Evolution and Learning in Heterogeneous Environments

Author:
Daniel Jones

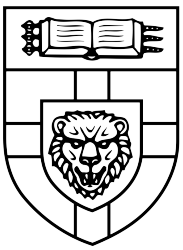

Doctor of Philosophy
Department of Computing
Goldsmiths, University of London

September 2015

Abstract

A real-world environment is complex and non-uniform, varying over space and time. This thesis demonstrates the impact of such environmental heterogeneity upon the ways in which organisms acquire information about the world, via a series of individual-based computational models that apply progressively more detailed forms of environmental structure to understand the causal impact of four distinct environmental factors: temporal variability; task complexity; population structure; and spatial heterogeneity.

We define a baseline model, comprised of an evolving population of polygenic individuals that can follow three learning modes: innate behaviour, in which an organism acts according to its genetically-encoded traits; individual learning, in which an organism engages in trial-and-error to modify its inherited behaviours; and social learning, in which an individual mimics the behaviours of its peers.

This model is used to show that environmental variability and task complexity affect the adaptive success of each learning mode, with social learning only arising as a dominant strategy in environments of median variability and complexity. Beyond a certain complexity threshold, individual learning is shown to be the sole dominant strategy. Social learning is shown to play a beneficial role following a sudden environmental change, contributing to the dissemination of novel traits in a population of poorly-adapted individuals.

Introducing population structure in the form of a k-regular graph, we show that bounded and rigid neighbourhood relationships can have deleterious effects on a population, diminishing its evolutionary rate and equilibrium fitness, and, in some cases, preventing the population from crossing a fitness valley to a global optimum. A larger neighbourhood size is shown to increase the effectiveness of social learning, and results in a more rapid evolutionary convergence rate.

The research subsequently focuses on spatially heterogeneous environments, proposing a new method of constructing an environment characterised by two key metrics derived from landscape ecology, "patchiness" and "gradient". We show that spatial complexity slows the rate of genetic adaptation when movement is restricted, but can increase the rate of evolution for mobile individuals. Social learning is shown to be particularly beneficial within heterogeneous environments, particularly when mobility is restricted, suggesting that phenotypic plasticity may act as a substitute for mobility.

Declaration

I hereby declare that the work presented in this thesis is entirely my own.

Daniel Jones
London
7 September 2015

Acknowledgements

I would like to thank Tim Blackwell, whose supportive and rigorous guidance was fundamental in shaping this research; and Mark d'Inverno, whose encouragement in the precursors to this work rendered it possible.

Thanks to all of the faculty and staff at the Department of Computing, whose ecosystem provided a thriving backdrop to work within. I would like to particularly thank Christophe Rhodes, for insights into some of the more intractable problems; and Spiros Andreou, for providing the cluster computing infrastructure that powered the simulations used in these experiments. Many thanks to my examiners, James Marshall and Peter Bentley, whose detailed comments have strengthened the dissertation immensely, and to Liam McNamara for feedback on earlier drafts. This research would not have been possible without the financial support provided by the Engineering and Physical Sciences Research Council (grant number EP/P503418/1).

I must acknowledge my gratitude to the Santa Fe Institute Complex Systems Summer School, for seeding the impetus behind these lines of thought, and to Richard Goldstein and Kyriakos Kentzoglanakis at the National Institute for Medical Research, whose passion and acuity provided a further intellectual catalyst.

Many thanks to my parents, for imparting curiosity from the very start, and in continuing to provide encouragement through the winding routes that led to this point.

Finally, to Julia, for her patience, support and perspective throughout the duration of this work, for which I owe a debt of boundless gratitude.

Contents

List of Figures

Chapter 1: Introduction

1.1 Overview

This research presented in this dissertation concerns the impact that environmental factors have on the ways in which an organism acquires information. Specifically, it proposes that environmental variability and heterogeneity affect the modes of learning that natural selection will give rise to, demonstrating that different forms of environmental structure produce quantifiably distinct combinations of innate behaviour, individual learning, and social learning. We adopt a computational approach, producing a series of individual-based models of increasing complexity, exploring incrementally more detailed types of environmental heterogeneity to understand how each class of heterogeneity modifies the population's baseline dynamics. These results are used to make predictions about the impacts of environmental heterogeneity in empirical contexts.

This chapter provides an introductory overview, beginning with a broad outline of the concepts under examination (§1.2). A statement of the overarching hypotheses is given (§1.3), outlining the particular environmental properties that are to be investigated. The research methods are summarised (§1.4), including the motivations, scope and limitations of the computational models adopted. Finally, the structure of the dissertation is described (§1.5), giving an overview of each chapter and its role in the context of the overall investigation.

1.2 Conceptual Context

The world is an uncertain place. An animal in the wild inhabits a landscape that is subject to continuous, unpredictable change (Wiens, 2000; Tilman et al., 1997); it may awake to find its world transformed by a major natural event. Even its own activities within the world cause persistent and long-lasting consequences, reshaping its own selective pressures and those of its neighbours (Odling-Smee et al., 2003). Environmental heterogeneity can manifest itself in many different ways: as gradual changes, sudden transformations, spatial patterns, and the social context constituted by an individual's kin. All of these factors have ramifications for the survival of a species in the wild.

The reproductive success of an individual is determined by how well-adapted it is to its environment, following the functional blueprint of its genetic material. As this material remains invariant across an individual's lifetime, the variability and heterogeneity of a real-world ecosystem poses a challenge for evolution, which must continue to maximise an individual's fitness under uncertain circumstances.

The solution that occurs in nature is for organisms to have the potential to produce a range of different phenotypes in response to changing environmental pressures (West-Eberhard, 1989; Scheiner, 1993). This phenotypic plasticity enables an individual to adapt its functional properties throughout its lifetime, improving its fitness across uncertain landscapes. The long-term dynamics of evolutionary development are thus buffered by short-term feedback through an organism's growth and development, providing adaptive robustness against change (Dukas, 2013).

The most sophisticated example of phenotypic plasticity is the capacity to learn novel behaviours. By trying out different approaches to interacting with the world and sampling their outcomes, an organism can gain insights into the costs and benefits of novel environmental resources (Thorpe, 1956). Rather than hard-code reaction norms to environmental stimuli within an individual's genotype, the evolutionary process may defer to uncertainty by providing the capacity to modify behaviours based on prior experience.

Learning in isolation is vital for the discovery of new traits, but poses additional costs, in the form of risks to the individual and time taken to accumulate experience. To mitigate these costs, many organisms engage in social learning, observing the behaviour of their peers and imitating novel traits (Heyes, 1994; Galef and Laland, 2005). Alongside reducing the risks of individual learning, social learning renders it possible for learned traits to be transmitted across generations, enabling useful behaviours to be ingrained within a population's collective memory (Laland and Hoppitt, 2003; Heyes and Galef Jr, 1996). However, social learning has additional cost/benefit trade-offs, including the potential for scrounging behaviours (Vickery et al., 1991; Galef and Laland, 2005) and the risk of acquiring outdated or inaccurate information (Giraldeau et al., 2002).

Each of these three modes of information acquisition – innate behaviour, individual learning, and social learning – possesses adaptive benefits and drawbacks in different situations. Each can arise at various degrees, with animals exhibiting a spectrum of behaviours from the thoroughly instinctive to those that are progressively learned (Thorpe, 1956; Lorenz, 1971). The research that follows seeks to understand how the balance between the three is affected by the ecological context that a species evolves within.

1.3 Statement of Thesis

The thesis proposed in this dissertation is that environmental heterogeneity has a measurable impact on the ways in which organisms acquire information about the world. Four types of environmental factors are identified: temporal variability; task complexity; population structure; and spatial heterogeneity. Each factor is studied in turn by constructing a series of computational models that exemplify their properties, thus quantifying how each factor is likely to influence the optimal balance of innate behaviour, individual learning, and social learning,

Specific hypotheses are given under each of the four environmental factors below. Some effects on evolution and learning are self-evident. We should, of course, expect to see general behavioural plasticity increase in environments of greater heterogeneity (Pigliucci, 2001), a prediction implicit in theoretical work as early as that of Wright (1931) and supported by the empirical work of (e.g.) Baythavong (2011). More details on the expected effects of each factor are given below.

The objective is to understand the impact of four different environmental factors, each of which will be represented and tested within the model. Summary definitions are given below; these will be expanded upon when they are defined within the models in Chapters 3, 4, 5 and 6.

- **Variability**: The rate of change that the environment exhibits (studied in §4), modelled as stochastic alternations to the environmental fitness function that are shared across the population. Particular focus will be given to the case of **sudden environmental change**, individual events that substantially alter the fitness landscape, a subclass of variability that has distinct ramifications and is often the subject of separate empirical study.

  Environmental variability is likely to select for modes of information acquisition that act on shorter timescales. Anderson (1995) considers learning to be of transient benefit in a static environment, but persists in one that varies; Feldman et al. (1996) and Kendal et al. (2009) conclude that social learning is likely to outcompete asocial learners when environmental variation is low. In an empirical study, Hallsson and Björklund (2012) show that a gradual change in conditions leads to an increase in phenotypic plasticity and in genetic variance, but that subsequent selection leads to a *decrease* in later plasticity.

- **Task complexity**: The degree of difficulty (and hence learning time) of the survival pressure posed by the environment (§4), modelled as the number of subtraits required to reach the global fitness optimum.

  A higher task complexity determines the learning time of a trait. Some particularly sophisticated natural traits (Dukas and Visscher, 1994; Marler and Slabbekoorn, 2004; Tebbich et al., 2001) are shown to be learned over extensive proportions of the individuals' lives. This predicts that task complexity may predict the level of plasticity and learning exerted in performing a task.

- **Population structure**: The structure of the society that an inhabitant resides within, bounding its interactants and defining its social environment (§5), modelled as a

fixed social graph.

Evolutionary graph theory (e.g. Taylor et al. (2007), Lehmann et al. (2007)) predicts that population structure can transform the equilibrium behavioural modes, fostering cooperation as a dominant strategy. In this context, it may suggest a higher prevalence of social learning.

- **Spatial heterogeneity**: The variance and patterning of selection across distinct areas of the environment, creating a disparity in adaptive requirements between individuals (§6), modelled with a spatially-extended 2D plane with fitness landscape imposed by a novel generator function. This will also encompass **task heterogeneity**, in which a number of different selection pressures are imposed concurrently.

  Spatial heterogeneity is known to produce plastic phenotypic responses (Baythavong, 2011), dependent on the type of heterogeneity displayed: Baythavong (2011) identifies a continuum between fine-grained varying environments, which result in raised phenotypic plasticity, and coarse-grained heterogeneity, which produces an adaptive genetic response; we hope to investigate this continuum and extend its predictions to distinct forms of heterogeneity. Furthermore, Rainey and Travisano (1998) concludes that environmental heterogeneity produces and sustains genetic polymorphism and population structure via studies on bacterial cultures.

These are introduced sequentially through this thesis. Chapter 4 begins by studying the effect of **variability** (including sudden environmental change) and **task complexity**, environmental factors that are experienced uniformly across the population, with every individual subject to the same fitness demands. Chapter 5 introduces **population structure**, investigating the effect that social environment has upon modes of learning. Chapter 6 introduces **spatial heterogeneity** (and task heterogeneity), siting the population on a spatially-extended 2D plane and examining the effect of varying fitness demands based on patchiness and fragmentation metrics drawn from landscape ecology.

For simplicity, we will henceforth use the term "heterogeneous" as shorthand for "spatiotemporally heterogeneous": that is, varying over space and/or time.

1.4 Motivations, Methods and Scope

This work seeks to address the under-representation of environmental factors within models of behavioural ecology. Many authors (Gordon, 2011; Laland et al., 2012; Marquet et al., 1993) have called for a closer integration between the fields of ecology (and landscape modelling) and behavioural studies, capturing both the environmental variation studied within ecology and the patterns of behaviour that this variation selects for. Levin (1992) identifies pattern and scale as the key problem in ecology, yet it is omitted from most theoretical models. This thesis demonstrates different types of impact that environmental heterogeneity can have upon patterns of evolution and learning.

It does so via computational simulation, creating a series of individual-based models that are intended to capture general evolutionary dynamics rather than the behaviours of a particular species. Each of these decisions is explained below.

Methods: Simulation

Simulation is an increasingly popular approach in computational biology and behavioural ecology, allowing for the prediction of real-world behaviours via the abstraction of observed processes (Bentley, 2009). There are multiple motivations for selecting modelling above empirical studies. The prime reason is that it allows for experiments that are difficult or impossible to carry out *in vivo* (Hartmann, 1996; Rohrlich, 1990) due to prohibitive timescales or other pragmatic issues. A limited number of empirical studies have been performed on learning-evolution interactions across multi-generational timescales (Dukas, 2008), and on the impact of environmental heterogeneity on evolutionary trajectories (Rainey and Travisano, 1998). However, to explore both factors simultaneously, across a range of quantifiable environmental factors, in a species sophisticated enough to also engage in social learning, would require a great deal of time and resources. This scope of research is thus an appropriate candidate for simulation.

Modelling these experiments in computer simulation allows us to ask what-if questions about states and processes in a less constrained manner than empirical studies, acting as a heuristic tool (Hartmann, 1996) that can potentially lead to unforeseen generalities underlying a system prior to discovering empirical or analytic answers (Wolfram, 1994). Particularly in cases when causal factors are complex or not well-understood, factors can be removed or isolated to better understand the causal pathways at play (Hartmann, 1996); here, we hope to better understand how complementary properties of environmental variability may affect learning and evolutionary tendencies.

In general, simulation can play a useful role bridging the gap between the purely empirical and purely theoretical (Humphreys, 1994; Winsberg, 2001), acting as an "opaque thought experiment" (Di Paolo et al., 2000) that can provide explanatory power when a systemic behaviour does not obviously follow from the interactions between its elements, closing the cognitive gap between model and scientist, and fostering creativity within the scientific process (Jones and d'Inverno, 2011).

Methods: Individual-Based Modelling

Individual-based models are commonly used to model populations in ecology and evolutionary studies (Grimm and Railsback, 2013) because they can accurately capture the dynamics of bottom-up, distributed phenomena such as genetic evolution and social behaviour, in which systemic properties arise from numerous individual interactions (Grimm et al., 2005). Unlike a mean-field model, which make assumptions of uniformity of behaviour, adopting the individual-based metaphor allows each simulated organism to possess a distinct set of states, a crucial property in discovering how differential types of behavioural tendencies can arise in a mixed population.

This thesis is also concerned with modelling particular types of landscape heterogeneity. Although the effect of space and environmental pattern is known to be a major factor in ecology (Levin, 1992; Fogarty et al., 2012), it is typically omitted from theoretical models of evolutionary and learning (Volterra, 1928; Hinton and Nowlan, 1987). Some analytical models have shown evolutionary consequences of general environmental variability (Borenstein et al., 2008; Feldman et al., 1996; Wakano and Aoki, 2006) and heterogeneity (McNamara et al., 2011; Boyd and Richerson, 1988; Zhivotovsky et al., 1996), but typically

as a single discrete category or axis. More sophisticated analytic approaches to modelling heterogeneity are possible, cf. (Taylor et al., 1993), but would not support the bottom-up social dynamics that we seek to explore.

To examine distinct classes of environmental heterogeneity, with properties determined by continuous variables (for example, testing a range of environmental task complexity or rates of change), requires a more detailed conception of space, in which individuals may experience different selection pressures following quantifiable structures. This is addressed by extending the individual-based model population with structure and an explicit form of space; initially placing the population on a graph structure, and subsequently extending the population across a 2D environment, with spatial patterning properties determined by quantifiable factors.

Methods: Species Agnosticism

The models described in this thesis are not representative of any particular animal species or environment, and have not been parametrised or calibrated as such. The objective is to create a model framework that makes as few assumptions as possible about the world whilst still demonstrating the phenomena in question, abstracting from micro-level details in order to be applicable to as wide a range of scenarios as possible (Bedau, 1999). This follows the principle of Occam's razor, in which an explanation should always be sought that makes no more assumptions than are necessary.

A number of other assumptions and simplifications are made in order to be maximally parsimonious with the assumptions and complexity encoded within the model. These include adopting a naive trial-and-error approach to learning, assuming full learnability (and evolvability) of all traits, simple unimodal fitness landscapes, and not imposing any explicit costs on learning activities. A critical evaluation of the work's limitations is given in Section 7.4.

1.5 Dissertation Overview

Chapter 2 presents a review of existing research within the field. We look first at ecological models which incorporate evolution, learning and social processes, and subsequently at those which involve a heterogeneous, fluctuating environment. We review metrics used within landscape ecology to quantify and model environmental heterogeneity.

Chapter 3 describes a novel theoretical model produced in support of this thesis, modelling a population of agents that are capable of engaging in a combination of innate behaviour, individual learning, and social learning. This baseline model is developed incrementally in subsequent sections, with the introduction of new structural and spatial properties.

In Chapter 4, an initial set of experiments are performed using this baseline model, exploring the impact of differing task complexity and environmental rates of change on the population's learning dynamics.

Chapter 5 extends the baseline model with the notion of population structure. We model social interactions by siting our population first on 1-dimensional and 2-dimensional lattices, and subsequently on regular graphs of arbitrary degree k. This enables us to rigorously

interrogate the interactions between population structure and learning behaviours: are different forms of population structure likely to foster different modes of learning?

Chapter 6 introduces environments that are spatially explicit and spatio-temporally heterogeneous. We define a two properties of heterogeneity, with reference to literature from landscape ecology and spatial analysis, and describe a new approach to constructing a spatially heterogeneous environment. We explore the interactions between learning and evolution across environments of different types of heterogeneity.

Chapter 7 summarises the results and predictions produced through the thesis, considers their limitations and threats to validity, and provides some suggestions for further work.

Finally, Appendix A presents further experimental results which are out of scope of the main body of the text; Appendix B publishes the parameter values used in each of the experiments detailed within; and Appendix C documents the simulation and analysis implementation details.

Chapter 2: Background and Related Work

2.1 Overview

In the following chapter, we review the current state of research informing this thesis, both empirical and theoretical. This includes ecological and zoological studies, computational and mathematical models, and theoretical discussions from the philosophy of biology.

We are concerned with how environmental variance affects the way that information is transmitted in biology. We begin by omitting environmental factors and looking at existing research on the baseline interactions between evolution and learning (§2.2), distinguishing between individual and social modes of learning. We discuss the Baldwin effect, a proposed mechanism in which learning can guide evolution. Results from empirical and theoretical studies are reviewed.

We then proceed to look at existing literature on the impact of environmental factors on learning and evolution (§2.3), focusing on each of our four key environmental properties: variability; task complexity; population structure; and spatial heterogeneity.

We finally review the analysis and modelling of environmental structure in ecology (§2.4), which will be relevant when constructing our own environmental models later in this thesis. We evaluate common landscape metrics, drawing on research from landscape ecology and spatial analysis.

2.2 Evolution, Learning and Sociality

2.2.1 Learning and Plasticity

The title of this thesis uses the term 'learning': the acquisition and modification of behavioural traits during an organism's lifetime, usually driven by evaluating alternative behaviours through trial or observation.

This is part of a wider class of biological capabilities under the umbrella of 'phenotypic plasticity' (Bradshaw, 1965; West-Eberhard, 1989), a responsive adaptability that was recognised as a crucial buffering factor in the models of evolution as far back as Wright (1931). As described by West-Eberhard (1989):

> "*Phenotypic plasticity* is the ability of a single genotype to produce more than one alternative form of morphology, physiological state, and/or behaviour in response to environment conditions." (p249 West-Eberhard, 1989, author's emphasis)

An example of mechanistic plasticity is the claw strength of certain crabs, which alters

adaptively based on the shell thickness of the mussels that they consume (Smith and Palmer, 1994; Agrawal, 2001). Consequently, mussel populations may begin to accordingly select for increased shell thicknesses, resulting in an evolutionary "arms race".

Those classes of change that act within the timespan of an organism's lifetime are known as *labile* plasticity (Scheiner, 1993). Conversely, changes which occur once during development, and remain constant for the rest of an organism's lifetime, are termed *fixed* plasticity. Here, variance occurs between, rather than within individuals (Gupta and Lewontin, 1982). The focus of this thesis is upon the former case, on traits and individuals that are able to continually adapt within their lifespan.

It is now understood that phenotypic plasticity is not a marginal concern, but a property that is universal across living creatures (Pigliucci, 2001). Moreover, plasticity itself evolves to maximise fitness in variable environments (West-Eberhard, 1989), acting as an adaptive trait that can be selected for (Bradshaw, 1965; Agrawal, 2001): the greater the degree of change or heterogeneity imposed by the environment, the greater the degree of plasticity that should result (Lefebvre and Palameta, 1988; Wcislo, 1989; Lewontin, 1991).

As well as environments predicting plasticity, understanding the dynamics of the plasticity exhibited by a species can tell us many things about the properties of the environment that it has evolved within. Charmantier (2008), for example, indicates that the phenotypes of avian populations have been reliably and rapidly correlated with changing environments, and suggests that they could be used as an early indicator of climate change.

Learning is a particularly powerful case of phenotypic plasticity in that it allows for the encoding of more complex stimulus-effect links (Thorpe, 1956; Pearce, 2013), allowing an individual to modify its behaviour over short timespans based on conditioning developed from previous encounters. This is in contrast to the gradual changes that typify morphological phenotypic plasticity (Pigliucci, 2001), which are often one-dimensional and sometimes irreversible. Moreover, individuals can engage in trial-and-error trials to innovate novel traits, selectively adopting those that result in some fitness payoff (Pearce, 2013) and thus creating the potential for new behavioural discoveries.

Miller and Todd (1991) identify three adaptive functions of learning. It can enable adaptations that respond to environmental changes at timescales quicker than evolution would otherwise be unable to track, expressing flexibility and reversibility over rapid periods (Pigliucci, 2001), and so providing fitness advantages over periods of major environmental upheaval. It can enable the organism to overcome the size limitations of the genotype by exploiting environmental regularities, extending phenotypic reach in ways that may otherwise not be evolutionarily possible, including acting as the basis of knowledge development via cumulative culture (Boyd and Richerson, 1985). It may also act to help and guide evolution, escaping local fitness optima and potentially changing the evolutionary trajectory of a species (Nolfi and Parisi, 1996).

The first of these capacities – learning's affordance of responsivity on short timescales – is key to its adaptive importance, and will be discussed in the following section. We will return to learning's potential for guiding evolution in Section 2.2.3.

2.2.2 Behavioural Optimality

Evolution operates by selecting on variance between generations, refining behaviours and optimising towards maximum net reproductive success, or fitness. Natural selection tunes an individual's traits to optimise the cost/benefit tradeoff that is experienced in developing and exhibiting a trait (MacArthur and Pianka, 1966; Davies et al., 2012); the thickness of a lobster's shell is selected to provide the optimal balance of defensive strength versus ontogenic and metabolic expenses.

For mechanical traits that may have a relatively small set of interactions with other facets of behaviour, optimal selection is a relatively uncomplicated process. A trait should express a mean value that most closely resembles the survival challenges likely to be imposed by the environment. This can be extended via phenotypic plasticity to allow for a trait to change its morphology based on experience, such as in the development of crab claw strength in response to mussel shell thickness, and vice versa (Smith and Palmer, 1994) – a phenomenon described as adaptive phenotypic plasticity (Bradshaw, 1965)

The optimality argument also predicts that learning capacities should likewise be optimised for the uncertainty that is likely to be experienced within the environment. This optimality is exploited throughout contemporary models of animal ecology and evolution: a behaviour that is shown to provide optimal net fitness within an evolutionary model is likely to have the same benefits within a real-world scenario, and so we can predict that behaviour to arise in similar contexts. Thus, if a model shows that a bird should spend a particular length of time nesting to optimise metabolic payoffs, it predicts that evolution will have selected for the same behaviour. Optimality has been successfully used to analyse decisions about foraging (MacArthur and Pianka, 1966) and mating (Davies et al., 2012).

We exploit this assumption of optimality in the models developed within this thesis, which make claims to predictive power by demonstrating the optimum balance of learning modes under particular sets of environmental contexts. However, there are limitations and caveats implicit when assuming optimality (Davies et al., 2012). Having incomplete knowledge of a scenario's cost/benefit tradeoffs may result in important factors being omitted, skewing the predictions of the model. This is particularly critical when considering behavioural traits, which can involve the nonlinear interactions of many different costs and considerations, including life history variance; a forager may make riskier decisions when it is desperate for food. It is also possible that particular behavioural traits may not be particularly well-tuned by the process of natural selection, or unable to reflect rapid changes in environmental requirements.

2.2.3 The Baldwin Effect

When evolutionary systems are extended with behavioural plasticity, we should expect some interesting interactions to arise. One that came to the attention of the first generation of evolutionary theorists after Darwin (Baldwin, 1896; Morgan, 1896) is the "Baldwin effect" (Simpson, 1953), a proposed mechanism in which learning guides the direction of evolution.

The general pattern encapsulated within the Baldwin effect is as follows.

1. A population arises in which some trait P becomes beneficial.

2. Some individuals arise which, through their phenotypic plasticity, are able to learn P.

3. In some of these individuals, the trait P becomes innate (*genetic assimilation*).

With the assumption that innate behaviours are less costly than those which are plastic, we would then expect selective pressure to lessen on these particular learning capabilities: if we can accomplish a task innately, we no longer need to be able to maintain the potentially expensive metabolic apparatus needed to learn it (West-Eberhard, 2003). This weakening of genetic selection is known as *shielding*, and is another phenomenon that we should expect to see in a simulation, if these hypotheses hold.

It is not obvious that such a process should occur; why should a population which is disposed to be able to learn a trait be disposed to acquire genetically? The machineries for learning and those for genetic acquisition may be mutually distinct, and there is no clear reason why a behavioural adaptation should have a causal impact on genetic change.

Godfrey-Smith (2003) argues that such a scenario is likely to arise in the simple case that, if the selection pressure for trait P becomes so great that a population will not survive without it, being able to learn P may buy the organisms enough time to develop it genetically. This is termed the "breathing space" scenario.

Papineau (2005) elaborates with a potential formalisation of the genetic assimilation process:

1. Suppose we have an advantageous trait P, which requires subtraits I. Every one of these subtraits is individually necessary and jointly sufficient to produce P. Without the other Is, however, each individual subtrait does not provide a selective advantage.

2. Suppose that each subtrait can either be performed by a learned mechanism (I_L), or by genetic instinct (I_G).

3. Each subtrait is difficult to obtain by learning (I_L), but even more unlikely to occur through genetic mutation (I_G). Conversely, once I has been acquired, it is *cheaper* to express through I_G, and more costly via I_L (through the aforementioned costs of learning: maintenance of machinery, time taken over trial and error, etc).

4. Without any Is at all, the selective advantage of any individual new I remains zero.

5. Now, assume that this probability differential renders the likelihood of obtaining *all* I_Gs is effectively nil, but the likelihood of obtaining all I_Ls is merely very small. Given the cheaper cost of innate ability versus learning, an individual who does succeed in learning P will suddenly experience a selective advantage of developing *any* I_G.

6. This I_G can then be passed to offspring. Now, a second benefit of I_G emerges: as one less subtrait is required, the process of learning the remaining I_Ls becomes cheaper.

The steps described above can be summarised as a kind of probabilistic piggybacking: a highly beneficial but negligibly unlikely scenario can be obtained via one that is less beneficial but more likely to occur. An example is the innate ability of woodpecker finches to utilise tools to collect grubs from the holes bored by woodpeckers (Tebbich et al., 2001). This requires the ability to *(i)* obtain a suitable tool, *(ii)* observe the creation of a hole, or at least its location, and *(iii)* apply it to scoop out said grubs. It is unlikely that each of

these genetic traits would arise independently, but plausible that, through individual or cumulative imitative learning, they could collectively be transmitted to the same individual.

Further specific examples of Baldwin-like phenomena have been discovered in nature (Wcislo, 1989; Agrawal, 2001; Simpson, 1953). But why should this seemingly obscure phenomenon have broader significance? One reason is that it can be seen as giving foresight to the typically blind process of natural selection, by permitting an organism to explore the region around the normal space delineated by its genotype. As John Maynard Smith (1987) observes,

> "...finding the optimal [solution] in the absence of learning is like searching for a needle in a haystack. With learning, it is like searching for the needle when someone tells you when you are getting close." (Smith, 1987, p762).

In other words, learning can act as a dowsing rod that points evolution in the direction of likely opportunities. We shall proceed to examine how this is further developed by the introduction of social learning.

2.2.4 Social Learning and Genetic Assimilation

Social learning (Zentall and Galef, 1988) is defined as the acquisition of behaviours from other individuals; usually members of an organism's species, though occasionally observed in inter-specific interactions (Heyes and Galef Jr, 1996; Seppänen and Forsman, 2007). Predominantly applied to behavioural traits, it may take place through mimicry, goal-based conditioning, cognitive reasoning (in the case of humans), or other distinct mechanisms (Lefebvre and Palameta, 1988; Papineau, 2005).

Historically thought to be a relatively infrequent curio (Galef, 1976; Boyd and Richerson, 1983, 1985), social learning has been demonstrated to empirically occur in a range of species, from chimpanzees, rodents and finches (Galef and Laland, 2005) to multiple different species of insect, both colonial (Lancet and Dukas, 2012; Sarin and Dukas, 2009) and noncolonial (Coolen et al., 2005).

Notable in many of these cases is the potential for culture to develop cumulatively, spreading from individual to individual in a chain- or web-like fashion. A further delineation is necessary to distinguish between these two separate forms of information transmission. Where the passage of genetic information can be described as *vertical* (from parent to child), the peer-to-peer exchange of traits within a generation is described as *horizontal* transmission (Boyd and Richerson, 1985; Cavalli-Sforza, 1981).

Boyd and Richerson assert that this persistence across generations is the factor that distinguishes cultural learning from other forms of phenotypic plasticity; it has "population-level consequences" (Boyd and Richerson, 1985, p4), persisting beyond the death of an individual and potentially resulting in traditions which are passed along to endure for generations (Avital and Jablonka, 2000; Galef, 1990).

A number of theoretical studies have been carried out on cultural evolution (Boyd and Richerson, 1985; Cavalli-Sforza, 1981; Lowen and Dunbar, 1996; Best, 1999), in which cultural information spreads either horizontally, or obliquely: with traits transmitted from the previous generation, whether related or unrelated. Predominantly, these focus on transmission of social traits in the form of directional bias to individual tendencies

(Cavalli-Sforza, 1981; Belew, 1990), taking mean-field, population-level approaches (Boyd and Richerson, 1985).

However, social learning comes with its own cost/benefit tradeoffs. Giraldeau et al. (2002) discusses the potential for ingrained social learning to lead to suboptimal behaviour via information "cascades", in which outdated information percolates across the population, giving rise to maladaptive behaviours. There is also risk posed by the potential for cheaters to arise in the population, creating a producer-scrounger dynamic (Barnard and Sibly, 1981; Vickery et al., 1991; Kurvers et al., 2012; Arbilly and Laland, 2014; Dyer et al., 2008; Kameda and Nakanishi, 2002)

Papineau (2005) continues his analysis of the Baldwin effect by turning to behavioural traits that are mediated specifically by social learning. The argument is encapsulated by the inequality:

$$p(G) < p(L) < p(S) \tag{2.1}$$

Where $p(G)$ is the probability of exhibiting a trait innately, $p(L)$ is the probability of learning it through exploration, and $p(S)$ is the probability of acquiring it through social learning. In a scenario wherein it is effectively impossible to obtain all of our aforementioned subtraits I through evolution (I_G), and unlikely to obtain them through individual exploration (I_L), it may be significantly more probable that a suite of subtraits may be picked up through social imitation (I_S). Returning to the case of the woodpecker finches, it seems plausible to imagine a finch duplicating the tool-finding and tool-usage behaviour of its peers. Through cultural spread, therefore, the chances are higher of entering into a situation in which all of our I_S subtraits are in place, primed for genetic assimilation.

To verify the theory proposed within this heuristic model, we will proceed to look to mathematical and computational models that provide concrete implementations of these ideas.

2.2.5 Theoretical Models

Hinton and Nowlan and its derivatives

The watershed model of learning and evolution is Hinton and Nowlan's 1987 paper (Hinton and Nowlan, 1987) (henceforth "HN"), which extends the notion of the genetic algorithm to support a highly simplified form of learning. The motivation is to demonstrate the idea that acquired characteristics can affect the trajectories of gene distributions over time. They observe that this is biologically preferable because "a learning trial is much faster and requires much less expenditure of energy than the production of a whole organism". As this model is critical to the field, we will summarise it here.

- A population of agents is created, each with a genome $G \in g^N$, comprising N traits $g \in \{1, 0, ?\}$, and a target phenotype $G_T = 1^N$. Any phenotype other than G_T has an effective fitness of zero, whereas G_T has a higher fitness. ? genes are interpreted as being undefined, and subject to lifetime learning.

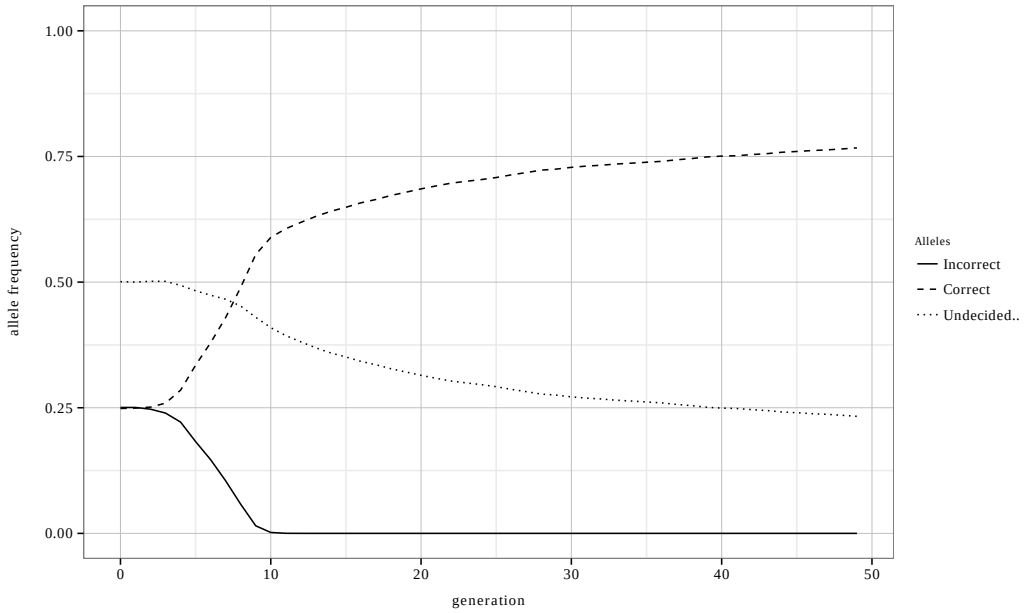

FIGURE 2.1: Gene frequencies in Hinton and Nowlan's learning model (1987)

- Each generation, every agent "learns" by repeatedly tossing q coins (1000 attempts) where q is the number of ? genes. If every trait $g = 1$, the agent has found the optimum and halts.
- Agents are selected to breed for the following generation with a fitness $1 + \frac{19n}{1000}$, where n is the number of learning trials remaining (out of 1000) after the optimum was encountered. This is thus equal to 1 for all agents which didn't locate the optimum, and 20 for those which found the optimum immediately.

The behaviour of the model is shown in Figure 2.1. A Baldwin-like effect should be evident. For the first few generations, the relative frequencies of each allele remains roughly constant: a few agents exist with genotypes solely made up of 1 and ? agents, with these distributions spreading via one-point recombination. After several more generations, recombination begins to replace these agents with those which have a higher relative proportion of 1s, due to the fitness differential caused by their more rapid learning rate (due to the fitness function's dependency on the number of trials). Ultimately, fixation occurs with a majority of alleles reaching 1. Some ?s do still exist; statistical analysis reveals (Belew, 1990) that this is because the *relative* fitness of a mostly-innate agent versus a completely-innate agent is small enough for evolution to be relatively neutral.

Though elegant, the simplicity of HN is also its weakness. The same analysis indicates that the model is highly parameter-dependent: if we double the number of ? alleles per agent, or halve the learning time, the optimum is never discovered, and the population languishes. The model still, however, succeeds in illustrating the general form of dynamics, and this model has acted as a springboard for numerous extensions.

Belew (1990) extends HN by introducing a basic form of social interchange. Successful agents can transmit a "cultural advantage" to their offspring – a purely vertical transmission – by inducing a bias to the likelihood of performing successful learning trials. A typical

member of the population has a 50% chance of guessing each of its learning alleles correctly, as in HN. Here, an individual that is the offspring of an individual that has managed to attain the optimum fitness (a "Winner": $G = 1^N$) has a higher probability of a correct guess. With a competitive advantage constant $CA = 0.1$, each allele is toggled to a 1 with a probability of 0.6.

The outcome is that convergence to optimality occurs much sooner. However, the resultant population harbours a greater number of ?s, indicating that the selection pressure towards full genetic assimilation has decreased in the face of a higher incidence of plasticity.

Best (1999) builds on these cultural ideas further by adopting overlapping generations, a truly horizontal mode of transmission, and a form of mimesis that does not always imply movement towards an optimum: in real-world scenarios, he argues, organisms are likely to mimic behaviours rather than strategies, and so are equally likely to learn deleterious traits. The outcome is that convergence accelerates once more, though with a similarly sub-optimal outcome.

Connectionist Models

A different approach is taken by Ackley and Littman (1991), who make use of neural networks and reinforcement learning to introduce a more continuous adaption process. In this model:

- a population of agents interacts with a hostile environment, populated with simple carnivorous enemies and metabolic resources
- each agent has a pair of neural networks: an *action* network A, to determine responses to stimuli, and an *evaluation* network E, to give responses to previous actions
- the weights of both are genetically determined, though E is fixed through the agent's lifetime, whereas A can be subsequently modified based upon the results of E's evaluations

To begin with, the majority of agents are eliminated as their random innate abilities lead them to deleterious behaviour: for instance, moving away from food or towards enemies. The few which survive proceed to replicate, with slight mutations in their descendants' connection weights producing modified behaviours. Following natural selection, latter generations thus have genetic capabilities that better reflect the world around them. However, plots indicate that behaviours continue to fluctuate, as the world and its selection pressures continue to change based on resource availability.

The genetically hard-coded E reflects an agent's evolved conception of how to evaluate its present state, with a positive value resulting in the back-propagated reinforcement (in A) of the previous timestep's action. This has been summarised as a scalar 'goodness' rating, roughly equivalent to an axis of pleasure versus pain.

An appeal of this strategy is that fitness cannot be derived directly from an organism's genotype, but emerges from a series of applications of A and E, effectively distinguishing between the choices that an individual makes within its lifetime, based on metabolism and learning, and the longer-term developmental cycles of innately programmed evaluative capacities in E.

Two major results arise from Ackley and Littman's work:

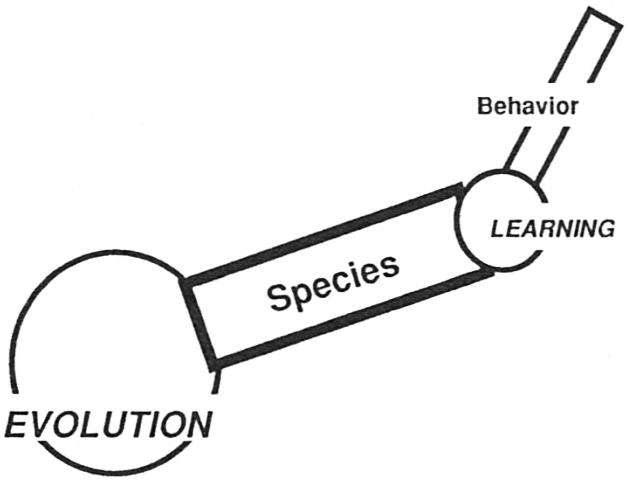

FIGURE 2.2: The "phenotypic limb", reproduced from (Belew, 1990, p15). The first joint is moved by the process of evolution, in which a species or gene pool becomes better adapted to an environment. The second joint is driven by learning and behaviour, enabling an individual to responsively alter its character through its lifetime, to adapt to changing environmental demands.

- that a population capable of both learning and evolutionary adaptivity greatly outlives one which has only one or the other
- that a learning function which is not directly coupled to fitness (that is, with a different internal criterion of success) can still serve to guide species-level evolution

However, perhaps due to computational limits of the time, the accompanying analysis based on a small set of experiments, and is brief and heuristic in its scope.

A more rigorous approach is taken by Nolfi et al. (1994), who use a similar connectionist-evolution framework with distinct behavioural and evaluative components. By applying analysis of variance to parallel situations with and without learning, the ability to learn certain traits is shown to correlate with a higher rate of genetic assimilation. Moreover, an evolved food-finding ability leads to an inherited predisposition towards learning.

Multi-Peaked Fitness Landscapes

Learning and evolution are known to have particularly complex non-linear interactions when acting upon structured, multi-peaked fitness landscapes. Borenstein et al. (2006) summarise that, in general, phenotypic plasticity serves to smooth a fitness landscape by enabling an individual to search the nearby fitness landscape area and escape local optima. Belew (1990) uses the metaphor of a phenotypic "limb" (Figure 2.2): a fixed joint leads to a specific location in functional space, with a flexible forearm of learned behaviour enabling flexibility around this axis.

Developing Hinton and Nowlan's (1987) single-optimum model, this notion is generalised to arbitrary multipeaked landscapes through analytic and numeric analysis (Borenstein et al., 2006). Three conclusions are drawn:

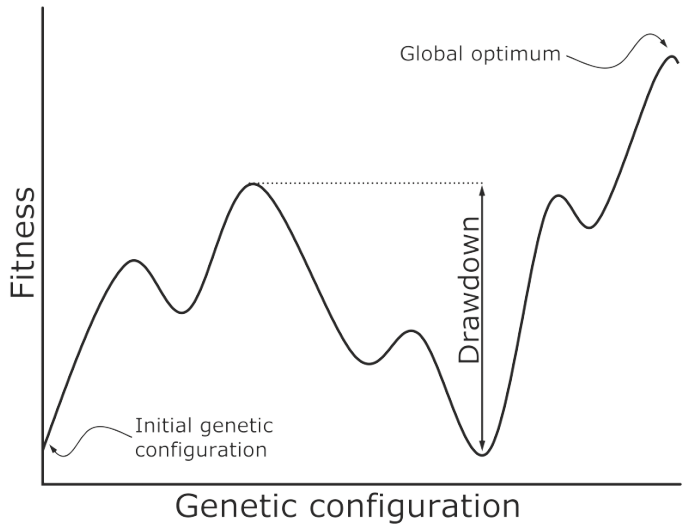

FIGURE 2.3: Fitness landscape and 'drawdown'.

- in a relatively rugged, multi-peaked (or rapidly changing) landscape, this smoothing effect serves to accelerate the movement towards optima
- in a generally flat (or static) landscape, behavioural plasticity actually slows genetic convergence to an optimum, by decreasing the relative fitness of higher points
- learning aside, the prime determinant of genetic convergence on a complex landscape is the maximal "drawdown", or the largest difference between a peak and trough encountered en route to an optimum (Figure 2.3). With learning, this effect is mitigated through smoothing.

Although the adaptive landscape metaphor is one of the dominant paradigms within population genetics, there are limitations to the metaphor. A fitness landscape is dependent on its neighbourhood operator, and the fitness function that generates the landscape, which must capture the complexity of interactions that take place within a real-world genome such as pleiotropy and other epistatic effects. Using a multi-valued function is generally insufficient to capture these real-world complexities. A full survey of this model should account for such interactions, and the neighbourhood relationships found in empirical gene interactions.

2.3 Environmental Heterogeneity: Types and Consequences

In Section 2.2, we looked at models and studies of evolution and learning that do not include environmental factors, or predominantly treat them as steady-state background values. This section proceeds to review existing research on the behavioural effects of the four major environmental factors identified as the focus of this thesis:

- **Variability**: environmental change over time, incorporating both gradual change and sudden transformative events

- **Environmental complexity**: the complexity and learning time of the selective pressures posed by the environment
- **Population structure**: bounding an individual's social environment based on a social interaction graph
- **Spatial heterogeneity**: extending the population into a spatially varying environment, in which individuals experience different selection pressures

We will look at the literature corresponding to each of these factors in turn, later returning to their results in the concluding chapter.

2.3.1 Variability

Environmental change has long been understood as a key factor that drives evolution (Levins, 1968), setting the pace of adaptive change as well as producing and maintaining phenotypic plasticity (Pigliucci, 2001). Much theoretical and empirical work has been done to understand the effects of temporal variance on evolution and learning.

Hallsson and Björklund (2012) perform a detailed study of the evolution of plasticity within insects, with rapid, continuous and fluctuating temperature changes imposed upon habitats over several generations. Their experiments confirm that a continuous change in temperature leads to the expected increase in genotypic variance and phenotypic plasticity, as does sudden fluctuation in temperatures. After 18 generations of selection, lineages that had been exposed to continuous temperature changes showed no change in their level of plasticity. Surprisingly, those lineages that had been exposed to major fluctuating changes had a *decreased* level of plasticity, perhaps indicating that the earlier plastic responses in fluctuating environments were actually maladaptive to the long-term survival of the species.

In an avian study, Dingemanse et al. (2004) show that fluctuating pressures in environments give rise to higher genetic variance, and that males individuals with a more exploratory propensity were frequently rewarded, though this result itself fluctuated between years.

Templeton and Rothman (1978) model evolution in changing environments, considering temporary variance to be a proxy for general fine-grained environmental heterogeneity. They claim that irregular fluctuations and "runs" of environmental states may be important in maintaining genetic polymorphism. However, their model does not incorporate plasticity, which is the typical evolutionary response to short-term change, and the predictions of the model have been disproven by more recent empirical work (e.g. Hallsson and Björklund (2012)).

One of the first models of to explicitly incorporate social learning within a fluctuating environment was created by Boyd and Richerson (1988), who show that social learning requires two conditions to dominate: (1) individual learning is inaccurate, and (2) peers experience a similar environment (spatially or temporally). The two factors are frequently correlated, with a low rate of change resulting in higher predictability and a better reliability of social learning.

Feldman et al. (1996) uses a single-locus population genetics model of individual versus social learning to show that greater probabilities of environmental change give a lower

adaptive success to social learning: if the environment changes state every generation, social learning is never seen to evolve, as the information received by a social learner becomes outdated. However, for longer periods of environmental fluctuation, social learning is seen to arise as a stable strategy in combination with a lower level of individual learning. This result is supported by Kendal et al. (2009)'s models of social learning strategies in changing environments, showing that social learning is likely to outcompete asocial learners when environmental variation is low.

Wakano et al. (2004) describe a mathematical model that allows any of innate behaviour, individual learning and social learning to arise. In static, unchanging environment, innate behaviour is the dominant strategy; in changing environments, either social or individual learning dominates. They later refine this work (Wakano and Aoki, 2006) to show that social learning is likely to arise in environments with a median rate of change, with individual learning prevailing in fast-changing environments.

Similarly using a two-state, stochastically fluctuating model of the environment, Borenstein et al. (2008) finds partial support for the three-regime result of Wakano and Aoki (2006). When these two states are switched between discretely, with no overlap, the model shows that a mix of individual and social learning is preferred for survival, in order to innovate and then disseminate the traits required to respond to the changing environmental demands. This appears to outcompete innate behaviour even in a relatively stable subsequent environment. However, if the two environmental states overlap, a pure social learning strategy can dominate.

Dyer and Bentley's (2002) PLANTWORLD model simulates the behaviour of a large population of plants within stable and varying environments. It demonstrates that environmental variability can be mitigated by the evolution of behavioural responses, preventing population decline in a suboptimal environment by switching to a strategy of "dormancy", coupling a population to its environment and therefore decoupling it from environmental variability.

Borg and Channon (2012) develop the baseline version of our model (Jones and Blackwell, 2011) to explore the effects of changing levels of environmental variability on population dynamics. They conclude that social learning is particularly advantageous when variability fluctuates, but that population collapse is likely to arise when extreme fluctuations take place within homogeneous populations.

Nolfi and Parisi (1996) uses a connectionist approach to model adaptation to changing environments, showing that neural networks evolved within changing environments show a greater propensity for learning than those that evolved in static conditions. This type of learning is out of the scope of the models described in this thesis, in which we will concentrate on single trial-and-error learning, with evaluation by comparison to a fitness function.

2.3.1.1 Sudden Environmental Change

Not all environmental change takes place gradually. Major events such as volcanic eruptions can impose a dramatic, unforeseen change on the demands and fitness pressures imposed by the environment; the current changes taking place to the environment are already causing rapid changes in selective pressures, evidenced from genetic and phenotypic

studies (Charmantier et al., 2008; Matesanz et al., 2010).

Various bodies of research have come to different conclusions about the impact of an unfavourable environment upon genetic diversity, with some indicating that a sudden change to a suboptimal environment will cause a decrease in genetic variability, and other studies showing an increase in variability (Hoffmann and Merilä, 1999). Several empirical studies of invasive species have shown variance in plasticity when the population enters an unfamiliar environment (Chapman et al., 2000; Lee et al., 2003); in general, species seem to vary in their response to novel and traumatic environmental changes, with some exhibiting increased plasticity, and others showing a more rapid genetic variance and selection.

A key theoretical study by Lande (2009) uses a quantitative genetics model to introduce a sudden major environmental change on a population, against a background of a low level of variance. The population is initially canalised, showing a high degree of specialisation to its environment. After the perturbation, plasticity increases to accommodate the change, allowing individuals the "breathing space" needed to survive the novel conditions, predicting an optimal level of plasticity that is proportional to the amount of change experienced. After the mean phenotype reaches the optimum, the required genes are slowly re-acquired, restoring fitness to its previous level.

Referring back to Hallsson and Björklund (2012), a secondary interesting result of major environmental fluctuations is that they have been shown to cause long-lasting effects on a lineage's phylogenetic development: selecting on the initial plasticity required to cope with the change can actually result in a subsequent, maladaptive reduction in plasticity.

2.3.2 Task Complexity

Many selective pressures entail an individual being able to fulfil multiple discrete traits simultaneously, or to be able to gradually develop increasingly fine-tuned responses to complex behavioural demands.

An example is the tool use of New Caledonian crows (Kenward et al., 2006; Hunt and Gray, 2003), who are able to acquire twigs and other such tool-making materials from the surrounding environment, fashion them into probes by removing excess, and use them to pry grubs and other prey from tree crevices.

In order to do so, these crows must possess a number of abilities:

- to locate source material for a tool;
- to fashion it into a suitable form;
- to locate a potential grub-dwelling crevice;
- to apply the tool appropriately and retrieve a grub

Each of these skills is of little use in isolation, and so it is unlikely that any of them would arise innately as there would be no positive selection pressure. In conjunction, however, they provide a significant competitive advantage.

Another example is the complex vocalisation of birdsong, whose forms exhibit a spectrum of complexity (Marler, 1970; Marler and Slabbekoorn, 2004). Some aspects of vocalisation appear to be encoded innately, but more sophisticated repertoires of note density and structure appear to be the result of practice and mimicry (Marler and Slabbekoorn,

2004). Similarly, the foraging behaviours of honeybees can take an entire lifetime to fully develop (Dukas and Visscher, 1994; Dukas, 2008), with evidence that efficiency continues to increase throughout the individual's lifespan.

In this thesis, we will describe the degree of behavioural requirements as an environment's "task complexity". From another perspective, a simpler task is one that can be learned in a shorter period of time; a more complex task could take an entire generation to learn. This may pose novel pressures on the learning modes that an individual should adopt to optimally respond to the task.

2.3.3 Population Structure

A substantial amount of empirical research has been carried out on the population structure of animal ecosystems. In a meta-analysis, Kudo and Dunbar (2001) review the findings of 43 studies on the social network size of various different primate species. A key factor when seeking to understand the nature of an animal's relationships with its community is its clique size, defined as the number of individuals that it regularly interacts with directly – its primary social partners. Over this comparative research, the mean clique size is 2.57, meaning that an agent's number of regular interactants is a small proportion of the collective population.

Beyond this immediate clique, links to other members of the species take place through a transitive chain of relationships. An individual animal is connected to the rest of its population via an indirect chain of transitive relations (Kudo and Dunbar, 2001, p717), forming a sprawling social network with heterogeneous connectivity. This can results in geographically bounded areas in which cultural traditions arise and persist, such as the regional dialects of passerine birds (Freeberg, 2000) and region-specific variance between the foraging tools constructed by New Caledonian crows (Hunt and Gray, 2003). These local traits have been found to not necessarily be the product of local adaptation, but can be result from arbitrary path-dependent lock-in: culturally ingrained traditions (Galef, 1990).

Even in an environment that is spatially homogeneous, the effects of a network of neighbours should not be underestimated. Evolutionary graph theory provides numerous examples of how population structure affects collective dynamics via local network effects such as frequency-dependent selection (Lieberman et al., 2005; Szabó and Fáth, 2007; Nowak, 2006). Here, heterogeneity is manifest in the differences between neighbours. Evolutionary games on a graph have quite different outcomes to those in a well-mixed environment. Most strikingly, in the prisoner's dilemma, population structure makes it possible for coooperation to evolve and persist as an evolutionarily stable strategy (Smith, 1982).

Lehmann et al. (2007) show that structuring an evolving population on a graph supports the predictions of inclusive fitness theory, creating kin selection benefits. A further recent result is that of Taylor et al. (2007), who use inclusive fitness analysis to show that cooperative alleles are capable of invading a structured population regardless of the particular form of structure used, demonstrating that altruism is likely to arise on a wider class of structured populations than previously believed. This has important implications for the understanding of population structure in modelling, as it implies that structure has significant impact particular as regards the evolution of social interaction.

2.3.4 Spatial Heterogeneity

Structuring a population through spatial locality can have significant effects on its dynamics. As Lowen and Dunbar (1996) observe:

> "In other areas of science where spatial models are in common use, it is now well known that mean-field models exhibit significantly different properties than their counterparts which explicitly account for local interactions. Indeed, it has been demonstrated that remarkably complex behaviour (such as that of the Ising model of magnetism or various population models) can emerge from extremely simple rules of local interaction." (Lowen and Dunbar, 1996)

Early work in optimal foraging theory showed that spatial heterogeneity could maintain genetic polymorphisms within a population (Brown, 2000; MacArthur and Pianka, 1966); subsequently, this would give rise to Maynard Smith and Price's (1973) work on evolutionarily stable strategies (ESS), in which multiple distinct behavioural groups can support one another. Wiens (2000) describes several origins and types of environmental heterogeneity, including spatial variance (pointwise autocorrelation), patterned variance (wider relationships and regional patterns), compositional variance (in which patches differ qualitatively as well as quantitatively), and locational variance (accounting for subjective spatial location to determine property neighbourhoods).

Where culture is concerned, spatial structure is particularly important. Hunt and Gray (2003) point to the spread of various forms of tools over the geography of the New Caledonian crow population, observing that this is likely mediated through the cultural transmission of traits from neighbour to neighbour. Likewise, theoretical investigations by Boyd and Richerson (1985; 1988) have shown that the dynamics of cultural transmission are markedly different when locality is introduced.

An explicit spatial context has been shown to favour the evolution of local structures which reinforce evolutionary benefits (Silver and Di Paolo, 2006), enabling the survival of traits that would otherwise not be supported by the environment. Odling-Smee et al. (2003) similarly emphasise the importance of spatial structure for niche construction, in the case that groups of individuals produce locally-clustered resource availability, creating a local selective advantage for other species.

There are multiple factors present that determine the composition of an environment: its elements, topology and patterning, and the availability of resources therein (Wiens, 2000). The majority of theoretical models, for simplicity, treat a population's environment as being homogeneous and unchanging. In many cases, this simplification is acceptable (and indeed necessary for tractability).

However, spatial patterning matters. Indeed, Levin (1992) argues that "the problem of pattern and scale is the central problem in ecology".

The relevance of spatial heterogeneity in ecology was recognised by Gause (1934), who conducted a series of lab experiments with microbial predator/prey duo *Didinium nasutum* and *Paramecium caudatum*, accompanied by differential equations to support his empirical findings. Growing populations of these microbes in a test tube, Gause found that they would not continue to coexist indefinitely within a homogeneous substrate; either predator or prey would eventually become extinct. When a spatial "refuge" was added

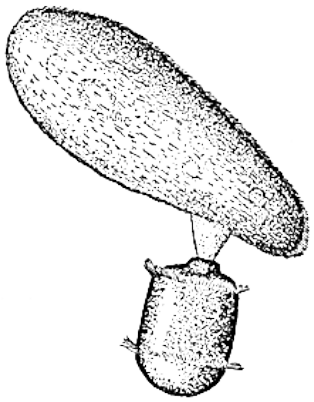

FIGURE 2.4: From Gause (1934): *Didinium nasutum* devouring *Paramecium caudatum*

– a compartment of different nutritional makeup – this heterogeneity would enable both populations to persistence in survival concurrently, but with a behaviour that is no longer broadly predictable from the oscillatory predictions of the Lotka-Volterra predator-prey formula. As Gause observes, "the struggle for existence is [in a heterogeneous environment] affected by a multiplicity of causes" (Gause, 1934, Chapter VI), irreducible to differential equations. He identifies the phenomenon of path-dependence as being critical to the system's subsequent trajectory; minor differences in initial conditions can lead to radically different outcomes due to small fluctuations being amplified and "locked in".

These themes were taken up by Huffaker (1958), whose lab experiments involved constructing a series of heterogeneous environments by distributing obstacles and resources over a short distance, over which mites and their predators were left to breed. In a homogeneous environment, populations quickly went to extinction as the prey was consumed and the predator had no further food source. In a heterogeneous environment, the existence of spatially distributed compartments with relatively slow dispersal between them allowed for the continued persistence of small pockets of predators or prey, resulting in multiple repeated oscillations of population growth and contraction.

The converse situation – heterogeneity of species within a homogeneous, wild environment – was the focus of Hutchinson (1961), whose observations of phytoplankton contradicted the knowledge of the field. Despite the apparent homogeneity of the open seas, there is a vast spectrum of plankton species across the oceans, in environments that are compositionally indistinguishable. Yet the principle of competitive exclusion formulated by Gause (1934) argues the converse: that a homogeneous environment would result in a complete reduction of evolutionary diversity. The initial coexistence of two species competing for the same resources, goes the argument, will eventually result in one species gaining the upper hand, and driving the other to extinction.

Why is this not happening in the case of plankton? Hutchinson offers a number of speculations as to why this is taking place: gradients of light causing subtle variation in the vertical plane; turbulence within the water; and differential predation in different regions, itself a consequence of environmental heterogeneity further afield (that is, at the boundaries of water and land). These theories have recently gained experimental support,

with Károlyi et al. (2000) discussing the encouraged coexistence of species in turbulent environments and Descamps-Julien and Gonzalez (2005) experimentally demonstrating the possibility of coexistence within an environment which is subject to temporal fluctuations, which may be driven by (for example) changing weather systems.

A number of recent bacterial studies have shown the impact of environmental heterogeneity over evolutionary timescales. The rapid reproductive life cycle of bacteria allows for experiments that can span thousands of generations, showing long-term patterns that would otherwise not be possible. Korona et al. (1994) shows that structured habitats can give rise to greater phenotypic plasticity than those that are homogeneous; subcolonies of variable mean fitness persisted, suggesting that they were trapped on suboptimal peaks in the fitness landscape, supporting the existence of multi-peaked landscapes in real-world scenarios. Rainey and Travisano (1998) conduct an elegant study showing that environmental heterogeneity can trigger and sustain evolutionary divergence, breeding bacteria in two types of culture: one with multiple heterogeneous compartments, and one whose culture was shaken, destroying any spatial structure. The spatially heterogeneous environment was shown to produce and support genetic polymorphism, extending the previous work of Gause et al onto evolutionary timescales.

Baythavong (2011) shows that, in plants, selection favours adaptive plasticity in fine-grained, heterogeneous environments. In homogeneous environments, however, the resulting plasticity is low. This highlights two different solutions to environmental heterogeneity: adaptive phenotypic plasticity Bradshaw (1965) and local adaptation. In fine-grained habitats, such that an individual's expected environment may differ from that of its parents, plasticity should be expected; in coarser-grained habitats, in which an individual's environment may remain the same, we should expect an adaptive specialism to arise.

A number of theoretical models have been used to explore the effects of "niche construction", in which an individual and its environment co-evolve simultaneously (Odling-Smee, 1988; Odling-Smee et al., 2003; Laland et al., 2000), dynamically modifying the selection pressures. Laland et al (2000) use a two-locus population genetic model to demonstrate that the ability of a population to alter its environment can dramatically alter selection processes, enabling otherwise deleterious alleles to reach fixation, producing new polymorphisms, and altering the timescales of selection.

Taylor (2004) uses niche construction as a basis to explore the drive towards complexity in natural systems. This model incorporates a more sophisticated conception of an organism, and the functional traits that it must exhibit within its environment: both positive and negative niche construction can take place, with true selective 'niches' opening up in the interactions between organisms. The motivation behind this model is to determine whether the process of niche construction drives an increase in the complexity of the organisms it inhabits, which it answers in the affirmative. This result throws up the question of whether other forms of dynamic – social interactions, for instance – are also under pressure to increase in complexity within the context of niche construction.

Borenstein (2005) explores a specifically cultural form of niche construction. Here, a structured series of populations (a "metapopulation") with a pair of transmissible cultural artefacts. They note that this configuration exhibits markedly different dynamics from those seen in a well-mixed model, with variance between and within populations). Silver and Di Paulo (2006) likewise adopt a spatial niche scenario, locating their agents on a finite,

toroidal lattice, with the capability of engineering resource availabilities. A significant conclusion here is that spatial effects were found to significantly increase the success of niche-constructing alleles, with local clustering resulting in positive feedback loops of selection. They conclude that "the stability of multi-species webs in natural populations may increase as the complexity of species-environment interactions increases.".

Comprehensive studies of the effects of spatial heterogeneity on learning-evolution interactions are rare, due to the complexity and long timescales required to conduct such studies, particularly on species sophisticated enough to support social learning. This thesis hopes to extend the work described in this chapter with results that bring together particular classes of environmental heterogeneity, making empirical predictions about the dynamics that are likely to occur in real-world scenarios.

2.4 Landscape Metrics

Spatial analysis is the pursuit of insight into spatial ecological processes (Fortin et al., 2006). It seeks to evaluate spatial structures upon well-defined axes by categorising and quantifying their spatial properties. Ecological forces operating on biotic or abiotic matter give rise to spatially-correlated effects on its organisation and organisation; matter may be eroded, consumed, moved or repurposed (Wiens, 2000). Many forces at play are geophysical – wind, rain, tides, earthquakes, volcanic eruptions – which can give rise to massive, discontinuous disturbances that can have substantial impact on the environment and its biotic inhabitants. Further environment heterogeneity is caused by the impact of the animals that inhabit the land, with grazing, foraging and "ecosystem engineering" (Jones and Lawton, 1995) causing changes to the landscape which can likewise have adaptive ramifications for surrounding species (Odling-Smee et al., 2003).

In a natural ecosystem, place matters. Transplanting a colony of termites ten metres to the north-east may maintain its internal population structure, but could move it into an environment which is devoid of food sources or inhabited by predators. It may also situate it in a landscape which is more variegated – possessing a wider array of different substrata, and imposing different fitness demands – or into a region that is uniform across its spatial extent. These different spatial formations can have significant ramifications for the evolutionary and behavioural trajectories of a population (Rainey and Travisano, 1998; Baythavong, 2011; Keymer et al., 2006).

The scale and extent of heterogeneity matters too; as it were, the heterogeneity of heterogeneity. What impact does this have on evolution? In a study of plant lineages, Baythavong (2011) characterises evolutionary trajectories over landscapes which have different levels of heterogeneity in different areas, observing that:

> "In fine grained environments, progeny are likely to experience an environment different from that of their maternal parent, and selection should favour the expression of adaptive phenotypic plasticity." (Baythavong, 2011)

In the landscape ecology literature, there is a spectrum of different approaches to classifying and describing empirical landscape heterogeneity. McGarigal (2006) distinguishes between *compositional* metrics, which quantify features related to the variety of different

resource types independent of spatial structure; and *spatial configuration* metrics, which describe the spatial characteristics and arrangement of resources. We will limit our scope to spatial configuration properties, as we will be restricting our investigation to one type of resource.

The two dominant spatial configuration qualities appearing throughout the literature (Riitters et al., 1995; Li and Reynolds, 1995; McGarigal, 2006) are *fragmentation* and *patchiness*.

Fragmentation (or its inverse, contagion) is the most common landscape metric used in quantifying the heterogeneity of discrete- space systems (Gustafson, 1998). It describes the tendency of patches to take the form of large, singular clumps (Frohn, 1998; Li and Reynolds, 1993) versus a dispersed distribution of smaller areas. An environment with a small fragmentation value may be made up of one or two monolithic zones of similar land usage. If fragmentation is high, there is likely to be a large number of small distinct areas. This classification is the inverse of *contagion*: a landscape with high contagion is grouped into a few large areas. As with many aspects of landscape classification (Frohn, 1998) this distinction is a matter of scale. At high magnification, or for a small organism, a particular factor may appear to have a low fragmentation value (or high contagion), whereas at low magnification it will appear to have high fragmentation (low contagion).

Patchiness[1] measures the degree of dissimilarity between neighbouring patches, determined by Romme's relative patchiness index, or RPI (Li and Reynolds, 1994; Romme, 1982). The greater the number of disparate patch types neighbouring each other, and the greater the degree of disparity, the higher the resultant RPI value. In a context characterised by high patchiness, there is a greater chance of moving from a familiar environment to one which is markedly less familiar. With a low RPI, conversely, we should expect to move relatively smoothly across terrain without major discontinuities.

Summary

This chapter has reviewed the background literature on evolution, learning and social behaviour, subsequently examining the impact that spatio-temporal heterogeneity has on the evolutionary dynamics of a population, and reviewing key environmental metrics drawn from landscape ecology and spatial analysis.

The following chapters will proceed to unify these ideas, describing a model that represents the evolutionary and learning capacities described in §2.2. This will be extended with progressively more detailed form of environmental structure to understand and quantify the impact of each of the four key environmental factors described in §2.3, latterly by constructing artificial environments following the landscape metrics described in §2.4.

[1] It should be noted that many writers use "patchiness" as a term to denote general spatial heterogeneity (Marquet et al., 1993; Grünbaum, 2012). We intend it here in its more technically specific sense; see Romme (1982), Li and Reynolds (1995).

Chapter 3: Baseline Model

3.1 Overview

In this chapter, we shall describe a computational individual-based model that allows us to explore the dynamics of a population that is able to evolve, learn, and socially interact. This baseline model acts as the foundation for a series of experiments in the following chapters (§4, §5 and §6), which progressively introduce novel elements to the model to pose specific questions and understand the impact of particular environmental factors on the optimal modes of information acquisition. This will include introducing population structure and locality, between-agent environmental heterogeneity, motion through space, and multiple environmental fitness objectives.

The present chapter begins by establishing a baseline model specification, Model ∅, which incorporates the minimal set of properties needed within an individual-based simulation in which evolution and learning interact. The objective is to introduce as few variables as possible, resulting in a model with a maximal level of abstraction, thus applicable to the widest possible range of real-world scenarios. Model ∅ does not incorporate any ideas of topographical structure. Each individual can interact with any other within the population.

In the baseline case, the model is limited in its scope. It is comprised of a population of agents, each of which possesses multiple properties: a genotype, which determines the individual's innate fitness; a phenotype, determining its subsequent fitness, mutable throughout in its lifetime; and behavioural characteristics which determine individual tendencies towards innate behaviour, individual learning and social learning. By differential selection and asexual reproduction, the population evolves to a steady state. The object of study of this thesis is the dynamics of interaction between these variables and the regimes that they pass through, as a consequence of the particular qualities of the environment that they reside within.

We begin in Section 3.2 by formally defining Model ∅ and its constants, variables, and state change equations. We provide a schematic overview of the model in pseudocode (Section 3.3), and proceed to comment on the implications, assumptions and limitations of this baseline model (Section 3.4).

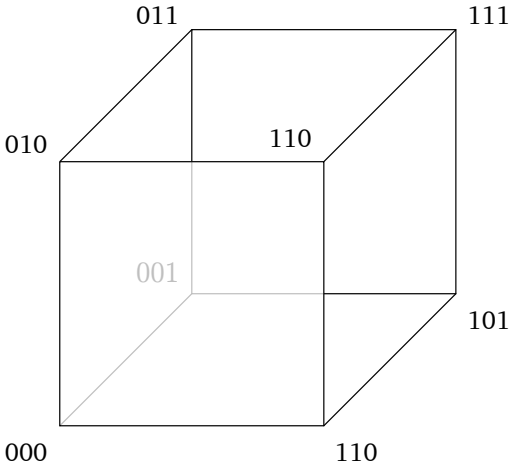

FIGURE 3.1: Representation of the space of B-bit fitness functions as a B-dimensional hypercube, for $B = 3$

3.2 Model Specification

We will now describe the components of the agent-based model used to explore these ideas[1]. An *environment* E consists of a single B-bit string, representing the 'objective' task that an individual must achieve to obtain maximal fitness: $E \in \{0,1\}^B$. The current environmental state can therefore be considered as a vertex on an B-dimensional hypercube (Figure 3.1).

The environment is inhabited by a population of N agents, each of which has the following properties:

- $b_{evo}, b_{ind}, b_{soc} \in [0,1]$ – *behavioural traits* determining the propensity towards genetically innate behaviour, individual learning, and social learning. These are collectively normalised to sum to unity, and remain fixed throughout an agent's lifetime.
- $g \in \{0,1\}^B$ – *genotype*, a B-bit string that establishes the agent's capability to fulfil the environment's objective. The genotype remains fixed throughout an agent's lifetime.
- $p \in \{0,1\}^B$ – *phenotype*, a B-bit string that is initialised equal to g when an agent is born, and subsequently subject to modification through individual and social learning. If p is equal to E then the agent's fitness is equal to 1.
- ϕ – current fitness, determined based on the agent's activity in the previous timestep.

An agent's current phenotype determines how well it fulfils the environment's objective task, based on its Hamming distance from E.

[1] For all subsequent parameter values, see *Table 3.1.*

3.2.1 Actions and Learning

Every timestep, each agent selects a behavioural mode according to a weighted random choice from $\{b_{evo}, b_{ind}, b_{soc}\}$:

- b_{evo} – **innate behaviour**: act according to the agent's current phenotype
- b_{ind} – **individual learning**: act according to the agent's current phenotype, with a single bit toggled at random
- b_{soc} – **social learning**: act according to the agent's current phenotype, with a single bit copied from a neighbour using fitness-proportionate (roulette wheel) selection, weighted by their fitness ϕ. With a probability p_{noise}, the bit may be copied erroneously (toggled from $0 \rightarrow 1$ or $1 \rightarrow 0$). This models the imperfection and noise present in real-world imitative learning: a behaviour may be only partially observed, or reproduced incorrectly.

If b_{ind} or b_{soc} is employed *and* the resultant action gives a higher payoff than the agent's own current phenotype, the corresponding bit in p is replaced by the new action: discovering (or imitating) a successful new trait results in its being incorporated into the agent's behavioural roster. This reflects the effects of phenotypic plasticity, or the incorporation of new traits during an organism's lifetime.

In the case of b_{soc}, weighting the exemplar by their ϕ value reflects a tendency towards mimicking those organisms which are perceived as being fittest (a "*copy-successful-individuals*" strategy (Laland, 2004), as observed in avian, chimpanzee and bat societies). At present, the agent's interaction neighbourhood comprises of the entire population. This will be refined in subsequent iterations of the model.

The agent's fitness is determined according to the following rule:

$$\phi = \left(1 - \frac{H(p, E)}{B}\right)^{\alpha^{-1}} \tag{3.1}$$

where H denotes the Hamming distance between two bit strings. Dividing by B and subtracting from 1 normalises the fitness value to $[0, 1]$, and transforms it from a distance to a proximity measure. The exponential of α is used to determine the fitness differential between perfect and almost-perfect task fulfilment (Figure 3.3): a lower value of α means that payoffs fall more rapidly with distance. With $\alpha = 1$, scaling is linear in distance. As α tends to zero, fitness becomes negligible unless every one of the traits is realised, similar to Hinton and Nowlan (1987) and Papineau (2005).

In general, if an agent's g precisely matches the environmental objective E, its fitness becomes the maximal value of 1. If g is precisely the complement of E, its fitness is 0.

3.2.2 Reproduction

Each timestep, reproduction occurs following a birth-death process. A single agent is selected using fitness-proportionate selection based upon fitness values ϕ, and reproduces

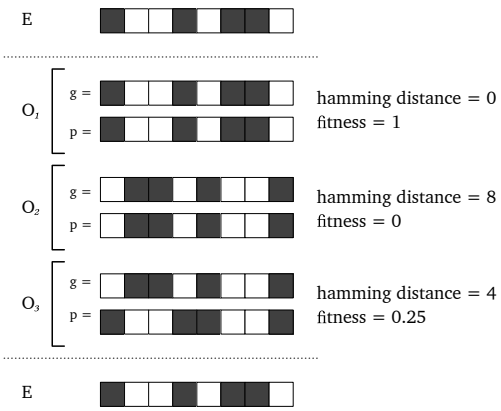

FIGURE 3.2: Deriving fitness from an individual's phenotype. Fitness is derived by first calculating a normalised proximity value to the bit string of the current environment E. This value is raised to the power of α^{-1}, representing an exponential benefit in possessing multiple traits. In the above example, $\alpha = 0.5$.

Variable	*Value*	*Comments*
N	64	Population size
B	16	Number of bits in E
α	0.01	Rate of fitness dropoff based on task proximity (Figure 3.3)
μ	0.05	s.d. of mutation as applied to $b_{evo}, b_{ind}, b_{soc}$
p_{switch}	0.01	Probability of a single environmental fluctuation
p_{noise}	0.25	Probability of incorrect observation during mimicking
p_{mut}	0.01	Probability of sustaining a mutation per gene
trials	50	Number of repeats of each trial, before averaging results

Table 3.1: Standard parameter values used in the baseline model. Parameter values per trial can be found in Appendix B.

asexually[2]. Its offspring has an identical genotype, subject to each bit of g mutating (that is, flipping from $1 \to 0$ or $0 \to 1$) with small probability p_{mut}. Behavioural traits $b_{evo}, b_{ind}, b_{soc}$ are modified by a zero-mean Gaussian noise function, standard deviation μ, and clipped to $[0, 1]$. These are again collectively normalised to unity. The child replaces another member of the population, selected uniformly randomly.

Discussion of Model Parameters

Some of the model parameters justify some further commentary and contextual explanation.

[2]Sexual recombination was considered as a reproductive strategy, as an effective method of finding 'middle ground' locations between points on a complex fitness landscape. Given the single-peaked landscape adopted in this model, we focus on clonal reproduction for the sake of simplicity. A number of recombinative trials indicated that the results would not be qualitatively different.

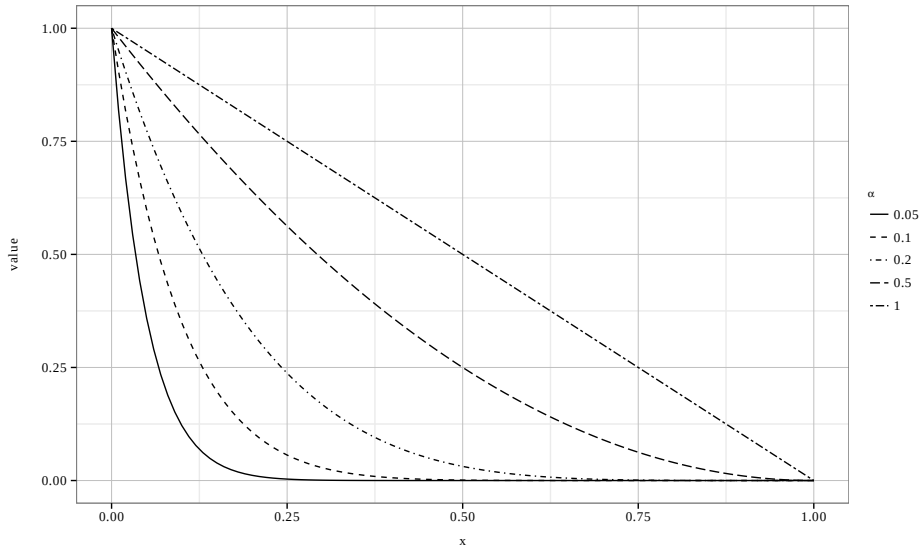

FIGURE 3.3: Proximity function $x^{\alpha^{-1}}$ for differing values of α. Smaller α gives a sharper fitness differential at larger Hamming distances.

N: Population size

As a Moran process, the population size represented in the model is fixed. See §7.4 for discussions of how this could be relaxed for future work; Borg and Channon (2012) extends our model with variable population size to make further conclusions about population collapse in variable environments.

B: Task complexity

B denotes the number of bits within the environment, and correspondingly the number of distinct alleles within an individual's genotype. This can be seen as modelling "task complexity", the number of distinct subtraits that must be acquired to reach maximal fitness. A higher level of B means that a greater learning time is required to achieve maximal fitness.

α: Fitness dropoff

$\alpha \in \{0, 1\}$ determines the steepness of the fitness function experienced by an individual; as α tends to zero, fitness becomes negligible unless every one of the traits is realised. This is similar to the approach taken in Hinton and Nowlan (1987).

p_{noise}: Social learning error

Social learning is an imperfect process and may result in maladaptive mimicry, either by individuals copying a non-beneficial trait or by failing to mimic a trait properly (Galef and Laland, 2005). This variable models the latter, imposing an implicit cost on social learning by adding a degree of noise to the copying process such that an individual may mimic a trait incorrectly, inverting the corresponding bit within its phenotype.

p_{switch}: Rate of environmental change

Models the stochastic environmental change needed to represent changing environments. This variable introduces gradual variability to the environment, determining the probability that a trait required within the environmental fitness function may invert, changing the selection pressure by a single bit.

3.3 Pseudocode

Below is a pseudocode formulation of the model described in §3.2.

```
environment ← 1^B
agents ← create_agents(N)

foreach agent in agents:
  agent.randomise_behaviour_weights()
  agent.randomise_genotype()
  agent.phenotype ← agent.genotype
  agent.current_fitness ← calculate_fitness(agent.phenotype, environment)

foreach timestep in 1 ... steps:
  foreach agent in agents:
    behaviour ← agent.roulette_wheel(b_evo, b_ind, b_soc)
    if behaviour = b_evo:
      agent.new_fitness ← calculate_fitness(agent.phenotype, environment)
    else if behaviour = b_ind:
      agent.toggle_random_phenotype_bit()
      agent.new_fitness ← calculate_fitness(agent.phenotype, environment)
      if agent.new_fitness < agent.current_fitness:
        agent.revert_to_previous_phenotype()
    else if behaviour = b_soc:
      neighbour ← agent.select_neighbour_by_fitness()
      agent.copy_random_phenotype_bit(neighbour)
      if coin_toss(p_noise):
        agent.toggle_last_copied_bit()
      agent.new_fitness ← calculate_fitness(agent.phenotype, environment)
      if agent.new_fitness < agent.current_fitness:
        agent.revert_to_previous_phenotype()
    agent.current_fitness ← agent.new_fitness

  parent ← agents.select_weighted_by_fitness()
  child ← parent.reproduce()
  foreach behaviour in child.behaviour_weights:
    behaviour <- behaviour + gaussian(μ)
  foreach gene in child.genotype:
    if coin_toss(p_mut):
      gene.toggle()
  child.phenotype ← child.genotype
  agents.remove_random()
  agents.add(child)

  if coin_toss(p_switch):
    environment.toggle_random_bit()
```

3.4 Commentary

There are a number of qualities to observe about Model ∅ and its relationship with existing theoretical and empirical work.

The present model takes the form of a Moran process (Moran, 1958; Lieberman et al., 2005; Nowak, 2006): its population size is constant, with one asexual (haploid) reproduction action each generation. Birth-death replacement is used to select a parent agent according to proportionate fitness, creating a new agent which displaces a random peer. Though a clear simplification of the variable population structures and reproductive schemes of a biological scenario, this mode of operation is common within the literature to simplify analysis whilst holding other conditions steady (Taylor and Fudenberg, 2004).

Genetic inheritance occurs vertically, from parent to offspring. Phenotypic modifications made throughout an agent's lifetime are not passed onwards at reproduction. This assumes that any epigenetic inheritance through the development period – such as in the case of DNA methylation (Jablonka and Lamb, 2005) – is negligible. Horizontal and oblique (that is, cross-generational (Best, 1999)) transfer of functional information, however, is possible through the mechanism of cultural inheritance. This is a crucial part of the model design, enabling the possibility of a cumulative culture which is able to act as a collective memory for functional traits (Avital and Jablonka, 2000; Laland and Hoppitt, 2003; Galef, 1990, 1976).

This model imposes no explicit cost for the maintenance of individual or social learning apparatus. The cost is incurred purely through the learning noise implicit in each of these mechanisms. With trial-and-error individual learning b_{ind}, a fit individual risks toggling a "correct" phenotypic bit and descending the fitness landscape. In the case of social learning b_{soc}, the chance of selecting a correct bit may be higher, given a sufficiently fit peer group; however, the error probability p_{noise} introduces the possibility of mimicking a trait incorrectly and decreasing one's own fitness. This models the observational imperfection present within animal social learning (Heyes and Galef Jr, 1996). In the wild, phenotypic plasticity does impose energetic and evolutionary costs in terms of the cellular machinery needed to maintain it (Scheiner, 1993; Moran, 1992). We avoid costs here partly for intellectual parsimony, and partly to demonstrate that the interesting range of reactions described in the below results *can* occur simply with learning error as their sole cost.

There is no scarcity of resources in this model. Resources are, at first pass, infinite and homogeneous; the payoff for attaining the peak of the fitness function remains identical for every agent. The between-agent competition is therefore purely a matter of maximising one's own fit to the environmental objective.

Each individual can engage in individual and social learning alongside innate behaviour, statistically modulated by their behavioural tendencies. This means that behavioural tendencies are not mutually exclusive but are always exhibited in different strengths. Game-theoretical approaches frequently pit discrete strategies against one another (Axelrod and Hamilton, 1981) and observe the evolution of evolutionarily stable strategies (Smith and Price, 1973). Here, the strategy emerges within individuals of the population, so we may expect to see individuals evolve to equilibria involving primary and secondary modes of learning. That is to say, if every individual can engage in a low-level of individual

learning, we may not expect to see such adversarial oppositions as producer-scrounger dynamics cropping up (Barnard and Sibly, 1981; Vickery et al., 1991).

This motivation also justifies the disparity between behavioural tendencies, which are modelled as continuous traits, and genotype/phenotype, which are modelled as discrete units. In the interests of parsimony, we would prefer to model all ecological qualities as discrete units where possible. However, the present research is particularly examining *tendencies* towards innate behaviour and modes of learning. If each individual only possesses one fixed behavioural mode throughout its lifetime – to act according to its genes, or to conduct trial-and-error learning – it is likely that a population would rapidly sweep to a single-behaviour equilibrium which it would never be able to escape.

These modelling paradigms were simply seen as more fitting for each case: genes require individual inheritance, whereas, in this model, we want to test a continuum of learning behaviours, from pure-inheritance to pure-learning and every combination between.

3.5 Summary

In this chapter, we described a novel individual-based model that models evolution, learning and mimicry, in a polygenic population with evolvable tendencies towards each behavioural mode (§3.2). We formulated the execution of the model in pseudocode (§3.3), and commented on its assumptions, relationships, and limitations (§3.4).

Chapter 4: Unstructured Populations

4.1 Overview

In Chapter 3, we formally defined a model of evolution and learning.

In this chapter, we begin by examining the behaviours of this model in static, unstructured environments, to establish a baseline set of dynamics that later experimental results can be compared against (§4.3.1). We then vary the first two environmental factors defined in §1.3 to understand their likely effects on optimal learning strategies: **variability**, in which the fitness pressures imposed by the environment gradually change over time (§4.3.2); and **task complexity**, the number of subtraits required to reach optimal fitness (§4.3.3).

We then look at how these two environmental factors interact, conducting a two-dimensional parameter sweep that enables us to show the dominant learning regimes across different levels of environmental complexity and rates of environmental change (§4.3.4).

4.2 Key Questions

Prior to exploring more complex environments in later chapters, our focus will begin on unstructured populations in spatially homogeneous environments. The primary objective is to gain an understanding of the dynamics of learning, evolution and sociality within a static environment, and subsequently within environments which are subject to different rates of change. This includes environments in which a sudden extreme perturbation takes place, emulating traumatic natural events such as an earthquake or volcanic eruption.

Alongside environmental fluctuations, we will examine a population's response to different levels of task complexity, measured in the number of phenotypic bits required to perform a task at optimal fitness. A simple environment may be comprised of a few bits of complexity, with more complex environments requiring many phenotypic bits to attain peak fitness.

We will finally explore how assumptions and strategies implicit within this model affect its macro dynamics, including normative learning modes and trait-selective learning regimes.

This chapter covers the first two environmental factors defined in §1.3: **variability** and **task complexity**.

Variability

A common view is that genetic evolution, social learning and individual learning are adaptively advantageous because they operate at different timescales within the evolutionary and ontogenic process (Cavalli-Sforza, 1981; Boyd and Richerson, 1985; Henrich and McElreath, 2003; Borenstein et al., 2008; Wakano et al., 2004).

We seek firstly to corroborate or challenge these results, and secondly to interrogate how they are affected by different types of environment.

Do these regimes hold in general, and how robust are they under different levels of environmental fluctuation and task complexity? Do transient behaviours occur which recede to a new equilibrium state? How do the interactions between individual and social learning affect these timescales and environment-specific trends? Do we expect to see multiple different learning strategies persist concurrently, or for strategies to reach extinction in certain settings?

This line of inquiry responds to empirical studies (Galef, 1976) and prior theoretical models that study the regimes of environmental change fostering different types of plasticity, whether behavioural (Levins, 1968; Templeton and Rothman, 1978; Lande, 2009) or cultural (Kendal et al., 2009; Feldman et al., 1996; Boyd and Richerson, 1988; Wakano et al., 2004).

Task Complexity

The second environmental explored in this chapter is the degree of task complexity, as exemplified by the examples of tool usage (Kenward et al., 2006; Hunt and Gray, 2003) and foraging tasks (Dukas and Visscher, 1994; Dukas, 2008) discussed in §2.3.2.

By imposing a range of complexity values, we hope to understand how the degree of difficulty posed by an environmental task – and thus the learning time that it requires to perfect – have on learning and evolution.

4.3 Results

As a starting point for investigating the interactions between the elements of the model, this analysis proceeds by introducing concepts gradually, with the intention of understanding fundamental interactions. Each experiment is introduced with its structure and motivation, and followed by a brief interpretation of results. Experiments are headed by their key findings.

In each of the following experiments, we begin by initialising the environment's objective to 1^B (following Hinton and Nowlan (1987)). Behavioural traits and genotypes are initialised to uniformly random values for each agent. Each experiment is left to unfold until its collective traits have reached stability. This is akin to introducing a mixed, unadapted population to a novel environment and allowing it to collectively evolve until it reaches a functionally adapted state.

Error bars indicate 95% confidence intervals ($p < 0.05$), with the number of repeat trials

per experiment given in Appendix B[1] . Default parameter values are given in Table 3.1; the precise values used in each experiment are also given in Appendix B.

4.3.1 Baseline Dynamics

We begin by exploring the dynamics of the model in a static environment, in which the environmental objective remains unaltered through the entire life cycle of the experiment. An organism that possesses the 'perfect' phenotype will be able to maintain maximal fitness, at least as long as it continues to act according to this phenotype.

The scenario of a completely static environment is one that is implausible in the wild, given that any organism modifies its environment in each interaction with it; real-world environments are subject to continuous change. It should be considered here as demonstrating potential behaviours within laboratory conditions, or within an environment whose rate of change is imperceptible relative to the short lifespan of its inhabitants.

4.3.1.1 Social learning initially dominates, subsequently giving way to innate behaviour

The changing distribution of behavioural traits over time in a static environment is shown in Figure 4.1. At initialisation, each trait – b_{evo}, b_{ind} and b_{soc} – has a mean value of $\frac{1}{3}$ across the population, as they are assigned uniformly random value per organism. As the population evolves, the mean behavioural traits evolve based on the relative fitness values that they confer.

The dynamics can be divided into three phases. Between steps **1–50000**, the population is dominated by social learners, rapidly rising to a peak at $t = 10000$, and then gradually yielding to innate behaviour. From steps **50000–150000**, b_{evo} continues to rise in dominance. Beyond step **150000**, a stable optimum has been reached, with innate behaviour reaching a peak mean value of 0.75. Analysis of the population within these experiments reveals that a single unified group emerges with behavioural values distributed around these means.

Parameter values (**4.3.1.1**)	N	B	α	μ	p_{switch}	p_{noise}	$trials$
	64	32	0.01	0.05	0.00	0.25	100

The trend of social learning giving way to innate behaviour is a instance of genetic assimilation (Waddington, 1953; Crispo, 2007; Pigliucci et al., 2006) taking place. In an unfamiliar environment, a population initially resorts to phenotypic plasticity (in this case, social learning) to accumulate information about the world that has not been encoded into their genotype. An individual copies the actions of its more successful peers to gain an immediate fitness boost. As generations progress, some surviving organisms develop the underlying set of genetic mutations needed to succeed in the environment innately. As innate behaviour does not include the noise-induced costs of learning, the predisposition towards genetic instinct is less costly, and gradually takes hold over the population.

Put differently, an evolutionary progression takes place towards individuals who are predisposed towards phenotypically accommodating the adaptive trait. Like the heat-shocked *drosophila* of Waddington's (1953) experiments, individuals are selected whose

[1]Typically 100 for single experiments.

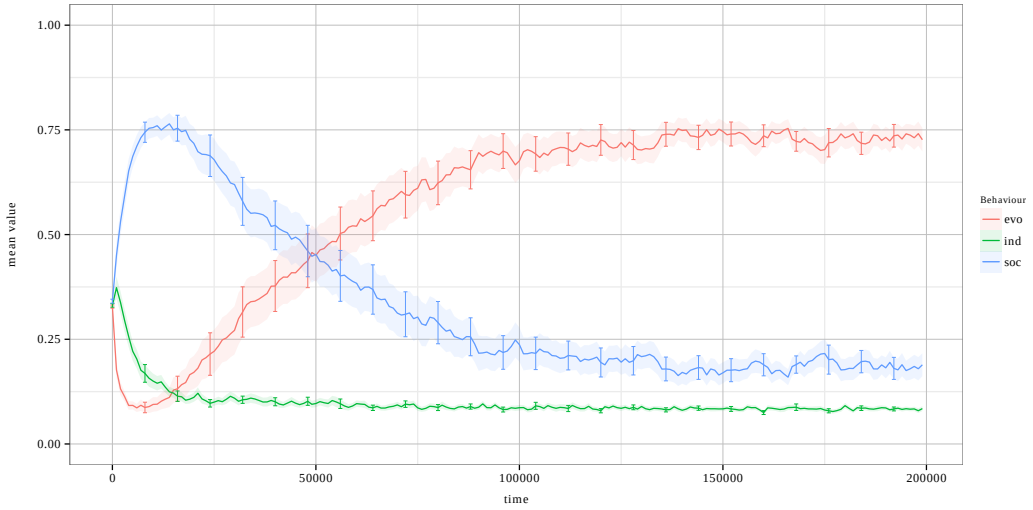

FIGURE 4.1: Static environment: Behaviour distributions over time

phenotypic reaction norms require the smallest employment of plasticity to exhibit the trait in question (Pigliucci et al., 2006).

Two properties of these dynamics are unexpected on first glance: the extended period of time taken to reach the population optimum, and the continued presence of low-level b_{soc} and b_{ind} traits. Both of these artefacts are a consequence of the relatively small cost of social and individual learning in this model. Neither bear an explicit constitutive cost, but have the caveat that, if an agent is already close to the optimum, testing a random new task or copying a peer may result in an act that is detrimental to the agent's native phenotype.

Sharpening these costs by reducing payoff scaling factor α results in a more rapid convergence to a predominantly b_{evo} population (see A.1.2 in Appendix A). However, this simultaneously reduces the selection pressure on lifetime learning, which introduces a penalty in fluctuating environments.

Within this static environment, there is clear evidence of a genetic assimilation process: phenotypically plastic individuals first out-compete their peers as they scramble to higher fitness through learning and social exchange, and are subsequently replaced by innate mutants, who do not bear the costs of exploration. It should be noted that, by initialising genotypes to random values, this is strictly an artificial example: in a natural ecosystem, genetic information would *begin* in a state that is correlated with the environment.

It may also seem surprising that there is little trace of individual learning in this experiment. However, given each agent's uniformly random genotype, there is no need for innovation within this context as all of the successful traits (1-bits) are already present within the population, ready to be socially transmitted and adopted.

4.3.1.2 Phenotypic fitness rapidly converges, temporarily shielding genotypic fitness

Figure 4.2 illustrates the mean Hamming distance between E and the population's genotypes and phenotypes. We see evidence for genetic shielding taking place (Wright, 1931; Anderson, 1995; Mayley, 1997), masking genotypic selective forces; the population's ability to phenotypically adapt to a niche reduces the pressure on genetic evolution. The

prevalence of learning behaviours means that the population's innate (genetic) fitness lags significantly behind the fitness that they express throughout their lifetimes, by utilising phenotypic plasticity to override their less-fit innate states. Eventually, as genetic assimilation completes, genotype and phenotype almost coincide. A small, statistically significant difference does remain (ANOVA, $p < 0.05$), indicating that a low level of genetic shielding persists even within a static environment.

Parameter values (**4.3.1.2**)	N	B	α	μ	p_{switch}	p_{noise}	*trials*
	64	32	0.01	0.05	0.00	0.25	100

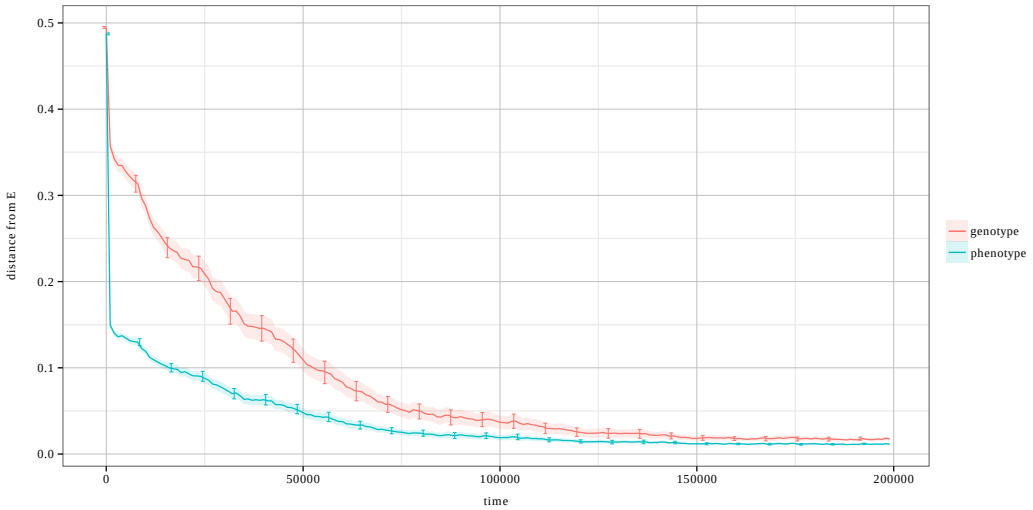

FIGURE 4.2: Static environment: Genotypic and phenotypic distance from E

It is interesting to note that there is an immediate leap in genotypic fitness at the very start of the experiment, dropping in a few generations from its starting value of 0.5 to around 0.35. This is because genetic shielding has not yet kicked in; at the start of the experiment, many individuals have a larger b_{evo} value, which imposes increased pressure on attaining a better genotypic value. As soon as social learning prevails, this pressure drops, which rapidly slows the rate at which genotypic assimilation occurs.

We can therefore infer that it is the differential between b_{evo} and b_{ind}/b_{soc} that gives rise to the selective pressure on the individual's genotype. An individual that exercises significant behavioural plasticity has a lowered genotypic selective pressure, which slows the rate of genetic convergence (Mayley, 1996).

4.3.1.3 Environmental perturbation results in a double Baldwin effect

We repeat the experiment as per the previous section, introducing an environmental perturbation after equilibrium has been attained. This mimics the effect of a traumatic change in the organism's surrounding landscape, in which a population is ill-adapted to their new environmental context (Lande, 2009). In the model, this is achieved by inverting E so that fit individuals immediately become maladapted.

As shown in Figure 4.3, this results in a temporary burst of individual learning followed by a longer wave of social learning, bringing up phenotypic fitness through plasticity

whilst evolution takes time to work out the necessary series of mutations. In a sense, this is a double Baldwin effect: the traits adopted by individual learning are assimilated into the cultural memory; and subsequently, socially-learned traits are assimilated into the long-term memory of the genotype.

Parameter values (**4.3.1.3**)	N	B	α	μ	p_{switch}	p_{noise}	$trials$
	64	32	0.01	0.05	0.00	0.25	100

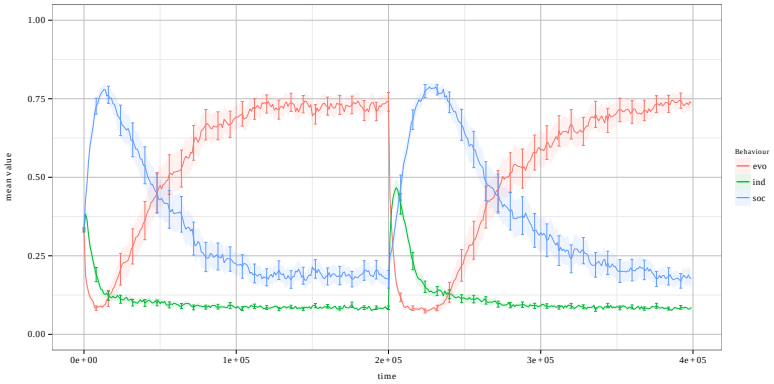

FIGURE 4.3: Introducing an environmental perturbation at $t = 2e + 05$. Subsequently, organisms are selected for increased learning tendencies.

This change takes place in three phases. Individual, exploratory learners first take precedence ($t = 200,000 \rightarrow 210,000$); at this early stage, mimicking one's neighbours is just as likely to result in learning incorrectly, as they are equally likely to be maladapted and possess outdated information about the environment.

After this peak, beneficial traits have become sufficiently prevalent to promote the use of social information. Social learning dominates for a long period ($t = 210,000 \rightarrow 270,000$), and finally genetic assimilation takes place, returning to evolutionary stability ($t = 400,000$).

Figure 4.4 demonstrates an accentuated version of the genetic shielding shown previously. Again, a rise in learning frequency means that phenotypic fitness improves more rapidly than in the genotype.

Continuing to examine Figure 4.4, two more observations may be made. The immediate, biphasic drop in genotypic distance observed in the first half of the experiment (and in Figure 4.2) is not visible in the second half. This supports the hypothesis stated in Section 4.3.1.2 that this drop is due to population's initial propensity towards innate behaviour, and the positive selection pressure that it exerts towards beneficial genetic traits.

Likewise, the rapid, biphasic phenotypic adaption at the beginning of the experiment is not seen following the perturbation, with the population taking a longer time (and following a smooth, exponential curve) to return to optimal fitness. Successful *phenotypic* traits are also missing from the population after the perturbation, meaning that the rapid behavioural adaptation offered by social learning is not possible. We can thus conclude that the genetic and behavioural homogeneity produced by a period of extended environmental stasis gives rise to maladaptivity in the population's ability to recover from a major environmental trauma.

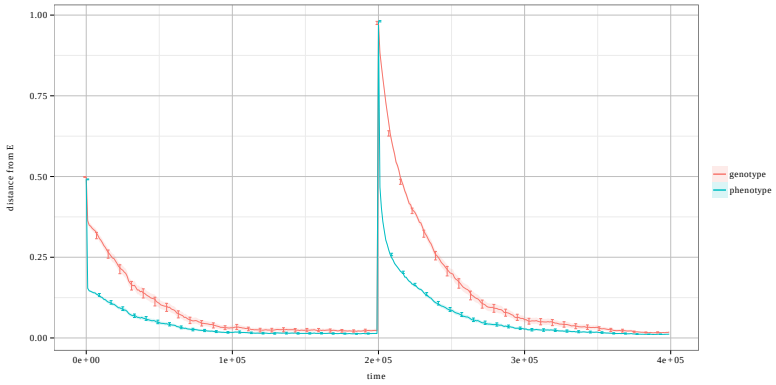

FIGURE 4.4: Genotypic and phenotypic distances from optimum following a perturbation, where 0 represents a perfect match with the environmental target. Phenotypic fitness returns to a stable level faster than genotypic fitness, reflecting the more rapid adaptive rate of lifetime learning.

These conclusions support the theoretical results of Lande (2009), who demonstrates that the slope of a reaction norm increases after a major environmental change, increasing plasticity for a period to accommodate the unpredictable environmental demands, before plasticity gradually falls away in the process of genetic assimilation. Our results interpose social learning interposes itself as a third adaptive recovery mechanism after a major environmental change, bootstrapping genetic evolution by allowing positive traits to spread horizontally and increasing the population's mean fitness.

4.3.2 Factor 1: Variability

We now extend the above by introducing regular environmental fluctuations. Each time step, a single bit of the environmental task may be toggled, according to a small probability p_{switch}. A value of $p_{switch} = 0.01$ reflects an expected period of 100 timesteps between 1-bit fluctuations. Ecologically, this models a gradual change in environmental conditions, altering the selective landscape within the lifetime of an individual.

4.3.2.1 Increasing instability encourages social learning at low rates of change, and individual learning at high rates of change

With a moderate fluctuation rate ($p_{switch} = 0.01$), the optimal combination of strategies is markedly different to that in a fixed environment (Figure 4.5). Social learning dominates, reflecting the benefit of lifetime plasticity in an uncertain environment. In this instance, genetic shielding appears to be highly evident (Figure 4.6), with the genotype's hamming distance from E vacillating between 0.44 and 0.49. This suggests that changes in the genotype are occurring almost solely due to genetic drift.

In this population, we can infer that the predominant method of information transfer is horizontal: the cultural transmission of traits from peer to peer via social learning. However, this is insufficient to acquire new traits alone; a low, baseline level of individual learning continues to exist, which acts as a source of functional innovation after a change in the environment does occur.

There is little pressure on innate behaviour to rise in dominance, as acting innately will result in functional behaviours which soon become outdated; even for an organism whose genotype currently matches E precisely, the environment will have altered in an expected time of 100 steps (as $p_{switch} = 0.01$, the expected time of next fluctuation is $\frac{1}{0.01} = 100$), at which point the genotype again lapses into suboptimality.

In a population size of $N = 64$, the expected lifespan of an organism is 64 steps, meaning that an environmental fluctuation should be expected within two generations.

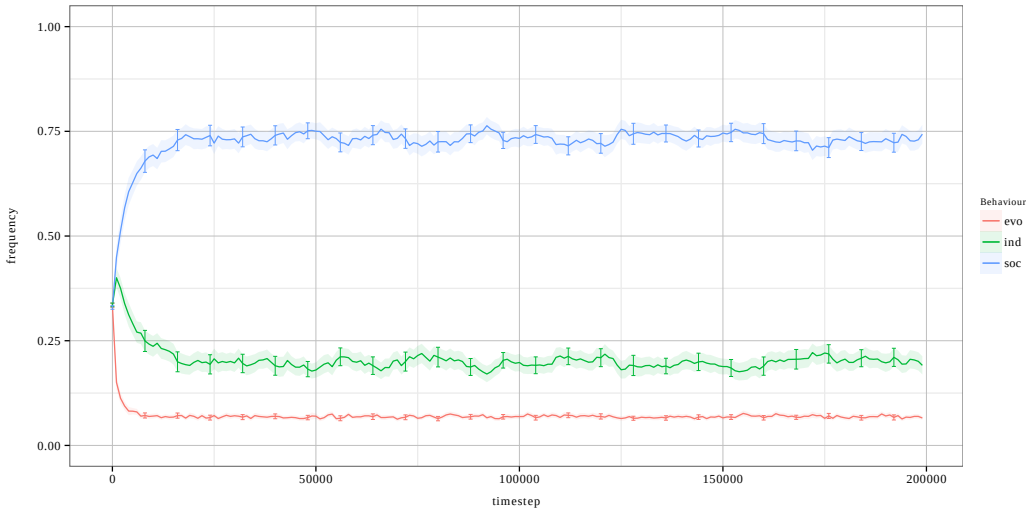

FIGURE 4.5: Fluctuating environment: Behaviour distributions over time

Parameter values (**4.3.2.1**)	N	B	α	μ	p_{switch}	p_{noise}	*trials*
	64	32	0.01	0.05	0.01	0.25	100

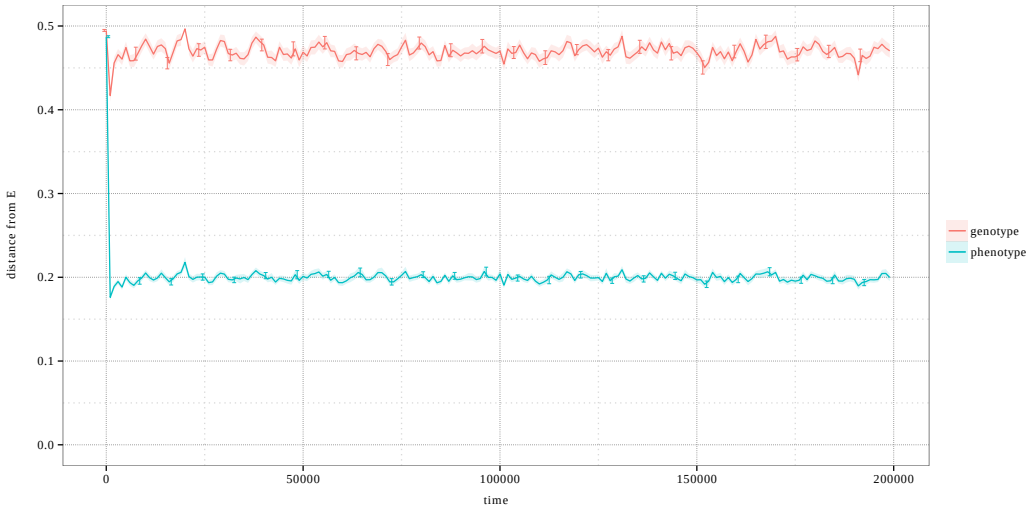

FIGURE 4.6: Fluctuating environment: Genotypic and phenotypic fitness over time

Figure 4.7 indicates the learning modes taking place at equilibrium over a range of fluctuation rates. It is interesting to note that three distinct regimes take place, without

substantial gradation between them:

1. at **slow** rates of environmental change ($p_{switch} \leq 1e-04$: less than 1 expected fluctuation per 150 generations), **innate behaviour** is dominant.

2. at **moderate** rates of environmental change ($1e-04 < p_{switch} < 0.05$: less than 1 expected fluctuation per generation), **social learning** is dominant; genetic evolution is not sufficiently fast to keep up with the changing fitness function, but there is a reasonable chance that surrounding peers will possess the correct traits (having also been learned via social transmission.)

3. at **fast** rates of environmental change ($p_{switch} \geq 0.05$: at least 1 expected fluctuation per generation), **individual learning** is dominant; at rates of change this fast, it becomes likely that shared social information will be outdated, and so it is optimum to resort to individual trial-and-error learning.

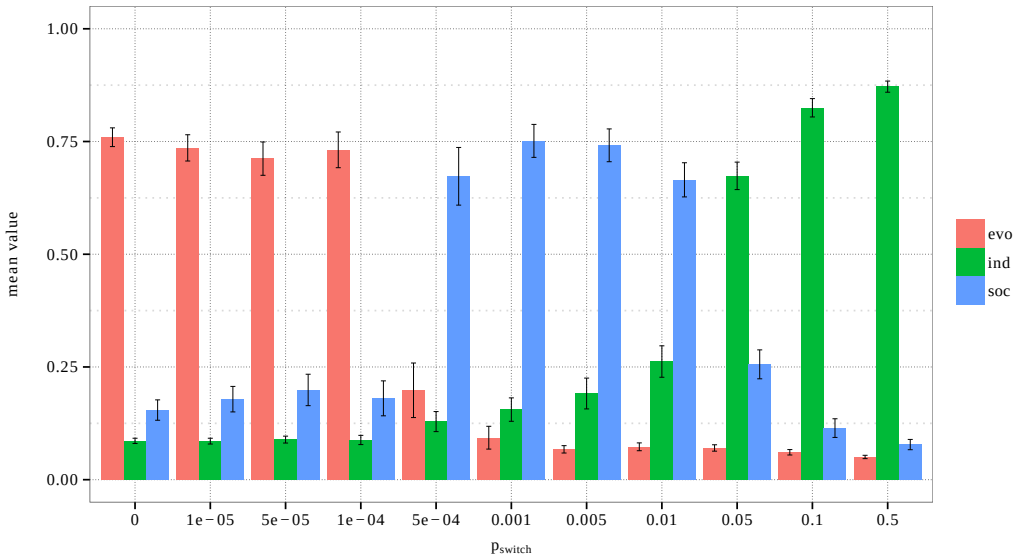

FIGURE 4.7: Learning modes at equilibrium across a range of p_{switch} values.

This result broadly supports the results of previous theoretical work that has described analogous learning regimes at slow, moderate and fast timescales (Cavalli-Sforza, 1981; Boyd and Richerson, 1985; Henrich and McElreath, 2003; Wakano and Aoki, 2006; Borenstein et al., 2008; Wakano et al., 2004). Our results highlight a novel aspect of this trend which is that these three phases are not discrete but there is a continuous overlap between them: note that social learning gradually diminishes in dominance between p_{switch} values of $0.01 \rightarrow 0.5$.

Having corroborated an existing result from the literature, we will now explore the ways in which these regimes are modified by different levels of environmental change and task complexity.

4.3.3 Factor 2: Task Complexity

The previous experiments all assume a constant environment task complexity, or B value: the number of phenotypic bits that must be correct to reach the peak of the unimodal fitness function. We can simulate modifying the complexity of the environmental objective by changing the value of B. Increasing B means that a greater number of subtraits must be obtained simultaneously to reach the peak of the fitness landscape. Conversely, an unfit agent can be so many bits away from the objective that any differential that it gains by obtaining one more bit becomes vanishingly small.

4.3.3.1 Increasing task complexity causes a phase shift from innate behaviour to individual learning

Parameter values (**4.3.3.1**)	N	B	α	μ	p_{switch}	p_{noise}	$trials$
	64	-	0.01	0.05	0.00	0.25	50

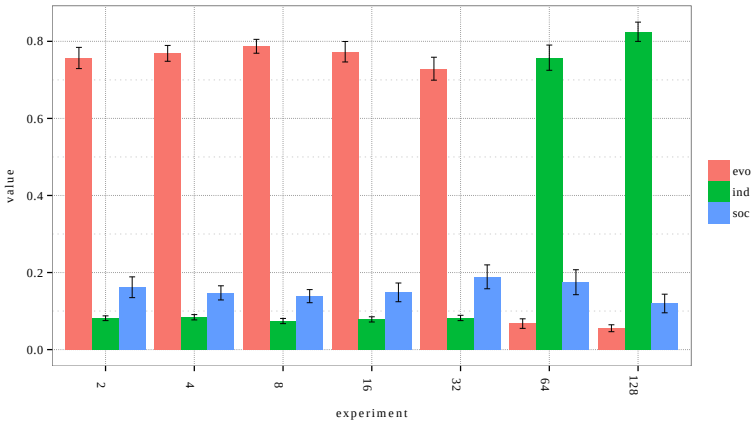

FIGURE 4.8: Changing B: Behavioural distributions at equilibrium

As seen in Figure 4.8, the results are surprising. At low task complexity ($B < 64$), innate behaviour consistently dominates. Beyond $B \geq 64$ is a second phase in which individual learning takes its place as the prevalent mode of learning. At equilibrium, social learning remains relatively infrequent at all levels of task complexity in a static environment.

Although the equilibrium states are similar for the low-B cases, the dynamics prior to this point differ substantially. Just as in Figure 4.1, social learning subsides and innate behaviour proliferates, with a determinate point at which mean $b_{evo} > b_{soc}$ within the population. We shall call the point at which this crossover occurs the *assimilation time*, as it represents the point at which genetic assimilation can be said to have taken place. In Figure 4.1, this takes place around $t = 50000$.

4.3.3.2 Learning Time Analysis

The reason for this two-phase distribution requires some further analysis. For the below discussion, we set aside social learning altogether and restrict our discussion to the pure interactions of evolution and individual learning.

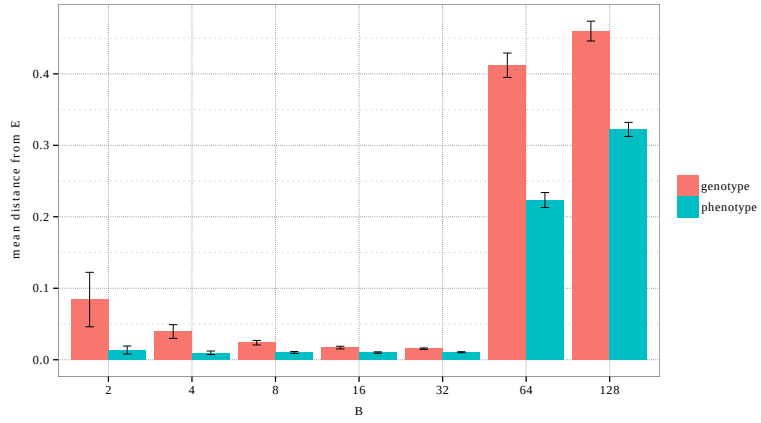

FIGURE 4.9: Changing B: Genotypic and phenotypic fitness at equilibrium

Consider an individual learner with a strongly unimodal fitness function of B bits, beginning from a genotype of minimal fitness. The learning process is a case of sampling repeatedly with replacement from $D = 1, 2, ..., B$. This is a case of a known problem in probability theory known as the Coupon Collector's Problem.

Let V_n denote the number of *distinct* selections in the first n learning attempts: that is, the number of different bits successfully learned. We are then interested in the random variable W_k, which denotes the distribution of samples needed to learn k bits. To reach optimal fitness (that is, setting all B bits of the phenotype to 1), let $k = B$.

W_B must equal the minimum number of trials needed for all B bits to be learned:

$$W_B = min\{n \in \mathbb{N}_+ : V_n = B\} \tag{4.1}$$

W_k is equivalent to a sum of independent random variables $Z = (Z_1, Z_2, ..., Z_B)$, where Z_i is the number of samples needed to go from $i-1$ distinct bits learned to i distinct bits learned. Each Z_i is a Bernoulli process whose probability is determined by the number of bits left to learn. It follows the geometric distribution:

$$Z_i \sim Geo\left(\frac{B-i+1}{B}\right) \tag{4.2}$$

From the moment generating functions of the geometric distribution, the expected (mean) value of Z_i can be derived. This gives the expected number of trials to learn the i^{th} bit of the genotype.

$$\mathbb{E}(Z_i) = \frac{B}{B-i+1} \tag{4.3}$$

By the linearity of expectations, the time taken to learn the entire genotype can be calculated by summing these independent variables.

$$\mathbb{E}(W_B) = \sum_{i=1}^{B} \frac{B}{B-i+1} \tag{4.4}$$

For an agent which already possesses A correct bits, to learn the remaining bits $(B - A)$, this becomes:

$$\mathbb{E}(W_{A\to B}) = \sum_{i=A}^{B} \frac{B}{B-i+1} \tag{4.5}$$

Example: An agent begins an experiment possessing 16 bits out of a total $B = 32$. To obtain the remaining bits, the mean time (in steps) is as follows:

$$\mathbb{E}(W_{16\to 32}) = \sum_{i=17}^{32} \frac{32}{32-i+1} \tag{4.6}$$

$$= \frac{32}{16} + \frac{32}{15} + ... + \frac{32}{1} \tag{4.7}$$

$$= 108.1833 \tag{4.8}$$

A single agent starts of an experiment ($B = 32$) with an average of 16 correct genotypic bits. We should thus expect each agent, if they are always engaged in individual learning, to take around 108 steps to attain maximal fitness.

Monte Carlo trials on 16-bit strings show that the expected time to learn the remaining 16 bits drops rapidly as population size N rises:

$$N = 1 \to \mathbb{E}(W_{16\to 32}) = 108.18 \tag{4.9}$$

$$N = 2 \to \mathbb{E}(W_{16\to 32}) = 87.55 \tag{4.10}$$

$$N = 4 \to \mathbb{E}(W_{16\to 32}) = 74.2 \tag{4.11}$$

$$N = 16 \to \mathbb{E}(W_{16\to 32}) = 57.65 \tag{4.12}$$

$$N = 64 \to \mathbb{E}(W_{16\to 32}) = 48.27 \tag{4.13}$$

For these trials, we use $N = 64$. This means that, according to the above analysis, the phase change does not take place at $B = 64$ (as inferred from the previous graphs) but at $B = 48$.

We shall proceed to examine the different regimes, where $B < 48$ and $B >= 48$.

The innate behaviour regime ($B < 48$)

Consider the expected lifespan of an agent. In a population of $N = 64$, given the birth-death scenario, each agent can expect to live for an average of 64 steps.

Based on the above experiments, the expected time for *some* agent to learn the remaining 16 bits is 48 steps.

This means that in a population of 64 agents, if all are engaged purely in individual learning, we can expect at least one agent to reach the fitness peak within its lifetime (the initial "optimal agent"). Any further individual learning trials will be deleterious and decrease the agent's fitness by increasing its Hamming distance to the fitness peak.

Meanwhile, other peers will have attained the same peak fitness also via having engaged primarily in individual learning. Multiple competitors will share the same maximal fitness as the optimal agent. However, they may – by chance mutations, and the same probability structure as described above – have also developed optimal traits in their genotype. Any offspring inheriting these bits will gain an adaptive advantage by exhibiting a higher tendency towards innate behaviour. Thus, the slow trend towards genetic assimilation begins, with the noisiness of individual learning giving way to reliable instinct.

The individual learning regime ($B \geq 48$)

As B rises, so too does the time taken to learn the remaining bits:

$$N = 64, B = 32 \rightarrow \mathbb{E}(W_{16\rightarrow 32}) = 48.27 \tag{4.14}$$

$$N = 64, B = 48 \rightarrow \mathbb{E}(W_{24\rightarrow 48}) = 87.64 \tag{4.15}$$

$$N = 64, B = 64 \rightarrow \mathbb{E}(W_{32\rightarrow 64}) = 132.4 \tag{4.16}$$

As we can see, even for a substantial population all of whom are engaging in learning, it **now takes longer than the agent's lifetime to reach optimal fitness**. This means that the population spends their entire lifetime trying to close the gap between possessing 50% to 100% of the environmental traits.

Here, the leading individual always has a fitness far below optimal, but still significantly greater than its peers. The only viable way to reach this fitness is to engage solely in individual learning, so learners proliferate rapidly. At this point, any tendency towards innate behaviour will be strongly deleterious as genotypes are purely the product of genetic drift, meaning that innate behaviour is penalised yet further. The greater the positive pressure on learners, the greater the negative pressure on individuals obeying their genetic instincts.

This cycle continues for the duration of the experiment; an agent spends its lifetime striving towards higher fitness by individual learning, but never reaching the top. This means that it never develops the breathing space needed to evolve other optimal behavioural modes. Every individual within the population is trapped in an evolutionary cycle of suboptimal behaviour. It is as if the entire population is attempting to climb up a slippery slope; the top can never be reached, but any momentary effort diverted from ascending the slope will lead to even more deleterious effects.

4.3.3.3 Task complexity ceiling, revisited

To confirm this theoretical analysis, this experiment was repeated with a focus on the band of B values between 32 and 64, and omitting social learners as in the analysis above. The results of this experiment are shown in Figure 4.10.

The predicted transition takes place at $B = 48$. Supporting the predictions above, B values less than or equal to 48 exhibit a dominance of innate behaviour. For $B > 48$, individual learning is the prevalent mode of transmission.

A small degree of noise is still present within the experimental outcomes, due to the stochastic nature of the simulations.

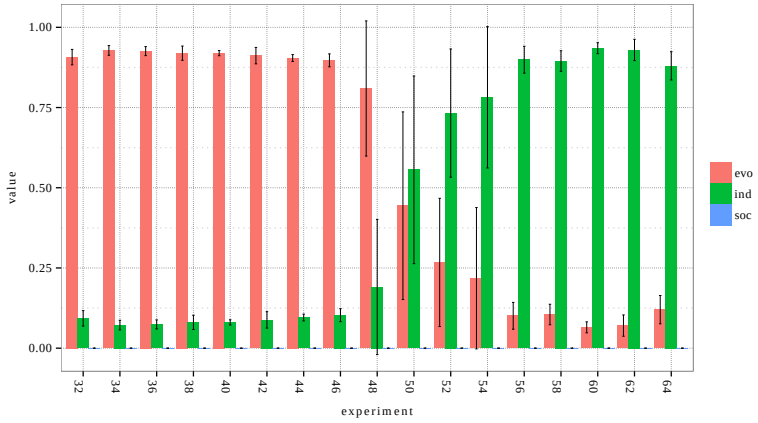

FIGURE 4.10: Changing B: Behavioural distributions between $[32, 64]$.

Parameter values (**4.3.3.3**)	N	B	α	μ	p_{switch}	p_{noise}	$trials$
	64	-	0.01	0.05	0.00	0.25	20

4.3.3.4 High task complexity locks a population in to a cycle of absolute genetic shielding, in which innate behaviour becomes deleterious

Much literature has been devoted to the concept of genetic shielding, in which behavioural plasticity reduces the strength and speed of genetic evolution by providing a timely phenotypic response to the environmental challenges that are posed (Wright, 1931; Ackley and Littman, 1991; Anderson, 1995; Mayley, 1997). If an organism is capable of altering its behaviour to adapt to a new selection pressure within its lifetime, the need to develop an innate tendency for the same trait is lessened.

We argue that the "slippery slope" cycle described above is a particularly insidious, absolute form of genetic shielding, in which the task complexity is so great that total plasticity arises to maximise fitness throughout the individual's lifetime. Yet in doing so, this absolute plasticity penalises innate tendencies. Whereas typical genetic shielding slows down the rate of genetic evolution, this absolute shielding actually *impedes* evolution, ultimately to the detriment of the individual's fitness.

4.3.4 Variability & Task Complexity

To establish a complete picture of the relative strengths of innate behaviour, individual learning and social learning in fluctuating environments, an array of experiments was executed over a range of rates of change ($p_{switch} \in [0.00001, 0.5]$) and degrees of environmental complexity ($B \in [2, 1024]$).

Each permutation of p_{switch} and B was executed for 5×10^5 timesteps, and a snapshot taken of the final distribution of behavioural traits. These are graphed in Figure 4.11, with red, green and blue segments corresponding to the proportional presence of b_{evo}, b_{ind} and b_{soc} respectively.

Parameter values (**4.3.4.0**)	N	B	α	μ	p_{switch}	p_{noise}	$trials$
	64	-	0.01	0.05	-	0.25	10

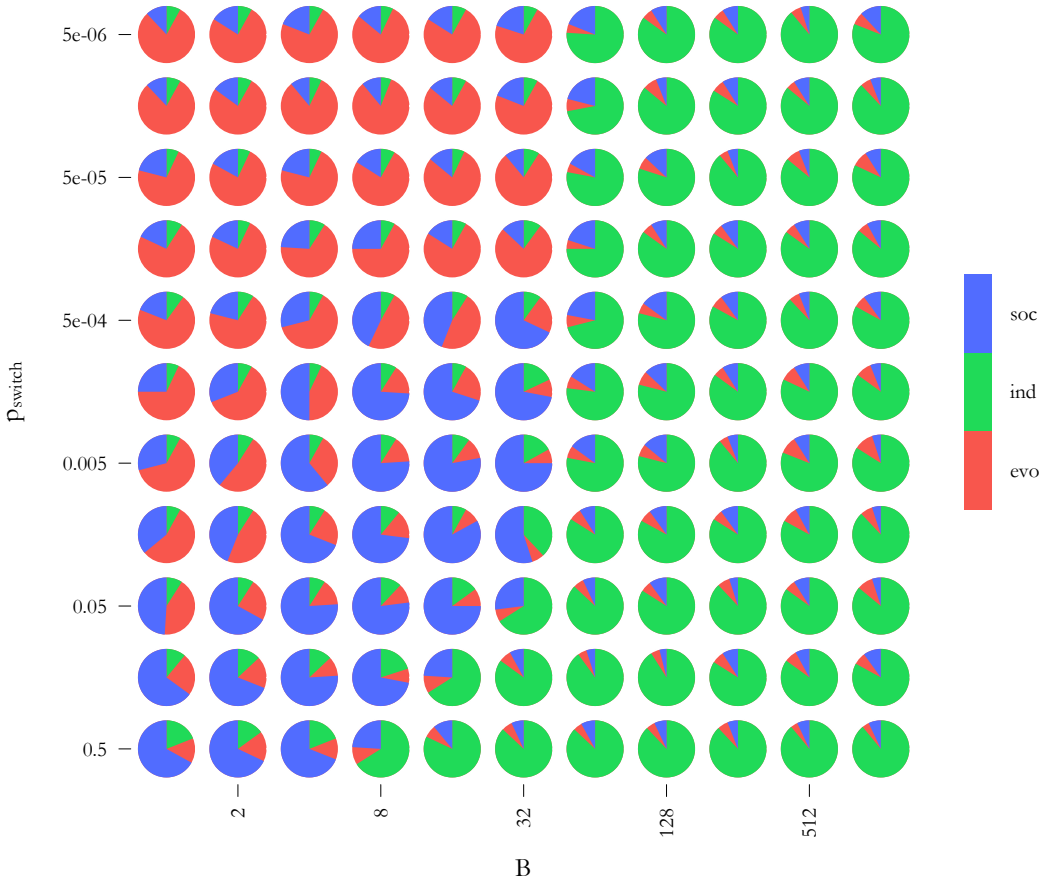

FIGURE 4.11: Learning modes at equilibrium across an array of p_{switch} and B values.

There are distinctive regimes in which each mode of learning dominates, with interactions between these two variables. At low rates of change and in simple environments, the population exhibits a prevalence of b_{evo}: individuals tend towards innate behaviour in contexts wherein an optimal genotype is simpler to obtain or maintain.

As either fluctuation rate or task complexity increase, strategies become more mixed, with a trend towards social learning at median values of each. A greater amount of noise in this regime suggests that selection pressures are weaker, leading to more vulnerability to stochastic variation. As per the analysis in Section 4.3.3.2, experiments with task complexity beyond $B \geq 48$ are uniformly dominated by individual learning.

4.3.4.1 Three-phase learning regime is bounded by task complexity

In Section 4.3.2.1, we showed that slow-, moderate- and fast-changing environments lead to dominant regimes of innate behaviour, social learning and individual learning respectively. From Section 4.3.4, we can see that these results are bounded by the level of complexity of

the environment.

In very simple environments ($B < 8$), individual trial-and-error is marginalised in favour of innate behaviour and social learning, and never becomes the dominant paradigm, even at a high level of p_{switch}. In increasingly complex environments, the window of social learning becomes narrower until individual learning dominates entirely.

Beyond an environment complexity bound of $B = 48$, following the analysis in Section 4.3.3.2, individual learning dominates entirely at all rates of environmental change.

The three-phase regime shown in §4.3.2.1, therefore, bounded by the complexity of the set of traits that an individual seeks to achieve: in particularly low-complexity and high-complexity environments, learning modes are constrained to b_{evo} and b_{ind} respectively.

4.4 Summary

The research goal in this chapter was to evaluate the different learning methods that individuals use to maximise their fitness in stable and changing environments, based on the model defined and described in §3.2.

In stable environments (§4.3.1), we showed that a new population initially adopts social learning (§4.3.1.1), subsequently undergoing genetic assimilation and reaching an equilibrium in which innate behaviours dominate. This is a socially-mediated form of the Baldwin effect. The genetic and phenotypic fitness curves of this population (§4.3.1.2) demonstrate substantial genetic shielding, in which phenotypic plasticity reduces genetic selection pressure by adapting to the environment within the individual's lifetime.

After a major environmental perturbation (§4.3.1.3), we see a series of three sequential regimes: individual learning immediately increases on prevalence, enabling individuals within the population to acquire the new successful traits, which then propagate via social learning. Finally, the traits are genetically assimilated, resulting in an equilibrium dominated by innate behaviour. This equilibrium behaviour has a price, however, which is that a homogeneously-adapted population has a diminished adaptive response to future perturbations.

The introduction of variable task complexity to a static environment demonstrates a phase change: beyond a certain threshold of task complexity, innate behaviour is superseded by individual learning (§4.3.3.1). A mathematical analysis shows that this is due to a "slippery slope" paradox (§4.3.3.2), in which a population becomes trapped at a suboptimal level of fitness when faced with a task whose subtraits they cannot acquire within their lifetime. The need to continually engage in learning activities causes innate tendencies to be completely eliminated, which causes an absolute form of genotypic shielding (§4.3.3.4).

In fluctuating environments (§4.3.2), the behavioural landscape features a much higher rate of social learning (§4.3.2.1). At slow rates of environmental change, innate behaviour is advantageous. At moderate and fast rates of change, social and individual learning respectively dominate. Some overlap is evident between regimes, with tendencies gradually diminishing in prominence.

Varying the environment's fluctuation rate and task complexity over a spectrum of values, we showed a series of analogous regimes for task complexity (§4.3.4), in which simple environments reward innate behaviour even at moderate rates of change. We infer

that the three timescale-based learning regimes (individual/social/innate) only takes place within a certain sweet spot of task complexity (§4.3.4.1), outside of which either social or innate behaviours are not found. Individual learning remains prevalent in all environments characterised by a complex task.

Chapter 5: Structured Populations

5.1 Overview

Chapter 4 explored the interactions between learning modes and environmental properties within a well-mixed population: that is to say, one in which population structure does not exist. The population was treated as unstructured, interacting with arbitrary neighbours. This is not representative of a real-world ecosystem, wherein an individual is embedded in a bounded spatial context, typically interacting with a regular selection of conspecifics.

In this chapter, we incrementally build upon the results of our numerical model with a series of experiments that introduce population structure via individual-specific locality. Agents can no longer interact with any other peers within the population, but are constrained to interact only with those neighbours nearby to them. Only these adjacent peers can be treated as exemplars (or subjects) for social learning activities.

This more closely resembles the structure of a real-world environment. In a real biological system, an organism has a limited range of perception, and can only accurately observe a phenomenon within a given range. Moreover, attention is a limited resource, and can only be devoted to a small number of targets at once (Heyes and Galef Jr, 1996).

5.2 Key Questions

In this chapter, we address the third of our major environmental factors: **population structure**.

In stable and unstructured populations, we have seen that individual learning becomes the dominant strategy in populations that are very fast-changing or characterised by complex tasks (§4.3.4). Innate behaviour proliferates in simple or static contexts. Do these results hold in a structured population, and those with variable neighbourhood size?

To respond to this question, we will investigate the learning trends that take place in transient and equilibrium states over a series of different discrete, fixed population structures, varying topology and neighbourhood size.

We will also visit some additional questions: Does population structure change the predicted outcome of the Baldwin effect? How does population structure affect the rate and success of adaptation?

5.3 Discrete Grid Structures

We place our agents on a one-dimensional lattice, with periodic boundary conditions. The general model structure remains the same as Model ∅, with no changes to genetic evolution, individual learning, or fitness parameters. Every agent continues to be awarded fitness based on its Hamming distance from the environment's bit string E.

Further properties are introduced:

- Social learning can only take place within the von Neumann neighbourhood (range $r = 1$) of each agent; that is, only the 2 adjacent agents are used as potential exemplars. Roulette wheel selection is still used to select an exemplar, with a preference for the fittest neighbour.
- When an agent reproduces, its offspring replaces one of its neighbours, selected uniformly randomly.

One agent inhabits each cell. The CA is therefore initialised with a width equal to the population size N, currently 256.

5.3.1 Commentary

We have now introduced a population structure to our model. How should we interpret this in biological terms?

It should be noted that our population is *structured* but not *located*. There is no sense of spatial embodiment within this system; though an agent has neighbours, it does not matter where on the plane the grouping sits, as the environment is invariant across its surface. An agent's spatial environment is constituted solely by its peer group. The peer group size is fixed for analytical simplicity. This should be understood as modelling the mean number of neighbours within a population.

The population structure also remains static, with no changes to network structure within or between generations. This implies that an individual's peer group remains fixed through its lifetime, and that circulation between new peers plays a negligible role.

A further spatial assumption introduced by this notion of space is that of exact neighbourhood regularity. In a 2-dimensional von Neumann neighbourhood, each organism has precisely four neighbours, with no long-range social links to join up distant clusters. In animal groups, social structures take a range of forms, often exhibiting scale-free patterns (Wey et al., 2008) with a long-tailed power law distribution of node degree (Barabási and Albert, 1999). Moreover, social networks themselves evolve substantially over time as peer groups change (Kossinets and Watts, 2006).

However, Flache and Hegselmann (2001) demonstrate that the qualitative behaviours of a range of theoretical models are robust to different grid formations: switching from regular to irregular neighbourhood structure did not have significant qualitative effects on the models' outcomes. Their models, like ours, concern social transmission and dispersion of traits, and should thus predict similar properties to hold within the model proposed here.

5.3.2 Results

5.3.2.1 Rigid environmental structure causes behavioural lockin, preventing the Baldwin effect from taking place

We begin by situating our population within a range of different environment topologies – the well-mixed, numeric setting of Chapter 4, and discrete grids in 1 and 2 dimensions – and allowing a series of trials to run to equilibrium. The equilibrium states are shown in Figure 5.1.

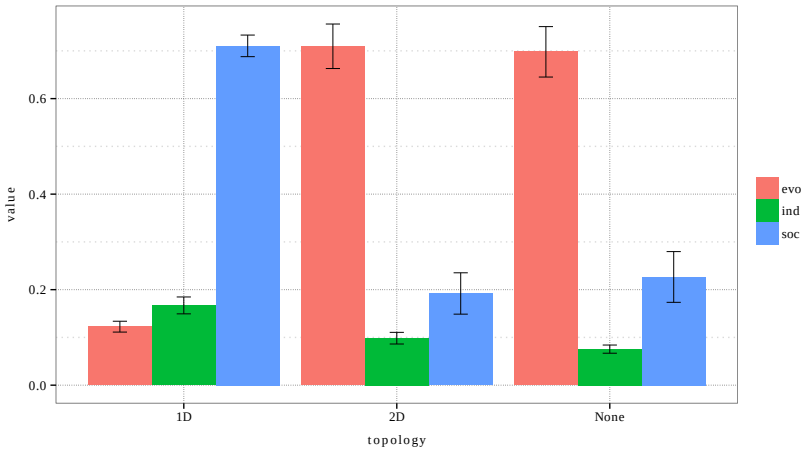

FIGURE 5.1: Population structure: Learning modes in well-mixed, 1D and 2D structures

.

Parameter values (**5.3.2.1**)	N	B	α	μ	p_{switch}	p_{noise}	*trials*
	64	32	0.01	0.05	0.0	0.25	50

We immediately see a major change in learning trends. Interestingly, the dividing line does not lie between unstructured and structured environments. The 1D structure results in a significant dominance of social learning behaviours, yet the 2D and well-mixed populations are both dominated by innate behaviour.

This seems counterintuitive. Individuals within the 2D and well-mixed populations possess larger neighbourhood sizes, meaning they have a greater number of exemplars to learn from. This would lead to the natural hypothesis of a larger population size *improving* the success of a social strategy. What is the cause behind this contradictory result?

Examining the time-series behavioural trends of a 1D-structured population (Figure 5.2), we can see that the familiar initial rise of social learning (*cf* Figure 4.3.1.1) never subsides to give way to innate behaviours; genetic assimilation does not take place within this 1D-structured population.

A closer inspection of the experiments' dynamics reveals that this is because the population becomes trapped in a behavioural "lock-in". A successful individual – in this case, one who obtains a reasonably fit genotype and a high propensity towards social learning – proliferates repeatedly, replacing one or other of its neighbours. These neighbours also possess a high fitness and are likely to replace one of *their* neighbours.

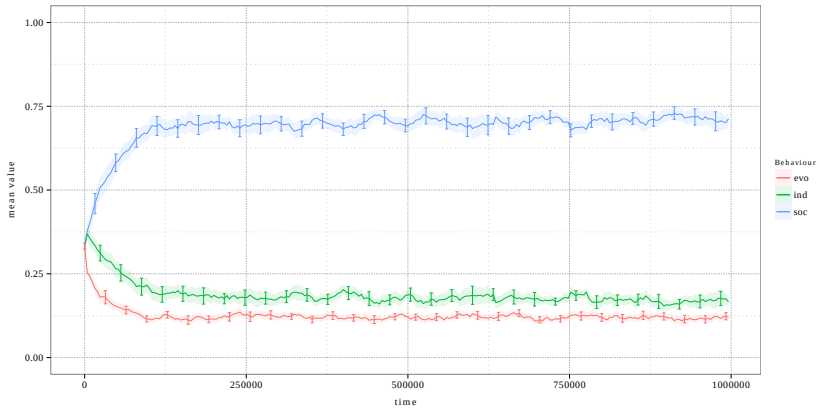

FIGURE 5.2: Population structure: The rigid structure of a 1D environment prevents the Baldwin effect from arising

.

Chains of reproductive events take place which ripple along the population, substantially reducing the chance that a suboptimal cluster of individuals will survive for long enough to develop the optimal genotype needed to supplant their learning peers.

Like the hidden "refuges" described by the empirical studies of Gause (1934) and Luckinbill (1973) – spatial pockets which allow a predator to remain concealed from prey and survive over a long period – the amorphous, less rigidly-structured populations in our well-mixed and 2D environments can foster small collections of individuals which are given the breathing space to evolve better-adapted genotypes (and the propensity to act upon them, in the form of b_{evo}).

The linear cascades of birth-death reproduction that take place in our 1D scenario, conversely, mean that there is effectively no place to hide for a suboptimal individual. Successful behaviours are reinforced and proliferate cyclically. This is evidenced by the relatively small behavioural variance shown in Figure 5.2, which is substantially lower than the variance in the 2D and well-mixed cases. Viewed from another perspective, it is this behavioural noise that is needed to give rise to the conjunction of genetic fitness *and* innate tendency (b_{evo}) that are required for genetic assimilation to occur.

This mechanism also gives rise to the following result.

5.3.2.2 Rigid population structure prolongs the life of successful individuals

Figure 5.3 illustrates the mean age of individuals within each of these population structures; that is, the number of timesteps that they have survived for. In a numeric population, a reproducing individual replaces a randomly-selected member of the population, meaning that every individual has an equal chance of being replaced every timestep. This means that the mean age will always maintain a value around the population size N (in this case, 64).

In the 2D and 1D structures, the picture is quite different. The rigid structure introduces a clustering effect as described in the previous section, meaning that a successful individual may survive and proliferate for a great many generations, unchallenged. In the 2D case, this occurs for a long time, until the population eventually converges to a unified state of

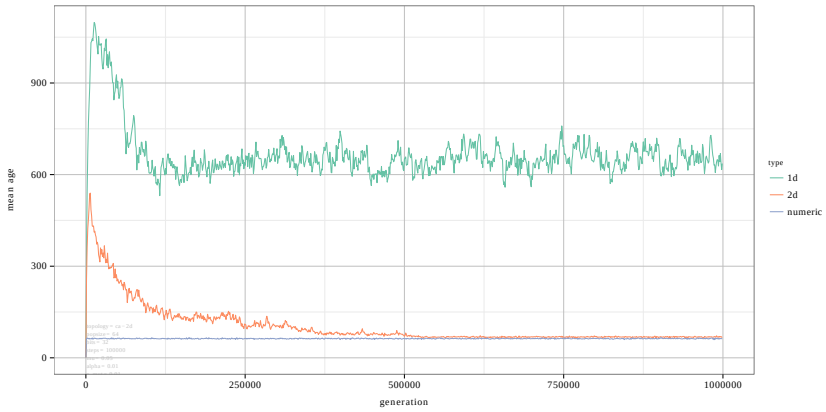

FIGURE 5.3: Population structure: Rigid population structure prolongs the life of successful individuals

.

innate behaviour.

In the 1D case, conversely, this global optimum is never reached; due to the ossified success of the social learning individuals, the rigid structure maintains the dominance of this suboptimal behaviour.

This has similarities with the "slippery slope" paradox described in §4.3.3.2; a learning regime maintains dominance at a local optimum because the population does not have the evolutionary breathing space needed to cross the fitness valley to the global optimum.

5.3.2.3 Social learning in a changing environment is less effective with a smaller neighbourhood

The subsequent experiments have all taken place in an unchanging world, in the absence of environmental changes. We now re-introduce instability to the population via the stochastic update process described in the previous chapter ($p_{switch} = 0.01$).

In a fluctuating environment, the equilibrium states across well-mixed, 1D and 2D population structures are given in Figure 5.4.

Parameter values (**5.3.2.3**)	N	B	α	μ	p_{switch}	p_{noise}	*trials*
	64	32	0.01	0.05	0.001	0.25	50

All three population structures now exhibit the same ordering of learning modes; social learning prevails, with a smaller degree of individual learning, and an even lesser degree of propensity towards innate behaviour.

Statistically significant differences do remain between the learning regimes in each of these contexts. Most prominently, social learning is less successful within a 1D population than in 2D or well-mixed environments. This seems surprising, given the previous section's conclusion that, in a stable environment, social learning is *more* successful in a 1D structure.

The reason for this differential is that offered in our original hypothesis; that the smaller neighbourhood size in a 1D population means that the efficacy of social learning is lessened. In a well-adapted population within an environment that is regularly fluctuating, an individual must engage on trial-and-error learning to acquire a newly-required trait.

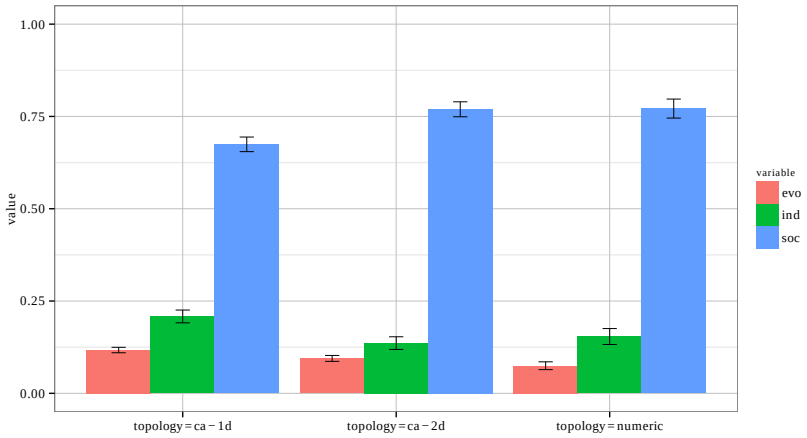

FIGURE 5.4: Population structure: Social learning in a changing environment is less effective with a smaller neighbourhood

This requires the correct trait to be selected by chance, a statistically infrequent occurrence. Alternatively, an individual can await one of their neighbours learning this trait and then use their social learning propensity to imitate. The latter is a less risky proposition. However, in a 1D environment, an individual only has two exemplars to copy from. A 2D structure gives twice the number of potential exemplars, meaning less time to wait for a nearby individual to have acquired the trait.

Put differently, a better-connected population means that there are shorter distances between any pair of individuals, meaning that social information can propagate much more quickly through a population that is well-mixed or has a wider neighbourhood.

To explore these issues further, we shall proceed to expand our study of neighbourhood sizes to populations structured on graphs. This will enable us to investigate populations that do not need to sit on a spatially-explicit field, allowing us to test arbitrary neighbourhood structures.

5.4 Regular Graph Structures

Grid-based lattice structures, whether 1D or 2D, are a subset of the class of regular graphs, with degree of 2 and 4 respectively. In this section, we will explore the wider class of regular graph structures, a standard approach to investigating the effect of population structure on collective dynamics (Newman, 2010; Lieberman et al., 2005; Szabó and Fáth, 2007).

5.4.1 Construction

The object is to create a k-regular social graph, in which every individual has precisely k connections to others. Varying k consequently allows us to manipulate the population's neighbourhood structure, with larger k equating to a larger clique size.

To construct k-regular graphs, we use the pairing method described by Bollobás (1985):

1. Create nk points, distributed equally across n buckets.
2. For each point, create a link between it and another randomly-selected point, until nk pairs are obtained.
3. Use these buckets to create a graph in which each bucket corresponds to a vertex, with links between points mapped to edges between vertices.
4. If the resulting graph is simple (that is, contains no self-connected nodes or repeated edges between nodes), terminate; otherwise, restart the process.

This gives a population in which every individual is connected to precisely k others. Such regular graph structures will be used throughout the rest of this section's results.

5.4.2 Results

5.4.2.1 Larger neighbourhood size speeds up evolutionary convergence

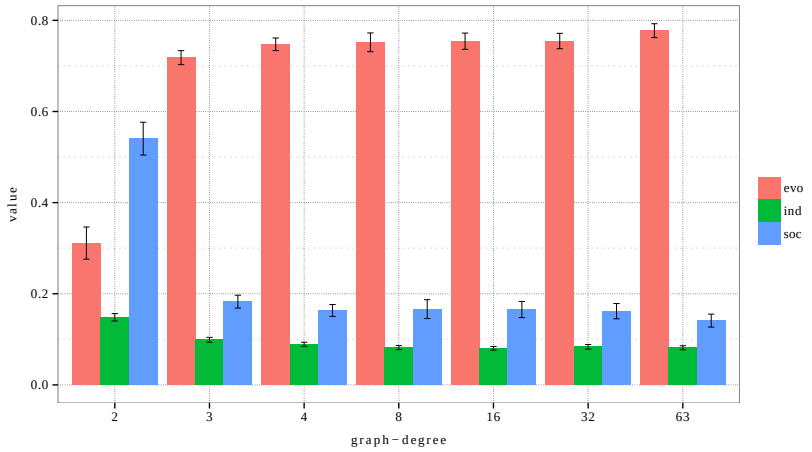

FIGURE 5.5: Regular graphs: Behavioural modes against k

Parameter values (**5.4.2.1**)	N	B	α	μ	p_{switch}	p_{noise}	*trials*
	64	16	0.01	0.05	0.00	0.25	100

We begin by carrying out a series of experiments on varying neighbourhood sizes, with a population of $N = 64$. The experiment runs until the population's behavioural traits reaches an equilibrium state. A neighbourhood size of $k = 63$ is equivalent to the well-mixed case, in which any individual is in the neighbourhood of any other individual.

Results are shown for the logarithmic series of k values: $k = 2, 4, 8, 16, 32, 63$ (with 63 being the maximum possible value). The result is also shown for the intermediate case $k = 3$, which was discovered to have an interesting behaviour.

The equilibrium behavioural traits at each of these k-values are shown in Figure 5.5. This is a continuation of the results demonstrated in the previous section. With a k of 2, the population fails to genetically assimilate the successful traits and remains dominated by social learning. For $k > 2$, genetic assimilation occurs, and innate behaviour dominates

(ANOVA, $p < 0.0001$). No statistically difference is found for any pairwise set of results beyond $k > 2$ (Tukey HSD post-hoc test: 3 vs 63, $p > 0.1$). Although level of adaptation in isolation is not a meaningful figure, as fitness only has relevance in a competitive situation, it is useful to note the disparities in genetic shielding and plasticity that take place as k rises.

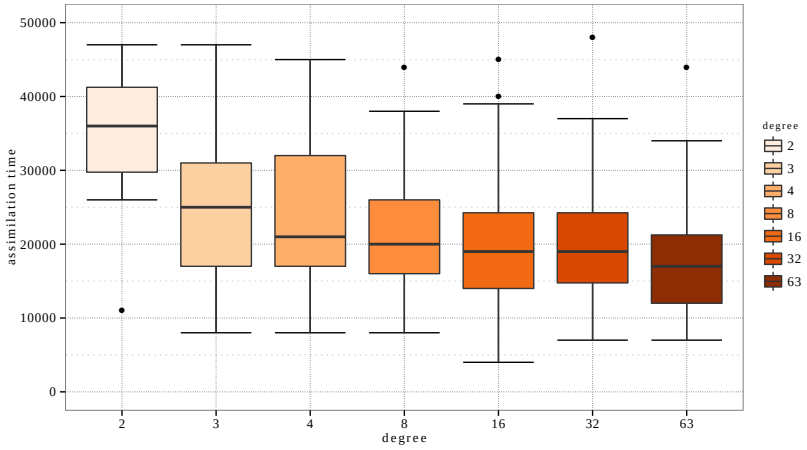

FIGURE 5.6: Regular graphs: Assimilation time against k

The assimilation times of these populations, as defined in §4.3.3.1, are shown in Figure 5.6.

A greater neighbour size results in a shorter assimilation time, with large-k populations undergoing genetic assimilation faster than those of a smaller neighbourhood size (ANOVA, $p < 0.0001$). Neighbourhoods of $k = 2$ converge significantly slower than larger k values ($p < 0.0001$, Tukey HSD post-hoc test); $k = 3$ converges marginally slower than those of $k \geq 16$ ($p = 0.05$ for $k = 16$).

Graph behavioural modes with non-adjacent birth

To verify whether the behaviour shown in Figure 5.5 is resilient to different modelling assumptions, we repeated this experiment whilst relaxing the constraint that children are created in neighbouring vertices from their parent individual; instead, offspring are positioned at a random location on the graph.

The results are shown in Figure 5.7. The behaviour shown in the above result no longer takes place. This suggests that the "breathing space" hypothesis was, in fact, an artefact of the model in question.

5.4.2.2 Larger neighbourhood sizes results in increased social learning propensity

We finally subject a graph-structured population to a changing environment, across a range of p_{switch} values. The resulting learning modes at equilibrium are shown in Figure 5.8.

Parameter values (**5.4.2.2**)	N	B	α	μ	p_{switch}	p_{noise}	*trials*
	64	16	0.01	0.05	-	0.25	100

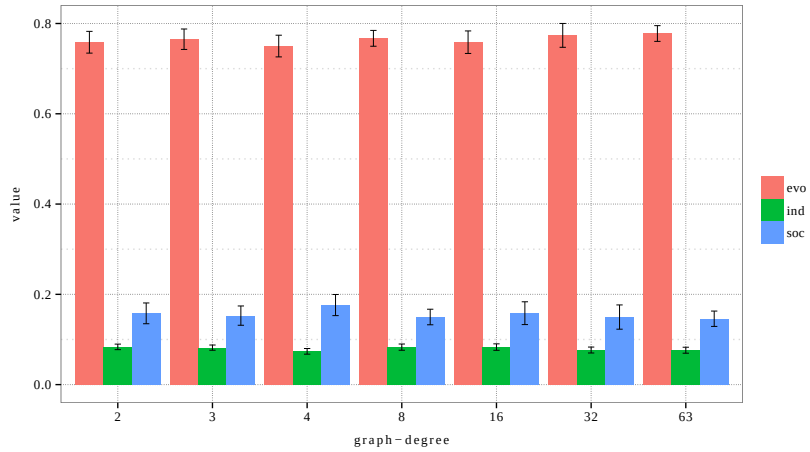

FIGURE 5.7: Regular graphs: Behavioural modes against k, random offspring placement

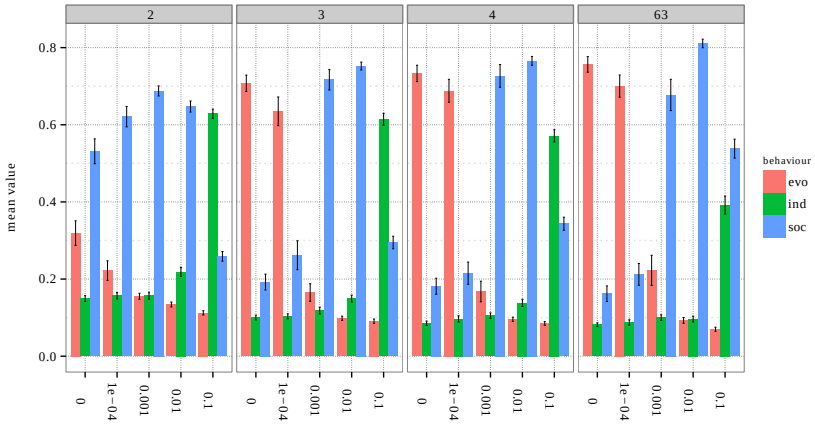

FIGURE 5.8: Regular graphs: Behavioural modes against p_{switch} and k

This is a rich and somewhat complex set of results that is not initially straightforward to digest. In general, individuals with a larger neighbourhood size demonstrate a higher level of innate behaviour, courtesy of the predisposition towards genetic assimilation described in the previous section. The mean value of b_{evo} in high-k environments is greater than those in low-k for the equivalent environmental rate of change.

However, at higher rates of p_{switch}, individuals within the high-k populations demonstrate a greater degree of social learning, courtesy of the superiority of social learning with a larger neighbourhood.

Indeed, a larger neighbourhood size results in better-adapted individual at all environment variabilities. Figure 5.9 depicts genotypic and phenotypic distance from the environmental objective E across each of these p_{switch}/k contexts. Individuals within the high-k population show improved genotypic and phenotypic fitness for all values of $p_{switch} \geq 0.001$ (ANOVA, $p < 0.0001$). For higher variability rates, no significant difference is observed (comparing genotypic fitness for $k = 2$ vs $k = 63$, Tukey HSD post-hoc test

indicates $p > 0.5$).

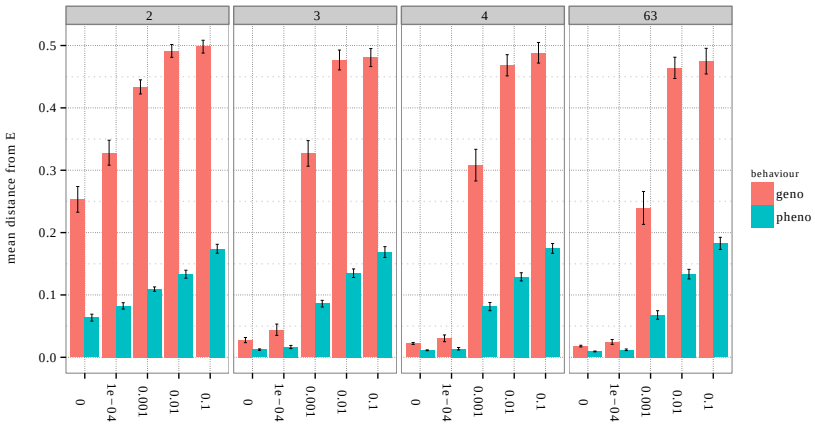

FIGURE 5.9: Regular graphs: Genetic and phenotypic fitness against p_{switch} and k

Returning to Figure 5.8, an interesting parallel emerges when these results are compared with the three-stage "double Baldwin effect" of Figure 4.3.

Recall that two transient stages occur after a significant perturbation: firstly, of individual learning, needed to acquire the novel environmental traits; and subsequently, of social learning, as these new culturally-ingrained traits proliferate via social transmission. It appears to be the case that smaller values of k remain trapped in these transient states. At low values of p_{switch}, low-k populations are dominated by social learning, whereas high-k populations assimilate this knowledge into their genotype and develop the preferential innate behaviour. Similarly, at high p_{switch}, low-k populations are dominated by individual learning, whereas high-k populates are able to disseminate information via social transmission, a cheaper and ultimately more successful approach. In summary, a larger neighbourhood structure enables a population to escape from suboptimal transients, giving the breathing space needed to reach a global optimum – be it social learning in changing environments, or gene-mediated instinct in static contexts.

5.5 SUMMARY

In this chapter, explored a range of different population structures, giving order to a community by limiting each individual to a bounded neighbourhood of interactants.

We began (§5.3) by focusing on 1D and 2D discrete lattice structures. We saw that the rigid structure of a 1D population prevents the Baldwin effect from taking place (§5.3.2.1), by suppressing the 'refuge'-like pockets of evolutionary breathing space needed to facilitate genetic assimilation. This results in a much wider distribution of individual survival times (§5.3.2.2), increasing the disparity between fit and unfit agents.

Introducing environmental fluctuation, we showed that social learning in a changing environment is less successful in 1D populations due to the smaller neighbourhood of exemplars available to these individuals (§5.3.2.3).

We then shifted to examine populations on regular graphs (§5.4.1), enabling us to

explore neighbourhood sizes of an arbitrary k. This produced two key findings: firstly, that a larger neighbourhood size speeds up evolutionary convergence (§5.4.2.1), allowing successful traits to propagate more rapidly through the population; and secondly, that larger k results in a generally better-adapted population (§5.4.2.2), over a range of rates of environmental change.

Chapter 6: Spatial Heterogeneity

The previous chapters have demonstrated the impacts of environmental change, task complexity and population structure on the optimal modes of learning arising within a population. However, selective pressures have remained uniform across the population, with every individual attempting to optimise the same fitness function. This means that a given adaptation – flipping phenotypic bit x from $0 \rightarrow 1$ – will be beneficial either for everybody, or for nobody. This correspondingly means that successfully copying a behaviour which is adaptively beneficial for my neighbour will also be adaptively beneficial for me, a constraint that is likely to impact upon preferential learning behaviours.

This chapter introduces spatial structure to the population, experienced as variation in fitness demands between individuals, or "individual stochasticity" (McNamara et al., 2011; Lande, 1988). Siting the population on a two-dimensional grid with patterned variance defined by a pair of quantifiable metrics, we will investigate the effects that spatial heterogeneity has on the evolution of a learning population.

6.1 Overview

We briefly summarise two major metrics that will be used to quantify spatial heterogeneity (§6.3), *fragmentation* and *gradient*, and outline a method of constructing a heterogeneous 2D landscape with spatial variance governed by these parameters (§6.3.1). We define three different levels of spatial heterogeneity: *uniform*, denoting completely homogeneous environments; *random*, denoting those which are have no correlation across pairwise points, and *structured*, which follow the two landscape metrics of fragmentation and gradient.

We subsequently describe an approach to modelling movement across the environment (§6.3.2), adding evolvable traits which determine an individual's mobility and tendency towards social clustering.

We then present a series of experimental results from these heterogeneous environments (§6.4). We begin by investigating learning regimes within uniform, structured, and random environments (§6.4.1.1), for populations that are alternately sessile (fixed) and mobile. We repeat these experiments over the same set of environments with a level of fluctuation over time (§6.4.1.4), and investigate the changes in mean fitness that each types of learning mode produce.

6.2 Key Questions

This chapter concerns the fourth of this thesis' overarching environmental factors: spatial heterogeneity. To investigate the behavioural effects of heterogeneity within our model, we

first need to be able to construct a spatial environment whose heterogeneity is determined by a minimal set of ecologically relevant properties.

This is a methodological issue, concerned with extending our current model into the realm of spatial variance in a way that is simple, quantifiable and extends from the existing research.

We will first define the set of metrics that we seek to model: properties of spatial arrangement and relationships between, and within, discontinuous clusters of resources. We then need to define an approach to creating an environment which exhibits the full range of values of these metrics. Once we have done so, we can carry out experiments at different locations in metric space. For example, we might seek to understand how evolution progresses in an environment with low patchiness and high fragmentation, or to carry out a series of experiments between low and high fragmentation.

We will define two key metrics to determine an environment's heterogeneity: *gradient* and *fragmentation*.

Field studies demonstrate correlations between environmental heterogeneity and tendencies towards phenotypic plasticity (Baythavong, 2011). Using this model, we can predict exactly how particular types of heterogeneity might result in tendencies towards different kinds of plasticity.

We are particularly interested in whether social learning may be seen to particularly arise in specific kinds of environment, as this has not previously been given much depth of study.

We will finally introduce mobility to our population. By moving through an environment, an organism may be able to reduce its experienced heterogeneity (by moving towards similar environments when an environment changes to become unfamiliar) or increase its experienced heterogeneity (by moving into new and unfamiliar terrain). We will introduce movement abilities to investigate interactions between movement, learning and evolution in heterogeneous spaces.

6.3 Constructing Heterogeneous Environments

Before we can conduct experiments within heterogeneous environments, we first need to determine what exactly we mean by heterogeneity. We will define two key properties to characterise a heterogeneous environment, and describe a method of constructing environments which possess these properties based on the common single-resource metrics found within landscape ecology (see Section 2.4 for background).

The metrics we define are *fragmentation* and *gradient*.

fragmentation, F – are resources aggregated in few large bodies, or are there multiple, smaller aggregations of resources?

gradient, G – are resource deposits surrounded by gradients of lesser abundance, so that organisms can follow a resource gradient to a local maximum? When a single resource is available, this is determined by adjacent cells having a high pairwise correlation in their resource abundance, so is equivalent to having a low patchiness. Gradient is the inverse of *patchiness*: an environment with a low G value has high patchiness.

We adopt this nomenclature partly to disambiguate with previous literature which uses "patchiness" in a looser sense (Marquet et al., 1993; Grünbaum, 2012), and partly because "gradient" is a more intuitive term in a single-resource environment.

6.3.1 Method

The present model differs from the landscape ecology literature in one key fashion. Unlike in a real-world ecosystem, which might be made up of segmented areas of qualitatively different kinds or niches, we will begin by describing landscapes containing one type of resource, with "heterogeneity" measuring differences of resource plenitude in neighbouring locations; that is, differences in payoffs to individuals inhabiting each location. Certain properties described in Riitters' and Li's analyses are therefore not applicable, and others require minor modification to apply in a single-niche world. For example, *NTYP*, one of the 6 dominant landscape characteristics in Riitter's factor analysis (Riitters et al., 1995), corresponds to the "number of attribute classes" (or resources) designated within that landscape, and so plays no role here.

Romme first describes patchiness as the "contrast between adjacent communities" Romme (1982), typically determined by the boundary relationships between different types of land cover. In our single-factor approach, we recast patchiness as a relationship between adjacent cells of the same resource type, wherein high patchiness is equivalent to having a small degree of autocorrelation between adjacent cells.

Objective: A method of stochastically constructing a two-dimensional spatially-varying landscape, whose heterogeneity properties are determined by the *fragmentation* and *gradient* indices.

Approach: The method we propose is similar to that described by Perlin (2002) in his work modelling fractal-like landscapes for computer graphics applications. The general approach is to sum together several N-dimensional interpolated noise functions at different spatial scales, or "octaves", where a lower octave has a wider period and a greater amplitude. The amplitude of each noise function can be thought of as corresponding to a height value. The amplitude at higher octaves falls off exponentially, according to a *rolloff* constant. This means that higher octaves are more detailed spatially but have less 'height' impact on the landscape.

To construct a 1-dimensional landscape with k octaves, we begin with a series of stochastic noise functions ζ_i, each of which takes a single position parameter x and returns a uniformly random number within $[-0.5, 0.5]$. The *period* of ζ_i is the distance in x between values, with the period halving every octave:

$$\text{Period}(\zeta_0) = 2^0$$
$$\text{Period}(\zeta_1) = 2^{-1}$$
$$\text{Period}(\zeta_2) = 2^{-2}$$
$$\ldots$$
$$\text{Period}(\zeta_i) = 2^{-i}$$

Between values, linear interpolation is performed, creating the interpolated layers of the landscape (Figure 6.1).

The aggregate height value of the landscape at a given point x is found by summing these noise functions, cumulatively multiplied by the *rolloff value* ($r \in [0, 1]$) which determines the ruggedness of the landscape:

$$\zeta(x) = \sum_{i=0}^{k} r^i \zeta_i(x) \tag{6.1}$$

The resultant output is shown in Figure 6.1 ($x \in [0, 4], k = 5, r = 0.75$). After 5 additive octaves, the landscape is fairly rugged, with several local optima. Increasing the value of r would increase this ruggedness proportionately.

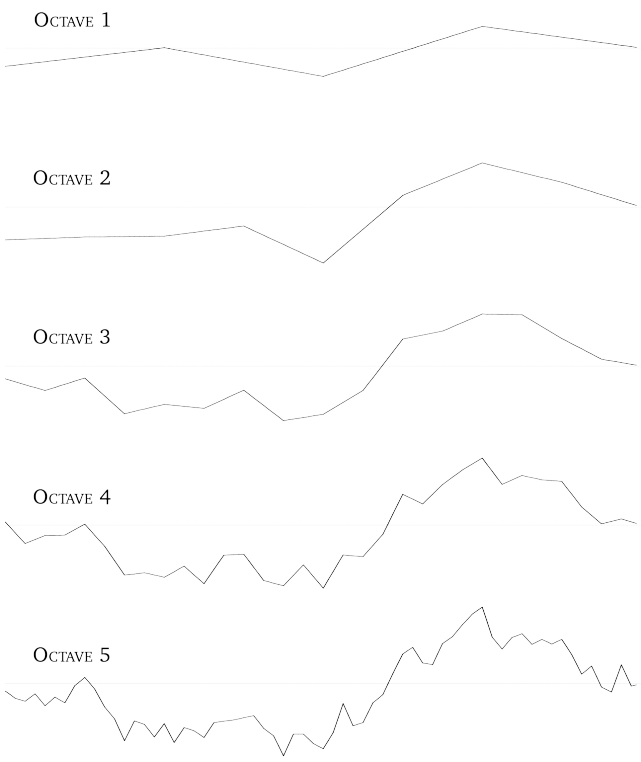

Figure 6.1: Constructing a landscape by summing cumulative octaves (where each octave is summed with those prior to it, thus "Octave 3" shows Octaves 1, 2 and 3 summed.). Each octave has finer detail and lower amplitude than the last.

6.3.1.1 Extending Into 2D

The same approach can be applied in a 2D (or higher-dimensional) landscape simply by scaling the dimensionality of the noise function. Instead of $\zeta_i(x)$, whose values are determined by a sequence of random values, we introduce $\zeta_i(x, y)$, whose values take the form of a grid. Interpolation is performed over two dimensions.

$$\zeta(x, y) = \sum_{i=0}^{k} r^k \zeta_k(x, y) \tag{6.2}$$

An example is shown in Figure 6.2, over $x \in [0,4], y \in [0,2]$. Lighter cells denote higher amplitude values.

In an ecological context, these amplitude values could represent the quantity of a given resource available in a cell, where a greater amplitude means a greater fitness payoff to an organism inhabiting that cell. Alternatively, they could represent the balance between an A/B duo of resources, in which lighter cells are more abundant in A. We will explicitly link these ideas to our model shortly.

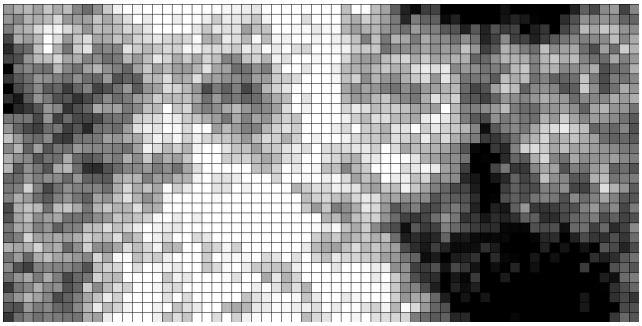

FIGURE 6.2: Landscape construction An example of a 2D landscape ($F = 0.5$, $G = 0.5$).

To introduce our landscape metrics, we extend by Equation 6.2 in two ways: by scaling the positional (x, y) parameters by F; and by wrapping it in a sigmoid *tanh* function, multiplied by G, to accentuate or lessen the gradient between adjacent cells. The resulting generator function is given in Equation 6.3.

$$\zeta(x,y) = \tanh\left(G\sum_{i=0}^{k} r^k \zeta_k(Fx, Fy)\right) \tag{6.3}$$

F is used to scale the generated landscapes to provide the desired granularity. A large F value is effectively the same as zooming out on the landscape, resulting in many smaller resource clusters.

G is used to adjust the step sharpness of the sigmoid function, serving to increase or decrease the differential between adjacent values. With a G value much greater than 1, $\zeta_(x,y)$ tends towards becoming a step function and gradients in the landscape are eradicated.

A low-F, high-G environment has no discontinuities and large areas of homogeneous resources, similar to a completely uniform environment. Conversely, in a high-F, low-G environment, neighbouring cells will exhibit almost no correlation, very close to a uniformly random distribution. The structured heterogeneity produced by this construction method is therefore bounded by completely uniform and completely random landscapes. See Figure 6.3 for example landscapes generated by this algorithm.

6.3.1.2 Measuring Landscape Properties with Moran's Index

To verify that the Perlin landscapes produced by this method reflect qualities typical of the landscape ecology literature, we subject them to Moran's I value, a measure of spatial

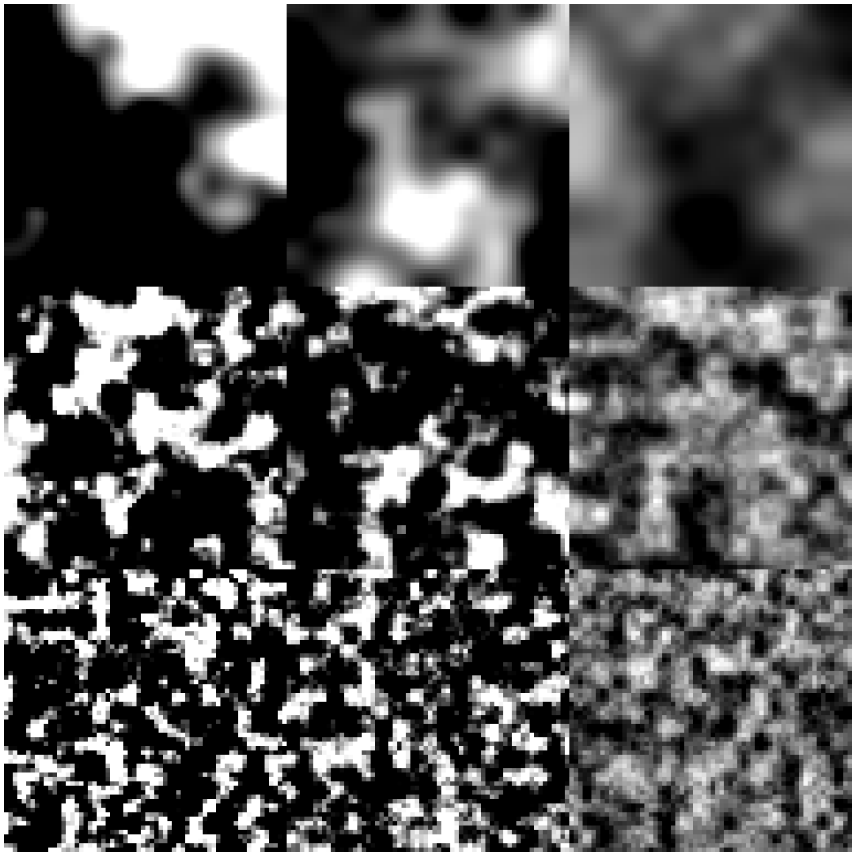

FIGURE 6.3: Landscape construction: Variable values. F increases from top to bottom, and G decreases from left to right.

autocorrelation. It varies from -1 to +1, where +1 denotes a perfect correlation between adjacent, and -1 denotes perfect negative autocorrelation (that is, all pairwise neighbouring cells have values each equal to the negation of the other). A value of 0 is expected when cells have uniformly random values.

Based on the definitions of our F and G values, we would expect that a higher F parameter should result in a lower mean of I, reflecting the lower autocorrelation between adjacent locations as the granularity of the landscape increases. Conversely, G should have no overall effect on the mean I value; despite the fact that a lower gradient value introduces sharp discontinuities at the borderline between patches, this should be counterbalanced by the increased homogeneity within a resource patch.

Applied to our spatial matrix, Moran's I is defined as below

$$I = \frac{N}{\sum_i \sum_j w_{ij}} \frac{\sum_i \sum_j w_{ij}(X_i - \bar{X})(X_j - \bar{X})}{\sum_i (X_i - \bar{X})^2} \tag{6.4}$$

We treat each pairwise case as equally weighted, with $\bar{X} = 0$, thus this becomes simply:

$$I = \frac{\sum_i \sum_j X_i X_j}{\sum_i X_i^2} \tag{6.5}$$

Figure 6.4 shows the changing Moran's I indices for a number of landscapes over a range of F and G indices. This confirms our predicted evaluations, from which we can conclude that the F and G matrix are having the intended consequences.

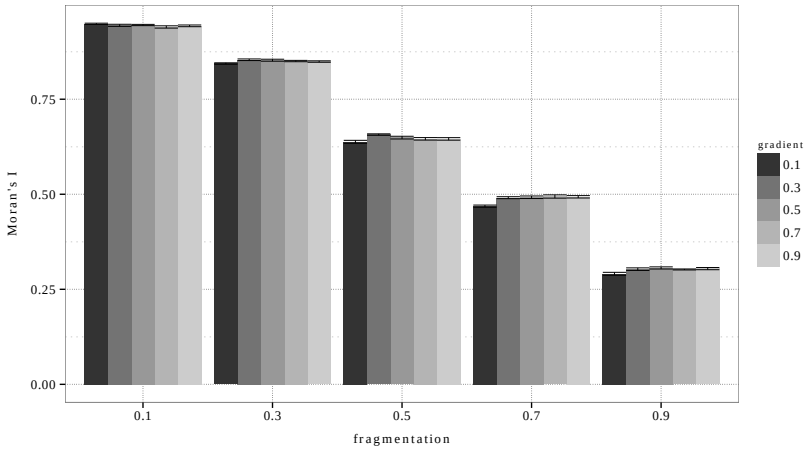

FIGURE 6.4: Moran's I for varying fragmentation (detail) and gradient measures

6.3.1.3 Integration with Baseline Model

Integration of this Perlin landscape construction method with Model ∅ is straightforward. On a 2D grid, we generate a landscape and use its intensity value to multiply the payoff of the fitness function as first described in 3.1 (with $H(p, E)$ denoting the Hamming distance between a phenotype and the environmental task). The fitness payoff for a given position $\{x, y\}$ becomes:

$$\phi_{x,y} = \zeta_{x,y} \left(1 - \frac{H(p, E)}{B}\right)^{\alpha^{-1}} \tag{6.6}$$

This introduces a structured heterogeneity to the fitness landscape which allows us to construct and interrogate the effects of different types of spatial variance.

6.3.1.4 Landscape Extremas: Uniform and random

The structured landscapes generated by this method have a range of Moran's I values which describe their heterogeneity. At either end of the I scale, we have two exceptional types of landscape which will be used to demonstrate the extreme cases.

Uniform landscape is the wholly homogeneous case used in previous chapters: each cell has an identical payoff of 1. This has a perfect autocorrelation of $I = 1.0$, meaning that every pair of neighbouring cells will be precisely the same.

Random landscape has a uniformly random payoff value per cell. It is neither correlated nor anticorrelated, with a mean autocorrelation of $I = 0.0$. There is no correlation between neighbouring cells.

6.3.2 Location and Movement

6.3.2.1 Stimulus Enhancement vs Local Enhancement

In a real-world environment, a learning organism typically has mobility, and the ability to navigate and relocate based on its drives. For most animals, a significant proportion of the daily time and energy budget is spent foraging for nutrition sources (Verner, 1965), searching out new spatial locations to improve fitness.

Social learning occurs frequently in foraging and other spatial exploration tasks (Heyes and Galef Jr, 1996). The socially-mediated learning activities performed by animals in this context may be divided into two forms in a spatial context.

In *stimulus enhancement*, an individual observes a demonstrator's interactions with a physical object and uses them to form a generalisation about other objects of that type. This may be used to ascertain whether certain classes of food are edible, modifying subsequent feeding patterns (Galef, 1976).

In *local enhancement*, an individual forms an association with a specific spatial location, learning tasks such as foraging that they associate with this particular space. This creates a useful distinction between different types of learning, which are exercised when an individual observes and mimics the spatial location of others.

To explore spatially-located properties of social learning, we will now add mobility to our population, with tendencies for movement and social cohesion.

We introduce a series of new properties of each agent.

- L – **location** $(x, y), \{x, y : \mathbb{N} | 0 \leq x < W \wedge 0 \leq y < W\}$: a pair of integers which locates the agent in 2D space
- m_{rate} – **movement rate** $[0, 1]$: the probability of movement in each timestep. A m_{rate} of 1 means that the agent moves every timestep.
- m_{coh} – **movement cohesion** $[-1, 1]$, the attraction or repulsion towards other agents. $m_{coh} = 1$ means that the agent moves towards groups of peers as often as possible, using roulette-wheel selection to weight in favour of moving towards peers; $m_{coh} = -1$ means that the weighting is inverted, so the agent moves away from groups of peers.

Movement occurs when an agent moves to an adjacent cell based on its current occupancy. The m_{coh} gene controls whether the agent should be attracted or repelled by the presence of neighbouring agents. We thus need an operator which allows this weighted movement to occur based on a linear value ranging from $[-1, 1]$.

We define the **weighted roulette-wheel selection** operator, with the requirement that:

- $m_{coh} = 1$: Selection operates according to standard roulette-wheel weighting, wherein an empty cell has a weight of 1, and a cell with population N has weighting $1 + N$. In this way, an empty cell can still be moved towards, but with 50% of the probability of a cell with occupancy 1.
- $m_{coh} = 0$: Selection operates uniformly randomly; any cell is as likely to be moved towards as any other.
- $m_{coh} = -1$: Selection operates according to inverse roulette-wheel weighting; a cell with occupancy 1 is 50% less likely to be moved towards than one of occupancy 0.

Mathematically, this function is defined as follows, with weight w:

$$\begin{aligned} w >= 0 : p(v_i) &= 1 + (wv_i) \\ w < 0 : p(v_i) &= 1/(1 + wv_i) \end{aligned} \tag{6.7}$$

Continuous boundary conditions are adopted; an agent moving beyond the right-hand side of the environment wraps onto the left. This effectively creates a toroidal structure.

6.4 Results

6.4.1 Uniform, Random and Structured Environments

6.4.1.1 Genetic convergence is slower in spatially complex environments

To understand the general impact of environmental heterogeneity, we repeat experiment 4.3.1.1: evolution to equilibrium in a static environment. This experiment is performed across three different types of environment: those that are **uniform**, **random** and **structured**, the latter with $F = 0.5$ and $G = 0.5$.

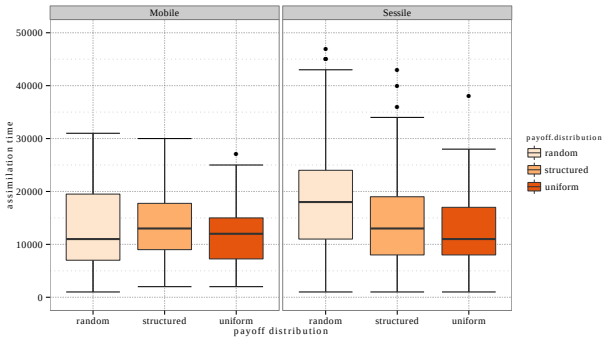

Figure 6.5: Genetic convergence is slower in more heterogeneous environments. With movement, genetic convergence is faster in spatially complex environments.

Parameter values (**6.4.1.1**)	N	B	α	μ	p_{switch}	p_{noise}	*trials*
	64	16	0.01	0.05	0.00	0.25	50

The trajectory of the simulation is similar, with genetic assimilation taking place across the population. However, the rate at which this occurs varies between different environment types, with innate behaviour taking longer to exceed social learning in more heterogeneous environments.

Figure 6.5 shows the generation at which this crossover takes place. Genetic assimilation takes place approximately 35% faster in a uniform environment, taking place after 13×10^3 generations, compared to 19×10^3 generations in a structured environment, and 20×10^3 generations in a random environment.

The source of this disparity is that, in spatially uncorrelated environments, a parent is more likely to experience a different set of selection pressures than its children. In

terms of genetic evolution, this is experienced as environmental noise; an individual that is perfectly adapted in a particular area of the landscape may produce identical offspring, which develop in an adjacent area with different adaptive requirements. This slows the process of convergence to peak fitness in spatially-complex surroundings.

The results (Figure 6.5) show an interesting difference when movement is present. When a population becomes mobile, the effect of landscape heterogeneity on assimilation times is reversed: convergence to optimal fitness occurs faster within a spatially complex environment than in one that is spatially uniform.

In a more spatially heterogeneous context, movement allows an individual to rapidly locate areas to which they are well adapted. This is a quick fix to improve an individual's fitness, consequently increasing their reproductive rate and ultimately speeding up the process of evolutionary convergence.

In a uniform environment, there is no adaptive difference between neighbouring cells. Movement, therefore, has no impact on evolutionary rate, as we would expect.

6.4.1.2 Social learning is more evident within complex environments and sessile populations

We have seen that evolutionary rates are markedly more rapid in environments that are more spatially homogeneous, and those in which the population is mobile and thus able to locate regions with more favourable selection pressures.

The distinction between learning trends at equilibrium is less significant (Figure 6.6). In a mobile population, no significant difference can be see between equilibrium learning modes. This reflects the fact that mobility enables a population to buffer against the effects of spatial uncertainty.

In a sessile population, changing spatial location is not possible and a small but statistically significant (95% CI) behavioural differential can be seen. Individual learning remains at a low level. Social learning is more evident across all environments, and is significantly more prevalent in those that have a random resource distribution.

This reflects the utility of specialised, local knowledge within the population. In a uniform environment – and, to a lesser extent, one that is structured – an individual's offspring is likely to experience a similar set of selective pressures. In an environment with low correlation between neighbouring regions, the small spatial dispersion of an individual's offspring means that they will experience a moderately different environment, reducing their innate fitness. By relying on the shared knowledge of individuals sharing that environment, however, fitness can be increased via phenotypic plasticity.

This is the spatially-analogous case to the uncertainty induced by temporal fluctuations, as described in §4.3.2.

6.4.1.3 Genetic shielding is most prevalent in sessile populations inhabiting random environments

The genetic and phenotypic fitnesses of the populations in these experiments is shown in Figure 6.7. Both innate and phenotypic fitness are significantly lower within a mobile population. Although convergence takes place more rapidly within these populations (Figure 6.5), the equilibrium state is inferior to those of a sessile population. This is because

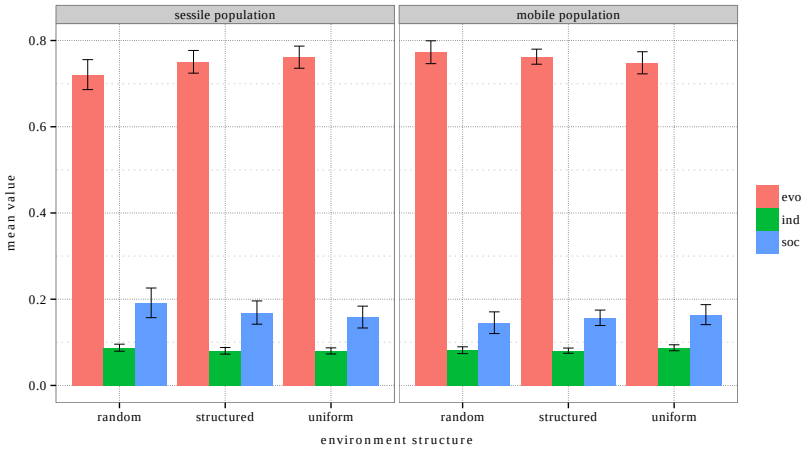

FIGURE 6.6: Social learning is more evident within complex environments and sessile populations

the concept of movement encoded in this model can often lead to an individual moving away from a region in which it is adaptively successful.

As described previously, no significant difference is found between genetic and phenotypic fitness in a mobile population, reflecting the fact that these populations rely less on phenotypic plasticity to improve their functional adaptation.

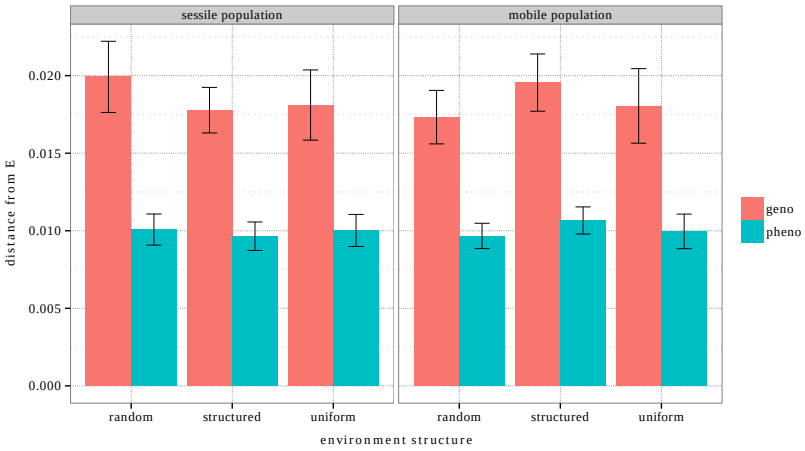

FIGURE 6.7: Genetic shielding is most prevalent in sessile populations inhabiting random environments

Sessile populations demonstrate a higher level of innate and phenotypic fitness, with the highest mean fitness in uniform environments, which make for a predictable adaptive context. A notable distinction can be seen in random environments, in which the population has an inferior level of innate fitness with a higher level of variance. We can infer from this that the population in these environments are engaged in a greater level of phenotypic plasticity which has a significant impact on their lifetime fitness.

6.4.1.4 In an unstable environment, phenotypic plasticity is particularly pronounced in sessile populations

We have seen that static and spatially complex environments foster a greater degree of phenotypic plasticity mediated by social learning, an adaptive advantage that can be equivalently countered (though at some fitness cost) by spatial mobility.

Real-world environments are not just spatially complex but exhibit change over time. This experiment re-introduces fluctuation to the environment, using the same stochastic change process described in §4.3.2 ($p_{switch} = 0.01$).

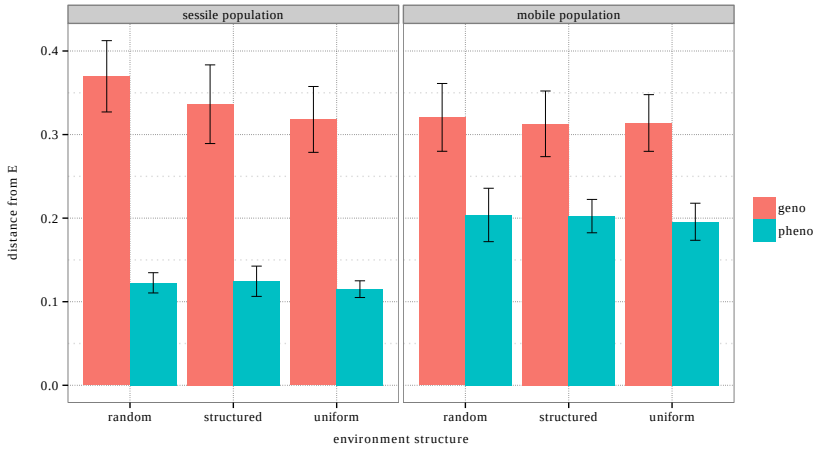

FIGURE 6.8: Genetic and phenotypic fitness in unstable environments

Parameter values (**6.4.1.4**)	N	B	α	μ	p_{switch}	p_{noise}	$trials$
	64	16	0.01	0.05	0.001	0.25	30

The results (Figure 6.8) show a disparity between mobile and sessile populations. Both show a significantly lower of mean genetic fitness than in the stable case (Figure 6.7), with a higher level of variance. However, the phenotypic fitness levels are only marginally lower in these unstable environments. Behavioural plasticity is applied to raise an individual's fitness within its lifetime.

This distinction is particularly evident within the sessile group, which have a greater genotype-phenotype differential, amplified moderately further within a random environment.

6.4.1.5 In a heterogeneous composite environment, individual learning is beneficial, leading genotypic evolution

This experiment introduces a composite environment, characterised by a pair of environmental tasks $T = 2$. Here, an individual has a pair of genotypes and phenotypes, evolved as before, which are exposed to both environmental tasks simultaneously. Fitness is obtained by summing the phenotypic fitness of both phenotypes.

To explore the impacts of different learning modes when multiple environmental tasks are present, we perform a number of experiments in which particular learning modes are suppressed: one series in which both b_{ind} and b_{soc} are present; one series in which each

of the learning modes is available in isolation; and one in which neither learning mode is available, restricting the population to only develop via genetic evolution. The results of these experiments are shown in Figure 6.9.

With this landscape configuration, we see some novel dynamics arising in the interaction between learning and evolution. The first notable property is that individual learning plays an unusually prominent role, improving phenotypic fitness in both random and uniform environments.

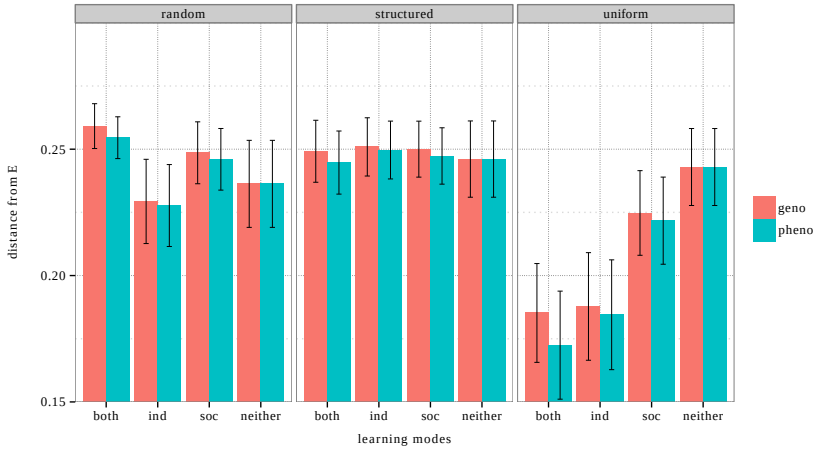

FIGURE 6.9: Spatial heterogeneity: Genetic and phenotypic fitness, restricted learning modes

Parameter values (**6.4.1.5**)	N	B	α	μ	p_{switch}	p_{noise}	$trials$
	64	16	0.01	0.05	0.00	0.25	100

A particularly interesting and unexpected property arises in uniform environments. As per previous experiments, genotypic and phenotypic fitnesses are higher in a uniform context than those which are more spatially heterogeneous, due to the lower level of uncertainty exemplified within this setting.

When individual learning is introduced in a uniform environment, fitness improves. The novel characteristic is that genetic fitness *also* improves when learning is introduced. This is a behaviour that has not arisen in any previous context; typically, learning will shield the emergence of innate behaviours, actively reducing genotypic adaptation. However, in this case, genetic fitness is lowest when neither learning capacity is available. Nor is this quite like the behaviour of genetic "guiding" described by Mayley (1997) and Maynard Smith (1987), which improves the *rate* of genetic evolution but can still have negative impact on a population's ultimate fitness. Moreover, we did not see this behaviour emerge with a single environmental task, meaning that it is a propensity specific to multi-task environments, and it likewise does not arise in spatially heterogeneous landscapes.

The cause behind the behaviour lies in the fact that two different sources of fitness are simultaneously available to the population at any point in the spatial environment. An individual can increase its fitness by improving its task adaptation towards E_1 or E_2. The effect is that an individual's behaviour diverges to specialise in either one of these tasks,

using individual learning to make the rapid transition towards fitness in E_1 or E_2 and to subsequently maintain its fitness within this task.

6.4.2 Landscape Metrics

6.4.2.1 Evolutionary convergence is delayed when gradient is low and fragmentation is high

The following results explore the impact of changing the landscape metrics which characterise a structured environment: fragmentation F and gradient G. As an environment is quantified by each of these properties independently, we will vary each parameter separately to examine its effects on learning and evolution behaviours.

The effects of altering the fragmentation and gradient of a heterogeneous environment are shown in Figure 6.10, with the Y-axis depicting the time at which genetic assimilation takes place (as per the definition in the previous section). The effects of heterogeneity are surprising. Convergence takes considerably longer in environments which are discontinuous (low value of G), an effect that is amplified by the degree of patchiness in the environment (ANOVA, $p < 0.05$). In clumped environments (low F), a population on a discontinuous landscape takes 150% of the time to converge than on a smooth landscape. In fragmented environments, a discontinuous landscape leads to convergence times that are 2.5x that of smooth landscapes.

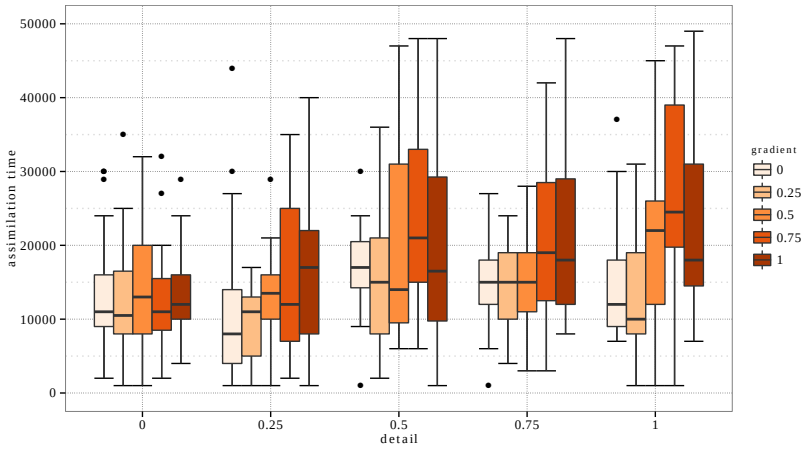

FIGURE 6.10: Convergence is slowest with high values of F and high values of G

Parameter values (**6.4.2.1**)	N	B	α	μ	p_{switch}	p_{noise}	$trials$
	64	16	0.01	0.05	0.00	0.25	50

6.5 SUMMARY

This chapter has investigated the dynamics of a learning and evolving population in environments which exhibit a range of different kinds of environmental heterogeneity. Environment types have featured those which are uniform, displaying no change across

their spatial extent; those which are random, in which adjacent pairs of cells have no correlation in their resource type or fitness payoff; and those which are structured, with clumps of resources separated by discontinuities of varying smoothness. We described a novel method of constructing a discrete two-dimensional landscape which can be characterised by two metrics drawn from the landscape ecology literature (§6.3.1): *fragmentation*, which describes whether the environment contains a large number of small resource patches or a small number of large patches; and *gradient*, determining the level of discontinuity between adjacent cells. We also introduced movement capabilities to our population (§6.3.2), with an individual's innate movement tendencies subject to evolution.

We have used this framework to generate a number of novel results.

In spatially complex environments, genetic assimilation takes substantially longer to complete due to the noise introduced by spatial uncertainty (§6.4.1.1). Introducing movement, however, causes the opposite result, with evolution operating more quickly over a mobile population in a complex environment, due to the population's ability to locate locations with more favourable selection pressures. Complex environments also foster a higher level of social learning, particularly within sessile populations (§6.4.1.2), whose fixed locations give rise to a greater demand for behavioural plasticity (§6.4.1.3). Introducing instability to a spatially heterogeneous environment, we saw that the tendency towards phenotypic plasticity in sessile populations became pronounced yet further (§6.4.1.4).

In composite heterogeneous environments, with more that one environmental task to attend to, individual learning plays a more prominent role, serving to increase both phenotypic and – unusually – genetic fitness (§6.4.1.5).

In structured environments, characterised by gradient and fragmentation values, we saw that evolutionary convergence is delayed significantly in gradients that are detailed and discontinuous; that is, low G and high F (§6.4.2.1).

Chapter 7: Conclusions

7.1 Overview

This section begins by summarising the key results obtained within this dissertation (§7.2), reviewing each chapter's results in context, and summarising the major research contributions (§7.3). We critically evaluate the limitations of the research and threats to validity posed towards it, and discuss future work that goes beyond that presented here (§7.4). We conclude with some closing remarks (§7.5).

7.2 Key Results

This dissertation has demonstrated the impact of different classes of environmental heterogeneity on the optimal modes of information acquisition that we should expect to see arise within a population. We defined four types of environmental factors under investigation: variability; task complexity; population structure; and spatial heterogeneity.

We described a novel individual-based model, representing a population of polygenic individuals that are able to engage in innate behaviour, individual learning, and social learning. We proceeded to introduce incrementally more detailed forms of environmental structure to understand how each of these four environmental factors affects the equilibrium behavioural modes that are likely to be produced by natural selection.

Predictions and comparisons with existing results are included within discussion of each factor below.

Baseline Results

Beginning from well-mixed populations in a static population, we proceeded to examine the baseline set of interactions generated by the simplest iteration of this model. In a stable environment, we saw that two distinct phases of behaviour take place. Initially, social learning dominates, as individuals adopt traits in a relatively unfamiliar environment by imitating those around them. This creates the breathing space needed to accumulate these traits innately, which brings about the genetic assimilation of the traits in question, allowing the population to exercise the less-costly facility of innate behaviour. This entails a transient period of genotypic "shielding", in which behavioural plasticity slows down evolutionary development by enabling individuals to adapt to an environment within their lifetime.

Environmental Factor 1: Variability

Environmental variability was introduced in the form of small, stochastic changes to the shared fitness function. We saw that this gradual type of environmental change leads to a different equilibrium set of states; social learning prevails at a moderate rate of change, and individual learning in fast-changing regimes. We diagrammed the complete relationship between task complexity and rates of change, showing that innate behaviour also dominates in very simple tasks. Social learning fills the middle ground, in environments which are moderately fast-changing *or* moderately complex.

These results corroborate those of Wakano and Aoki (2006), Boyd and Richerson (1985) et al, who predict the same three-phase regime of learning modes, offering support for this model. They also offer a new constraint on these dynamics, and a prediction for empirical studies: that even at moderate rates of change, social learning is only likely to arise as a dominant strategy when the environment's demands are of moderate complexity.

In contrast to gradual environmental change is the phenomenon of a rapid environmental shift (Lande, 2009; Hallsson and Björklund, 2012), modelled here as a complete change to the environment's fitness function. This was shown to bring about *three* phases of behavioural dominance. Firstly, individual learning proliferates, as individuals scramble to acquire the new traits required in this unfamiliar circumstance (phase I). These traits then circulate socially (phase II), before being assimilated into the genotype (phase III). With cultural assimilation followed by genetic assimilation, we described this as a "double Baldwin effect".

This temporary rise in plasticity supports Lande (2009)'s quantitative genetics model, but is extended with the addition of social transmission. Our results make a further prediction, namely that the social assimilation and subsequent genetic assimilation will operate on distinct timescales, with overlap between each: social assimilation takes place rapidly, with a fast decline in individual learning even before social learning has peaked. The process of genetic assimilation then takes place substantially more slowly.

In our introduction, we discussed results by Feldman et al. (1996) and Kendal et al. (2009) who show that social learning outcompetes asocial learners when environmental variance is low. Our results support these predictions, albeit with a continued, persistent low level of individual learning. This is necessary to support the continued innovation of traits in the population, in order to more rapidly acquire a trait after a fluctuation takes place, which can subsequently be disseminated.

Our results partly support the empirical work of Hallsson and Björklund (2012), in that phenotypic plasticity and genetic variance increase after environmental fluctuations are introduced and subsequently decreased. This rise in genetic variance is induced by the lower dependence on innate behaviour, particularly in the parts of the phenotype in which successful traits have been phenotypically assimilated via social learning. However, they do not support the observation that, after selection on these plastic individuals, subsequent plasticity is lower. This is likely due to the limitation of the simplicity of the model used, omitting the mechanisms that dictate plasticity in different areas of behaviour. An extension to this work could use a connectionist model to create a more sophisticated model of adaptive learning (similar to Nolfi et al. (1994)) which may result more nuanced plastic behaviours.

Environmental Factor 2: Task Complexity

We explored the consequences of changing the level of complexity of the task that an individual needs to fulfil to reach optimal fitness in its environment. The low-complexity case, where $B = 1$, is equivalent to the single-locus quantitative genetics models used in prior work (Ancel, 2000). It is interesting to note that, at this minimum level of task complexity, individual learning never appears as the dominant mode of behaviour; social learning dominates, even at high rates of environmental fluctuation. This is because, in well-mixed populations, there is virtually always a small number of individuals who possess the correct trait (either 1 or 0) and so a stronger exemplar can always be found and reliably mimicked.

Beyond a certain threshold level of task complexity, a phase transition to a population of individual learners comes about. A mathematical analysis revealed that this was due to the task's learning time exceeding the lifetime of an individual, giving rise to a virtually paradoxical state we described as a "slippery slope" scenario: the population becomes wholly reliant on individual learning to try to attain an unattainable goal, thereby causing any type of instinctual behaviour to have deleterious effects.

This result has interesting implications in an evolutionary context. It suggests that an extreme pressure on phenotypic plasticity may correspondingly result in an absolute form of extreme genetic drift, as genetic pressure falls to zero.

In our introduction, we discussed work by Dukas and Visscher (1994) and Marler (1970) who indicate that complex tasks may require lifetime learning to achieve. Our results are theoretical support that, for tasks beyond a given complexity, we would indeed expect to see a strong tendency towards individual learning. Moreover, they predict that innate disposition should fall away for these tasks.

Environmental Factor 3: Population Structure

This question is tackled in Chapter 5, by introducing two different forms of discrete structure to a population. Initially, the population is sited on one- and two-dimensional lattice structures similar to a cellular automata. Latterly, it is modelled as a k-regular graph, with an individual interacting with a neighbourhood of size k.

We encountered a surprising result. Unlike the well-mixed and 2D-structured population, the 1D population never reaches the stage of genetic assimilation, thus not demonstrating the Baldwin effect. We determined that this was because the rigid population structure prevents the refuge-like breathing space required for innate behaviour to emerge, thus locking in the population to a maladaptive behaviour. We suggested that the more rigid population structure locks a population in to the earlier transients of the three-phase learning process described in the double Baldwin effect.

However, further experiments showed that relaxing the model's representation of reproduction – in which offspring always replace a neighbour of their parent – eradicates this behaviour. It is unlikely that such a constraint would be imposed within a real-world model, so the predictions generated may be attributed to a modelling artefact.

Switching to k-regular graphs for finer control over neighbourhood structure, we saw that, in general, larger neighbourhood size accelerates convergence to equilibrium, and results in a better-adapted population.

In our introduction, we discussed work from evolutionary game theory that suggests that social structure can give rise to altruistic behaviour. Due to the form of our model, altruism cannot truly arise as a strategy, although we did demonstrate a substantial increase in the success of social learning as a dominant mode of learning.

Environmental Factor 4: Spatial Heterogeneity

The final chapter of results concerns environments which are spatially explicit and demonstrate more complex spatial patterning. We began by describing a novel method of constructing spatially heterogeneous environments that can be characterised by two qualities drawn from the landscape ecology literature. *Fragmentation* describes whether an environment is comprised of a large number of small areas of resources, or a small number of large areas. *Gradient* describes whether adjacent areas of resources are joined by smooth gradients or sudden discontinuities. These "structured" environments are supplemented by two extreme types of spatial environment: uniform environments, which are spatially homogeneous over their extent; and random environments, in which pairwise regions have no correlation.

This spectrum of heterogeneity – from random to uniform – was used to respond to the above question, via a series of experiments in which populations are sited on environments exhibiting the properties under investigation.

We first saw that evolution occurs more slowly in spatially-complex environments, due to the inconsistency in selective pressures between the environments of a parent and its offspring. However, this result was then shown not to hold when mobility is introduced to a population; instead, convergence occurs *faster* when a mobile population is within a complex environment.

Social learning was shown to succeed more widely within spatially complex environments, particularly when inhabited by sessile (static) populations. We observed that this spatial discontinuity gives rise to uncertainty analogous to the temporal change described in previous sections. We saw that genetic shielding is particularly prevalent in sessile populations in complex environments, in which innate behaviour recedes and is replaced by phenotypic plasticity; likewise in faster-changing environments. Mobile populations, conversely, exhibit substantially less phenotypic plasticity, instead able to move to locations with more favourable selection pressures.

In composite environments, with multiple environmental tasks to address, we showed that individual learning is beneficial. Surprisingly, individual learning can act to lead the population to a more successful genotypic state – a case of learning not just accelerating evolution, but increasing a population's mean fitness.

We finally turned to our landscape metrics, exploring the ways in which fragmented and graded environmental qualities impact upon learning and evolution. We saw that environments with high fragmentation and low gradient give rise to significant delays in convergence rates, with the low level of spatial autocorrelation adding noise to the evolutionary process.

Our theoretical results support those of Baythavong (2011), who showed that phenotypic plasticity is preferred in heterogeneous environments that are fine-grained (equivalent to our "fragmentation"). Baythavong's work focused on sessile organisms; our models extend

this theory to offer the new prediction that this conclusion would not hold as strongly for mobile organisms, in which we would instead predict a *lower* level of phenotypic plasticity and a higher level of genetic specialism, with motion enabling the individual to target the area of the environment to which they are best adapted.

7.3 Research Contributions

This thesis includes a number of results and approaches that are new to the field.

In Chapter 4, we describe a novel minimal individual-based framework for modelling the interactions between evolution, individual learning and social learning in environments of arbitrary task complexity and rates of change. This in turn extends the approaches of Hinton and Nowlan (1987), Belew (1990), Best (1999) and others, with properties which allow for the investigation of continuously-changing learning trends, gene-culture co-evolution, and – as later demonstrated – the introduction of arbitrary environmental structures.

In Section 4.3, we used this model to make a substantial series of novel predictions. We demonstrated that social learning gives rise to a particularly potent, cultural form of the Baldwin effect, as predicted by Papineau (2005). We then proceeded to show that, after a major environmental perturbation, this mutates into a "double Baldwin effect", with a brief burst of individual learning followed by a longer phase of social learning, before genetic assimilation takes place, a testable prediction in the effects of major environmental change.

In Section 4.3.3.1 and the subsequent mathematical analysis (§4.3.3.2, we discover a vicious evolutionary circle that arises within environments characterised by a task beyond a certain, learnable level of complexity; a population can become collectively trapped in a suboptimal state, in which individual learning becomes dominant and innate behaviour becomes deleterious. To our knowledge, this situation of "absolute genetic shielding" – and the phase transition that it gives rise to – has not been encountered or discussed in the existing literature, though it is likely to emerge in other situations.

A further new result in this chapter (§4.3.4) is the demonstration of the interactions between task complexity and rates of change visualised in Figure 4.11. This suite of experiments corroborates the consensus view that evolution, individual and social learning succeed respectively in slow-, moderate- and fast-changing environments; it then expands these results by demonstrating that task complexity is a further vital characterising factor in the interaction of these dynamics. Innate behaviour should not only be expected in slow or static environments, but also those in which a task requires an exceptionally small number of subtraits; social learning should be expected in environments of a moderate rate of change *or* those of moderate complexity.

In Section 6.3, we take two of the dominant metrics from landscape ecology and describe a novel method of constructing artificial 2D landscapes that can be parametrised by this model, an approach that could be used in the context of other simulations predicting the impact of quantifiable landscape properties on behavioural traits.

A number of new predictions are also offered. One key conclusion is that, in spatially complex environments, phenotypic plasticity is likely to be particularly pronounced in sessile populations (reflecting the results of Baythavong (2011)), yet *reduced* in mobile

populations. Plasticity thus acts as a replacement for mobility, in which a mobile population can alter its selection pressures by relocating to a more favourable location.

7.4 Limitations and Future Work

The models described in this thesis are, by design, unrealistic. They are designed to be maximally parsimonious in their means and, by extension, applicable to a wide range of scenarios. As a consequence, in limiting the scope of the research, it has many omissions imposed by design. This section addresses these limitations and omissions, and identifies further work that could be performed.

Prior to this, a couple of specific discrepancies and shortfalls should be remarked upon.

The results shown in Section 5.4.2.1 (and by extension, Section 5.3.2.2) seems to demonstrate that small neighbourhood size results in a behavioural lock-in, in which the population remains trapped in a suboptimal behavioural mode. However, as described later in the section, this behaviour disappears when one of the model constraints is relaxed: the placement of offspring in neighbouring vertices. Treated empirically, this constraint would be as if an individual had a highly constrained number of offspring, which were competing for the same small set of resources. It is not clear that this is a plausible situation, and so the result should be treated as an artefact of the model selection.

A result we have failed to reproduce is that of Rainey and Travisano (1998), who demonstrated that a bacterial population within a heterogeneous environment would adaptively diverge to form multiple spatial compartments, genetically demarcated by the survival demands of their respective habitats. Even in heterogeneous environments, the population converged to a single lineage. This is almost certainly an effect of the model's fixed population size and birth/death reproduction structure, which could result in a pocket of survival being stochastically eradicated by a spatially distant competitor. Extending the model to allow for a variable population size with a limited resource structure would be an interesting development, creating the potential for multiple distinct evolutionary lineages corresponding to environmental heterogeneity.

Varying fluctuation and population collapse (Borg and Channon, 2012)

This model and part of the results included in Chapter 3 were published in Jones and Blackwell (2011), encompassing many of the results obtained in a well-mixed population. Some small differences are evidenced, most notably the inclusion of metabolic state in an individual agent's state, used to determine its reproductive rate. This was subsequently removed as the model was refined and simplified.

An extension of Jones and Blackwell (2011) was published by Borg and Channon (2012), who develop our basic model to draw new conclusions in environments which are increasingly variable. By introducing periodic levels of variability to the environment, they demonstrate that major evolutionary adaptations are particularly likely occur at times of higher variability. Further conclusions are made about the likelihood of population collapse. Although small departures are made from our model, including the incorporation of a population-wide learning rate, the fundaments of the model remain the same, including the bit-wise genotype/phenotype structure and the processes determining their development.

We believe that this simple but powerful and intuitive framework will continue to bear fruit in other such novel contexts.

Trait learnability

An implicit assumption of the model is that every genetically-encoded trait possessed by an individual can be modified by learning, meaning that the optimal phenotype can always be obtained by sampling distant parts of the fitness landscape. Such a powerful form of learning is less affected by the structure of the fitness landscape Sznajder et al. (2012), enabling the individual to cross fitness valleys and escape local maxima.

To address this limitation, an extension could follow the approach taken by Anderson (1995) or Ancel (2000) in bounding the level of learning that can take place. Exploring the interaction between bounded learning on more sophisticated fitness landscapes could reduce the predicted benefits of learning, as well as making further predictions relating to gene epistasis, which is disregarded by our assumption of a monotonically-increasing fitness function.

Moreover, the level of plasticity itself should ideally become an evolvable trait, perhaps as segregated into different areas of behaviour. It may be that the demands posed by foraging are likely to vary over time, or over different areas, requiring a greater level of plasticity than mating behaviours. In real-world genotypes, different faculties show different degrees of innateness vs learnability, even between species (Marler, 1970; Lorenz, 1971; Thorpe, 1956). An interesting development of this work would be to segregate the genotype and phenotype into different behavioural areas, with evolvable innate/learning traits for each. This would be a suitable platform to investigate the ways in which particular behavioural groups tend towards the innate or learned.

Trait evolvability

Similarly, this model makes the assumption that all traits can be evolved innately, regardless of their complexity. Sophisticated behaviours such as avian tool-use have been shown to have substantial innate foundations (Tebbich et al., 2001; Kenward et al., 2006) but appear to be relative naive in execution without practise and social exposure, suggesting that instinct only provides part of the picture.

Deferring the fine-tuning of such behaviours to learning may be an adaptive choice due to unpredictability survival challenges, granting a degree of generalism to allow instinct to be modified to novel scenarios; or it may be a bound on the precision of coordination that can be encoded innately. It would be an interesting complement to this work to explore the bounds of innateness, although these bounds are most likely connected with metabolic and biophysical costs, which would require a more mechanistic approach to modelling.

Explicit costs of learning

No constitutive cost of learning is applied within these models. Recent empirical work has demonstrated that explicit costs are imposed for the apparatus involved in learning Mery and Kawecki (2003), paid regardless of whether the ability is used. This would alter the preferential balance between innate and acquired behaviours. A fixed cost could

be introduced for learning, imposed as a reduction in fitness proportional to the level of phenotypic plasticity. Again, this would most likely be more illuminating in the context of a mechanistic model which reproduces the biophysical substrates of behaviour.

Fixed population size

Many recent models have looked at the effect of environmental fluctuation (Lundberg et al., 2000; Boyce et al., 2006) and spatial variation (Schreiber, 2010) on changes to population size, considering the interactions between internal processes (such as competition and sociality) and external processes of environmental variability. A simple extension of the model would be to relax the fixed population constraint, enabling the population size and structure to vary in response to environmental factors.

Borg and Channon (2012) have done some initial work in this area by developing an earlier version of the baseline model described in this thesis (Jones and Blackwell, 2011) with variable population structure to explore the effects of changing environmental variability, showing that increasingly variable environments select for social learning capacities, and demonstrating the risk of population collapse introduced by population conformity.

Simplistic fitness landscape

The fitness landscapes adopted throughout these models are unimodal. This implies that fitness increases monotonically as successful traits are acquired; there are no local maxima, and no modelling of complex inter-trait relationships. Actual genetic interactions are far much complex and interdependent, with combinations of alleles resulting in nonlinear and contradictory patterns of effects: epistasis, in which two or more genes contribute in different ways to one phenotype, with effects that are not simply additive; and pleiotropy, in which one gene has an impact on multiple phenotypes. Work by Borenstein et al. (2006) shows that rugged fitness landscapes can alter the dynamics of learning and evolution.

7.5 Concluding Remarks

We began this thesis by drawing up an operational definition of heterogeneity as the uncertainty experienced by an individual, both in time and space. By systematically exploring the learning dynamics in populations in both of these dimensions, we have demonstrated that each of our four aspects of environmental heterogeneity – variability, task complexity, population structure, and environmental heterogeneity – have substantial and measurable impacts on the optimal modes of information acquisition that we would expect to see arise from natural selection.

We have also demonstrated that behavioural plasticity is not simply a deficiency of evolution, or an omission to be corrected during an organism's lifetime. Plasticity – and the uncertainty that it engenders and embodies – is itself an active part of the evolutionary process, capable of optimising and leading evolution through unforeseeable terrains.

Appendix A: Further Model Results

The body of this thesis omits certain results which are not central to its argument, particularly those pertaining specifically to properties of the models in question (rather than the systems whose behaviours are under investigation). These include parameter sweeps and sensitivity analyses, and certain negative or uninteresting results.

This appendix summarises these peripheral results, referring back to the main text when particular iterations of the model are used.

It is the author's belief that, even in an abstract artificial life model, parameter values should not simply be plucked from thin air, but justified either by biological plausibility or by demonstrating that values within a particular range are necessary to exhibit certain effects. Moreover, it is vital to establish confidence in a model by showing that it is not particularly brittle to a magical set of parameter values. Equally, it should be shown that a parameter does have some impact within a model, and is not simply there for decoration or other voodoo purposes.

A.1 Model Parameter Sweep

A.1.1 Larger population size slows convergence time

Population size can have a substantial effect on the evolutionary dynamics of an ecosystem (Reed, 2005; Jain et al., 2011; Rozen et al., 2008). It can affect the likelihood that a genetic trajectory is "locked in" to a particular vector, and the chance of a beneficial mutation sweeping to dominance.

We can trivially explore different population sizes by altering our model's N value. Behavioural trends follow the same pattern as those in §4.3.1.1. Convergence rates, however, are much slower for large N (Figure A.1); on a log/log scale, assimilation time varies linearly with N.

Parameter values (**A.1.1.0**)	N	B	α	μ	p_{switch}	p_{noise}	*trials*
	-	16	0.01	0.05	0.00	0.25	8

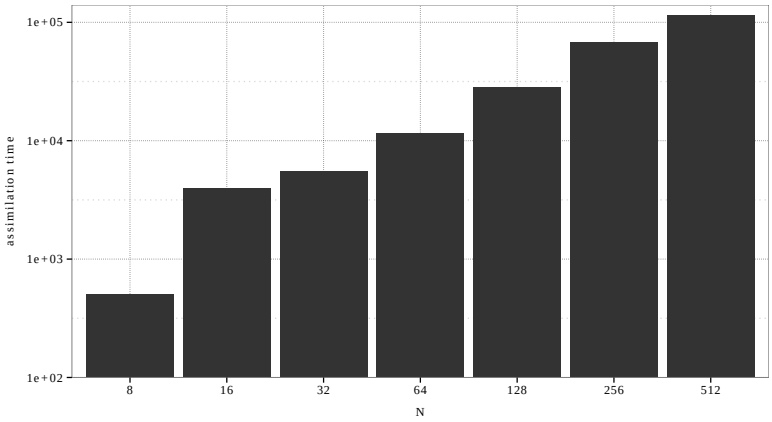

Figure A.1: Convergence times for varying N. Axes on logarithmic scale.

A.1.2 Larger α weakens evolutionary pressure

The α scaling value affects the environment's fitness curve. A smaller α creates a steeper falloff for agents with a greater Hamming distance from the environment's objective; a larger α means that the fitness differential decreases, lowering the selective pressure for agents to attain a precise match to the environment's demands.

Figure A.2 shows the effect of varying α values for a fixed population size in a static environment. For α of 0.05 or higher, the model follows a noisy trajectory and with less clear-cut equilibrium state. This is due to the fact that, when reproduction occurs, the roulette-wheel selection has a greater chance of selecting an agent with a poorer fit to the environment. Weaker agents are less heavily penalised, and genetic drift can occur.

Figure A.3 shows two different α values over a range of population complexities and rates of change. The boundaries between different equilibrium regimes are less clear-cut, with a series of noisy outcomes around $B = 32$. The overall results do not substantially differ, but a lower α value of 0.01 was selected to maximise the clarity between results.

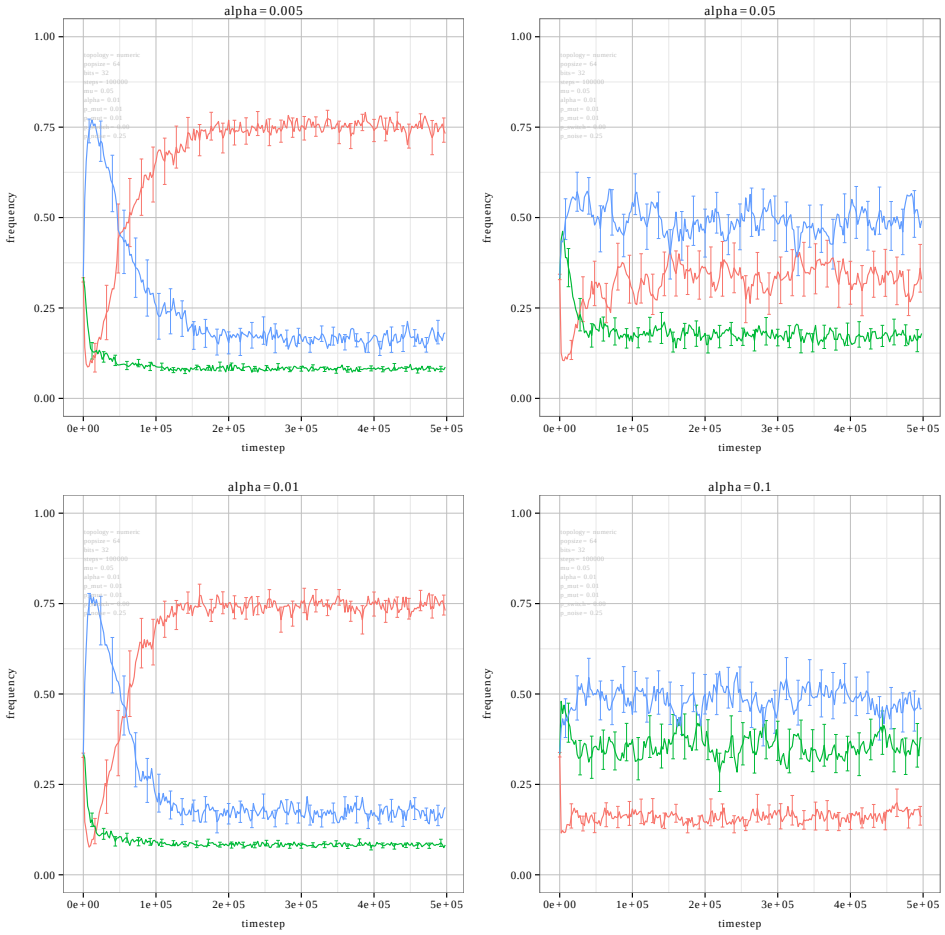

FIGURE A.2: Effects of varying values of α. A larger α decreases the fitness differential and introduces noise to the system's stable states.

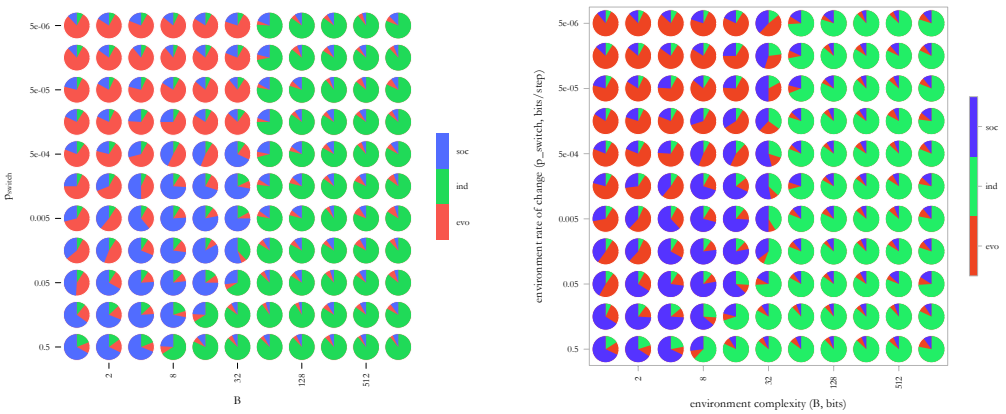

(a) $\alpha = 0.01$ (default)

(b) $\alpha = 0.05$

FIGURE A.3: Two values of α

A.1.3 Larger p_{noise} decreases the efficacy of social learning

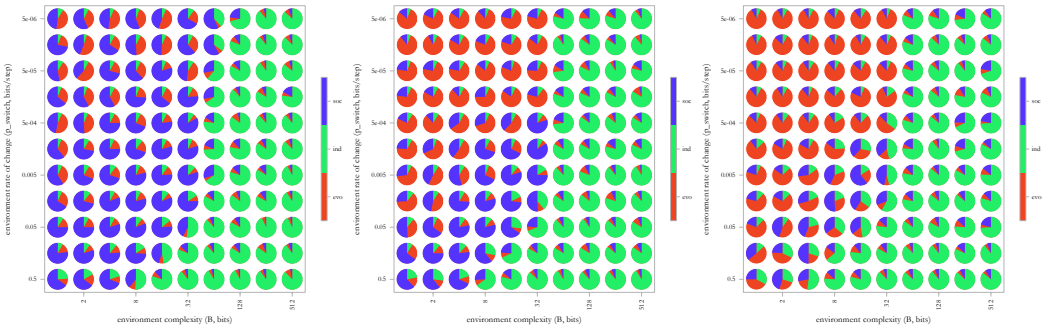

(a) $p_{noise} = 0.0$ (b) $p_{noise} = 0.25$ (c) $p_{noise} = 0.5$

FIGURE A.4: Varying p_{noise}. At very p_{noise} values, social learning operates accurately, with individuals able to copy from their fitter peers with high fidelity. Within this regime, p_{soc}e entirely supersedes innate behaviour. At $p_{noise} = 0.5$, a social learner is effectively mimicking values with a white noise distribution, and so are equally likely to attempt a deleterious action as one that is successful. Values of p_{noise} beyond 0.5 indicate that the individual copies from their peers, but is more likely to perform the inverse action (probability of flipping the operative bit is greater than 0.5). This could be viewed as functionality analogous to non-conformity or anti-conformity (Efferson et al., 2008), but has no direct correlate within our model as we are treating this as a noise term, not a behavioural tendency.

A.1.4 Larger μ increases the rate of genetic assimilation, adding genetic noise

μ is the rate at which behavioural trends mutate between generations. A greater μ value means that an agent's offspring may have a significantly larger tendency to engage in different kinds of learning.

For the Baldwin effect to occur, two steps must take place: firstly, a beneficial trait must arise in an individual through genetic mutation; and subsequently, its offspring must evolve a greater tendency to behave innately.

Thus, μ affects the rate at which genetic assimilation reaches fixation in the population. Figure A.5 shows the effects of different μ values. A very small value ($\mu = 0.005$) takes a very long time to reach fixation; even after 10^6 generations, equilibrium is not reached. Moreover, there is great variance between the point and rate at which fixation occurs, as indicated by the large standard error bars.

As the value is increased ($\mu = 0.01, \mu = 0.05$), fixation time decreases, as does the quantity of noise. However, after a certain threshold ($\mu = 0.1$), the inter-generational variance is so large that the equilibrium behaviours start to converge together.

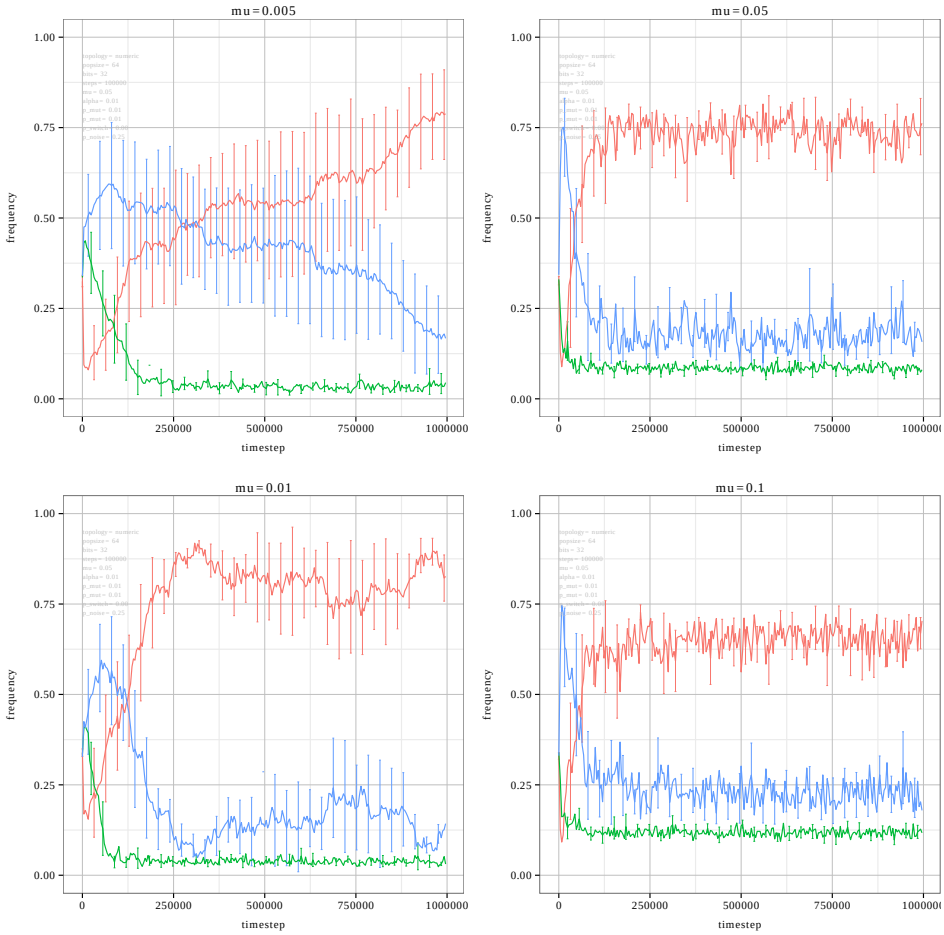

FIGURE A.5: Effects of varying values of μ. A larger μ increases the rate at which genetic assimilation sweeps to dominance within a population.

A.1.5 Larger p_{mut} increases rate of genetic assimilation, but can encourage social learning

p_{mut} determines the rate at which individual genes mutate between generations. A small value of p_{mut} decreases the rate at which genetic assimilation takes place, as it takes more generations for the beneficial trait mutation to occur (Figure A.6).

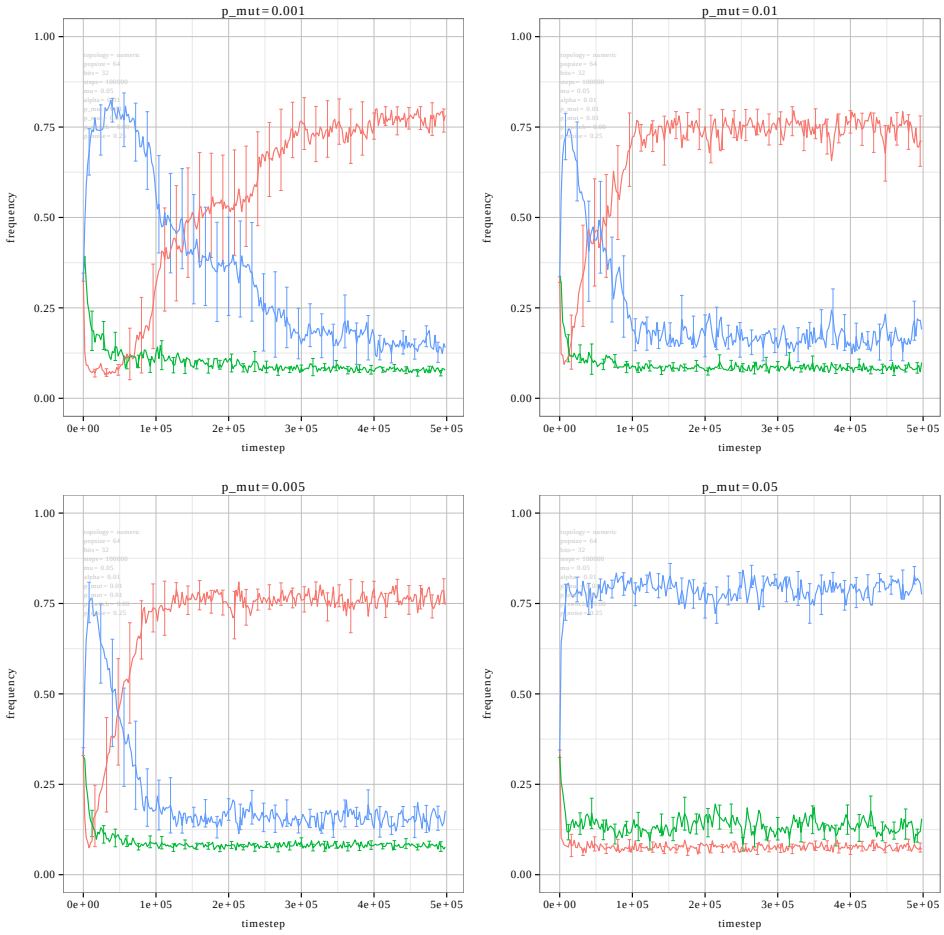

FIGURE A.6: Effects of varying values of p_{mut}. At larger values, acting innately becomes deleterious.

After p_{mut} passes beyond a certain threshold ($p_{mut} = 0.05$, Figure A.6), it becomes detrimental to act innately as there is such a high chance of possessing a deleterious gene through inter-generational mutations, thus social learning dominates.

A.2 Learning Strategies

We examine a number of discrete strategies, in which entire behavioural qualities are switched on and off.

A.2.1 Copy Fittest Neighbour *vs* Copy Random Neighbour: Fitness-weighted exemplar selection has little effect on dynamics

Laland (2004) discusses a number of learning strategies apparently adopted by real species, placing emphasis on a handful, notably *copy the majority* (that is, normative behaviour) and *copy if better* (that is, comparative learning). In the case of this model, we are interested in determining to what extent these particular strategies affect the success of social learning.

At present, agents use a *copy fittest* strategy; observing their fittest neighbour, weighted by roulette wheel selection. What happens if we disable this constraint and copy an arbitrary neighbour?

In the standard model: When engaged in social learning, an exemplar to copy from is selected using roulette-wheel selection, weighted towards fitter neighbours.

In *Copy Random Neighbour*, an exemplar is selected uniformly randomly. We would expect this to significantly weaken the power of social learning as it would increase the probability of copying an incorrect trait. However, the impact is insignificant (Figure A.7)

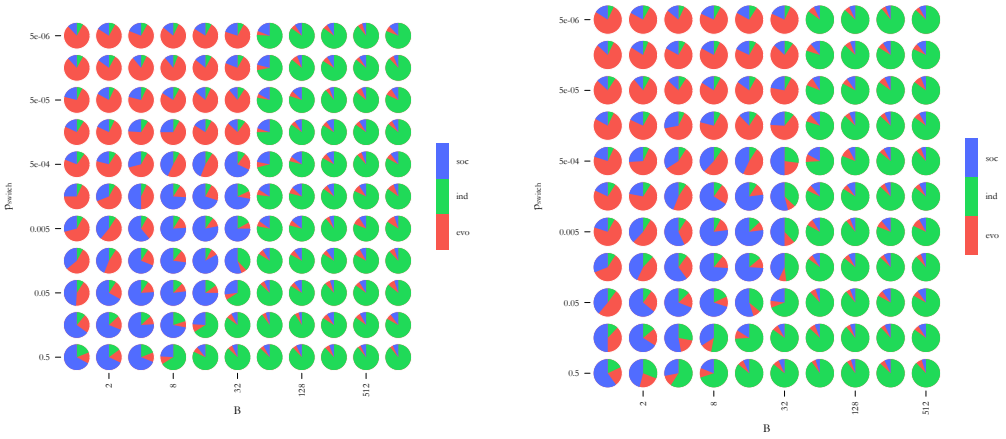

(a) Copy Fittest Neighbour (default)

(b) Copy Random Neighbour

FIGURE A.7: Copying a random neighbour, rather than weighting towards fitter peers, has an insignificant effect on model dynamics

A.2.2 Copy Random Trait *vs* Copy Novel Trait: Copying only novel traits significantly increases the advantage of social learning

In the standard model, social learning entails an agent observing a neighbour and mimicking a single trait selected uniformly randomly (that is to say, a single bit of their phenotypic bitstring).

It may be argued that real creatures are only likely to mimic behaviours that differ from their existing behaviours. We call this strategy *Copy Novel Trait*.

The results of adopting this strategy are shown in Figure A.8. It substantially increases the adoption of social learning, as it means that less learning trials are needed to discover successful new traits.

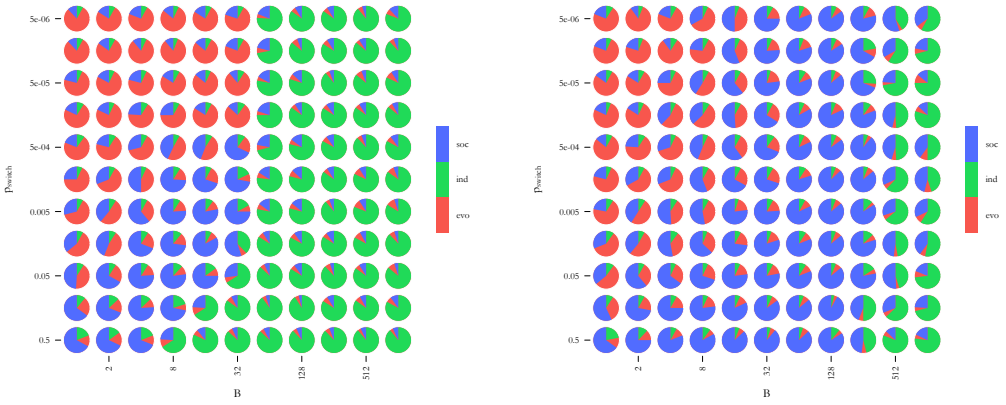

(a) Copy random trait (default)

(b) Copy novel traits only

FIGURE A.8: Adopting a strategy of only copying novel traits, rather than selecting a trait at random, causes social learning to dominate through a wider parameter space.

A.2.3 Assimilate If Advantageous *vs* Always Assimilate: Always Assimilate weakens individual and social learning

In the standard model: When engaging in a learning activity, either individual or social, an agent attempts to toggle a single bit of its phenotype and evaluates whether the fitness payoff would be higher than using its current base phenotype. If so, it incorporates the new bitstring to become its phenotype for subsequent timesteps.

With the strategy *Always Assimilate*, an agent toggles a bit permanently; that is, it always incorporates the newly-learned bit.

This mode of operation is severely detrimental to learning activities (Figure A.9). Individual learning is never engaged, as attempting to toggle an erroneous bit will be retained for the individual's entire lifetime. Social learning is now only seen in environments that are very fast-changing and simple, as the risk of incorporating an incorrect trait are outweighed by the fitness costs of remaining genetically outdated in a fast-paced environment.

A.2.4 Modelling learning modes with fixed, thoroughbred behaviours obstructs the Baldwin effect

Models of evolutionary games (Axelrod and Hamilton, 1981; Nowak, 2006; Rendell et al., 2010) typically give agents one discrete choice of action, which remains fixed throughout their lifetime: they exhibit strategy A *or* strategy B. The results given above are produced by a population whose learning modes are continuous, with a spectrum of behavioural tendencies that an agent can select from, proportionately to the tendency value $(b_{evo}, b_{ind}, b_{soc})$ at each timestep.

Here, we modify the model by allocating each agent a fixed and singular behavioural mode at the start of the trial, which remains fixed throughout its lifetime. This disallows the notion of mixed strategies; an agent can be *either* following its genotype *or* engaging in individual learning, but not both.

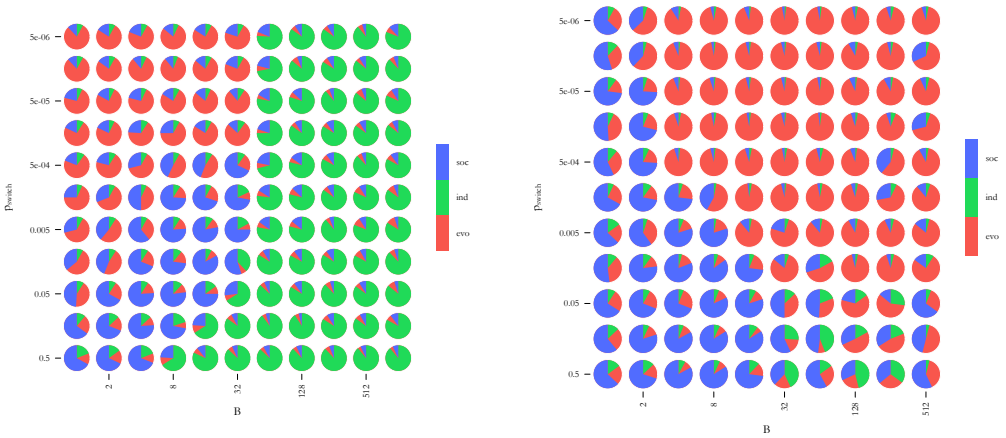

(a) Assimilate If Advantageous (default)

(b) Always Assimilate

FIGURE A.9: Adopting a strategy of always assimilating learned traits, rather than critically assimilating based on the trait's success, dramatically alters the behavioural landscape.

The results are shown in Figure A.10 (left), across the same lattice of parameter values as in Section 4.3.4; task complexity B is varied over the x-axis, and environmental rate of change p_{switch} over the y-axis.

Compared to the standard results (right), we see an distinct dominance of social learning over regimes wherein genetic assimilation would normally take place, displacing social learners with their less costly genetic counterparts. This is due to the fact that this new model obstructs the very behaviour that allows the Baldwin effect to take place; here, the initial fitness benefits of plasticity cause social learners to sweep to dominance, and eradicate any innate tendencies from the population. The Baldwin effect relies on the possibility of genotypic assimilation supplanting a population of learning agents. Yet, if innate behavioural capacity is eliminated from the population entirely, genetic assimilation can no longer take place.

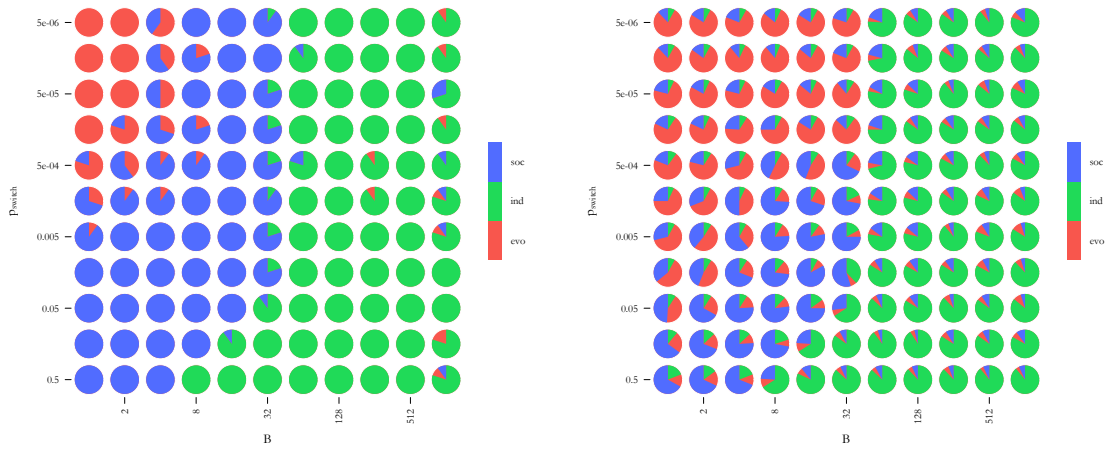

FIGURE A.10: Left: Learning modes in a thoroughbred population across values of p_{switch} and B. (Learning modes using normal behavioural traits are shown on the right for comparison.) Social learning appears to be a much more successful strategy in this context (cf. Figure 4.11), but is in fact costlier than innate behaviour; the population becomes locked in to a suboptimal solution.

Appendix B: Simulation Parameters

This appendix details the model parameters used in each of the experiments detailed within this dissertation. Parameter values marked with a hyphen ("-") are not applicable because they are varied within the experiment.

Experiment	N	B	α	μ	p_{switch}	p_{noise}	*trials*
4.3.1.1 numeric-static	64	32	0.01	0.05	0.00	0.25	100
4.3.1.2 numeric-static	64	32	0.01	0.05	0.00	0.25	100
4.3.1.3 numeric-perturbation	64	32	0.01	0.05	0.00	0.25	100
4.3.2 numeric-fluctuating	64	32	0.01	0.05	0.01	0.25	100
4.3.3.1 numeric-static-B-sweep	64	-	0.01	0.05	0.00	0.25	50
4.3.3.3 numeric-static-B-sweep-focused-no-soc	64	-	0.01	0.05	0.00	0.25	20
4.3.4 numeric-lattice	64	-	0.01	0.05	-	0.25	10
5.3.2.1 topology-sweep-static	64	32	0.01	0.05	0.0	0.25	50
5.3.2.3 topology-sweep-fluctuating	64	32	0.01	0.05	0.001	0.25	50
5.4.2.1 graph-k-sweep-initial-fixed	64	16	0.01	0.05	0.00	0.25	100
5.4.2.2 graph-pswitch-sweep	64	16	0.01	0.05	-	0.25	100
6.4.1.1 spatial-structure-distribution-sweep-movement	64	16	0.01	0.05	0.00	0.25	50
6.4.1.4 spatial-structure-distribution-sweep-fluctuating	64	16	0.01	0.05	0.001	0.25	30
6.4.1.5 spatial-structure-distribution-sweep-T=2-restricted	64	16	0.01	0.05	0.00	0.25	100
6.4.2.1 spatial-structure-metric-sweep	64	16	0.01	0.05	0.00	0.25	50

All error bars on results are shown at a 95% confidence interval ($p < 0.05$).

Appendix C: Implementation

The implementation of the simulations used to realise this research was performed as a two stage process.

Prototypes were first developed in the Python[1] programming language, with a novel agent-based evolutionary modelling framework using elements from NumPy and SciPy.

An optimised version was subsequently developed using C++, with the GNU Science Library[2] for pseudo-random number generation and other probability functions. This provided over two orders of magnitude performance increase from the prototype. This was used to carry out distributed experiments in conjunction with a novel Python framework for distributed processing, using the Xgrid cluster framework and GNU Parallel[3](Tange, 2011).

Analysis was performed using the R[4] environment, with most graphs produced using grammars from `ggplot2`.

The source code for all of these simulations can be obtained from the author's source code repository:

`https://github.com/ideoforms/phd`

[1] http://www.python.org
[2] http://www.gnu.org/software/gsl/
[3] http://www.gnu.org/software/parallel/
[4] http://www.r-project.org

Bibliography

Ackley, D. and Littman, M. (1991). Interactions between learning and evolution. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 487–509. Addison-Wesley, Reading, MA.

Agrawal, A. A. (2001). Phenotypic plasticity in the interactions and evolution of species. *Science*, 294(5541):321–326.

Ancel, L. W. (2000). Undermining the Baldwin expediting effect: does phenotypic plasticity accelerate evolution? *Theoretical population biology*, 58(4):307–19.

Anderson, R. W. (1995). Learning and evolution: a quantitative genetics approach. *Journal of theoretical biology*, 175(1):89–101.

Arbilly, M. and Laland, K. N. (2014). The local enhancement conundrum: In search of the adaptive value of a social learning mechanism. *Theoretical population biology*, 91:50–7.

Avital, E. and Jablonka, E. (2000). *Animal traditions: Behavioural inheritance in evolution*. Cambridge University Press, Cambridge.

Axelrod, R. and Hamilton, W. (1981). The Evolution of Cooperation. *Science*, 211(4489):1390–1396.

Baldwin, J. (1896). A new factor in evolution. *The American Naturalist*, 30(354):441–451.

Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):509–512.

Barnard, C. and Sibly, R. (1981). Producers and scroungers: A general model and its application to captive flocks of house sparrows. *Animal Behaviour*, 29(2):543–550.

Baythavong, B. S. (2011). Linking the spatial scale of environmental variation and the evolution of phenotypic plasticity: selection favors adaptive plasticity in fine-grained environments. *The American Naturalist*, 178(1):75–87.

Bedau, M. (1999). Can Unrealistic Computer Models Illuminate Theoretical Biology? *Proceedings of GECCO 1999*, pages 20–23.

Belew, R. (1990). Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems*, 4:11–49.

Bentley, P. J. (2009). Methods for improving simulations of biological systems: systemic computation and fractal proteins. *Journal of The Royal Society Interface*, 6(Suppl 4):S451—-S466.

Best, M. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, 1(617):1–31.

Bollobás, B. (1985). *Random Graphs*. Academic Press, New York, NY.

Borenstein, E., Feldman, M. W., and Aoki, K. (2008). Evolution of learning in fluctuating environments: when selection favors both social and exploratory individual learning. *Evolution*, 62(3):586–602.

Borenstein, E., Meilijson, I., and Ruppin, E. (2006). The effect of phenotypic plasticity on evolution in multipeaked fitness landscapes. *Journal of Evolutionary Biology*, 19(5):1555–70.

Borenstein, E. and Ruppin, E. (2005). The evolution of imitation and mirror neurons in adaptive agents. *Cognitive Systems Research*, 6(3):229–242.

Borg, J. and Channon, A. (2012). Testing the Variability Selection Hypothesis: The Adoption of Social Learning in Increasingly Variable Environments. *Artificial Life*, 13:317–324.

Boyce, M. S., Haridas, C. V., Lee, C. T., Boggs, C. L., Bruna, E. M., Coulson, T., Doak, D., Drake, J. M., Gaillard, J. M., Horvitz, C. C., Kalisz, S., Kendall, B. E., Knight, T., Mastrandrea, M., Menges, E. S., Morris, W. F., Pfister, C. a., and Tuljapurkar, S. D. (2006). Demography in an increasingly variable world. *Trends in Ecology and Evolution*, 21(3):141–148.

Boyd, R. and Richerson, P. (1983). The cultural transmission of acquired variation: Effects on genetic fitness. *Journal of Theoretical Biology*, pages 567–596.

Boyd, R. and Richerson, P. (1988). An evolutionary model of social learning: the effects of spatial and temporal variation. In Zentall, T. and Galef, Jr., B., editors, *Social Learning*, pages 29–48. Erlbaum, Hillside, NJ.

Boyd, R. and Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press, Chicago.

Bradshaw, A. (1965). Evolutionary significance of phenotypic plasticity in plants. *Advances in genetics*, 41(July):491–502.

Brown, J. S. (2000). Foraging ecology of animals in response to heterogeneous environments. *The Ecological Consequences of Environmental Heterogeneity*, pages 181–214.

Cavalli-Sforza, L. L. (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton University Press.

Chapman, L. G., Galis, F., and Shinn, J. (2000). Phenotypic plasticity and the possible role of genetic assimilation: Hypoxia-induced trade-offs in the morphological traits of an African cichlid. *Ecology Letters*, 3(5):387–393.

Charmantier, A., McCleery, R. H., Cole, L. R., Perrins, C., Kruuk, L. E. B., and Sheldon, B. C. (2008). Adaptive phenotypic plasticity in response to climate change in a wild bird population. *Science*, 320(5877):800–3.

Coolen, I., Dangles, O., and Casas, J. (2005). Social learning in noncolonial insects? *Current biology*, 15(21):1931–5.

Crispo, E. (2007). The Baldwin effect and genetic assimilation: revisiting two mechanisms of evolutionary change mediated by phenotypic plasticity. *Evolution; international journal of organic evolution*, 61(11):2469–79.

Davies, N. B., Krebs, J. R., and West, S. A. (2012). *An introduction to behavioural ecology*. John Wiley & Sons.

Descamps-Julien, B. and Gonzalez, A. (2005). Stable coexistence in a fluctuating environment: an experimental demonstration. *Ecology*, 86(10):2815–2824.

Di Paolo, E. A., Noble, J., and Bullock, S. (2000). Simulation models as opaque thought experiments. *The Seventh International Conference on Artificial Life*, pages 497–506.

Dingemanse, N. J., Both, C., Drent, P. J., and Tinbergen, J. M. (2004). Fitness consequences of avian personalities in a fluctuating environment. *Proceedings. Biological sciences / The Royal Society*, 271(1541):847–852.

Dukas, R. (2008). Evolutionary biology of insect learning. *Annual review of entomology*, 53:145–60.

Dukas, R. (2013). Effects of learning on evolution: robustness, innovation and speciation. *Animal Behaviour*, 85.

Dukas, R. and Visscher, P. K. (1994). Lifetime learning by foraging honey bees.

Dyer, D. and Bentley, P. (2002). PLANTWORLD: Population Dynamics in Contrasting Environments. In *GECCO Late Breaking Papers*.

Dyer, J. R., Croft, D. P., Morrell, L. J., and Krause, J. (2008). Shoal composition determines foraging success in the guppy. *Behavioral Ecology*, 20(1):165–171.

Efferson, C., Lalive, R., Richerson, P. J., McElreath, R., and Lubell, M. (2008). Conformists and mavericks: the empirics of frequency-dependent cultural transmission. *Evolution and Human Behavior*, 29(1):56–64.

Feldman, M. W., Aoki, K., and Kumm, J. (1996). Individual Versus Social Learning: Evolutionary Analysis in a Fluctuating Environment. *Anthropological Science*, 104(3):209–232.

Flache, A. and Hegselmann, R. (2001). Do Irregular Grids make a Difference? Relaxing the Spatial Regularity Assumption in Cellular Models of Social Dynamics. *Journal of Artificial Societies and Social Simulation*, 4(4).

Fogarty, L., Rendell, L., and Laland, K. N. (2012). The Importance of Space in Models of Social Learning, Cultural Evolution and Niche Construction. *Advances in Complex Systems*, 15(01n02):1150001.

Fortin, M., Dale, M., and Hoef, J. (2006). Spatial analysis in ecology. In El-Shaarawi, A. H. and Piegorsch, W. W., editors, *Encyclopedia of environmetrics*, volume 4, pages 2051–2058. John Wiley & Sons.

Freeberg, T. (2000). Culture and courtship in vertebrates: a review of social learning and transmission of courtship systems and mating patterns. *Behavioural Processes*, 51:177–192.

Frohn, R. C. (1998). *Remote sensing for landscape ecology: new metric indicators for monitoring, modeling, and assessment of ecosystems*, volume 6. CRC Press.

Galef, B. (1976). Social Transmission of Acquired Behavior: A Discussion of Tradition and Social Learning in Vertebrates. *Advances in the Study of Behavior*, pages 77–100.

Galef, B. (1990). Tradition in animals: field observations and laboratory analyses. In Bekoff, M. and Jamieson, D., editors, *Interpretations and explanations in the study of behaviour: comparative perspectives*, volume I, pages 74–95. Westview Press.

Galef, B. G. and Laland, K. N. (2005). Social Learning in Animals: Empirical Studies and Theoretical Models. *BioScience*, 55(6):489.

Gause, G. (1934). *The struggle for existence*. Williams & Wilkins, Baltimore, MD.

Giraldeau, L.-A., Valone, T. J., and Templeton, J. J. (2002). Potential disadvantages of using socially acquired information. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 357(1427):1559–66.

Godfrey-Smith, P. (2003). Between Baldwin skepticism and Baldwin boosterism. In *Evolution and learning: The Baldwin effect reconsidered*, pages 53–67. The MIT Press.

Gordon, D. M. (2011). The fusion of behavioral ecology and ecology. *Behavioral Ecology*, 22(2):225–230.

Grimm, V. and Railsback, S. F. (2013). *Individual-Based Modeling and Ecology*. Princeton University Press.

Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., Thulke, H.-H., Weiner, J., Wiegand, T., and DeAngelis, D. L. (2005). Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science*, 310(5750):987–91.

Grünbaum, D. (2012). The logic of ecological patchiness. *Interface focus*, 2(2):150–5.

Gupta, A. and Lewontin, R. (1982). A Study of Reaction Norms in Natural Populations of Drosophila pseudoobscura. *Evolution*, 36(5):934–948.

Gustafson, E. (1998). Quantifying landscape spatial pattern: what is the state of the art? *Ecosystems*, 1(2):143–156.

Hallsson, L. R. and Björklund, M. (2012). Selection in a fluctuating environment leads to decreased genetic variation and facilitates the evolution of phenotypic plasticity. *Journal of evolutionary biology*, 25(7):1275–1290.

Hartmann, S. (1996). The world as a process. In *Modelling and simulation in the social sciences from the philosophy of science point of view*, pages 77–100. Springer.

Henrich, J. and McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews*, 12(3):123–135.

Heyes, C. M. (1994). Social learning in animals: categories and mechanisms. *Biological reviews of the Cambridge Philosophical Society*, 69(2):207–31.

Heyes, C. M. and Galef Jr, B. G. (1996). *Social learning in animals: the roots of culture*. Elsevier.

Hinton, G. and Nowlan, S. (1987). How learning can guide evolution. *Complex systems*, 1:495–502.

Hoffmann, A. a. and Merilä, J. (1999). Heritable variation and evolution under favourable and unfavourable conditions. *Trends in Ecology and Evolution*, 14(3):96–101.

Huffaker, C. B. (1958). Experimental studies on predation: dispersion factors and predator-prey oscillations. *Hilgardia*, 27:343–383.

Humphreys, P. (1994). Numerical experimentation. *Scientific Philosopher*, pages 103–121.

Hunt, G. R. and Gray, R. D. (2003). Diversification and cumulative evolution in New Caledonian crow tool manufacture. *Proceedings of Biological Sciences*, 270(1517):867–74.

Hutchinson, G. E. (1961). The paradox of the plankton. *American Naturalist*, 95:137–145.

Jablonka, E. and Lamb, M. J. (2005). *Evolution in four dimensions: Genetic, epigenetic, behavioral, and symbolic variation in the history of life*. MIT press.

Jain, K., Krug, J., and Park, S. (2011). Evolutionary advantage of small populations on complex fitness landscapes. *Evolution*, pages 1–27.

Jones, C. and Lawton, J. H. (1995). *Linking species and ecosystems*. Springer.

Jones, D. and Blackwell, T. (2011). Social Learning and Evolution in a Structured Environment. In *Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems*, Paris, FR.

Jones, D. and d'Inverno, M. (2011). Amplifying scientific creativity with agent-based simulation. In *Proceedings of the International Conference on Complex Systems*.

Kameda, T. and Nakanishi, D. (2002). CostâĂŞbenefit analysis of social/cultural learning in a nonstationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior*, 23:373–393.

Károlyi, G., Péntek, A., Scheuring, I., Tél, T., and Toroczkai, Z. (2000). Chaotic flow: the physics of species coexistence. *Proceedings of the National Academy of Sciences*, 97(25):13661–13665.

Kendal, J., Giraldeau, L.-A., and Laland, K. (2009). The evolution of social learning rules: payoff-biased and frequency-dependent biased transmission. *Journal of theoretical biology*, 260(2):210–9.

Kenward, B., Rutz, C., a.S. Weir, A., and Kacelnik, A. (2006). Development of tool use in New Caledonian crows: inherited action patterns and social influences. *Animal Behaviour*, 72(6):1329–1343.

Keymer, J. E., Galajda, P., Muldoon, C., Park, S., and Austin, R. H. (2006). Bacterial metapopulations in nanofabricated landscapes. *Proceedings of the National Academy of Sciences of the United States of America*, 103(46):17290–5.

Korona, R., Nakatsu, C. H., Forney, L. J., and Lenski, R. E. (1994). Evidence for multiple adaptive peaks from populations of bacteria evolving in a structured habitat. *Proceedings of the National Academy of Sciences of the United States of America*, 91(19):9037–9041.

Kossinets, G. and Watts, D. (2006). Empirical analysis of an evolving social network. *Science*, 311(January):88–90.

Kudo, H. and Dunbar, R. (2001). Neocortex size and social network size in primates. *Animal Behaviour*, 62(4):711–722.

Kurvers, R. H. J. M., Hamblin, S., and Giraldeau, L.-A. (2012). The effect of exploration on the use of producer-scrounger tactics. *PloS one*, 7(11):e49400.

Laland, K. N. (2004). Social learning strategies. *Learning & behavior*, 32(1):4–14.

Laland, K. N. and Hoppitt, W. (2003). Do animals have culture? *Evolutionary Anthropology: Issues, News, and Reviews*, 12(3):150–159.

Laland, K. N., Odling-Smee, J., and Feldman, M. W. (2000). Niche construction, biological evolution, and cultural change. *The Behavioral and Brain Sciences*, 23(1):131–46; discussion 146–75.

Laland, K. N., Odling-Smee, J., Hoppitt, W., and Uller, T. (2012). More on how and why: cause and effect in biology revisited. *Biology & Philosophy*, 28(5):719–745.

Lancet, Y. and Dukas, R. (2012). Socially Influenced Behaviour and Learning in Locusts. *Ethology*, 118(3):302–310.

Lande, R. (1988). Genetics and demography in biological conservation. *Science*, 241(4872):1455–1460.

Lande, R. (2009). Adaptation to an extraordinary environment by evolution of phenotypic plasticity and genetic assimilation. *Journal of evolutionary biology*, 22(7):1435–46.

Lee, C. E., Remfert, J. L., and Gelembiuk, G. W. (2003). Evolution of physiological tolerance and performance during freshwater invasions. *Integrative and comparative biology*, 43(3):439–449.

Lefebvre, L. and Palameta, B. (1988). Mechanisms, ecology, and population diffusion of socially learned, food-finding behavior in feral pigeons. In Zentall, T. R. and Galef, B. G., editors, *Social Learning*. Erlbaum, Hillside, NJ.

Lehmann, L., Keller, L., and Sumpter, D. J. T. (2007). The evolution of helping and harming on graphs: The return of the inclusive fitness effect. *Journal of Evolutionary Biology*, 20(6):2284–2295.

Levin, S. (1992). The problem of pattern and scale in ecology. *Ecology*, 73:1943–1967.

Levins, R. (1968). *Evolution in changing environments: some theoretical explorations*. Princeton University Press.

Lewontin, R. (1991). *The Doctrine of DNA*. Penguin, London, UK.

Li, H. and Reynolds, J. (1994). A Simulation Experiment to Quantify Spatial Heterogeneity in Categorical Maps. *Ecology*, 75(8):2446–2455.

Li, H. and Reynolds, J. (1995). On definition and quantification of heterogeneity. *Oikos*, 73:280–284.

Li, H. and Reynolds, J. F. (1993). A new contagion index to quantify spatial patterns of landscapes. *Landscape Ecology*, 8(3):155–162.

Lieberman, E., Hauert, C., and Nowak, M. (2005). Evolutionary dynamics on graphs. *Nature*, 433(7023):312–316.

Lorenz, K. (1971). *Evolution and modification of behavior*. University of Chicago Press, Chicago, US.

Lowen, C. and Dunbar, R. (1996). Cultural evolution in a structured population. Technical report, ESRC Centre on Economics Learning and Social Evolution.

Luckinbill, L. (1973). Coexistence in Laboratory Populations of Paramecium Aurelia and Its Predator Didinium Nasutum. *Ecology*, 54(6):1320–1327.

Lundberg, P., Ranta, E., Ripa, J., and Kaitala, V. (2000). Population variability in space and time. *Trends in Ecology and Evolution*, 15(11):460–464.

MacArthur, R. and Pianka, E. (1966). On optimal use of a patchy environment. *American Naturalist*, 100(916):603–609.

Marler, P. (1970). A comparative approach to vocal learning: Song development in white-crowned sparrows. *Journal of Comparative and Physiological Psychology*, 71(2):1–25.

Marler, P. and Slabbekoorn, H. (2004). *Nature's Music: The Science of Birdsong*. Academic Press.

Marquet, P., Fortin, M., Pineda, J., and Wallin, D. (1993). Ecological and evolutionary consequences of patchiness: a marine-terrestrial perspective. *Patch dynamics*.

Matesanz, S., Gianoli, E., and Valladares, F. (2010). Global change and the evolution of phenotypic plasticity in plants. *Annals of the New York Academy of Sciences*, 1206:35–55.

Mayley, G. (1996). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, 4(3):213–234.

Mayley, G. (1997). Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In *Proceedings of the Fourth European Conference on Artificial Life*, volume 97, pages 135–144.

McGarigal, K. (2006). Landscape pattern metrics. In El-Shaarawi, A. H. and Piegorsch, W. W., editors, *Encyclopedia of environmetrics*. John Wiley & Sons.

McNamara, J. M., Trimmer, P. C., Eriksson, A., Marshall, J. a. R., and Houston, A. I. (2011). Environmental variability can select for optimism or pessimism. *Ecology Letters*, 14(1):58–62.

Mery, F. and Kawecki, T. J. (2003). A fitness cost of learning ability in Drosophila melanogaster. *Proceedings. Biological sciences / The Royal Society*, 270(1532):2465–2469.

Miller, G. F. and Todd, P. M. (1991). Exploring adaptive agency I: Theory and methods for simulating the evolution of learning. In *Connectionist Models, Proceedings of the 1990 Summer School*.

Moran, N. (1992). The Evolutionary Maintenance of Alternative Phenotypes. *American Naturalist*, 139(5):971–989.

Moran, P. (1958). Random processes in genetics. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 54, pages 60–71. Cambridge University Press.

Morgan, C. L. (1896). On modification and variation. *Science*, 4(99):733–740.

Newman, M. (2010). *Networks: an introduction*. Oxford University Press.

Nolfi, S. and Parisi, D. (1996). Learning to Adapt to Changing Environments in Evolving Neural Networks. *Adaptive Behavior*, 5(1):75–98.

Nolfi, S., Parisi, D., and Elman, J. L. (1994). Learning and Evolution in Neural Networks. *Adaptive Behavior*, 3(1):5–28.

Nowak, M. A. (2006). *Evolutionary dynamics: exploring the equations of life*. Harvard University Press.

Odling-Smee, F. J. (1988). Niche-Constructing Phenotypes. In Plotkin, H. C., editor, *The Role of Behavior in Evolution*, pages 73–132. MIT Press, Cambridge, MA.

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

Papineau, D. (2005). Social learning and the Baldwin effect. In Zilhão, A., editor, *Evolution, Rationality, and Cognition*, pages 40–60. Taylor & Francis, New York, NY.

Pearce, J. M. (2013). *Animal learning and cognition: an introduction*. Psychology Press.

Perlin, K. (2002). Improving noise. *ACM Transactions on Graphics*, 21(3):2–3.

Pigliucci, M. (2001). *Phenotypic plasticity: beyond nature and nurture*. Johns Hopkins University Press, Baltimore, MD.

Pigliucci, M., Murren, C. J., and Schlichting, C. D. (2006). Phenotypic plasticity and evolution by genetic assimilation. *The Journal of Experimental Biology*, 209:2362–7.

Rainey, P. and Travisano, M. (1998). Adaptive radiation in a heterogeneous environment. *Nature*, 32:69–72.

Reed, D. H. (2005). Relationship between Population Size and Fitness. *Conservation Biology*, 19(2):563–568.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., and Laland, K. N. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, 328(5975):208–13.

Riitters, K. H., Neil, R. V. O., Hunsaker, C. T., Wickham, J. D., Yankee, D. H., and Timmins, S. P. (1995). A factor analysis of landscape pattern and structure metrics. *Landscape Ecology*, 10(1):23–39.

Rohrlich, F. (1990). Computer Simulation in the Physical Sciences. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 2:507–518.

Romme, W. (1982). Fire and landscape diversity in subalpine forests of Yellowstone National Park. *Ecological Monographs*, 52(2):199–221.

Rozen, D. E., Habets, M. G. J. L., Handel, A., and de Visser, J. A. G. M. (2008). Heterogeneous adaptive trajectories of small populations on complex fitness landscapes. *PloS one*, 3(3):e1715.

Sarin, S. and Dukas, R. (2009). Social learning about egg-laying substrates in fruitflies. *Proceedings of the Royal Society B: Biological Sciences*, 276(1677):4323–4328.

Scheiner, S. (1993). Genetics and evolution of phenotypic plasticity. *Annual review of ecology and systematics*, 24(1993):35–68.

Schreiber, S. J. (2010). Interactive effects of temporal correlations, spatial heterogeneity and dispersal on population persistence. *Proceedings. Biological sciences / The Royal Society*, 277(1689):1907–1914.

Seppänen, J.-T. and Forsman, J. T. (2007). Interspecific social learning: novel preference can be acquired from a competing species. *Current biology : CB*, 17(14):1248–52.

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

Simpson, G. (1953). The Baldwin effect. *Evolution*, 7(2).

Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK.

Smith, J. M. (1987). When learning guides evolution. *Nature*, 329:761–762.

Smith, J. M. and Price, G. (1973). The Logic of Animal Conflict. *Nature*, 246:15–18.

Smith, L. D. and Palmer, a. R. (1994). Effects of manipulated diet on size and performance of brachyuran crab claws. *Science*, 264(5159):710–712.

Szabó, G. and Fáth, G. (2007). Evolutionary games on graphs. *Physics Reports*, 446(4-6):97–216.

Sznajder, B., Sabelis, M. W., and Egas, M. (2012). How Adaptive Learning Affects Evolution: Reviewing Theory on the Baldwin Effect. *Evolutionary Biology*, 39(3):301–310.

Tange, O. (2011). GNU Parallel: the command-line power tool. *The USENIX Magazine*, 36(1):42–47.

Taylor, C. and Fudenberg, D. (2004). Evolutionary game dynamics in finite populations. *Bulletin of Mathematical Biology*, 66(6):1261–1644.

Taylor, P. D., Day, T., and Wild, G. (2007). Evolution of cooperation in a finite homogeneous graph. *Nature*, 447(7143):469–472.

Taylor, P. D., Fahrig, L., Henein, K., and Merriam, G. (1993). Connectivity is a vital element of landscape structure. *Oikos*, pages 571–573.

Tebbich, S., Taborsky, M., Fessl, B., and Blomqvist, D. (2001). Do woodpecker finches acquire tool-use by social learning? *Proceedings of the Royal Society B: Biological Sciences*, 268(1482):2189–93.

Templeton, A. and Rothman, E. (1978). Evolution in fine-grained environments I. Environmental runs and the evolution of homeostasis. *Theoretical Population Biology*, 355:340–355.

Thorpe, W. H. (1956). *Learning and instinct in animals*. Harvard University Press, Cambridge, MA.

Tilman, D., Knops, J., Wedin, D., Reich, P., Ritchie, M., and Siemann, E. (1997). The influence of functional diversity and composition on ecosystem processes. *Science*, 277(5330):1300–1302.

Verner, J. (1965). Time Budget of the Male Long-billed Marsh Wren During the Breeding Season. *The Condor*, 67(2):125–139.

Vickery, W., Giraldeau, L., and Templeton, J. (1991). Producers, scroungers, and group foraging. *American Naturalist*, 137(6):847–863.

Volterra, V. (1928). Variations and fluctuations of the number of individuals in animal species living together. *Journal du Conseil Permanent International pour l'Exploration de la Mer*, 3(1):3–51.

Waddington, C. (1953). Genetic assimilation of an acquired character. *Evolution*, 7(2):118–126.

Wakano, J. Y. and Aoki, K. (2006). A mixed strategy model for the emergence and intensification of social learning in a periodically changing natural environment. *Theoretical population biology*, 70(4):486–97.

Wakano, J. Y., Aoki, K., and Feldman, M. W. (2004). Evolution of social learning: a mathematical analysis. *Theoretical population biology*, 66(3):249–58.

Wcislo, W. (1989). Behavioral environments and evolutionary change. *Annual Review of Ecology and Systematics*, 20(1989):137–169.

West-Eberhard, M. (1989). Phenotypic plasticity and the origins of diversity. *Annual Review of Ecology and Systematics*, 20(1989):249–278.

West-Eberhard, M. J. (2003). *Developmental Plasticity and Evolution*. Oxford University Press.

Wey, T., Blumstein, D. T., Shen, W., and Jordán, F. (2008). Social network analysis of animal behaviour: a promising tool for the study of sociality. *Animal Behaviour*, 75(2):333–344.

Wiens, J. A. (2000). Ecological heterogeneity: an ontogeny of concepts and approaches. *The ecological consequences of environmental heterogeneity*, pages 9–31.

Winsberg, E. (2001). Simulations, Models, and Theories: Complex Physical Systems and Their Representations. *Philosophy of Science*, 68(S1):S442.

Wolfram, S. (1994). *Cellular automata and complexity: collected papers*, volume 1. Addison-Wesley Reading.

Wright, S. (1931). Evolution in Mendelian Populations. *Genetics*, 16(2):97–159.

Zentall, T. R. and Galef, B. G. (1988). *Social learning: Psychological and biological perspectives.* Erlbaum, Hillside, NJ.

Zhivotovsky, L. A., Feldman, M. W., and Bergman, A. (1996). On the Evolution of Phenotypic Plasticity in a Spatially Heterogeneous Environment. *Evolution*, 50(2):547–558.