998). Gaia and natural selection. *Nature*, 394:439–447.

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

[Lenton and van Oijen M., 2002] Lenton, T. and van Oijen M. (2002). Gaia as a complex adaptive system. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 357:683–695.

[Lenton et al., 2004] Lenton, T. M., Caldeira, K. G., and Szathmary, E. (2004). What does history teach us about the major transitions and the role of disturbances in the evolution of life and of the earth system? In *Earth System Analysis for Sustainability*. Dahlem Workshop Report 91. H.-J.

[Lenton et al., 2008] Lenton, T. M., Held, H., Kreigler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J. (2008). Tipping elements in the earth's climate system. *Proceedings National Acadamies of Science*, 105:1786.

[Lenton and Lovelock, 2001] Lenton, T. M. and Lovelock, J. E. (2001). Daisyworld revisited: Quantifying biological effects on planetary self-regulation. *Tellus Series B-Chemical and Physical Meteorology*, 53:288–305.

[Lenton and Watson, 2000] Lenton, T. M. and Watson, R. A. (2000). Redfield revisited 1: Regulation of nitrate, phosphate and oxygen in the ocean. *Global Biogeochemical Cycles*, 14(1):225–248.

[Lewontin, 1982] Lewontin, R. C. (1982). Organism and environment. In Plotkin, H. C., editor, *Learning, Development and Culture: Essays in Evolutionary Epistemology*. Wiley, New York.

[Lewontin, 2001] Lewontin, R. C. (2001). *The Triple Helix*. Harvard University Press.

[Liss et al., 1997] Liss, P. S., Hatton, A. D., Malin, G., Nightingale, P. D., and Turner, S. M. (1997). Marine sulphur emissions. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 352:159–169.

[Lorenz, 2002] Lorenz, R. D. (2002). Planets, life and the production of entropy. *International Journal of Astrobiology*, 1(1):3–13.

[Lorenz et al., 2001] Lorenz, R. D., Lunine, J. I., and Withers, P. G. (2001). Titan, mars and earth: Entropy production by latitudinal heat transport. *Geophysical Research Letters*, 28(3):415–418.

[Lovelock, 1979] Lovelock, J. E. (1979). *Gaia: a new look at life on Earth*. Oxford University Press, Oxford.

[Lovelock, 1983] Lovelock, J. E. (1983). Daisy world - a cybernetic proof of the gaia hypothesis. *The Co-evolution Quarterly*, Summer:66–72.

[Lovelock, 1988] Lovelock, J. E. (1988). *Ages of Gaia*. Oxford University Press, Oxford.

[Lovelock, 1991] Lovelock, J. E. (1991). *Gaia: the practical science of planetary medicine*. Gaia Books Ltd, Stroud.

[Lovelock, 1992] Lovelock, J. E. (1992). A numerical model for biodiversity. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 338:383–391.

[Lovelock, 1997] Lovelock, J. E. (1997). A geophysiologist's thoughts on the natural sulphur cycle. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 352:143–147.

[Lovelock, 2006] Lovelock, J. E. (2006). *The Revenge of Gaia*. Allen Lane, London.

[Lovelock and Margulis, 1974] Lovelock, J. E. and Margulis, L. (1974). Atmospheric homeostasis by and for the biosphere. *Tellus Series B-Chemical and Physical Meteorology*, 26(4):299–327.

[Maddock, 1991] Maddock, L. (1991). Effects of simple environmental feedback on some population models. *Tellus Series B-Chemical and Physical Meteorology*, 43B:331–337.

[Mamontov, 2007] Mamontov, E. (2007). Modelling homeorhesis by ordinary differential equations. *Mathematical and Computer Modelling*, 45:694–707.

[Margulis and Lovelock, 1974] Margulis, L. and Lovelock, J. E. (1974). Biological modulation of the earth's atmosphere. *Icarus*, 21(4):471–489.

[Margulis and Sagan, 1997] Margulis, L. and Sagan, D. (1997). *Slanted Truths: Essays on Gaia, Symbiosis and Evolution*. Springer-Verlag, Berlin.

[Martyushev and Seleznev, 2006] Martyushev, L. M. and Seleznev, V. D. (2006). Maximum entropy production principle in physics, chemistry and biology. *Physics Reports*, 426(1-45).

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

[Maynard-Smith, 1989] Maynard-Smith, J. (1989). *Evolutionary Genetics*. Oxford University Press, Oxford.

[McCann, 2000] McCann, K. S. (2000). The diversity-stability debate. *Nature*, 405:228–233.

[McDonald-Gibson, 2006] McDonald-Gibson, J. (2006). Investigating gaia: A new mechanism for environmental regulation. Master's thesis, University of Sussex.

[McDonald-Gibson et al., 2008] McDonald-Gibson, J., Dyke, J. G., Di Paolo, E., and Harvey, I. R. (2008). Environmental regulation can arise under minimal assumptions. *Journal of Theoretical Biology*, 251(4):653–666.

[North et al., 1981] North, G. R., Cahalan, R. F., and Coakley Jr, J. A. (1981). Energy balance climate models. *Reviews of Geophysics*, 19:91–121.

[Novick and Sziliard, 1950] Novick, A. and Sziliard, L. (1950). Description of the chemostat. *Science*, 112:715–716.

[Odling-Smee et al., 2003] Odling-Smee, F. J., Laland, K. N., and Feldman, M. (2003). *Niche Construction: The Neglected Process in Evolution*. Princeton University Press.

[Odum, 1953] Odum, E. P. (1953). *Fundamentals of Ecology*. Saunders, Philadelphia.

[Ozawa et al., 2003] Ozawa, H. A., Ohmura, A., Lorenz, R. D., and Pujol, T. (2003). The second law of thermodynamics and the global climate system: A review of the maximum entropy production principle. *Reviews of Geophysics*, 41(4):1018.

[Paltridge, 1975] Paltridge, G. W. (1975). The steady-state format of global climate systems. *Quarterly Journal of Royal Meterological Society*, 104:927–945.

[Prigogine, 1962] Prigogine, I. (1962). *Introduction to non-equilibrium thermodynamics*. Wiley Interscience, New York.

[Pujol, 2002] Pujol, T. (2002). The consequences of maximum thermodynamic efficiency in daisyworld. *Journal of Theoretical Biology*, 217:53–60.

[Redfield, 1958] Redfield, A. C. (1958). The biological control of chemical factors in the environment. *American Scientist*, 46(3):205–221.

[Robertson and Robinson, 1998] Robertson, D. and Robinson, J. (1998). Darwinian daisyworld. *Journal of Theoretical Biology*, pages 129–134.

[Saunders et al., 1998] Saunders, P., Koeslag, J. H., and Wessels, J. A. (1998). Integral rein control in physiology. *Journal of Theoretical Biology*, 194:163–173.

[Saunders et al., 2000] Saunders, P., Koeslag, J. H., and Wessels, J. A. (2000). Integral rein control in physiology ii: A general model. *Journal of Theoretical Biology*, 206:211–220.

[Saunders, 1994] Saunders, P. T. (1994). Evolution without natural selection: further implications of the daisyworld parable. *Journal of Theoretical Biology*, pages 365–373.

[Schwartman and Volk, 1989] Schwartman, D. W. and Volk, T. (1989). Biotic enhancement of weathering and the habitability of earth. *Nature*, 340:457–460.

[Schwartman and Volk, 1991] Schwartman, D. W. and Volk, T. (1991). Biotic enhancement of weathering and surface temperatures on earth since the origin of life. *Palaeogeography*, 90:457–460.

[Silver and Di Paolo, 2006] Silver, M. and Di Paolo, E. (2006). Spatial effects favour the evolution of niche construction. *Theoretical Population Biology*, 70(4):387–400.

[Staley, 2002] Staley, M. (2002). Darwinian selection leads to gaia. *Journal of Theoretical Biology*, pages 35–46.

[Stocker, 1995] Stocker, S. (1995). Regarding mutations in daisyworld models. *Journal of Theoretical Biology*, 175:495–501.

[Sugimoto, 2002] Sugimoto, T. (2002). Darwinian evolution does not rule out the gaia hypothesis. *Journal of Theoretical Biology*, 218:447–455.

[Thomas et al., 2004] Thomas, R., Basios, V., Eiswirth, M., Kruel, T., and Rossler, O. E. (2004). Hyperchaos of arbitrary order generated by a single feedback circuit, and the emergence of chaotic walks. *Chaos*, 14(3).

[Thomas and D'Ari, 1990] Thomas, R. and D'Ari, R. (1990). *Biological Feedback*. CRC Press.

[Tonaizzo et al., 2004] Tonaizzo, T., Lenton, T. M., Cox, P. M., and Gregory, J. (2004). Non-equilibrium thermodynamics and the production of entropy: Life, earth and beyond. In Kleidon, A. and Lorenz, R. D., editors, *Entropy and Gaia: Is there a link between MEP and self-regulation in the climate system?*, pages 223–239. Springer, Berlin.

[Vernadsky, 1926] Vernadsky, V. I. (1926). *The Biosphere*. Copernicus, New York, English translation, 1998.

[Volk, 1998] Volk, T. (1998). *Gaia's Body: Toward a Physiology of Earth*. Springer-Verlag, Berlin.

[Volk, 2002] Volk, T. (2002). Toward a future for gaia theory. *Climatic Change*, 52:423–430.

[Volk, 2003] Volk, T. (2003). Natural selection, gaia, and inadvertant by-products. *Climatic Change*, 58:13–19.

[Volk, 2007] Volk, T. (2007). The properties of organisms are not tunable parameters selected because they create maximum entropy production on the biosphere scale: A by-product framework in response to kleidon. *Climatic Change*, 85:251–258.

[von Bloh et al., 1997] von Bloh, W., Block, A., H, and Schellnhuber, H. J. (1997). Self stabilization of the biosphere under global change: a tutorial geophysiological approach. *Tellus*, 49B:249–262.

[Waddington, 1957] Waddington, C. H. (1957). *The Strategy of the Genes. A Discussion of Some Aspects of Theoretical Biology*. George Allen and Unwin, London.

[Watson and Lovelock, 1983] Watson, A. J. and Lovelock, J. E. (1983). Biological homeostasis of the global environment: the parable of daisyworld. *Tellus Series B-Chemical and Physical Meteorology*, 35B:284–289.

[Weber and Robinson, 2004] Weber, S. L. and Robinson, J. M. (2004). *Scientists Debate Gaia: The Next Century*, chapter Daisyworld Homeostasis and the Earth System, pages 231–240. MIT Press, Cambridge MA.

[Wilkinson, 2003a] Wilkinson, D. M. (2003a). Catastrophes on daisyworld. *Trends in Ecology and Evolution*, 18(6):266–268.

[Wilkinson, 2003b] Wilkinson, D. M. (2003b). The fundamental processes in ecology: a thought experiment on extraterrestrial biospheres. *Biological Reviews*, 78(2):171–179.

[Williams and Lenton, 2007a] Williams, H. and Lenton, T. (2007a). The flask model: Emergence and disruption of nutriend recycling in a simulated microbial ecosystem. *Oikos*, 116(7):1087–1105.

[Williams and Noble, 2005] Williams, H. and Noble, J. (2005). Evolution and the regulation of environmental variables. In Capcarrere, M., Freitas, A., Bentley, J., Johnson, C., and Timmis, J., editors, *Proceedings of VIIIth European Conference on Artificial Life, ECAL 2005*, pages 332–342. Springer-Verlag.

[Williams and Lenton, 2007b] Williams, H. T. P. and Lenton, T. M. (2007b). Artificial selection of simulated microbial ecosystems. *Proceedings National Acadamies of Science*, 104:8918–8923.

[Williams and Lenton, 2008] Williams, H. T. P. and Lenton, T. M. (2008). Environmental regulation in a network of simulated microbial ecosystems. *Proceedings National Acadamies of Science*, 105:10432–10437.

[Wilson and Dugatkin, 1997] Wilson, D. S. and Dugatkin, L. A. (1997). Group selection and assortative interactions. *American Naturalist*, 149:336–351.

[Wood et al., 2008] Wood, A. J., Ackland, G. J., Dyke, J. G., Williams, H. T. P., and Lenton, T. M. (2008). Daisyworld: a review. *Reviews of Geophysics*, 46.

[Wood et al., 2006] Wood, A. J., Ackland, G. J., and Lenton, T. M. (2006). Mutation of albedo and growth response produces oscillations in a spatial daisyworld. *Journal of Theoretical Biology*, 242(1):188–198.

[Wood and Coe, 2007] Wood, A. J. and Coe, J. B. (2007). A fitness based analysis of daisyworld. *Journal of Theoretical Biology*, 249(2):190–197.

[Wright et al., 2002] Wright, J. P., Jones, C. G., and Flecker, A. S. (2002). An ecosystem engineer, the beaver, increases species richness at the landscape scale. *Oecologica*, 132:96–101.

[Wright, 1948] Wright, S. (1948). On the roles and directed and random changes in gene frequency in the genetics of natural poplulations. *Evolution*, 2:279–294.