UNIVERSITY OF SOUTHAMPTON

On the Origin and Evolution of Social Learning

Reducing Individual Requirements for the Emergence of Cultural Evolutionary Systems

by

Miguel González Canudas

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

in the
Faculty of Physical Sciences and Engineering
Electronics and Computer Science

August 2017

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING
ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Miguel González Canudas

The apparent adaptive value of culture was once assumed to be an explanation for the evolution of social learning. The original argument was that obtaining information from others, even if it was just by unbiased imitation, would save social learners the costs and risks of finding solutions as individual learners or instinctive actuators. These assumptions have been challenged by existing theoretical work where the evolution of social learning and the consequent emergence of culture are shown to require more than just unbiased imitation. The consensus is that an extra source of selection, apart from natural selection, is required for social learning to evolve. This consensus is based on two basic problems of unbiased imitation in the context of existing models. The first problem is that, when only natural selection is considered, unbiased imitation breaks the relation between phenotypic fitness and the frequency of phenotypic replication. The second problem is that perfect imitation can cause phenotypic stagnation as populations with a large proportion of social learners maintain low phenotypic variance, which in turn prevents the fixation of social learners under any rate of environmental change. Subsequently, researchers in the field have developed models that necessitate individual learners (i.e., individuals with the ability to improve on innate behaviours) and different forms of decision-making processes, like biases towards imitating fitter strategies or the ability to improve on the strategies already acquired, in order to explain the evolution of social learners and the emergence of cultural evolutionary systems.

Here we claim individual learning and decision-making processes are unnecessary for the emergence of cultural evolutionary systems. These systems can be explained because, when allowed to explore an extended strategy space, they are faster problem solvers than genetic evolutionary systems. We base this conclusion on the results of a series of individual based simulations, which show that under strong survival selection and while imitation error rates are higher than genetic mutation, culture can find solutions to environmental problems faster than genetic systems. Our work shows that unbiased horizontal imitation - a particular form of social learning that is generally considered a hindrance in classical models of the origin of culture - can also be adaptive under certain conditions of selection pressures and error rates. Furthermore, we show moderate environmental change increases the chances of culture emerging in a population of pure social learners; a process previously thought to require a mixed strategy of social and individual learning. Compared to existing models our approach makes fewer and more plausible assumptions, as it does not involve additional sources of selection in the form of individual learning or decision-making processes. This broadens the range of species for which culture could evolve as it reduces individual-level requirements for its evolution.

Academic Thesis: Declaration Of Authorship

I, Miguel Gonzalez Canudas declare that the thesis entitled

On the Origin and Evolution of Social Learning:
Reducing Individual Requirements for the Emergence of Cultural Evolutionary Systems

and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published or have been accepted to be published as:
    - Gonzalez, M., Watson, R. A., Noble, J., and Bullock, S. (2014). The origin of culture: selective conditions for horizontal information transfer. Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems. MIT Press. 408-414.
    - Gonzalez, M., Watson, R., and Bullock, S. (2017). Minimally sufficient conditions for the evolution of social learning and the emergence of non-genetic evolutionary systems. Artificial Life. MIT Press. (Accepted).
8. Parts of this work have been submitted for publication, or are being prepared for submission as:
    - Gonzalez, M., Watson, R., and Bullock, S. (2018). When viability selection is strong, horizontal imitation is dominant over vertical imitation in a purely Darwinian model of the evolution of social learning (in preparation).
    - Gonzalez, M., Watson, R., and Bullock, S. (2018). Under moderate environmental change unbiased social learning increases the overall mean fitness of a population (in preparation).

Signed: ……………………………………………………………………………

Date: ……………………………………………………………………………

Contents

List of Figures

Nomenclature

Cavalli-Sforza and Feldman's Model

v	Proportion of *skilled* social learners
u	Proportion of *unskilled* social learners
w	Proportion of instinctive actuators
b	Probability of vertical transmission
f	Probability of horizontal transmission
v_0	v after imitation
u_0	u after imitation
s	Fitness advantage of *skilled* individuals
v'	v after imitation and selection
u'	u after imitation and selection
w'	w after selection
ϕ	Normalisation term

Rogers' model

ω	Base fitness value
ω_i	Individual learner's fitness
ω_s	Social learner's fitness
b	Fit behaviour advantage
c	Individual learning cost
τ	Generations since last fit imitation
p	Proportion of social learners
μ	Environmental change rate per generation
$\hat{p}$	p at equilibrium

Simulation model and boundary estimations

N	Population size
L	Length of genotypes and phenotypes (bits)
L_1	Sum-of-1s

L_m	Sum of environment-matching bits
x_r	Reproductive selection exponent
x_d	Survival selection exponent
μ	General mutation/error rate
μ_g	Genetic mutation rate
μ_p	Imitation error rate
μ_v	Vertical imitation error rate
μ_h	Horizontal imitation error rate
μ_e	Environmental change rate (per bit)
μ_c	Imitation switch mutation rate
P_{ri}	Prob. of being selected to breed
P_{dj}	Prob. of being selected to die
α	Critical sum-of-1s value
δ_e	Env. change rate (for randomly selected bit)
$\omega_g(b, a)$	Genetic fitness for b in pop. of a
$\omega_p(b, a)$	Cultural fitness for b in pop. of a

Acknowledgements

I would like to thank my supervisors, Dr. Richard Watson, Dr. Seth Bullock and Dr. Jason Noble for their involvement, support and patience throughout my PhD. I would also like to thank Dr. Markus Brede, Dr. Simon Tudge, Lewis Brace, Dr. Chris Cox, Dr. Adam Jackson and Dr. Paul Ryan for engaging in numerous discussions that have benefited the development of this work. Finally, my partner, Mary Crossland; and family, Dr. Nieves Canudas, Vicente Gonzalez, Maria Fernanda Lopez, Dr. Jose Villamizar and many others; for their emotional and moral support.

To my family

Chapter 1

Introduction

Culture is one of the most important characteristics of our species as it allows us to obtain valuable information from others and from past experiences passed down generations (Mesoudi, 2011). Culture is a population-level process that emerges as a consequence of the evolution of social learning, a form of learning where individuals copy information or behaviour from one another (Richerson and Boyd, 2006). Culture is not only present in humans as other primates, mammals, birds and fish species have been found to have some basic form of collective information copying (Whiten and van Schaik, 2007; Laland and Galef, 2009; van de Waal et al., 2013). Whether these forms of animal cultures can be considered qualitatively the same as the human case is an on-going debate in the literature (Laland and Hoppitt, 2003; Laland and Janik, 2006; Whiten and van Schaik, 2007). In this thesis, I am interested in the common processes leading to its emergence across taxa. For this reason, I consider animal and human culture to originate by the same basic mechanism (Richerson and Boyd, 2006) and refer to the product of this mechanism simply as *culture*.

Developing culture requires individuals to be able to imitate one another, but current theory states that just imitating individuals at random is not sufficient to explain its origin, ergo, the inclusion of *individual learning* or some form of imitation bias in favour of fitter strategies is required (Richerson and Boyd, 2006; Jablonka and Lamb, 2006; Rogers, 1988). In all contexts within this work, unless stated otherwise, I define individual learning as the ability to improve on innate or acquired behaviours by processes which are not dependent on other individuals in the population (e.g., by trial-and-error). The requirement for individual learning or imitation bias reduces the range of species for which culture could be expected to arise to only those where individuals improve the behaviours that they already have acquired, or alternatively, where individuals have a direct or indirect ability to critically assess the value of the behaviours of others.

In contrast to individual learning, *social learning* is defined as the ability to learn from others by any process that results in the imitation of the behaviour of one individual (i.e., the model to imitate) by another individual (i.e., the social learner). When social learners imitate any member of the population at random (unbiased imitation) and are under natural selection (fitness is a function of reproductive output alone), they have been shown to be unable to fix and consequently originate a cultural evolutionary system. This is because the pool of strategies that social learners imitate from (i.e., the population) has, logically, an average fitness value. Since every social learner's offspring picks a strategy to imitate from this pool, their strategy is not defined by their genetic endowment and therefore the link between an increased reproductive output and the increase in frequency of the strategy that produced that output is broken. Here I call this process *phenotype-fitness de-correlation*. In existing models, introducing an imitation bias towards fitter strategies or, alternatively, improving on the strategies already acquired (individual learning/guided variation) are the preferred approaches to compensate for phenotype-fitness de-correlation. Only when adding these *decision-making abilities* can social learners outcompete genetically innate individuals (acting instinctively) (Henrich and Boyd, 1998; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007; McElreath and Henrich, 2007).

In this thesis, I present an alternative explanation for the evolution of social learning that does not require biased imitation or guided variation. I claim that social learning can evolve as an adaptation because culture outcompetes and then masks its genetic counterpart when finding optimal solutions to environmental challenges. I show that this is the case when a component of selection affects longevity, the imitation error rate is larger than genetic mutation, and the strategy space is large (i.e., when many strategies can be explored).

None of these requirements depend on the capabilities or complexity of organisms and are rather common evolutionary assumptions. My work is framed as a theoretical contribution that shows how simple cultural systems (i.e., non-genetic evolutionary systems) can arise and fix once unbiased imitation has evolved.

My work is structured in six chapters, in the current chapter I introduce and briefly describe the problems that existing models have encountered when trying to explain the evolution of unbiased imitation. This chapter also contains my main research questions, hypotheses, claims and contributions.

The second chapter comprises a short literature review that focuses on (1) describing culture as an evolutionary system, and (2) explaining two seminal models that expose the issues of phenotype-fitness de-correlation and *social learning stagnation* (i.e., the reduction of phenotypic variance in the population caused by social learning) present in previous theoretical work.

The core results of my thesis are explained in Chapter 3 with a simulation model that shows how culture can evolve under minimal assumptions (i.e., without decision making attributes or individual learners). Based on the conclusions of Chapter 3; Chapter 4 and Chapter 5 respectively, I explain: (i) the value of horizontal vs. vertical imitation in social learning and (ii) the effects of environmental change on the emergence of adaptive and unbiased social learning (i.e., unbiased social learning that increases the overall mean fitness of the population). Chapter 6 provides general conclusions for this thesis. An extended description and summary of the contributions of this thesis can be found in Section 1.8. This work uses subject-specific terms. In order to maintain consistent definitions, I have included a glossary of technical terms as an Appendix of this thesis (see Appendix A).

1.1 Disclaimer: Introducing My Working Definition of Culture

Culture is a difficult concept to grasp, yet it is easy to recognise its importance to our species. Robert Boyd and Peter Richerson argue that we owe to it our spatial expansion and environmental adaptation, our innovation and civilisation, all of which have increased in magnitude and efficiency over a time period that is short enough to remove genetic evolution as an explanation (Boyd and Richerson, 1988a; Richerson and Boyd, 2006; Perreault, 2012). The specific processes that gave rise to culture are still unclear, but we know that the ability to imitate behaviours (i.e., social learning) must have evolved genetically before culture could emerge (Richerson and Boyd, 2006).

Attempts to define culture are abundant in the literature (Kroeber, 1948; Ingold, 1986; Richerson and Boyd., 2001; Rogers, 1988; Richerson and Boyd, 2006). However, many of these definitions include characteristics that I do not consider in the work of this thesis as they are too specific and/or sophisticated. If I was to use these definitions, which claim that culture requires physical artefacts and/or basic forms of language, I would restrict the existence of any kind of cultural system to humans and a few other species of primates. My work, which shows how basic forms of culture can emerge with minimal individual requirements, is therefore much better suited by a general definition of culture that explicitly distinguishes it from a genetic evolutionary system but that does not include specific details that are unique to culturally advanced species. In this thesis, I define culture as a simple non-genetic evolutionary system maintained by social learning. More specifically, I adhere to the definition derived from the theory of cultural evolution developed by Boyd and Richerson, which defines culture as a set of behavioural traits that are not the direct result of genetic expression but the product of an evolving pool of variants (i.e., behaviours and ideas) stored and imperfectly transmitted within and/or between overlapping generations by means of social learning (Holden and Mace, 1997; Avital and Jablonka, 2005; Richerson and Boyd, 2006).

To frame my work within the scope of existing theory, I must first establish that the term “culture” is used here to refer to all forms of evolutionary systems that emerge as a consequence of non-genetic information transfer between individuals; regardless of how simple these systems are. In nature, the existence of non-genetic evolutionary systems driven by social learning has been identified in many species (Laland and Janik, 2006; Laland and Hoppitt, 2003; Avital and Jablonka, 2005). Evidence of these systems is found in birds, where song learning occurs by means of imitation (Jenkins, 1978; Heyes, 1994); chimpanzees, with well-known cases of tool use and advanced material culture (McGrew, 1998); Cetaceans, imitating cooperative hunting practices and mating calls (Rendell and Whitehead, 2001); and others (Laland and Galef, 2009).

Human culture is a sophisticated example of a non-genetic evolutionary system (Richerson and Boyd, 2006). The distinctions between human and animal cultures are outside the scope of this thesis, but I am aware that some schools of thought object to classifying non-human, socially learnt animal behaviours as examples of "real" culture (Laland and Galef, 2009). There is little contention, however, that all these instances share a general mechanism during their emergence, most certainly involving some form of social learning (Richerson and Boyd, 2006; Mesoudi, 2011; Laland and Galef, 2009). In this thesis, I am interested in explaining the evolution of this general mechanism.

1.2 Current Consensus: Unbiased Social Learning is Not Sufficient to Originate Culture

Early literature readily assumed that there was an adaptive value for social learning at the individual level and for culture at the population level (White, 1959; Durham, 1976). The argument for these assumptions was a combination of the seemingly evident advantages that culture offered our species and the notion that by imitating others, social learners would save on costs associated with trial and error during individual learning (Boyd and Richerson, 1988a; Pulliam and Dunford, 1980).

Formal mathematical analyses of this rationale questioned these assumptions and concluded that adaptive culture could not be explained by unbiased social learning alone (Rogers, 1988; Richerson and Boyd, 2006). The seminal models of Rogers (Rogers, 1988) and Cavalli-Sforza and Feldman (Cavalli-Sforza and Feldman, 1983) highlighted the reasons leading to this conclusion. In this chapter these seminal models are briefly described as example cases for two persistent problems in theoretical work.

Chapter 2 will include an in-depth explanation of these two problems and the assumptions recent models have made in order to solve them. Chapter 3 presents my model and explains why previous assumptions are unnecessary.

1.2.1 The First Problem: Horizontal Imitation Causes Phenotype-Fitness De-correlation

One of the first models to suggest that social learning is not adaptive was proposed by Cavalli-Sforza and Feldman in 1983 (Cavalli-Sforza and Feldman, 1983). In their model, the frequency of social learners in a population always converged to zero for all non-zero rates of *horizontal imitation*. In this form of imitation, information is copied from any individual in the population within the same generation, in contrast to *vertical imitation*, where it only occurs from parents to offspring. Cavalli-Sforza and Feldman showed that any degree of horizontal imitation in social learners introduces a cost in the form of *phenotype-fitness de-correlation*. This means, to the extent that social learners imitate random members of the population, the correlation between the fitness of a phenotype (i.e., the effect of a phenotype on genetic reproduction) and the change of frequency for that phenotype (i.e., the change in the proportion of individuals expressing the phenotype in the population) is broken.

In natural populations, social learners can imitate any other individual and not exclusively their own parents. It is then likely that some degree of horizontal imitation, however marginal, occurs in this form of learning. This causes social learners' offspring to effectively have fitness values closer to or equal to the population average regardless of how fit their parents were. How close to average these fitness values are depends on how frequent imitation is horizontal and/or random. In the case where social learning is entirely horizontal and random, the offspring of a social learner performing a highly fit behaviour will acquire on average the same phenotype as the offspring of a less fit social learner. Another way of looking at this issue is to consider the offspring of social learners as "anyone's offspring" because their phenotype can come from any model-to-imitate in the population. The causal link between horizontal imitation and the cost of phenotype-fitness de-correlation is not emphasised in current literature.

In this thesis, I do not claim to have eliminated the problem of phenotype-fitness de-correlation but have found conditions that attenuate its effects enough so that the explorative advantage of culture can outcompete genetic evolution. Chapter 3 contains a detailed explanation of how this is achieved by separating fitness into *survival* and *reproductive* components. In this context, the survival component of fitness corresponds to the effect of fitness on viability and lifespan, in contrast, the reproductive component of fitness corresponds to the effect of fitness on reproductive output (i.e., number of genetic offspring).

In contrast to reproductive fitness, the survival component influences both genetic and cultural transmission. When exclusively under survival selection, longer-living agents increase their chances to serve as models for unbiased imitation by the same degree that they increase their chances to be randomly selected for reproduction. In this scenario, genetic and cultural systems compete under the same selective pressures, and therefore the higher exploration capabilities of culture (due to the increased rate of transmission errors that are assumed to accompany social learning when compared with genetic transmission) have a chance to outcompete genetic evolution. Furthermore, I found conditions where culture continues to outcompete its genetic counterpart even when a small component of reproductive selection is present.

1.2.2 The Second Problem: Social Learning Stagnates When the Environment Changes

Another model produced by Alan Rogers in 1988 (Rogers, 1988) concluded that under environmental fluctuations, the average fitness of a population of social learners (i.e., the average fitness of its individuals) will converge to the same value of a *monomorphic* (i.e., homogeneous) population of individual learners (i.e., a population that consists entirely of agents that pay a cost in time and resources, ultimately expressed as fitness, to guarantee that their behaviours are advantageous in the current environmental state) (Rogers, 1988).

Rogers suggested that social learning fails at tracking and matching changing environmental states. In populations of pure social learners, individuals will keep copying one another and fix a particular behavioural strategy. When the environment changes, social learners maintain the previous strategy by conformity because they have no access to environmental-matching options to imitate. This mismatch produces a fitness cost. Mixed populations achieve equilibrium between social and individual learners when the cost of individual learning is equal to the cost of potential environmental mismatch by social learners (see Chapter 2 for a detailed explanation). This equilibrium is reached at the same fitness value as the one derived from individual learning. The mixed population cannot be, on average, more fit than a monomorphic population of individual learners.

As is the case for most of the existing theoretical work in the field, Rogers' used a simple dynamical systems approach to model social learning. In his model, environmental change is represented by the fluctuation of two states that can be matched respectively by one of two available behaviours (Rogers, 1988; Henrich and Boyd, 1998; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Borenstein et al., 2007). Existing models do not include an extended strategy space (i.e., a fitness landscape with many more than two genotypes) where strategies can be gradually evolved by either genetic or cultural evolution. The lack of models with this feature has created a void of explanations for the evolution of social learning focusing on the exploratory advantage of culture as an evolutionary system.

The theoretical scenario where pure populations of social learners have no imitation error, no strategy space to explore and no environmental tracking abilities is here referred to as *social learning stagnation.* I consider this problem to be an artefact introduced by existing models of cultural evolution rather than a problem present in nature. In Chapter 3, I show that culture, with its exploration advantage, can outcompete genetic evolution when including a strategy space and an imitation error.

It is not my intention to solve the two problems described here by including any particular single feature in my model but by considering the interaction of a few features in my simulation-based approach (see Section 1.7). In this work, I evaluate the capacity of culture to find solutions faster than its genetic counterpart and find that this capacity alone can overcome the problems mentioned without requiring individual decision-making capabilities. In this context, culture is simply an evolutionary system sustained by imitation (equivalent to reproduction and inheritance in genetic evolutionary systems), imitation error (equivalent to mutation in genetic evolutionary systems) and any form of selection that causes some behaviours to be imitated more frequently than others (equivalent to natural selection in genetic evolutionary systems).

1.3 Existing Solutions

Subsequent theoretical work aimed at explaining the evolution of adaptive culture grew after the publication of Rogers' model in 1988 (Boyd and Richerson, 1988b; Henrich and Boyd, 1998; Wakano and Aoki, 2006; Richerson and Boyd, 2006; McElreath and Henrich, 2007; Enquist et al., 2008; Borenstein et al., 2007; Rendell et al., 2010b). These models claimed to have solved the question of how culture could be adaptive, generally referred to as the "Rogers Paradox", by including decision-making processes (Henrich and Boyd, 1998; Richerson and Boyd, 2006; Borenstein et al., 2007; Enquist et al., 2008; Mesoudi, 2011). The specifics of all these models are too numerous to be described in this introduction chapter (see Section 2.5 for more details). However, I can divide them into two groups. In one group, the origin of culture is explained by imitation biases, also known as transmission biases or critical social learning. The second group comprises models with some degree of individual learning in their social learning strategies, and where this form of mixed learning drives a process called guided variation.

In models with imitation bias, social learners imitate behaviours with higher-than-average fitness more frequently. This bias can be driven directly by a fitness-assessing criteria or by a proxy criterion that results in the selection of higher-than-average behaviours to imitate (Richerson and Boyd, 2006; McElreath and Henrich, 2007; Mesoudi, 2011). Current literature refers to the former as *direct-bias.* The latter includes different forms of biases like *prestige bias* (when imitating according to social status) and *conformity bias* (when imitating the most frequent strategy). From here onwards, I group all bias-dependent processes under the label *critical social learning* or simply *biased imitation.*

Critical social learning models solve the problem of phenotype-fitness de-correlation by incorporating an additional source of selection in the form of bias. The behaviours that effectively increase fitness are then copied more frequently. If this bias is strong enough, imitating becomes the best available strategy compared to innate actuation or individual learning. To evolve culture, these models require the pre-existence of traits allowing, directly or indirectly, a bias between high-fitness and low-fitness strategies.

The second group of models rely on *guided variation* which consists of improving upon behaviours that have been acquired by imitation. Most guided variation models introduce it as a form of individual learning. Social learners imitate behaviours from others and then increase their value by a process of trial and error. This can be considered an additional form of selection, one that occurs within each individual after poor-performing variations of the original behaviour are discarded in favour of the fittest one (Kameda and Nakanishi, 2002; Aoki et al., 2005; Wakano and Aoki, 2006; Borenstein et al., 2007).

Both groups of models (critical social learning and guided variation) are encompassed here under the label decision-making models. These models form the core of existing theoretical work attempting to explain the origin of culture. In this thesis I claim decision-making processes are not required because the fixation of social learners can be explained by their explorative advantage alone.

While researching existing literature related to the subject of the origin of culture, I identified a recent “gap” in the last six years. That is, except for rare exceptions, most references of relevant theoretical models have been published before 2011. In general, this is a consequence of a shift in focus within this multi-disciplinary field. In the last few years, theoretical researches have moved their attention to different subjects, and focused their efforts on explaining macro-trends in human culture, and characterising processes that occur in established cultural systems, rather than explaining the origin of culture itself. In general, recent literature focuses on specific cultural dynamics in advanced cultural systems (i.e., systems with sophisticated tools and languages) and the effect of known historical trends on the evolution of human culture (Pfau et al., 2013; Acerbi et al., 2014; Kolodny et al., 2016; Henrich, 2015; Boyd et al., 2013). This gap has been produced because, by the end of the 2000s decade, all models that extended seminal work (i.e., work published in the 1980s), had reached a consensus that we will show to be incorrect: *social learning through unbiased imitation is not sufficient to explain the origin of culture.* As mentioned before, the models produced in the 2000s decade (i.e., between 2000 and 2010) explain the evolution of culture by incorporating decision-making abilities for social learners, or alternatively, by incorporating individual learners in the population. In Section 2.5 a brief description of these models is provided.

1.4 Minimal Assumptions to Explain Cultural Evolution

The work in this thesis takes a different approach to previous models, I take an approach based on dual-inheritance theory. This theory frames culture as an evolutionary system with its own information substrate, that is: ideas and behaviours instead of genes (McElreath and Henrich, 2007). Evolutionary systems can be explained in algorithmic terms as optimisation programs containing key functions of replication, variation and selection (DeJong, 2002). Replication and variation functions for cultural systems are implicit in the nature of its substrate: transmissible ideas and behaviours. These ideas pass from one individual to another and on each transmission some errors are introduced. For culture to evolve as an evolutionary system, it must outperform its genetic counterpart when finding solutions. In my model, the genetic system is exclusively formed by innate actuators that follow genetically-coded instructions. Individual learning is not included in any of my simulations. In Chapter 3, I further discuss how this assumption simplifies my work and creates the most challenging conditions for culture to evolve, which in turn adds to the robustness of my conclusions.

I aim to explore the role of replication, variation and selection on the origin of culture without invoking decision-making processes (i.e., without invoking the ability of individuals to evaluate and compare the potential fitness values of different behaviours). To understand these properties, I use the concepts of variants and memes. *Memes are discrete units of cultural information that are ultimately expressed as behaviours and are transmitted by means of social learning.* Each meme can have variations in a population and these variations are referred to as *cultural variants* (Richerson and Boyd, 2006; Dawkins, 2006). Variants and memes are respectively the alleles and genes of cultural systems. In these terms, my model explains culture by finding conditions of reproduction, variation and selection where cultural variants outperform *genetic alleles* (i.e., the set of existing genetic variants for any specific gene), for a given behavioural challenge.

As explained, the classical interpretation of natural selection in reproductive terms alone does not allow for culture to evolve due to phenotype-fitness de-correlation effects. In this scenario, behavioural traits always evolve through genes. When the effect of natural selection is separated in reproductive and survival components, we can start conceiving combinations of these components where the selective disadvantage of culture compared to its genetic counterpart is lessened, specifically when the survival component is substantially larger than the reproductive one.

Under this condition, social learners merely experience the same selective pressure as their instinctive counterparts. To further claim that culture is a better evolutionary system, they need an advantage. This advantage comes from cultural transmission having a higher transmission error rate than genetic reproduction (Henrich et al., 2008; Huber et al., 2009; Shea, 2009); a pattern that has been observed in several animal populations (Henrich and Boyd, 2002; Boyd and Richerson, 1988a; Brown and Laland, 2003; Monfardini et al., 2014). It is then a combination of increased variation and specific selection components that allows cultural variants to improve faster than genetic alleles, and hence culture to originate from an initially pure genetic system. Chapter 3, 4, and 5 results substantiate these claims.

1.5 The Research Questions

The research questions of this thesis need to be contextually framed. First, I assume a sole evolutionary system was in place before the origin of culture: a genetic system. This system had a reproduction function based on genetic replication, a variation function based on mutation, and a selection function represented by natural selection (Mesoudi, 2011; McElreath and Henrich, 2007). The chance for a second evolutionary system to be instantiated came when a mutation for acquiring information by copying others entered the population, and with it, the first social learner. Cultural and genetic systems use different substrates of information, and transmit this information with different error rates and different frequencies of horizontal and vertical transmission. Considering this, my first research question is:

Research Question 1

What are the conditions of variation and selection in evolutionary systems that lead to the invasion of social learners and the emergence of culture without invoking individual decision-making capabilities?

In nature, social learning has both horizontal and vertical transmission pathways. In populations with overlapping generations, this means social learners can imitate any member at random in the population including their parents. Existing literature considers horizontal imitation a hindrance for the evolution of culture and attaches the problem of phenotype-fitness de-correlation to this form of social learning. I challenge this conclusion and argue that when comparing the two competing evolutionary systems (i.e., the cultural evolutionary system against the genetic evolutionary system), horizontal imitation can be an exploration aid that imports fitter trending variants (i.e., variants increasing in frequency in the population) into the cultural system regardless of their origin as genetic mutations or imitation errors. This feature is not available to innate actuators due to their exclusively vertical/genetic inheritance. I thus investigate if this difference contributes to the evolution of social learners, or whether social learning evolves exclusively due to its higher mutation rate. My second research question is therefore:

Research Question 2

Does horizontal transmission enhance the competitive ability of cultural evolution relative to genetic evolution and vertical cultural inheritance?

In previous models, a moderate rate of environmental change (i.e., a moderate change of the optimal solution over time) has been shown to facilitate the origin of culture, as it increases the frequency of social learners compared to *instinctive actuators* (i.e., agents whose behaviour is the product of genetic expression). However, in these models culture emerges in a mixed population of individual and social learners (Kameda and

Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007), with the frequency of individual learners at equilibrium being proportional to the rate of environmental change (Kameda and Nakanishi, 2002; Borenstein et al., 2007). These models conclude that for culture to evolve there should always be a minimal proportion of individual learners. Individual learners have a fixed fitness value as they are guaranteed to learn (at a cost) the optimal behaviour regardless of environmental conditions, therefore the presence of social learners does not increase the overall mean fitness of the population in these models (i.e., at equilibrium they will have the same fitness as individual learners).

My simulation-based approach challenges this conclusion and considers that culture can: (1) emerge by pure social learning regardless of the rate of environmental change, and (2) increase the overall mean fitness of the population by tracking environmental fluctuations better than a monomorphic population of instinctive actuators. To test this hypothesis, I look for answers to the following two questions using my simulation framework:

Research Question 3

Does environmental change increase the chance of social learners fixing in a population without individual decision-making capabilities?

and

Research Question 4

Does the fixation of social learners increase the overall mean fitness of a population under moderate environmental change?

In this context, I define the process of *fixation* as the complete invasion of one type of agent over all others, in a population. Logically, this process leads to a monomorphic population of the invading type (e.g., the fixation of social learners in a population that was initially formed by instinctive actuators leads to a monomorphic population of social learners).

The following sections will explain the aims of my thesis and the a-priori hypotheses that arise when considering these features.

1.6 The Aims of this Thesis

1. To identify conditions under which social learning can evolve due to the properties of cultural replication, variation and natural selection, without invoking decision-making processes such as critical social learning or guided variation.

2. To explain the existence of horizontal imitation in cultural systems. Not only as a collateral evolutionary contingency but as a trait with adaptive value (i.e., over and above the value of error prone vertical imitation).
3. To explain the effects of environmental change on the origin of culture and find conditions (if any) where social learners can fix and increase the overall mean fitness of a population compared to a monomorphic population of instinctive actuators.

1.7 The Hypotheses

Explanations for the evolution of culture that do not invoke decision-making processes must overcome the two problems mentioned in this chapter, that is phenotype-fitness de-correlation and social learning stagnation, by only using the distinctive properties of variation, inheritance and selection in each evolutionary system. My hypothesis is that this could be achieved by including an extended strategy space, independent rates of genetic and cultural variation, and by splitting selection into survival and reproduction components.

Decomposing selection into survival and reproduction, rather than following the classical view of selection acting on reproduction alone, allows social learners to imitate longer-living strategies more frequently. This attenuates the effect of phenotype-fitness de-correlation for cases where survival selection is larger than reproductive selection. Including transmission error rates for both genetic reproduction and social learning imitation and allowing these systems to explore an extended strategy space solves the problem of social learning stagnation. In this scenario, both types of agents have their own source of variation to track environmental changes if the need arises.

Regarding my second research question: horizontal imitation can be an exploration aid during early stages of the evolution of culture, as this form of transmission allows social learners to imitate strategies that are increasing in frequency in the population, likely because of their adaptive value, regardless of their origin as genetic mutations or imitation errors.

The ability to imitate these trending behaviours is not available to instinctive actuators as they can only inherit information passed vertically from their parents. In this way, horizontal imitation gives social learners a tool to incorporate solutions found by any individual in the population, while instinctive actuators (i.e., agents whose behaviour is the product of genetic expression) can only exploit solutions found by mutations in their own genetic lineage. My hypothesis is that this asymmetry is an advantage, as long as selective conditions attenuate the negative effects of horizontal transmission (i.e., if survival selection is stronger than reproductive selection).

Regarding my third and fourth research questions: if culture evolves because social learners find solutions faster than their genetic counterparts (i.e., instinctive actuators), environmental change must facilitate the evolution of this trait. This is because a changing environment (i.e., an environment that changes its optimal solution over time) is better exploited by a system with an explorative advantage. This system must surely be culture, not genetic evolution, as long as the high imitation error rate does not cause an error catastrophe (i.e., a process where high error rates destroy adapted variants faster than the evolutionary process can construct them).

Furthermore, I hypothesise that a monomorphic population of social learners can sustain culture in a changing environment, and that this population can track environmental changes better than a monomorphic population of instinctive actuators, therefore proving that unbiased social learning can be adaptive (i.e., that it can increase the overall mean fitness of the population).

1.8 Core Chapters: Claims and Contributions

1.8.1 Chapter Three

This chapter claims social learners can fix in a population when behavioural traits have a strong effect on survival compared to reproduction (i.e., when behavioural traits increase the lifespan of individuals but not their reproductive output) and the error rate associated with imitation is greater than genetic mutation. This conclusion is derived from the results of an individual-based simulation where selection is divided into survival and reproductive components, and where mutation and imitation error rates can be adjusted independently. In this simulation, social learners can only imitate horizontally and the extended strategy space has a single optimal solution.

The model consists of a steady-state genetic algorithm where individuals are represented by two bit-strings respectively storing their genotype and phenotype. Agents also include a switch as part of their genotype. When a new individual is introduced to the population it inherits the genotype and switch state of its parent. If the switch is “off”, it expresses its own genotype as its phenotype in a one-to-one map; and when “on”, the individual gets its phenotype from a random agent in the population. This switch is considered part of the genotype during reproduction but it is not present on the phenotype during expression. The evolution of culture is signalled by the invasion and fixation of the “on” gene.

Fitness and selection decomposition into survival and reproduction components allows for the frequency of horizontal imitation to be proportional to the fitness value of the trait being imitated. This is because individuals executing fitter strategies live longer and, as a consequence, are more likely to be imitated by others throughout generations.

This process by itself counteracts the effect of phenotype-fitness de-correlation caused by horizontal imitation, but it does not offer an advantage for cultural transmission when compared to genetic reproduction; non-imitators (instinctive actuators) that live longer also produce a greater number of offspring.

Including relatively high imitation errors (relative to mutation in genetic replication) solves the issue of social learning stagnation while adding an exploratory advantage for social learners. High error rates increase the chance of social learners finding the fittest solution at an early point of the evolutionary optimisation. The early advantage of social learners allows them to rapidly invade since fitter variants live longer and therefore are copied more frequently. This process eventually originates culture.

In a population of social learners, genetic information is not expressed in the phenotype; this means selection cannot act on genes. The resulting genetic drift makes it less likely for mutant non-imitators (i.e., social learners' offspring that mutate into instinctive actuators) to reinvade by instinctively expressing genetic information. Cultural variants keep evolving by imitation while the fitness of the genotype strings drifts towards poor-performing solutions. This process is defined as *genotype-phenotype disengagement* and it is considered a clear sign of the evolution of culture throughout this work.

High error rates and fitness decomposition are more parsimonious assumptions than decision-making processes as they explain the origin of culture without invoking sources of selection that depend on individual capabilities. The claims of this chapter thus lower the theoretical requirements of species that could evolve culture.

In this chapter, I also discuss the potential limitations of the model as a tool to predict the emergence of culture in natural populations while emphasising its value as a theoretical contribution that brings insight into the exploration and exploitation advantages of social learning evolutionary systems when compared to genetic evolution.

1.8.2 Chapter Four

This chapter claims that horizontal imitation is adaptive under strong survival selection and challenges the idea that culture evolves in my model exclusively due to the high error rate of imitation. The results in Chapter 3 seem to suggest that social learning invades only because of its high error rate during imitation. If this were true, exclusive vertical imitators would invade when given the chance, as vertical imitation can exploit both fitness components, survival and reproduction, while still having an explorative advantage with its high error rate. In this chapter, I thus challenge the notion that horizontal transmission is just a hindrance compensated by the exploratory advantage of a relatively high imitation error rate.

To test my hypothesis, I extend the original model in Chapter 3 where the imitation switch could only code for two states: non-imitation (innate expression of genes) or indiscriminate horizontal imitation. This limited set of options does not include a frequently observed form of imitation in nature: vertical imitation (i.e., imitating the phenotype of ones' parent) (Heyes, 1994). The extended version consists of a model where the switch can take three different values corresponding to: non-imitation, horizontal imitation only and vertical imitation only. I show that conditions for the evolution and maintenance of culture by horizontal imitation exist even when vertical imitation is an option. Analyses of the results expose an asymmetry between horizontal and vertical imitators in terms of their access to fit models to imitate. This asymmetry is accentuated under strong survival selection. The asymmetry is explained by vertical imitators only obtaining trending variants (i.e., incipient, fitter than average and increasing in frequency) from their parents which are vertical imitators themselves. Horizontal imitators, on the other hand, quickly incorporate trending variants regardless of the type of imitator that expresses it, which gives them access to phenotypic variation across the population in mixed populations of vertical and horizontal imitators. When compared to instinctive actuators, this access is a secondary explorative advantage of social learners, different from that produced by a high imitation error rate.

A further extension of the model explores the effects of varying the error rate of vertical imitation compared to horizontal imitation. Reinforcement and local proximity between parents and offspring is considered as an explanatory argument for reduced error rates in vertical imitation. The effect of this extension is a sequential change of horizontal and vertical imitators that responds to different stages of the optimisation process. During the initial optimisation stage, solutions coded in the bit-strings are far from optimal and high error rates in horizontal imitation have an explorative advantage in terms of finding fitter ones. In a later stage, solutions are closer to the optimum and high error rates become destructive; vertical imitators are then able to invade given their low error rate (i.e., high fidelity). This change in error rate from horizontal imitation to vertical imitation cannot occur from horizontal imitation to non-imitation due to the aforementioned process of genetic drift and genotype-phenotype disengagement.

1.8.3 Chapter Five

In this chapter, I claim that environmental change has a positive effect on the evolution of unbiased social learning, as it increases the area of the parameter space of my model (i.e., combinations of survival selection coefficient, reproductive selection coefficient, imitation error rates and genetic mutation rates) where a cultural evolutionary system emerges. This claim is in agreement with existing work, however I explain it through a different process compared with other models in the literature.

In my model, social learners can completely fix under moderate rates of environmental change, and do not require individual learners to do so (i.e., they do not need to imitate individual learners' behaviours in order to invade instinctive actuators). This claim contrasts with existing models where the emergence of culture can only occur in populations with a mixture of social and individual learners (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

The claims of this chapter are based on a model extension where the optimum strategy can change with a given rate through a simulation run. This rate is controlled by a variable that quantifies the chance of a bit flip in the optimal string (i.e., the string against which all phenotypes are evaluated to calculate their fitness values) in each generation, where a generation corresponds to $N/2$ iterations, and where N is the population size. Analysis of the results shows how moderate environmental change slows down the rate of adaptation for both genetic and cultural evolution. This is beneficial for social learners, as genotype-phenotype disengagement can occur before phenotypes evolve genetically above the critical value where the high error rates of social learners are no longer beneficial (i.e., where high error rates produce more deleterious phenotypes than advantageous ones compared to genetic mutation).

Further analysis shows that moderate environmental change makes social learning adaptive at the population level (i.e., under environmental change, social learners increase the overall mean fitness of the population compared to a monomorphic population of instinctive actuators). This is an important finding, as it challenges the existing assumptions regarding the adaptive value of unbiased imitation (Richerson and Boyd, 2006; Rogers, 1988; Cavalli-Sforza and Feldman, 1983). I show that for a range of rates of environmental change, genetic evolution cannot evolve phenotypes beyond a given fitness value, but social learners can due to (1) high imitation error rates and (2) horizontal transmission. In this range, the average phenotype value maintained by cultural evolution (i.e., by a population of social learners) is higher than the value that could be maintained by a population of instinctive actuators. I therefore conclude that the evolution and fixation of unbiased social learners can be adaptive.

My analyses find culture to be the dominant evolutionary system for moderate rates of environmental change. Further increases of this rate make constructive evolution unviable, as rapidly changing environments are impossible to track by either evolutionary system (i.e., cultural or genetic). Under this condition, the population cannot produce consistent improvements of the phenotypes (i.e., adaptive evolution does not occur).

1.9 Publications From This Thesis

1.9.1 Publications From Chapter Three

Published:

1. Gonzalez, M., Watson, R. A., Noble, J., and Bullock, S. (2014). The origin of culture: selective conditions for horizontal information transfer. *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems.* MIT Press. 408-414.

Accepted:

1. Gonzalez, M., Watson, R., and Bullock, S. (2017). Minimally sufficient conditions for the evolution of social learning and the emergence of non-genetic evolutionary systems. *Artificial Life.* MIT Press. (Accepted for publication).

1.9.2 Publications From Chapter Four

In Preparation:

1. Gonzalez, M., Watson, R. and Bullock, S. (2018). When viability selection is strong, horizontal imitation is dominant over vertical imitation in a purely Darwinian model of the evolution of social learning.

1.9.3 Publications From Chapter Five

In Preparation:

1. Gonzalez, M., Watson, R. and Bullock, S. (2018). Under moderate environmental change unbiased social learning increases the overall mean fitness of a population.

Chapter 2

Cultural Evolution and the Challenges of Explaining its Origin

Culture is not easy to define. Discussions on what culture is and what it is not have been explored in countless articles in biology, anthropology, and the social sciences (Jahoda, 2012; Heyes and Galef, 1996). Some of these definitions consider culture as a unique feature of the human species:

"Culture consists of patterns, explicit and implicit, of and for behaviour acquired and transmitted by symbols, constituting the distinctive achievement of human groups, including their embodiments in artefacts; the essential core of culture consists of traditional (i.e. historically derived and selected) ideas and especially their attached values; culture systems may, on the one hand, be considered as products of action, on the other as conditioning elements of further action."(Kroeber and Kluckhohn, 1952).

"[We] think of culture as a dynamically changing environment that is transformed by the artefacts created by prior generations... an artefact is an aspect of the material world that has been modified over the history of its incorporation into goal-directed human thought and action."(Cole, 1996; Cole and Parker, 2011).

Here we explain why our working definition - which encompasses human and non-human "culture" - has been selected and how it differs from others, with the intention to frame our research within this broad range of fields of study. In this thesis, we use a formal definition based on the theory of cultural evolution developed by Robert Boyd and Peter Richerson (Richerson and Boyd, 2006). According to this theory, *culture* is:

A set of behavioural traits that are not the direct result of genetic expression but the product of an evolving pool of variants (i.e., behaviours and ideas) stored and imperfectly transmitted within and/or between overlapping generations by means of social learning.

To frame our work within the scope of existing theory, we must establish that the term "culture" is here used to refer to all forms of evolutionary systems that emerge as a consequence of social learning information transfer between individuals, regardless of how simple these systems are. We consider cumulative primate culture (including human culture) as a sophisticated example of a cultural system (Richerson and Boyd, 2006), but leave the distinctions between human and animal cultures outside the scope of this thesis.

We are aware that some schools of thought object to classifying non-human, socially learnt animal behaviours as examples of "real" culture (Laland and Galef, 2009). However, there is little contention that all these instances share a general mechanism during their emergence, most certainly involving some form of social learning (Richerson and Boyd, 2006; Mesoudi, 2011; Laland and Galef, 2009). In this thesis, we have selected our working concept of culture precisely because it is general enough to explore questions regarding the emergence of these first non-genetic evolutionary systems based on social learning.

2.1 Definitions of Culture and Social Learning

Historically one of the first formal definitions of culture comes from the late 19th century when the popular anthropologist Edward Tylor defined it as: *That complex whole which includes knowledge, belief, art, law, morals, custom, and any other capabilities and habits acquired by man as a member of society* (Tylor, 1920). This definition, like many from this period, excluded the possibility of culture in animals and failed to offer useful properties for modelling its origin and evolution.

Over time, definitions of culture became more refined and generalizable with the introduction of terms like "knowledge" and "information" (Mead and Metraux, 1953). According to these new definitions, culture, just like other bodies of information, could be encoded, transmitted and ultimately quantified (Mead and Metraux, 1953; Richerson and Boyd, 2006).

Nowadays, concepts of culture reference evolutionary aspects that explain its change over time (i.e., its evolution) (Avital and Jablonka, 2005; Richerson and Boyd, 2006). In our definition, we name the process of social learning in order to include the key evolutionary functions of reproduction, inheritance, variation, and selection (Richerson and Boyd, 2006). Social learning is a catchall term encompassing several mechanisms of information transmission between individuals. During these transmissions errors can occur, and behaviours (i.e., variants) may be transmitted at different rates (Richerson and Boyd, 2006; Mesoudi, 2011).

Social learning is formally defined as *the set of learning mechanisms that result in the effective transference of behavioural traits between individuals in a population* (Heyes, 1994; Richerson and Boyd, 2006). In other words, if culture is said to be a population level characteristic (i.e., a pool of information, or the sum of all behavioural variants), social learning is the individual trait that underpins its emergence and maintenance (Richerson and Boyd, 2006).

Different forms of social learning have been identified in nature (Heyes and Galef, 1996). Some examples are *local enhancement*, where individuals performing a behaviour create local conditions that facilitate the independent discovery of that behaviour by others (e.g., nut-cracking chimpanzees accumulate stones of the right size and shape to crack nuts, in a small area); *social induction*, where an individual causes a specific behaviour to develop in others even when he is not performing the behaviour himself; and *observational learning* where basic imitation between individuals happens through visual cues (Heyes, 1994; Heyes and Galef, 1996; Jablonka and Lamb, 2006). These forms of social learning are mechanistically different, and some of them require specific environmental interactions and cognitive processing of information (Jablonka and Lamb, 2006).

In our model, social learning mechanisms are only relevant to the extent that they involve imitation, which can be thought of as the most basic form of social learning. Here we define imitation as *all instances of model copying that have as a consequence one individual performing another's behaviour* (Richerson and Boyd, 2006; Mesoudi, 2011). We further assume that all forms of social learning are subject to a source of variation in the form of imitation errors caused by imperfect transmission of information (Jablonka and Lamb, 2006; Mesoudi, 2011).

2.2 The Evidence for Culture and Social Learning

Due to its complex nature and the way it interacts with genes, conclusive evidence for the existence of culture as an evolutionary system is difficult to obtain. In order to frame our definition of culture, and by extension our work, within the known body of experimental and empirical evidence, we here outline well-studied phenomena that support its existence as a body of evolving information independent from genes.

1. Twin and Sibling Studies:

   Genetically identical individuals have been used for research on culture for many years as they offer genetically-controlled cases where behavioural differences can be partially attributed to social learning (Vernon et al., 1997; Bouchard et al., 1990). By comparing twins and siblings who develop in similar physical environments (i.e., that have similar non-social environmental inputs), behavioural geneticists have estimated that most cognitive traits have a genetic heritability of only around

50%. Furthermore, large-scale studies agree with this estimation, showing that only 50% of the variation in the behaviour of individuals within a society can be attributed to genes (Plomin et al., 2003). Such a conclusion leaves 50% of unexplained variation that can be attributed to a combination of culture, driven mainly by social learning; and behavioural plasticity, as a product of environmental interactions.

2. Genetic and Behavioural Variation Patterns Don't Match:

   Around 94% of the genetic variation in humans comes from within societies rather than between societies. In other words, the average genetic distance between two societies is very small compared with the genetic distances that exists within those societies (Rosenberg et al., 2002). In comparison, most of the variation in social norms and practices exists between societies rather than within them; a pattern that is observed even when compared societies live under similar environmental conditions (Bell et al., 2009). The mismatch between these variation patterns shows evidence, albeit non-conclusive, for the existence of culture as a body of behavioural information residing outside of our genome.

3. Cultural Adoption by Immigrants

   Immigrants adopt cultural norms and practices from their host population within a generation, and their children are more culturally similar to the members of the hosting culture than the parental one (Norenzayan et al., 2002). This similarity is intensified in the second generation even when couples are formed by first generation members of the same immigrant group (Norenzayan et al., 2002). This transitional pattern points to cultural traits gradually permeating genetic lines through vertical and horizontal transmission processes.

4. Fast Changes in Opinion Formation

   Studies of cultural fads, fashion and marketing trends, show transitions in opinion formation can take as little as months or even weeks to occur (Mesoudi, 2011; Bentley et al., 2004). These opinions affect behavioural practices and are transmitted trough social networks in very short periods of time. Genetic inheritance is not considered a potential explanation, given that the rate of change for these opinions occurs in time scales between two and three orders of magnitude shorter than generational times (Bentley et al., 2004).

5. Innate Mechanisms

   Mirror neural systems are instinctively wired for imitation of actions performed by other members of the population (Huber et al., 2009). Studies in ethology, human psychology and neuroscience suggest they are part of our most primitive neural circuitry, and their presence in relatively simple species of fish and rodents suggests social learning mechanisms and perhaps primitive forms of culture can exist in more species than initially thought (Huber et al., 2009; Ferrari et al., 2003). The existence of mirroring does not directly prove the presence of culture as a population-level phenomenon, but it is an individual-level requirement for basic forms of social learning.

2.3 Human Culture and Non-Human Culture

According to certain schools of thought in anthropology and social science, "real" culture is exclusively found in the human species (Tylor, 1871; Case, 1927; Jablonka and Lamb, 2006). In these fields, culture is restricted to humans by establishing cognitively demanding requirements for its existence: selective biases, guided variation, mixed learning strategies, abstract symbolic representation (i.e., language) and the presence of transmission artefacts (i.e., material culture) (Baldwin et al., 2005; McGrew, 1998; Tomasello, 1999). In order to answer the research questions in our work, we have invoked a more general definition of culture, a definition we believe to be necessary if we are to explain its origin (or the origin of a primitive form of it). In other words, we consider that in order to explore the origin of something as complex as culture, one must have a definition that allows the existence of proto-forms of it.

Defining culture in simple evolutionary terms allows us to include many groups of animals in our example cases (Whiten and van Schaik, 2007; Laland and Galef, 2009; van de Waal et al., 2013; Laland and Janik, 2006; Laland and Hoppitt, 2003; Avital and Jablonka, 2005). Evidence of culture as information transmitted and maintained by social learning alone is found in birds, where song learning occurs by means of imitation (Jenkins, 1978; Heyes, 1994; Nelson, 2000); in chimpanzees, with well-known cases of tool use and primitive forms of material culture (McGrew, 1998); in Cetaceans, imitating cooperative hunting practices and mating calls (Rendell and Whitehead, 2001); and in others (Laland and Galef, 2009; Laland et al., 2011).

There is little contention that all instances of culture need to share some general mechanisms for their emergence (Richerson and Boyd, 2006; Mesoudi, 2011; Laland and Galef, 2009). Existing models tend to assume that in order for such a mechanism to lead to culture, decision-making attributes, such as learners choosing fitter individuals to imitate or expressing improved versions of the variants they acquire, are required (Rogers, 1988; Henrich and Boyd, 1998; Enquist et al., 2008). This is a reduced list of requirements compared with the requirements made by definitions of "real" culture in anthropology and social science, but we intend to show that even this reduced list is over specific and unnecessarily limiting.

Reducing the number of requirements that basic forms of culture must satisfy not only broadens the range of species for which we could expect culture (and explain it), but also contributes, by means of contrast, to a better understanding of the effects of individual cognitive traits in more sophisticated cases (i.e., in human culture).

The debate on what is true culture and what counts only as "animal traditions" has been an on-going discussion in the field (Laland and Hoppitt, 2003; Laland and Janik, 2006; Whiten and van Schaik, 2007). In this work, we aim to further expand the range of species that could develop the most basic form of culture and will assume that the evolutionary processes that gave rise to early versions of existing cultural systems are qualitatively the same. We also assume that further evolution (after the installment of a basic form of culture) is responsible for the cognitively demanding traits present in species with sophisticated cultural systems (i.e., humans and chimpanzees) (Jablonka and Lamb, 2006; Richerson and Boyd, 2006).

2.4 The Challenges of Explaining the Evolution of Social Learning: Seminal Models

Mathematical modelling has been key in the development of cultural evolutionary theory, from basic mathematical expressions that explain the evolution of social learning (Cavalli-Sforza and Feldman, 1983; Rogers, 1988; Henrich and Gil-White, 2001) to models and simulations of opinion formation in social networks (Moussaid et al., 2013; Holyst et al., 2000; Watts, 2007). These tools allow researchers to explore the dynamics of hypothetical scenarios and compare theories regarding adaptation, speed of change and the effects of population structure on the evolution of culture.

When explaining the origin of culture, current mathematical models state that simple unbiased imitation does not offer any advantage for social learners compared to instinctive actuation (i.e., where behaviours are the result of genetic expression) (Rogers, 1988; Richerson and Boyd, 2006). In this section, we deconstruct two seminal models to show why the current literature has reached this conclusion.

2.4.1 The Cavalli-Sforza and Feldman Model

Generally accepted as an adaptive trait at a population level (White, 1959; Durham, 1976), social learning passed under the radar of theoretical scrutiny until the early years of the 1980s. In 1983, Cavalli-Sforza and Feldman used a model to show that when competition takes place between instinctive actuators and social learners using unbiased imitation, the instinctive actuators would always fix in the population, driving social learners extinct.

The model assumes there are two types of behavioural strategies that may be carried out (one fitter than the other) and that there are two ways of acquiring these behaviours:

they can either be learned by observation or inherited genetically. Non-imitators (also referred to as *instinctive actuators*) inherit their behavioural strategy vertically from their parents in the form of a genetically encoded instinctive behaviour. By contrast, social learners obtain their behavioural strategy by imitating the behaviour exhibited by a member of the current population (i.e., potentially their own parents but possibly some other random member of the population). The probability of imitating a parent (i.e., vertical imitation) rather than an individual selected at random (i.e., horizontal imitation) is a model parameter. While social learners acquire their behaviour via imitation in this way, they inherit their acquisition strategy (to be an imitator) genetically from their parents. The successful evolution of culture in this model is defined as the fixation of social learners.

For cross-referencing purposes, the following model description will use the same symbols and notations as the ones used in the original paper (Cavalli-Sforza and Feldman, 1983) even when some of the notation is non-standard for a time-stepped model. We warn the reader that the following model description is mathematically dense and requires attention to the equations between paragraphs in order to appreciate the problems highlighted by the model.

2.4.1.1 Model Description

Let us assume a population entirely formed by social learners, use v to represent the proportion of these that has learned the fitter behavioural strategy, and use u to represent the remaining proportion of social learners (i.e., those that have not acquired the behaviour) so that $u = 1 - v$. Following the original paper's terminology, we call v the proportion of *skilled* social learners and u the proportion of *unskilled* social learners.

Let b represent the probability that social learners imitate their parents, and let f be the probability that an offspring that chooses not to imitate their parents imitates an individual chosen at random from the population. Note that with probability $(1 - b)(1 - f)$ an individual offspring will neither copy their parent nor a random member of the population and will therefore remain unskilled. We then calculate the respective frequencies for u and v after imitation as u_0 and v_0 in Equations 2.1 and 2.2. The next-generation frequencies are calculated in two steps; first imitation and then selection, the following equations only describe the first of these steps.

$$u_0 = [u + (1 - b)v](1 - fv) \tag{2.1}$$

$$v_0 = bv + fv[u + (1 - b)v] \tag{2.2}$$

In these equations u_0 is calculated as the existing proportion of unskilled social learners, u, plus the offspring of skilled parents that did not receive the information from their parents, $(1-b)v$. These individuals remain unskilled if they fail to copy skilled behaviour from the population, an outcome that occurs with probability, $1 - fv$. Similarly, the calculation of v_0 sums the proportion of social learners that received information from their skilled parents, bv, with the proportion that were unskilled ($[u + (1 - b)v]$) but acquired the behaviour from horizontal imitation $fv[u + (1 - b)v]$.

Let us now allocate a fitness advantage of s to the skilled behaviour so the relation between skilled and unskilled social learners' fitness is $(1 + s) : 1$. The next-generation frequencies after imitation and selection are then expressed as v' and u' in Equations 2.3 and 2.4.

$$v' = v_0 \frac{(1 + s)}{1 + sv_0} \tag{2.3}$$

$$u' = 1 - v' \tag{2.4}$$

These equations show two points of equilibrium (i.e., where $v' = v$): one in the case of extinction for skilled social learners $v_0 = 0$, and the other corresponding to the fixation of skilled social learners if the condition $(1+s)(b+f) > 1$ for $f < 1$ and $b < 1$ is fulfilled.

Equations 2.3 and 2.4 describe the simple dynamical model for the spread of cultural variants in a population of social learners. To compare this social learning system with its genetic (i.e., instinctive) counterpart, we add a third type of individuals: instinctive actuators.

Instinctive actuators are skilled on an instinctive level (innates) and do not depend on imitation to acquire advantageous behaviours (in this model unskilled instinctive actuators were not considered, as this type is trivially destined to extinction). We use w to represent the proportion of skilled instinctive actuators. Because this type does not imitate, we will skip the first of the two steps required to calculate the social learners' frequencies (i.e., imitation). If we then consider only the second step (i.e., selection) for the calculation of w' and also take into account that innate individuals can act as models for social learners to imitate, the recursion system for all three frequencies (i.e., skilled social learners, unskilled social learners and instinctive actuators) is:

$$\phi u' = [u + (1 - b)v][1 - f(v + w)] \tag{2.5}$$

$$\phi v' = (1 + s)\{bv + f[u + (1 - b)v](v + w)\} \tag{2.6}$$

$$\phi w' = (1+s)w \tag{2.7}$$

Here ϕ is a normalisation term defined by the sum of all the right sides in the system of equations:

$$\phi = 1 + s\{bv + w + f(1-u)[u + (1-b)v]\} \tag{2.8}$$

This system has a trivial equilibrium point ($w' = w$, $v' = v$, $u' = u$) at $w = 0$ which leads to the scenario described in Equations 2.3 and 2.4 where the population is only composed of social learners.

In all possible equilibria, the advantageous behaviour fixes. This can be proven using Equation 2.7, where $\phi w' = (1+s)w$, an expression that makes $1 + s = \phi$ a requirement for $w' = w$, hence the population is completely formed by skilled individuals at equilibrium. The following analysis assesses under which conditions this invasion is achieved by instinctive actuators or by skilled social learners (if any).

Considering the condition at equilibrium $1 + s = \phi$, Equation 2.8 becomes:

$$1 = bv + w + f(1-u)[u + (1-b)v] \tag{2.9}$$

To find the frequency of social learners at equilibrium, we then need to evaluate Equation 2.6 for $1 + s = \phi$ while setting $v' = v$:

$$v = bv + f[u + (1-b)v](v + w) \tag{2.10}$$

To satisfy Equations 2.9 and 2.10, v must be equal to $1 - w$ in equilibrium and therefore $u = 0$. From Equation 2.5, we can prove that $v = 0$ for $b > 0$ and $f < 1$ if $u = 0$. If $v = 0$, w must then be 1 in any equilibrium for $b > 0$ and $f < 1$.

In other words, instinctive actuators will prevail and fix in a population with any degree of horizontal imitation and only two conditions can sustain polymorphic populations of instinctive actuators and skilled social learners. The first one is when $b = 1$, which means that vertical transmission for imitators occurs in every reproductive event and therefore no skilled social learner will be involved in horizontal imitation. The second is when $f = 1$ and $u = 0$, a condition where every horizontal imitator is guaranteed to obtain the advantageous behaviour when selecting at random a model to imitate.

A further analysis for scenarios with $b < 1$ and $f < 1$ proves that instinctive actuators will also fix in the population. Using Equations 2.7 and 2.8 to calculate the change of frequency for instinctive actuators $\frac{w'}{w}$ we have:

$$\frac{w'}{w} = \frac{1+s}{1+s[bv+w+fu(1-u)+fv(1-b)(1-u)]} = \frac{1+s}{1+sB} \tag{2.11}$$

For conditions where $0 < b < 1$ and $0 < f < 1$, the previous equation results in a constant increase of w. Notice that the abbreviated term B ($B = [bv + w + fu(1-u) + fv(1-b)(1-u)]$) is always less than 1 for this condition.

In this model, instinctive actuators have guaranteed inheritance and fitness advantage s, hence none of the two types of social learners can match their growth. On one side, unskilled social learners do not possess the fitness gain of skilled individuals. On the other side, skilled social learners do not have guaranteed transmission of their behaviour, which translates into a fitness penalty proportional to the frequency of unskilled social learners in the population (i.e., potential unskilled models to imitate).

Cavalli-Sforza and Feldman go on to execute a local stability analysis for monomorphic populations at the corners of the trivariate system (u, v, w):

1. Local stability for $u = 1$

   Unstable equilibrium. Both instinctive actuators w and skilled social learners v increase for all values of $s > 0$.

2. Local stability for $w = 1$

   This is a locally stable attractor, even for small values of s and values of b and f close to 1.

3. Local Stability for $v = 1$

   This is a locally unstable equilibrium. In the neighbourhood of $v = 1$, w increases if:

$$bv + fv(1-bv) < 1 \tag{2.12}$$

   A condition that is always satisfied for $0 < b$ and $f < 1$.

Figure 2.1 shows four simplexes for the trivariate system (u, v, w) with their respective vector fields. Red dots and/or lines represent stable attractors in these simplexes. Each plot corresponds to a different combination of b and f values. In Figures 2.1(A) and 2.1(B), the system converges towards the w apex (representing the invasion of skilled instinctive actuators), this convergence occurs for all combinations of b and f where $b < 1$ and $f < 1$. In Figures 2.1(C) and 2.1(D), the two singular cases where all social learners have either guaranteed vertical transmission $b = 1$ or guaranteed horizontal transmission $f = 1$, a continuum of polymorphic equilibria along the $[w, v]$ edge is possible.

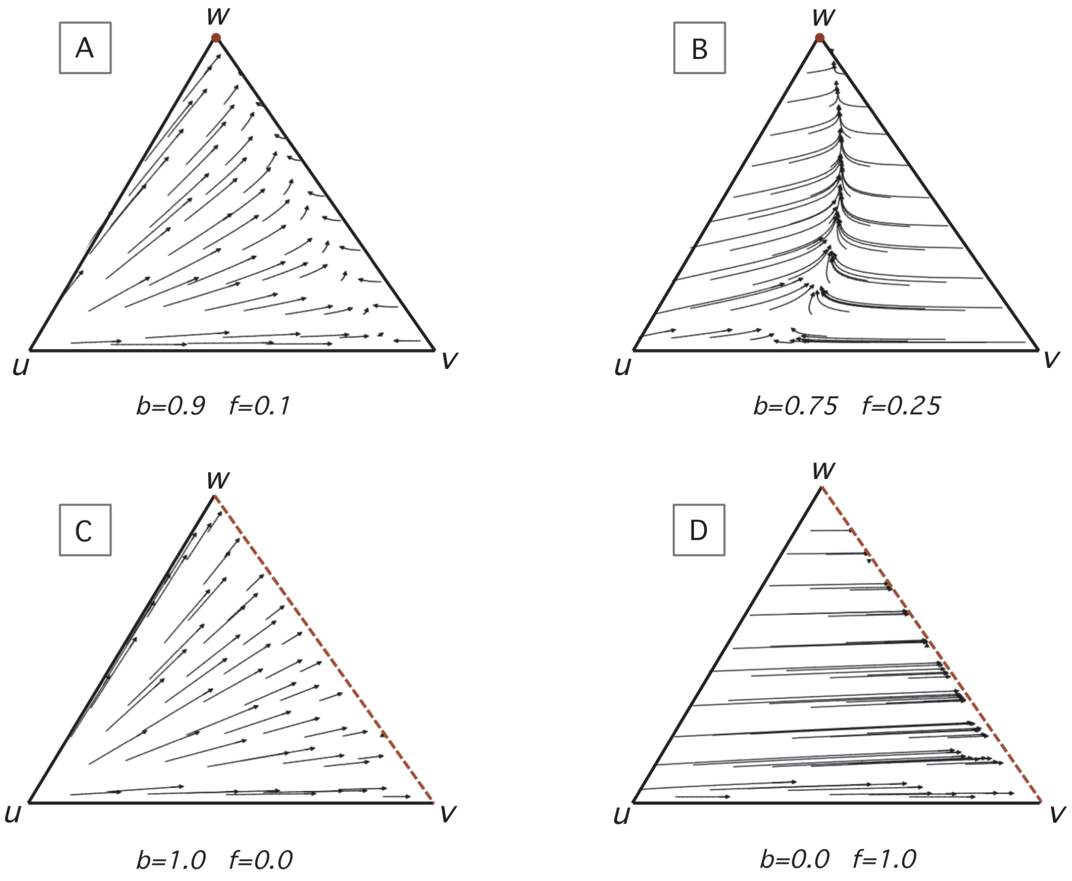

FIGURE 2.1: Four simplexes with their respective vector fields showing the Cavalli-Sforza and Feldman model dynamics for different combinations of vertical (b) and horizontal (f) transmission. The only stable attractor in simplexes A and B is denoted by a red dot at the vertex w. Simplexes A and B are representative cases of all conditions where $b < 1$ and $f < 1$. In simplexes C and D, the edge $[w, v]$, represented by a red dashed line, is an attractor. Simplexes C and D are singular cases. This is an original figure produced using the Cavalli-Sforza and Feldman model equations; this type of figure is not shown in their paper.

In the introductory chapter, we considered horizontal imitation to be a hindrance for cultural evolution as it implicitly includes a cost for social learners; we refer to this issue as *phenotype-fitness de-correlation*.

Phenotype-fitness de-correlation is perfectly exemplified by the Cavalli-Sforza and Feldman model. When the offspring of a social learner imitates other individuals, it risks imitating an unskilled member of the population. In other words, according to Cavalli-Sforza and Feldman, any degree of horizontal imitation will always induce a fitness value closer to the average of the population and hence a cost compared to fit instinctive actuators.

2.4.2 The Rogers' Paradox

In 1988, in a similar fashion to Cavalli-Sforza and Feldman, Rogers showed with a simple model that even if it is not wrong to assume that social learning could be adaptive, the reason for its adaptive value lies beyond the basic process of unbiased imitation. In contrast to previous models, Rogers focused on the cost of losing track of environmental conditions when social learning fixes in populations living in changing environments (Rogers, 1988).

Rogers' mathematical model describes a hypothetical population in an environment that may be in one of two possible states. These states alternate over time with a given frequency. The individuals in this population could either be social learners who imitate the behaviour of randomly selected members of the population, or individual learners who pay a small cost (e.g., the cost of learning by trial-and-error) to always acquire the appropriate behaviour for the current environmental state (Rogers, 1988).

For any initial non-zero frequencies of individual and social learners, and as long as there is some marginal alternation of environmental states over time, Rogers' model converges to a mixed-strategy population at equilibrium. Because the fitness of individual learners is not density dependent, the average fitness value of any social learner in this mixed-strategy equilibrium must therefore be equal to the fitness value of an individual learner. Rogers then concludes that social learning driven by unbiased imitation is not adaptive.

2.4.2.1 Rogers' Model

Let us consider a population of individuals that can each perform one of two behaviours and live in an environment which can change between two states, each behaviour being respectively advantageous for one of the two environmental states but not for the other. The following table contains the fitness value for each combination of behaviour and environmental state:

Environmental State	0	1
Behaviour 0	$\omega + b$	$\omega - b$
Behaviour 1	$\omega - b$	$\omega + b$

Here ω represents the base fitness value for every individual, regardless of their behavioural match with the current environmental state, and b and $-b$ represent the bonus or penalty associated with a match or mismatch between environment and behaviour.

Individual learners, with fitness represented by ω_i, are guaranteed to match their behaviour to the current environmental state, but in order to do so, they incur in a fitness cost c proportional to b:

$$\omega_i = \omega + b(1 - c) \tag{2.13}$$

To calculate the fitness of social learners, we need to calculate how likely it is that the current behaviours in the population (i.e., the models available to be imitated) will match the current state of the environment. As a start, let us agree that all imitated behaviours in social learners ultimately came from a chain of imitation that started with a social learner imitating an individual learner a number of generations ago.

Let τ be the number of generations ago that the first imitator in the chain got its then-environmentally-matching behaviour. The probability that a behaviour was obtained from an individual learner τ generations ago can be expressed as $p^{\tau-1}(1-p)$. With p representing the proportion of social learners in the current generation.

Changes in the state of the environment in the past τ generations will cause a reduction in the likelihood current behaviours matching the current environmental state. If the environment has not changed in the past τ generations, the fitness gained by imitating an individual that obtained its behaviour from an imitation chain starting τ generations ago will then be equal to b.

The fitness of a social learner in the current generation is therefore calculated by multiplying b by the probability of non-change for all values of τ (i.e., from one to infinity). The probability for a specific value of τ is expressed as $(1-\mu)^{\tau}$ where μ is the probability per generation that the environment undergoes a change in state. The fitness of social learners ω_s is then expressed as:

$$\omega_s = \omega + b\sum_{\tau=1}^{\infty} p^{\tau}(1-p)(1-\mu)^{\tau} \tag{2.14}$$

Which leads to:

$$\omega_s = \omega + \frac{b(1-p)(1-\mu)}{1-p(1-\mu)} \tag{2.15}$$

Having now constructed expressions for the fitness of both individual learners and social learners, the calculation of the proportion of social learners at equilibrium ($\hat{p}$) can be derived by setting $\omega_i = \omega_s$. By solving p for $\omega_i = \omega_s$, we arrive at $\hat{p}$ as:

$$\hat{p} = 1 - \frac{(1-c)\mu}{(1-\mu)c} \tag{2.16}$$

The proportion of social learners at equilibrium will therefore be greater than zero as long as $(1-\mu) > (1-c)$, which means that as long as the rate of environmental change per generation is less than the cost of individual learning, a non-zero proportion of social learners will remain.

In this mixed population, the fitness value of social learners is equal to the fitness value of individual learners. Individual learners' fitness is not density dependent, therefore the average fitness value in the entire mixed population is not different to the average fitness value of a monomorphic population of individual learners. In other words, not only is social learning not adaptive at the individual level (to the extent that it can achieve fixation), but also culture cannot be said to be adaptive at the population level.

The plot in Figure 2.2 shows the fitness values for each type vs. the proportion of social learners, p. A third curve represents the average fitness of the population. The point where all three curves cross denotes the equilibrium point of the system. Notice how this point has the same fitness value as the constant value for individual learners.

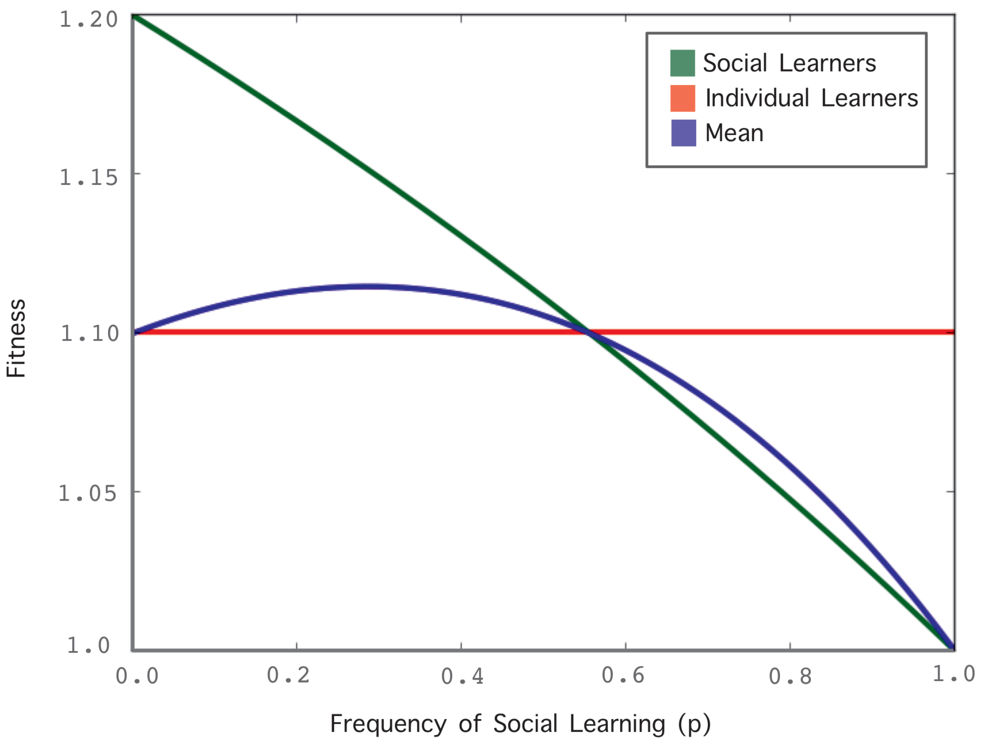

FIGURE 2.2: Individual learners' fitness, social learners' fitness and weighted mean fitness plotted for all frequencies of social learners (p). The point where the three lines cross represents the equilibrium point of the system. The weighted mean fitness (M) is calculated as $M = \omega_s p + \omega_i(1-p)$ being p the proportion of social learners. This figure was reproduced independently using Rogers' model, the result is identical to Figure 1 in his paper (Rogers, 1988) ($\omega = b = 1$, $c = 0.9$ and $\mu = 0.8$).

In a similar way to the Cavalli-Sforza and Feldman model, Rogers explains why culture is not adaptive when implemented via unbiased imitation alone. Both models have limitations as they only consider two behaviours and two environmental states, and do not include genetic mutation or imitation errors. We claim these simplifications create a restricted scenario where social learners cannot access environmentally matching strategies without the presence of individual learners. The lack of access to matching states in a monomorphic population of social learners causes what we've called *social learning stagnation*. We consider this process an artefact introduced by the simplifications of existing models.

In Chapter 3 we show that the conclusions of models like Rogers' do not hold when mutation rates for genetic replication and imitation error rates for cultural imitation are introduced. The model in this chapter also includes an extended strategy-space where the parallel evolutionary systems (i.e., genetic and cultural) can compete with one another to find fitter solutions. This feature contrasts with existing models where only two states represent the whole strategy-space (Rogers, 1988; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

The models of Rogers and Cavalli-Sforza and Feldman were not the only ones to highlight the issues of explaining the origin of culture in the 1980s (Henrich et al., 2008). In more recent models, these issues are also identified, but only solved when invoking decision-making processes (see next section for an extended explanation). In this thesis, we have selected Rogers' and Cavalli-Sforza and Feldman's models as clear and simple examples in order to frame the need for the simulation approach explained in Chapters 3, 4, and 5. When using this simulation approach, we show how the origin of culture can be explained without invoking decision-making processes.

2.5 Models That Explain the Evolution of Social Learning Using Decision-Making Processes

If the claims of Roger's and Cavalli-Sforza and Feldman's models are correct (i.e., social learning is not adaptive) then how did culture evolve?

Not every trait in a species is necessarily an adaptation, as some traits can simply be contingencies particular to a species' evolutionary history (Futuyma, 2009). However, the evidence for adaptive benefits derived from culturally acquired information makes that an implausible hypothesis for culture (Richerson and Boyd, 2006; Mesoudi, 2011). Recent models incorporate at least one of two assumptions in order to make social learning, and by extension culture, adaptive. Here we respectively refer to these assumptions as *critical social learning* (i.e., biased imitation) and *guided variation*.

2.5.1 Critical Social Learning Models

In models with critical social learning, imitators do not select a random individual in the population. Instead, imitation is biased towards those individuals with fitter behaviours (Richerson and Boyd, 2006; Enquist et al., 2008). Alternatively, critical social learning can be introduced as a conditional function of imitation itself. In that case, social learning becomes a facultative trait (i.e., a trait that is expressed depending on a situational context) that will derive in "actual" imitation only if the potential behaviour-to-imitate is fitter than the behaviour encoded in the imitator's genes (i.e., its own instincts) (Enquist et al., 2008). Other models that feature critical social learning include it in the form of "parental assessor behaviour" (i.e., parents approve or disapprove of their offspring's learned behaviours) (Castro and Toro, 2004).

For example, similar to other authors, Castro and Toro argue that the evolution of social learning through unbiased imitation cannot explain the emergence of adaptive culture (Castro and Toro, 2004). To explain the emergence of culture, the authors identify the need for individual learning, guided variation or some form of decision-making ability that allows social learners to imitate behaviours with higher-than-average fitness. However, in contrast to other authors, Castro and Toro claim that the explanation for the emergence of adaptive culture in nature is the evolution of "parental assessor behaviour" (i.e., the capacity for parents to approve or disapprove of offspring's learned behaviors and enforce their judgment on them). In their paper, the authors admit that this from of selection might only be observed in the hominid lineage and perhaps a few other species of primates (Castro and Toro, 2004). Nevertheless, they show how this approach could theoretically explain the emergence of culture. To do this, they developed a mathematical model composed of a set of differential equations that calculate the frequencies of "assessor parents" (i.e., parents that bias their offpring's behaviours) and "non-assessor parents" (i.e., parents that do not bias their offpring's behaviours). In their results, the evolution of assessor behaviour occurs in parallel with the emergence of adaptive culture in the population. This assessing/teaching behaviour is mechanistically different from standard biased imitation (i.e., critical social learning). However, the model is in essence a critical social learning model. That is, to explain the evolution of culture the authors introduced an internalised form of selection (e.g., a decision-making ability) that ultimately biases effective social learning transmission in favour of behavioural variants with higher-than-average fitness.

In contrast to Castro and Toro's work, and similar to other authors in the field, our model investigates the evolution of social learners, not assessor parents. However, in contrast to all authors' work, including Castro and Toro's, our model does not rely on any decision-making abilities. In Chapter 3, we provide a detailed description on how we achieve this.

More sophisticated models introduce the term "adaptive filtering". In models with adaptive filtering, several variants are imitated and evaluated before most of them are discarded for the one with the highest fitness value (Enquist and Ghirlanda, 2007). In essence, these models explain the adaptive value of culture over genetic systems (i.e., over instinctively evolved behaviours) by introducing a selection process that can iterate several times during the lifetime of an individual, while genes, on the other hand, can only do so once (i.e., fitness evaluation only occurs once per lifetime) (Enquist and Ghirlanda, 2007). We consider that this difference in time scales (between cultural and genetic systems) might be a realistic assumption although, as we will see in the next chapter, it is theoretically unnecessary in order to demonstrate the adaptive value of culture.

In their work, Enquist and Ghirlanda claim that current theory explains the evolution of culture with models in which social learners only imitate a single behavioural variant in their lifetime (Enquist and Ghirlanda, 2007). However, advanced cultural systems such as human culture involve social learners accumulating and discerning between several of these variants. Based on this argument, they introduce a mathematical modelling strategy that models the adaptive value of many cultural variants and the frequency of social learners in the population, using a system of differential equations. In their model, a behavioural variant can be adaptive or maladaptive. Also, an adaptive variant can become maladaptive after certain time (i.e., cultural variants have a rate of "corruption") (Enquist and Ghirlanda, 2007). Using this model, Enquist and Ghirlanda claim that unbiased social learning cannot explain the evolution of adaptive culture. Their work shows how the mean adaptive value of cultural variants in the population remains neutral (i.e., both adaptive and maladaptive variants are found at the same frequency) when social learners only assess/imitate one behavioural variant in their lifetime. In order to show conditions in which culture is adaptive, the authors gave social learners the ability to identify and discard maladaptive behavioural variants in favour of adaptive variants (i.e., the ability to discern between a maladaptive behaviour and an adaptive behaviour). They concluded that the evolution of *adaptive filtering* or any other form of biased imitation may have been a requirement for the evolution of culture in natural populations (Enquist and Ghirlanda, 2007). In our simulation model, explained in Chapter 3, no form of biased imitation is required for social learners to fix and for culture to emerge.

The mechanisms for what is effectively fitness-biased imitation differ drastically between critical social learning models. However, the adaptive value of social learning - and by extension culture - is explained in all of them by a single algorithmic addition to the standard seminal models: a secondary selection function. This function patches the effect of phenotype-fitness de-correlation and solves the problem described in the Cavalli-Sforza and Feldman model.

Fitness-biased imitation does not necessarily imply imitators being able to assess fitness values; in natural populations several properties of individuals are correlated with fitness.

For this reason imitation biases that act as critical social learning mechanisms can take many forms depending on what fitness-proxy is used. A description of the most popular forms of fitness-proxy imitation biases follows (McElreath and Henrich, 2007).

1. Direct Bias:

   Direct biases are defined by the *content* of the cultural variants. This form of bias is perhaps the most straight-forward to understand. Direct-biased imitation occurs because the information that defines the variant is perceived as advantageous. This may occur because the variant itself is associated with a positive physiological reward, like imitating high-calorie food consumption, or it might come as a result of pre-established cultural traits in the population (McElreath and Henrich, 2007).

2. Model-Dependent Bias:

   In this kind of bias, individuals choose models to imitate according to their status or structural identity within a social group (e.g., how many connections does a potential model have). This form of bias decreases cultural diversity within groups but has the potential to increase intergroup differences. Symbolic systems of communication (e.g., language) are assumed to have evolved with a model-dependent bias (Richerson and Boyd, 2006).

   Like all biases in this list, model-dependent bias can make social learning adaptive as long as it acts as a proxy for fitness bias (Richerson and Boyd, 2006; McElreath and Henrich, 2007). This means that adaptation (by model dependent bias) requires a correlation between the individual attributes that cause biased imitation and the attributes that confer a fitness advantage.

   This bias can in some cases lead to a runaway process that creates culturally acquired, extravagant and sometimes maladaptive behavioural traits analogous to the peacock tail in genetic evolution (Boyd and Richerson, 1988a).

3. Frequency-dependent Bias:

   With this bias, common behaviours (i.e., behaviours that are more frequent than average) are imitated with a greater probability than their representation in the population. Some authors refer to this bias as *conformity* or *hyper-conformity* (Richerson and Boyd, 2006). It follows that when the strength of hyper-conformity is zero, social learners will imitate randomly selected behaviours (i.e., they will select behaviours with equal probability to their representation)(Henrich and Boyd, 1998).

   In changing environments, frequency-dependent bias is favoured by natural selection as imitators can rapidly track environmental fluctuations by acquiring "trending" behaviours (i.e., behaviours that are fitter than average and increasing in frequency) (Henrich and Boyd, 1998).

4. Communicator Biases:

   This is a particular form of bias which depends on behaviours being expressed in front of imitators rather than imitators having access to all behaviours and then biasing their selection when picking a model to imitate. With this bias some cultural variants transmit faster than others simply due to the effect of communication restrictions (Richerson and Boyd, 2006; McElreath and Henrich, 2007). If the variants that transmit faster and are available to as many imitators as possible also confer higher-than-average fitness values, communicator bias can explain the origin of culture (McElreath and Henrich, 2007).

   In nature some forms of communicator bias occur when individuals change their behaviour depending on the identity of potential observers around them. Again, this bias can increase adaptation or decrease it depending on the criteria being used to adjust the behaviour (Richerson and Boyd, 2006).

In this thesis, critical social learning encompasses many mechanistically different processes with the same outcome, biasing the selection of behaviours-to-imitate towards the fittest variants. We claim that the assumption of pre-existent cognitive abilities that lead to biased imitation, or alternatively the assumption of pre-existent systems (e.g., spatial distribution, social structure, instinctive behaviours) that lead to biased imitation, is unnecessary (see Chapter 3).

2.5.2 Guided Variation

Models that are able to maintain culture by relying on a mixed form of learning strategies (social and individual learning) explain culture's adaptive value by a process called *guided variation* (Richerson and Boyd, 2006). In these models, imitators improve upon their acquired behaviours through additional individual learning (i.e., the performed behaviour is a fitter version of the one that was imitated originally) (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

Populations employing guided variation evolve adaptive culture by constantly improving acquired behaviours after every transmission event (Richerson and Boyd, 2006). In many ways, this mechanism can be thought of as a genetic system where mutations do not occur at random but are disproportionately fitness-positive instead (i.e., a system with *facilitated variation*). This process is fundamentally different from biased transmission in the sense that it is not dependent on the diversity of strategies existing in a population or the selection of fit variants over unfit ones (i.e., it is not density dependent).

A characteristic example of a guided variation model is found in Borenstein, Feldman and Aoki's work (BFA model) (Borenstein et al., 2007). In their model, the evolution of culture is evaluated in the context of a fluctuating environment (i.e., the fitness value associated with each behavioral variant changes in time). In the BFA model, social learners can evolve a propensity to perform individual learning. This means, social learners can stop expressing their socially learned behaviour in favour of an individually learned behaviour. Here, an individually learned behaviour is a behaviour learned with an associated trial-and-error cost, but also guaranteed to match the current environmental state. The BFA model is formed by a set of differential equations that include conditional probabilities (i.e., expressions that describe the probability of individuals obtaining behavioral variants that match the current environmental state, given probabilities of engaging in social or individual learning). Using this model, the authors show that culture could only be adaptive if a certain degree of individual learning exists in the population. In their model, the frequency of individual learning required for the evolution of adaptive culture increases with the rate of environmental change (i.e., to track high rates of environmental change a higher rate of individual learning is required). In Chapter 5, a similar scenario will be described in more detail using our simulation model. However, in contrast to the BFA model, our model does not require any form of individual learning for social learners to fix and for adaptive culture to emerge.

Like critical social learning models, guided variation models make the assumption that imitators have the ability to evaluate different strategies before deciding which one to perform. In these models, the evaluations take place by a trial-and-error process that improves on already learned variants, rather than a process of comparison between the behaviours of different potential models (Richerson and Boyd, 2006; Mesoudi, 2011).

As in the case of critical social learning, we will show, in the remainder of this thesis, how the assumptions required for guided variation are unnecessary to explain the adaptive value of culture.

2.6 Changing Environments and the Emergence of Cultural Systems

Environmental change is considered a key factor in the evolution and origin of human culture (Richerson and Boyd, 2006; Kameda and Nakanishi, 2002; Potts, 1996). In many models the fluctuation of environmental states is included as a “de-facto” part of the pre-existing scenario from which culture emerged (Rogers, 1988; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

Models attempting to explain the origin of culture explore the effect of environmental change using populations where instinctive actuation, individual learning and social learning are all possible strategies for individuals to adopt (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007). These models unanimously conclude that under moderate environmental change, social learning will be the dominant strategy at equilibrium. However, in order for social learners to invade, a considerable fraction of the population needs to be individual learners or alternatively social learners need to resort to individual learning with a given frequency (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007). The requirement that individual learners must be present in order to maintain a population of social learners means that a process similar to guided variation is being used to maintain phenotype-fitness correlation during the emergence -and further maintenance- of culture in these models.

In existing models, an increasing rate of environmental change initially makes social learning the best option when compared to instinctive actuation. This occurs because social learners' ability to adopt trending variants (i.e., by horizontal imitation) helps them track environmental changes better than instinctive actuators (i.e., by imitating individual leaners without paying the trial-and-error cost of individual learning). However, when social learners reach a critical frequency, the process of social learning stagnation (shown in Rogers' model) makes them dependent on individual learners in order to keep track of the changing environmental state (i.e., individual learners serve as environmentally updated models to imitate).

In agreement with these models, we show in Chapter 5 how increased rates of environmental change facilitate the emergence of culture. However, in contrast to these models, we show this emergence can occur within a monomorphic population of social learners and without the requirement of individual learning strategies.

2.7 Modelling Culture Through Dual Inheritance Theory

To explain our work we must understand culture as an evolutionary system in its own right. Dual-inheritance theory is a theoretical framework that considers culture as an independent evolutionary system parallel to genes. Like any other evolutionary system, culture must then have four internal functions: reproduction, inheritance, variation and selection. Before introducing our original model in Chapter 3, let us first establish these unique functions in cultural systems according to this theory:

1. Reproduction and Inheritance:

   Culture is formed by a set of transmissible behaviours (i.e., variants). These behaviours are units of information stored in a non-genetic substrate. Dual inheritance theory assumes this information can be transmitted vertically or horizontally between members of a population (McElreath and Henrich, 2007).

   When a unit is transmitted and/or stored, we assume it does so in a *particulate* form. We use the term "particulate" in contrast to "blended" to describe an inheritance system where units can combine and form sets (i.e., cultures) without losing their ability to later separate and recombine with another set in their original form (Blackmore, 1998, 2000). A "blended" cultural system, which we do not endorse in this thesis, assumes units of information can infinitely disaggregate. In blended inheritance, once units combine with a set, they are not able to recombine in their original form with a new set (i.e., they dissolve and loose the property of being a unit) (Blackmore, 1998).

The particulate inheritance system used in dual inheritance theory is also broadly used in the field of *memetics* where units of information are referred to as *memes* (Blackmore, 1998; Dawkins, 2006)

2. Variation

   In cultural systems, errors in transmission and storage can change units (i.e., a variant can turn into a different variant). These errors generally cause small changes that gradually accumulate (i.e., similar to the effects of mutation in genetic systems). Dual inheritance theory considers this form of low-level variation a fundamental part of culture as it introduces an explorative component in its evolutionary algorithm.

   Experimental evidence on information transmission errors and error propagation exists for human and animal species (Avital and Jablonka, 2005; Mesoudi, 2011; Massen and Prinz, 2009; Knuf et al., 2001). However, cultural information is not physically encoded in units with a finite number of states, like genes are in sequences of DNA bases (Richerson and Boyd, 2006). Precise estimations of how transmission error rates compare to genetic mutation rates are not easy to obtain (Mesoudi, 2011; Massen and Prinz, 2009; Knuf et al., 2001).

   Nevertheless, it is sensible to assume, as many authors have done in the past, that error rates for cultural transmission (occurring at the macro scale where animal species interact) are several times higher than the mutation rates for the highly evolved, biochemically-guided process of genetic replication (Richerson and Boyd, 2006; Futuyma, 2009; Mesoudi, 2011).

3. Selection

   Different cultural variants transmit at different rates; hence variants that transmit more often will eventually displace competing variants that fulfil the same role (e.g., a method to sharpen a stone tool or to dig a hole efficiently). Imitation biases are the most common explanation used to introduce selection in models of cultural evolution. However, other processes like natural selection or population structure can also play the role of selection according to dual-inheritance theory (McElreath and Henrich, 2007).

   Natural selection can play the role of selection in these systems only when cultural variants affect the longevity and/or fertility of their performers. For instance, variants that extend the life of their carriers make them more likely to serve as

> models-to-imitate and variants that increase the relative fertility of their carriers (and transmit vertically) increase their frequency in the next generation (McElreath and Henrich, 2007).
>
> In most existing models of the origin of culture, individual fitness (as a product of natural selection) is exclusively represented by fertility (Cavalli-Sforza and Feldman, 1983; Rogers, 1988; Henrich and Gil-White, 2001). Longevity is therefore considered only relevant to the extent a longer life leads to a higher number of offspring (i.e., as a proxy for fertility). This simplification creates a scenario where the previously explained effect of phenotype-fitness de-correlation prevents culture, with any degree of horizontal imitation, from emerging without the effects of individual learning, or the effects of decision-making processes like critical social learning or guided variation.
>
> As will be shown in the next chapter, the separation of these two types of natural selection is an important distinction of our model. When introducing longevity as a function of fitness, imitators can copy fitter individuals more frequently than unfit ones without the need for biased imitation.

Dual inheritance theory is concerned with the evolution of culture after it has emerged rather than the origin of culture itself (which is the focus of this thesis). However, the framework built by this theory establishes culture as a well-defined evolutionary system that can be compared with its genetic counterpart considering all of its four functions: reproduction, inheritance, variation and selection. Previous attempts to explain the origin of culture have overlooked the effects of considering differences in the variation functions (i.e., differences between imitation error rates and genetic mutation rates) along with the role of natural selection as a selection function that can be separated into survival and reproductive components.

Although these factors have not been taken into account all together in any existing model, we would like to make a special mention of McElreath and Strimling (2008) (McElreath and Strimling, 2008). In their work, the authors extend Rogers' model by dividing selection into survival and reproductive components, one of the two features that differentiates our model in Chapter 3 from others. However, their model does not incorporate an extended strategy space along with the respective mutation rates and imitation error rates for instinctive actuators and social learners. For this reason, McElreath and Strimling's model is unable to reach the same conclusions we arrive at in Chapter 3. With regard to the conditions under which social learners could invade a population, contributions from McElreath and Strimling's model conform with those found in the original Rogers' model. Their paper, however, concludes that vertical imitation (i.e., imitating ones' parents as opposed to any member of the population) is the favoured form of social learning when reproductive selection is stronger than survival selection. In Chapter 4, where we analyse the effect of vertical vs horizontal imitation, we will revisit and discuss this model and their conclusions.

Our model in Chapter 3 borrows from dual inheritance theory the assumption that culture is an evolutionary system in its own right and explores the emergence of this system in a situation where it competes with its genetic counterpart to find a solution for an environmental problem. Given the number of non-linear processes emerging in a population with two parallel evolutionary systems (i.e., cultural and genetic), we decide to build our model using a simulation-based approach (with further analytical estimations) rather than a purely analytical approach.

2.8 Using Individual-Based Simulations

Mathematical models are powerful tools used to gain insight into processes and mechanisms that explain the evolution of culture. However, they are in principle limited to rather simple scenarios and can get increasingly complicated when interdependencies and non-linearities are included. We claim that when analysing the emergence of culture, important interdependencies and non-linearities may arise during the transition from a genetic evolutionary system to a cultural one. Therefore, to study the origin of this cultural system we have resorted to an alternative tool: a simulation-based approach.

Individual-based simulations (i.e., numerical simulations of populations of agents) allow us to construct complex models with minimal time-cost in development and a broader range of analysis (Edmonds, 2001). Furthermore, this type of simulation approach can easily uncover unexpected emergent patterns in the systems to be modelled simply as an outcome of basic assumptions (Resnick, 1997). Unfortunately there is a trade-off when using individual-based simulation models rather than analytical models: the introduction of undesired artefacts (Axelrod, 1997b; Gilbert, 2007; Galán et al., 2009). The model in Chapter 3 and its subsequent extensions in Chapters 4 and 5 aim to be simple enough to avoid these unnecessary artefacts.

Some of the relevant limitations of existing analytical models have already been discussed, environmental matching being represented by a two-state variable (Roger's model) and a limited strategy space (Cavalli-Sforza and Feldman's model) being the main ones. We claim that culture may benefit from theoretical work which uses explicit strategy spaces where solutions are encoded, and both genetic and cultural systems compete to find the optima in these spaces. At present, we can only do this using simulation models.

Other uses of simulation methods include recent work in the field of artificial life where individual-based approaches to model culture have been used (Borg and Channon, 2012; Marriot and Chebib, 2014). These models were not mentioned in previous sections of this chapter because they are not relevant to the research questions of this work. In other words, these models do not explain the origin of simple cultural systems; instead, they explain the evolution of traits found in advanced culture. Uses of individual-based simulations in cultural evolution are also abundant in social science fields where authors explain cultural patterns in humans societies rather than the origin of culture itself (Dignum and Dignum, 2014; Morris et al., 2014). Other models that explore topics involving evolutionary dynamics in cultural systems have benefited in the past from a wide range of simulation approaches (Smith et al., 2003; Kameda and Nakanishi, 2002; Kohler and Gummerman, 2001; Axelrod, 1997b,a; Macy and Svoretz, 1998).

Chapter 3

The Origin and Evolution of Cultural Evolutionary Systems in the Absence of Decision-Making Processes

In this chapter we aim to identify conditions under which social learning can evolve due to the properties of cultural replication, variation and natural selection, without invoking decision-making processes (see first research question in Chapter 1). Here we explain the development of an individual-based simulation model and discuss its results. The results support the claim that social learning can evolve to displace its genetic counterpart (i.e., instinctive actuation) when natural selection on survivability is stronger than selection on reproduction and when imitation error rates are higher than genetic mutation rates. The model does not rely on critical social learning or guided variation to achieve these results.

3.1 The Model

The individual-based simulation model represents each individual by two strings of bits. The first string represents an individual's phenotype and the second represents its genotype. The fitness of an individual is determined by considering only the phenotypic information. Individuals also have a single additional genotypic bit which acts as a switch determining whether or not the individual engages in social learning. If the value of the switch is 0, the individual's phenotype is simply a copy of its own genotype (for the purposes of our research questions, a one-to-one identity mapping from genotype to phenotype is sufficient). However, if the value of this switch is 1, then rather than express a phenotype derived from its own genotype, an individual will express an "imitated" phenotype randomly selected from the population. These social learners will each imitate a randomly selected phenotype at birth and their phenotypes will then remain unchanged for the remainder of their lifetimes.

All the agents in the initial population have an imitation switch value of 0 and each bit of their genotype bit string set to 1 or 0 with uniform random probability per bit. Each initial agent's phenotype is set to be a perfect copy of its genotype. These initial conditions prevent mutation biases from imposing directed drift on the genotypes or phenotypes, even in the absence of selective pressure. For all the results shown here, unless stated otherwise, the length of the phenotype and genotype bit strings is 200 bits ($L = 200$), and the size of the population is 100 individuals ($N = 100$). We have selected these values considering, (1) the need for short computational times per simulation run (i.e., small populations and short bit strings require less computational time), and (2) the need for minimal bit string length and population sizes in order to show key processes described in this thesis (i.e., if bit strings and/or population sizes are too small some processes described in this thesis will not occur) (see Appendix C for more details). Selection is defined by a Boltzmann-weighted sum-of-1s in the string of phenotypic bits. A phenotypic string of all-1s represents the optimal solution.

Qualitatively similar results can be reproduced with larger populations and longer bit strings. Smaller populations produce variation between individual simulation runs, but on average they also agree qualitatively with the results described here. Shorter bit-strings can produce different results, as they represent simpler search problems, and a sustained search and optimisation period of several iterations is required in order to produce the results reported here. For example, a single-bit (two-strategy) model will not reproduce our results.

Selection is established by the joint action of a reproduction function and a death function, which impose *selection for reproduction* and *selection for survival*, respectively. Reproduction selects an individual i from the population with a probability P_{ri} using a Boltzmann-weighted function of its sum of phenotypic 1s (i.e., L_{1i}) normalised across the population. The Boltzmann exponent x_r in the function, controls the strength of selection for reproduction. High values of x_r reduce the effect of reproductive selection (Equation 3.1).

$$P_{ri} = \frac{e^{\frac{L_{1i}}{x_r}}}{\sum_{k=1}^{N} e^{\frac{L_{1k}}{x_r}}} \tag{3.1}$$

The death function selects an individual j from the population with a probability P_{dj} using a Boltzmann-weighted function of the complement of its sum of phenotypic 1s (i.e.,

$L - L_{1j}$). Thus, P_{dj} is the relative probability of dying, or anti-fitness. The Boltzmann exponent x_d in the function, represents the strength of selection for survival. High values of x_d reduce the effect of survival selection (Equation 3.2).

$$P_{dj} = \frac{e^{\frac{L-L_{1j}}{x_d}}}{\sum_{k=1}^{N} e^{\frac{L-L_{1k}}{x_d}}} \tag{3.2}$$

A key feature of our model is the ability to change the relative balance between reproductive selection and survival selection. Higher values for exponents x_r and x_d reduce the effect of the sum of phenotypic 1s on probabilities of reproduction and death, respectively.

The simulation is a steady-state model, which means that, in each iteration, a single individual is selected for reproduction by the reproduction function and a single individual is selected for death by the death function (i.e., the model produces overlapping "generations"; therefore, agents created at different iterations coexist at any given point of a simulation run). The genotype and imitation switch of the reproducing individual are mutated, and they then replace the genotype and imitation switch of the dying individual. Genotype mutation occurs as a uniform bit-flip probability of μ_g per bit, except for the imitation switch which has a uniform bit-flip probability of μ_c per bit. If the imitation bit of the new copy has value 0, the new phenotype bit string is a one-to-one copy of the new genotype (not including the imitation switch). If the switch has value 1, a random phenotype from the population is copied (unbiased horizontal imitation) with a bit-flip error rate of μ_p per bit. The imitation switch can only be passed genetically and is not part of any phenotype.

For simplicity, individuals reproduce asexually with no "crossover" during genetic replication or phenotypic imitation and there is no environmental change altering the optimum bit-string (i.e., the string where $L_1 = L$), although an extension on the effect of the latter is presented in Chapter 5.

Changing the values of μ_g and μ_p controls the rates of mutation and imitation error, respectively. We fix $\mu_c = 0.01$ for all results reported here and vary μ_g and μ_p relative to this value. Figure 3.1 illustrates an algorithmic implementation of the model.

As it will be discussed in Section 3.5, we chose to fix $\mu_c = 0.01$ to incentivise the formation of *chains of imitation* (i.e., sequential imitation among imitators) in the early stages of our simulation runs. For lower values of μ_c, the density of social learners produced by mutation alone (i.e., before any form of selection acts on the population), is too low for imitation among imitators to occur with high enough frequency as to kick-start an evolutionary system (i.e., culture). Values of μ_c that are larger than those minimally required to start chains of imitation, produce a large initial density of social learners without the intervention of any evolutionary process. For obvious reasons, a high initial density of social learners would make our model unsuitable to describe the emergence of culture, as this process is precisely the outcome of the evolution of social learning. The value of $\mu_c = 0.01$ was selected as a compromise between these two constraints. This caveat/limitation would be abstracted-away in a mathematical model with an infinite population because, so long as the frequency of social learners is non-zero, the frequency of imitation between social learners will also be non-zero. In contrast, any finite-population process instantiating our model, will have the same limitation as our simulation (i.e., the need for a minimal initial frequency of social learners, that guarantees a non-zero frequency of imitation between social learners). Note that both our simulation model and an infinite-population, mathematical model do not explain how a critical initial frequency, or a non-zero frequency, is attained - but both kinds of models do address the conditions under which this initial frequency takes-off. These points should be taken into consideration when interpreting our results relative to other models in the literature, in which the invasion of social learners can start with a differential proportion of social learners. Please see Section 3.5 and Appendix C for more details.

In this thesis we explore a wide range of parameter values in our simulation model, for more information on why this explored range was selected, please see Appendix C.

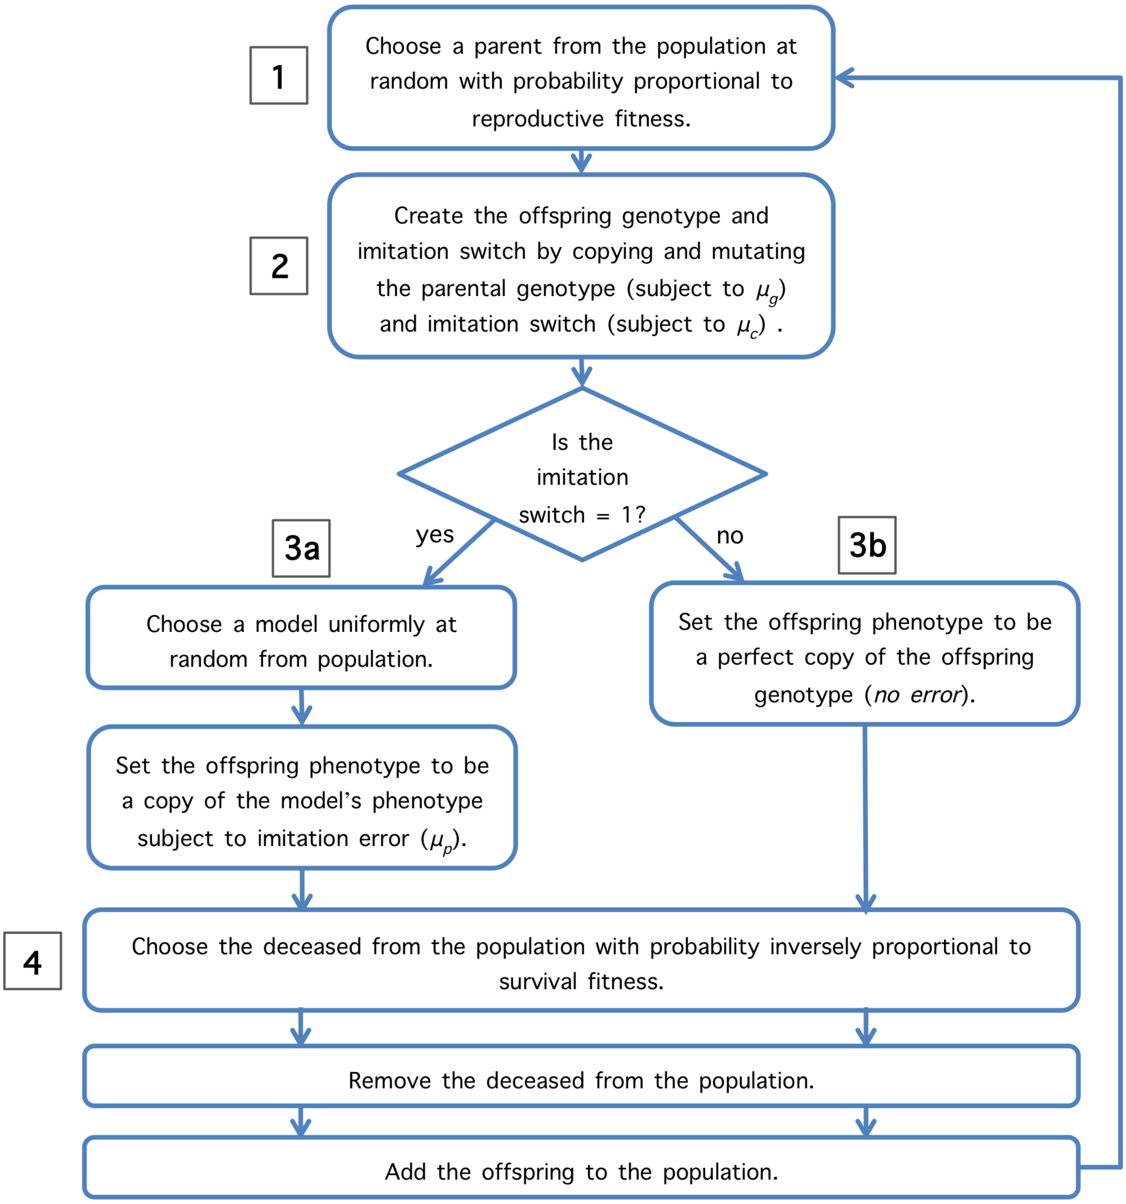

FIGURE 3.1: An algorithmic representation of the model's logic. (1) The selection function picks an individual from the population. (2) Its genotype string is copied including the imitation switch (mutations occur at a rate of μ_g per genotype bit and μ_c for the imitation switch). (3a) If the imitation switch has value 1, a random individual from the population is selected and its phenotype will be copied (imitation errors occur at a rate of μ_p per bit). (3b) If the imitation switch has value 0, the phenotype will be a perfect copy of the individual's genotype. (4) The resulting combination of phenotype and genotype will replace an individual selected by the death function.

3.1.1 Simulation Steps

In this subsection, the model presented in Figure 3.1 is explained in detail, with particular emphasis on clarifying the internal logic of each function and how it affects the bit strings that define individual agents in the simulation. This information can also be found in Appendix B, along with a description of the extensions explained in Chapters 4 and 5. To assess the internal logic of our model, before generating our analytical results, all individual functions described here were "unit tested". That is, for each function, the expected outputs for a range of possible inputs were evaluated.

3.1.1.1 Population Setup

To setup a simulation run, we initialise a population of 100 agents (i.e., $N = 100$) by allocating two bit strings of 200 bits per individual (i.e., $L = 200$). One string of 200 bits represents the phenotype, and the other string of 200 bits represents the genotype. For all simulations in this thesis, unless stated otherwise, the value of each bit is assigned at random during the setup (i.e., each bit has equal probability of being 1 or 0). Phenotype and genotype strings represent encoded solutions for a behavioural challenge (e.g., how to cross a river safely, build a nest or sing a song).

In addition to these two strings, each individual also has an extra "genotype" bit. This bit acts as a switch. If the switch has value 1, the individual will substitute, at birth, their entire 200-bit phenotype string for the phenotype string of a randomly selected individual from the population. If the switch has value 0, the individual will express an exact copy of its 200-bit genotype string as their phenotype (without including the imitation switch). Unless stated otherwise, all individuals in the population are initialised with their imitation switch in position 0. This means all individuals are *instinctive actuators* (i.e., their genotype string and phenotype string are exact copies). Figure 3.2 shows a diagram explaining this setup.

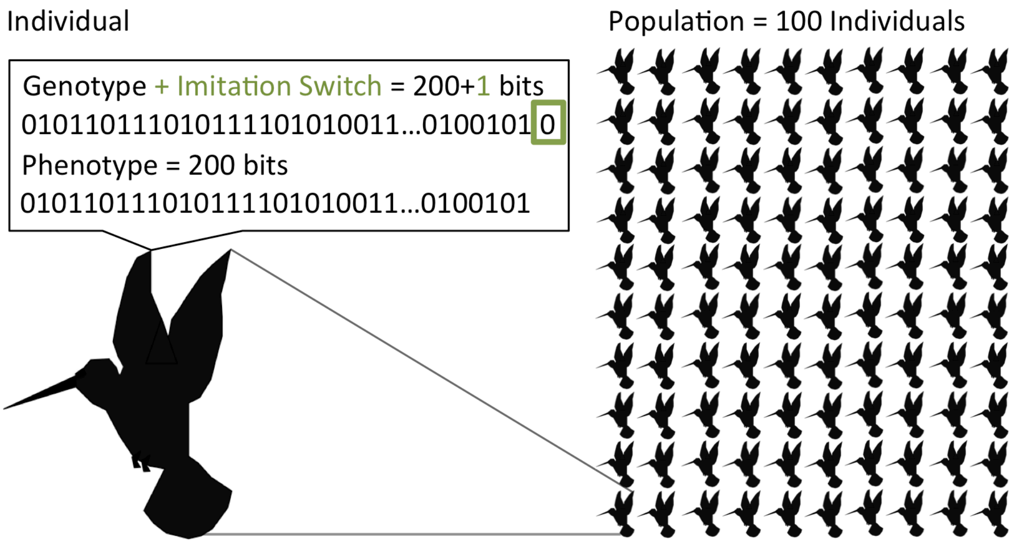

Figure 3.2: Simulation setup of a population with 100 individuals. Each individual is represented by two strings of 200 bits. One string represents the phenotype, and the other string represents the genotype. In addition to these strings, each agent has an imitation switch bit (framed in green) which is inherited during genetic reproduction but not expressed in the phenotype. All agents are initialised with their imitation switch set to 0.

For each individual we calculate *reproductive selection scores* as $e^{\frac{L_1}{x_r}}$, where L_1 is the sum-of-1s in the phenotype string and x_r is the Boltzmann coefficient that controls the strength of reproductive selection in the simulation run. In a similar manner, we calculate *death selection scores* for every individual as $e^{\frac{L-L_1}{x_d}}$, where $L-L_1$ is the sum-of-0s in the phenotype string, and x_d is the Boltzmann coefficient that controls the strength of survival selection in the simulation run.

Note that x_r and x_d are global parameters in the simulation (i.e., we use the same x_r and x_d to calculate the reproductive selection scores and death selection scores of all individuals, respectively). To avoid repeating computationally expensive operations, we store these selection scores and only recalculate them when necessary.

3.1.1.2 Step 1: Select a Parent With Probability Proportional to Reproductive Fitness

In the first step of our algorithm, we select an individual using the reproductive selection scores calculated during the setup. An individual i will be selected for reproduction with probability P_{ri} as described in Equation 3.1. This means, the probability of selecting any given individual for reproduction is equal to the reproductive selection score of that individual divided by the sum of the reproductive selection scores of all individuals in the population (i.e., the scores are normalised).

3.1.1.3 Step 2: Create the Offspring Genotype and Imitation Switch

In the second step of our algorithm, we create the genotype and the imitation switch of a new individual. This new individual will be introduced into the population in step 4, and represents the offspring of the parent selected in step 1. To produce its genotype, we copy the genotype string of the parent with an error rate of μ_g per bit. This means, each bit copied into the genotype of the new individual will change its value (i.e., changing from 0 to 1, or from 1 to 0) with probability μ_g. The imitation switch of the parent is also copied into the offspring. However, this switch will change its value with probability μ_c, instead of μ_g. Figure 3.3 shows a diagram explaining this step.

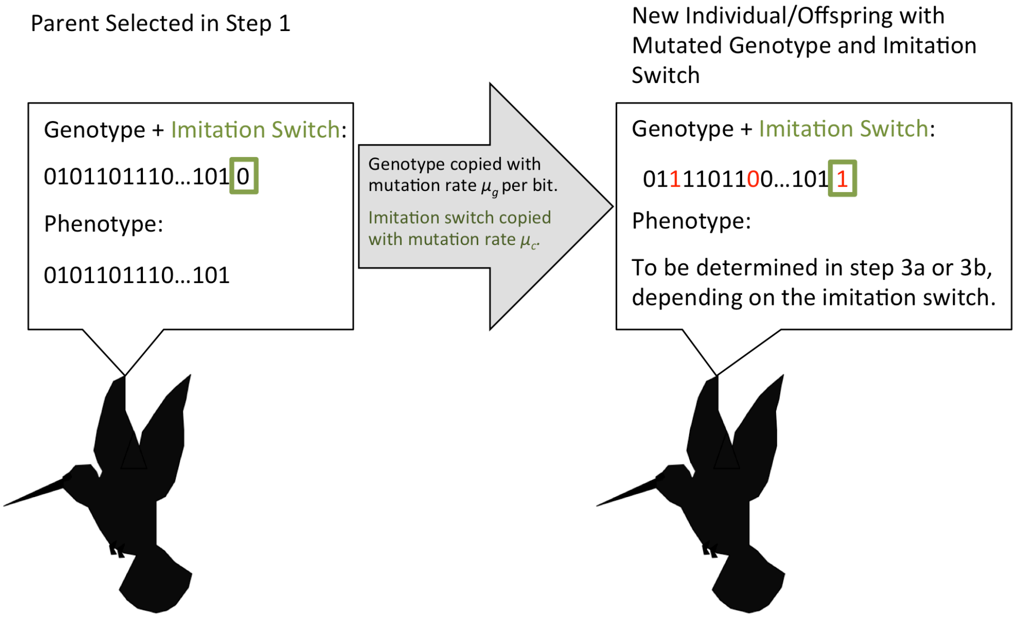

FIGURE 3.3: Diagram representing the second step of our simulation model. We create a new individual using the individual selected in the first step (i.e., parent). The genotype of the new individual is a copy of the parental genotype. Each genotype bit is copied subject to a mutation rate of μ_g. The imitation switch of the new individual is a copy of the parental imitation switch. This switch is copied subject to a mutation rate of μ_c. Mutated bits are shown in red.

3.1.1.4 Step 3a: If the Offspring is a Social Learner

In the previous step, we produced a new individual/offspring with a genotype string and an imitation switch, but without a phenotype string. In this step, we will create a phenotype for this individual assuming that its imitation switch is 1 (i.e., the individual is a social learner). When the imitation switch is set to 1, the phenotype of the new individual is a copy of the phenotype string of another individual, selected at random from the entire population (parent included). This process of phenotype imitation occurs once at birth. The phenotype string is copied with an error rate of μ_p per bit. This means, each bit copied into the phenotype of the new individual will change its value (i.e., changing from 0 to 1, or from 1 to 0) with probability μ_p. Notice that by imitating the phenotype of an individual selected at random from the population, social learners are effectively engaging in *unbiased imitation*. Figure 3.4 shows a diagram explaining this step.

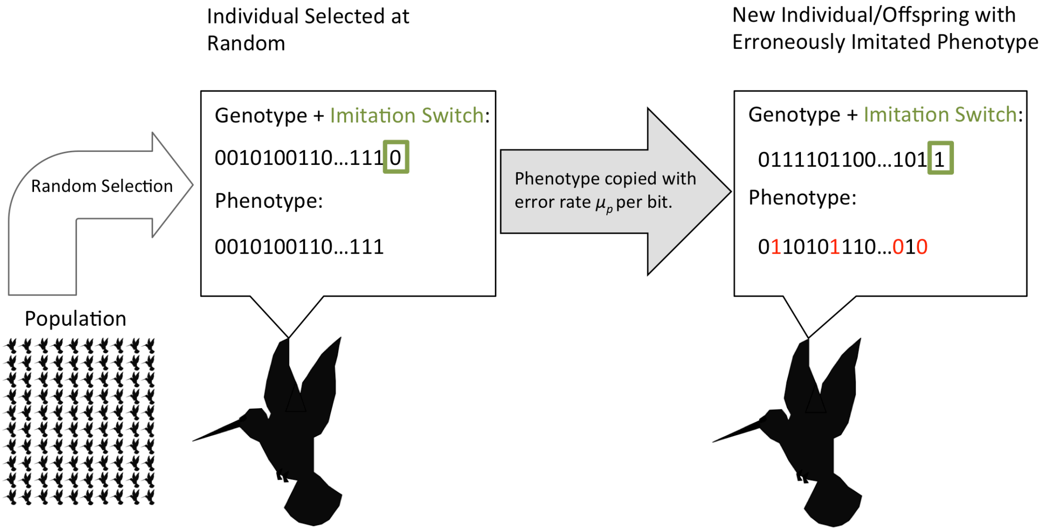

Figure 3.4: Diagram representing the third step of our simulation model (if the new individual is a social leaner). If the value of the imitation switch of the new individual is equal to 1 (i.e., the individual is a social learner), we create its phenotype by copying the entire phenotype string of an individual selected at random from the population. Each phenotype bit is copied subject to a mutation rate of μ_p. Notice that the imitation switch is not part of the phenotype. Mutated bits are shown in red.

3.1.1.5 Step 3b: If the Offspring is an Instinctive Actuator

In step 3a we created the phenotype of the new individual/offspring by copying the entire phenotype of a randomly selected agent. We used this method because the imitation switch of the new individual had a value of 1. In contrast, in this alternative step, which we labelled as step 3b, we will describe the method for creating the phenotype of the new individual when its imitation switch is set to 0. It is important to note, that new individuals will create their phenotypes using step 3a or step 3b, but not both.

If the imitation switch of the new individual is set to 0, an exact copy of its genotype string will be used to create its phenotype. This copying process represents genetic expression, and occurs without error. For the focus of our research questions, a 1-to-1 map between phenotype and genotype will suffice. The imitation switch is not part of the phenotype, and therefore is not copied. Figure 3.5 shows a diagram explaining this step.

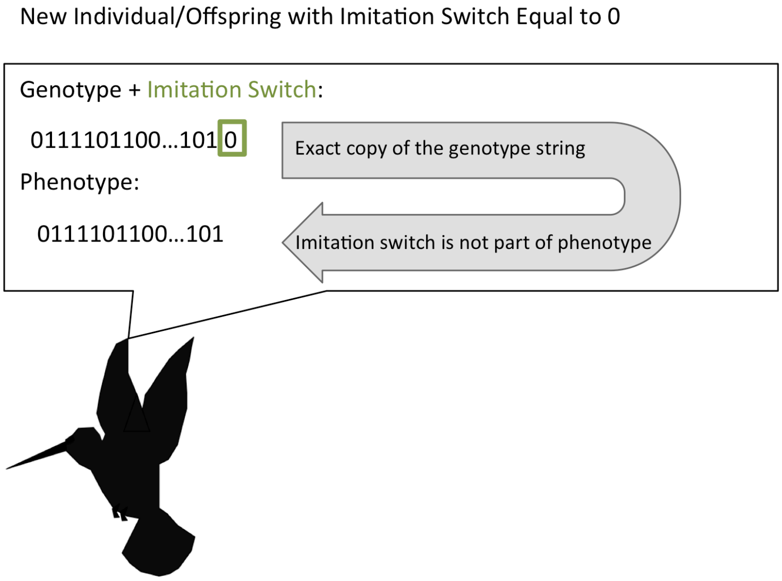

FIGURE 3.5: Diagram representing the third step of our simulation model (if the new individual is an instinctive actuator). If the value of the imitation switch of the new individual is equal to 0 (i.e., the individual is an instinctive actuator), its phenotype is an exact copy of its own 200-bit genotype string. Each bit is copied without error. This step represents genetic expression (i.e., the phenotypic expression of the behavioural variant encoded in the genotype). Notice that the imitation switch is not part of the phenotype; therefore it is not "expressed".

3.1.1.6 Step 4: Select a Replacee With Probability Inversely Proportional to Survival Fitness

In this step, the new individual created through step 3a or 3b is introduced in the population. However, in order to maintain a constant population size, another individual needs to be removed/killed/replaced. The replacee is selected using the death selection scores mentioned in our setup. An individual j will be selected for replacement/death with probability P_{dj} as described in Equation 3.2. This means, the probability of selecting any given individual for death is equal to the death selection score of that individual divided by the sum of the death selection scores of all individuals in the population (i.e., the scores are normalised). Note that the death selection scores use the sum-of-0s instead of the sum-of-1s. Therefore, the probability of being selected for death is inversely proportional to the total number of 1s. Also, note that the parent selected for reproduction in step 1 can also be selected as the replacee in this step.

After a replacee as been selected, the new individual takes its place in the population. The reproduction selection scores and death selection scores of the new individual are calculated and stored. These scores will be used in the next iteration of the model, starting on step 1.

3.1.2 System Details

As it can be appreciated, the core simulation of our work is very simple. During the research and development phase of this thesis, our model was independently implemented in three different computer languages: NetLogo 4.0, Python 2.7 and C. Results from each implementation were crosschecked to guarantee consistency and reproducibility. Single-run results, shown in Section 3.2, were produce using the implementation in python, on an Apple-MacBook Pro laptop (2.3 Intel Core i5 - 8 GB of RAM) and an Apple-IMac desktop (3.4 GHz Intel Core i7 - 8 GB of RAM). However, implementations in NetLogo and Python were deemed too slow to produce results that required a large number of replicated runs, in a sensible amount of time (e.g., Figures 3.8 and 3.9 in Section 3.2).

For results that required a large number of replicates, a trivially parallelised version of our implementation in C was executed using OpenMP (Open Multi-Processing) on the IRIDIS4 supercomputing cluster, at the University of Southampton. The IRIDIS4 supercomputing cluster is one of the largest computational facilities in the United Kingdom; it is professionally managed, and available to the university's entire research community. IRIDIS4 delivers a reliable batch service for distributed memory parallel jobs. For this thesis, a single node in the cluster was used. A node has 16 Intel Sandy Bridge CPUs at 2.6 GHz with 64 GB of RAM. All plotting and visualisations were produced using matplotlib 4: a python visualisation library.

3.2 Results

Simulations were run for a range of mutation rates, $\mu_g \in [0.0, 1 \times 10^{-4}, \ldots, 50 \times 10^{-4}]$; imitation error rates, $\mu_p \in [0.0, 1 \times 10^{-4}, \ldots, 50 \times 10^{-4}]$; survival selection strengths, $x_d \in [1, 5, 10]$; and reproductive selection strengths, $x_r \in [1, 5, 10]$. Each simulation run consisted of 5×10^4 iterations. The average proportion of 1s in the population's genotypes, the average proportion of 1s in the population's phenotypes, and the proportion of individuals with their imitation switch "on" (i.e., the proportion of social learners) were recorded. All simulation runs had a population size of $N = 100$ and a string length for phenotype and genotype strings of $L = 200$. For each combination of μ_g, μ_p, x_d and x_r, 100 replicate simulations were run. See Appendix C for more details on parameter selection.

Where imitation error rate is larger than genetic mutation rate, $\mu_p > \mu_g$, simulations converge to one of three distinctive patterns (Figure 3.6): A) Both genotypes and phenotypes converge on the single optimum, while the proportion of imitators remains close to zero; B) Phenotypes and genotypes initially improve together, but after around two-thousand iterations, imitators rapidly invade the population and only phenotypes continue to improve to the optimum; C) imitators increase in frequency slowly and inconsistently (i.e., as if the proportion of 1s in the imitation switches of the population was drifting), and eventually fix in the population; after fixation only the phenotypes continue to improve. In B and C the genotype drifts around the neutral average containing approximately equal numbers of 1s and 0s. The de-correlation of genotype and phenotype information, accompanied by the fixation of imitators, is here referred to as *genotype-phenotype disengagement*. This pattern is considered to indicate the emergence of a non-genetic evolutionary system, here referred to as “culture”. The stochastic nature of the model makes convergence probabilistic. Multiple replicates are hence required to analyse outcomes for each parameterisation. The chances of a given mutation rate and imitation error rate converging to pattern A, B or C are also affected by the weight of reproductive and survival selection. These weights are respectively controlled by the exponents x_r and x_d in Equations 3.1 and 3.2, with higher exponents leading to lower selection strengths. Hence, for each simulation we must specify selection coefficients, mutation rate and imitation error rate.

Scenarios in which the genetic mutation rate is larger than the imitation error rate, $\mu_g > \mu_p$, are not considered relevant to the evolution of non-genetic evolutionary systems in this work, since it is unrealistic that this incipient secondary system has more fidelity than the process of genetic replication (Dugatkin, 2000; Avital and Jablonka, 2005; Richerson and Boyd, 2006; Mesoudi, 2011).

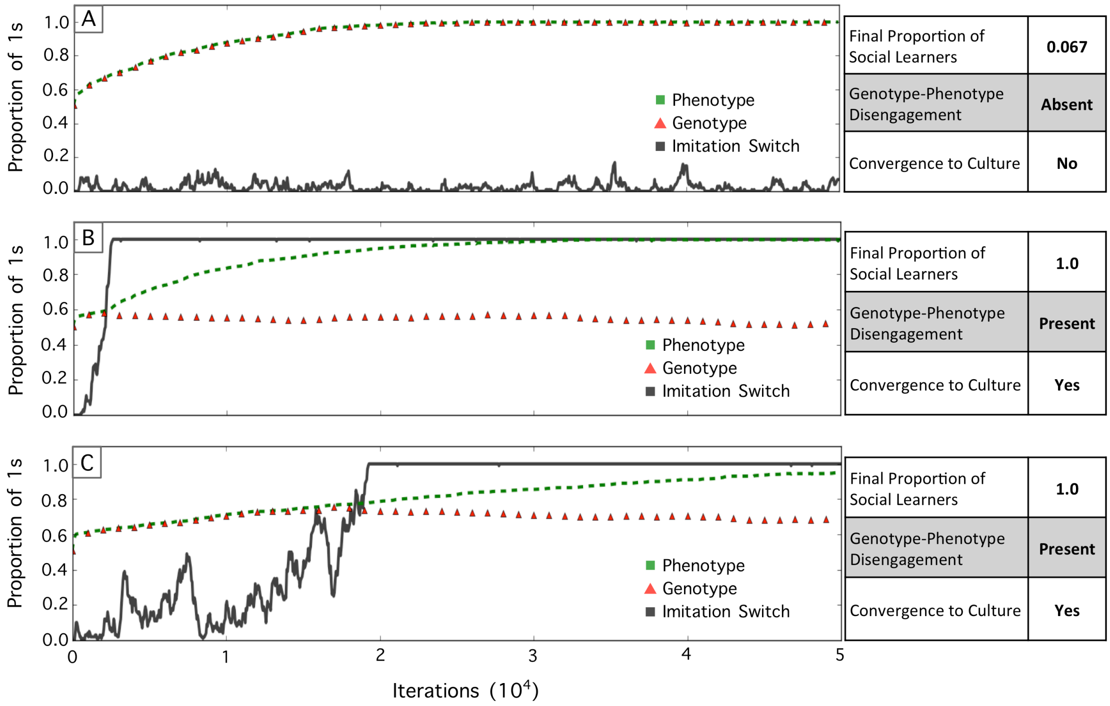

Final Proportion of Social Learners	**0.067**
Genotype-Phenotype Disengagement	**Absent**
Convergence to Culture	**No**

Final Proportion of Social Learners	**1.0**
Genotype-Phenotype Disengagement	**Present**
Convergence to Culture	**Yes**

Final Proportion of Social Learners	**1.0**
Genotype-Phenotype Disengagement	**Present**
Convergence to Culture	**Yes**

Figure 3.6: A single representative simulation run for each of three different parameter sets showing characteristically different convergence patterns A, B and C. Each plot depicts the change over evolutionary time in the average proportion of 1s in the population's phenotypes (green) and genotypes (red), and the proportion of imitators in the population (black). (A) No disengagement: phenotype and genotype improve together and imitation frequency remains low. This case is encountered when genetic mutation is high and the imitation error rate offers not enough exploration advantage for a given set of reproductive and survival selection strengths. (Parameter set: $x_d = 1$, $x_r = 10$, $\mu_g = 0.003$, $\mu_p = 0.005$). B) Early disengagement: the non-genetic evolutionary system emerges rapidly in the simulation, indicated by a sharp rise in imitators and genotype-phenotype disengagement. Imitators fix and phenotype improvement is due to non-genetic evolution from then on. (Parameter set: $x_d = 1$, $x_r = 10$, $\mu_g = 0.0005$, $\mu_p = 0.005$). (C) Late disengagement: non-genetic evolution takes over after several iterations. Imitators invade slowly and with large fluctuations in frequency. This pattern occurs when mutation and imitation error rates are very low and/or similar to one another. (Parameter set: $x_d = 1$, $x_r = 10$, $\mu_g = 0.0002$, $\mu_p = 0.0002$). For each simulation run, a summary table is presented. The table describes (1) the proportion of social learners observed in the last iteration of the simulation, (2) the presence or absence of phenotype-genotype disengagement, and (3) the emergence of culture. As we will explain in this chapter, these three variables are *not* independent. In our simulation, convergence to culture is defined by the fixation of social learners, which in turn, occurs inextricably in parallel with the process of genotype-phenotype disengagement.

Scenarios in which the genetic mutation rate is larger than the imitation error rate, $\mu_g > \mu_p$, are not considered relevant to the evolution of culture in this work, since it is unrealistic that this incipient secondary evolutionary system (i.e., culture by social learning transmission) has more fidelity than the process of genetic replication. It is the overall consensus in the field, that the error rate of social learning transmission (i.e., imitation) must be higher than the mutation rate of genetic replication (Dugatkin, 2000; Avital and Jablonka, 2005; Richerson and Boyd, 2006; Mesoudi, 2011). For completeness, we show a single simulation run in which $\mu_g > \mu_p$, in Figure 3.7. Under these conditions, social learners rapidly fixate after the genetic evolutionary system has reached sum-of-1s values close to the optima (i.e., a string of all-1s). This is trivially explained by the advantage of high-fidelity transmission that social learners have over instinctive actuators, in this scenario. Consider that once the optima has been reached, the best replication/transmission strategy would the one that offers higher fidelity, when $\mu_g > \mu_p$, such strategy is social learning.

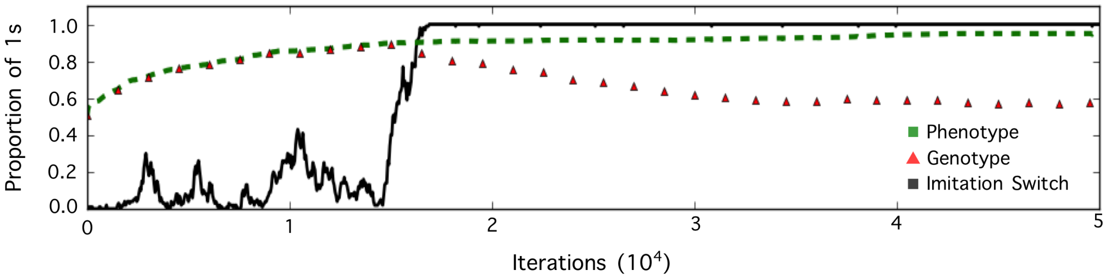

FIGURE 3.7: A single representative simulation run for a parameterisation in which $\mu_g > \mu_p$ is shown. In this work we consider cases in which $\mu_g > \mu_p$, unrealistic. In these cases, the incipient cultural evolutionary system is assumed to have more inheritance fidelity (i.e., a lower error transmission rate), than that of genetic replication. In these simulations, social learners tend to invade the population after the genetic evolutionary system, driven by genetic replication, has evolved the phenotype strings close to the optima (i.e., a string of all-1s). The green dashed line represents the average proportion of 1s in the phenotype strings of the population, the red triangles represent the average proportion of 1s in the genotype strings, and the black line represents the proportion of social learners (i.e., the proportion of 1s in imitation switches). (Parameter set: $x_d = 5$, $x_r = 5$, $\mu_g = 0.003$, $\mu_p = 0.0001$).

One hundred replicates for each combination of mutation rate, μ_g, and imitation error rate, μ_p, drawn from the set $[1 \times 10^{-4}, 2 \times 10^{-4}, \ldots, 50 \times 10^{-4}]$, and where $\mu_p > \mu_g$, were carried out with fixed values of survival and reproductive selection coefficients ($x_d = 1$, $x_r = 10$). The proportion of replicates that converge to non-genetic evolution (where the proportion of imitators is above 0.95 by the end of the run) is presented in a two-dimensional heat map in Figure 3.8. This map shows all combinations below the diagonal $\mu_p = \mu_g$, where imitation error is greater than genetic mutation.

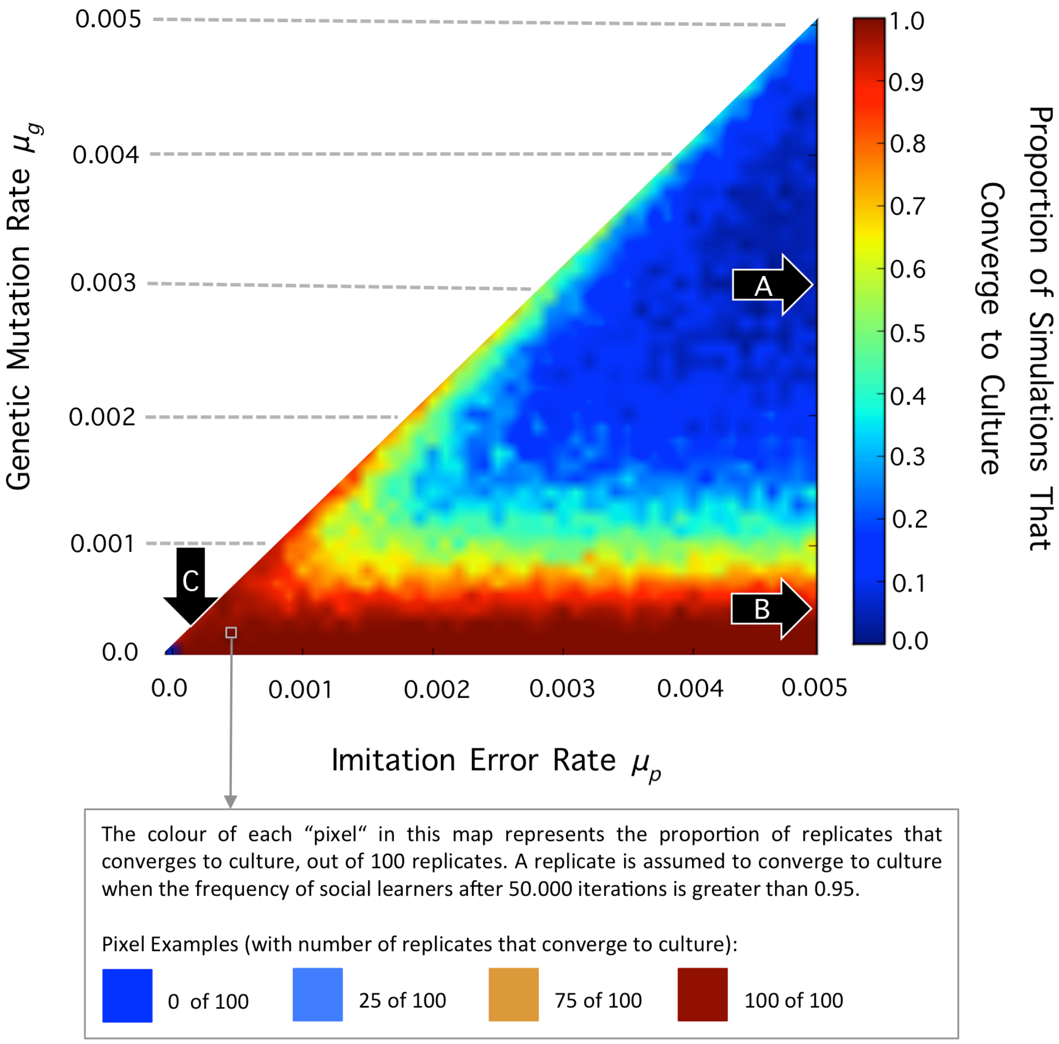

The colour of each "pixel" in this map represents the proportion of replicates that converges to culture, out of 100 replicates. A replicate is assumed to converge to culture when the frequency of social learners after 50.000 iterations is greater than 0.95.

Pixel Examples (with number of replicates that converge to culture):

🟦 0 of 100 🟦 25 of 100 🟧 75 of 100 🟥 100 of 100

FIGURE 3.8: Heat map depicting how the tendency for social learners to evolve varies with the genetic mutation rate, μ_g, and imitation error rate, μ_p. For all points in this plot, $\mu_p > \mu_g$. The colour of each pixel in the heat map, represents the proportion of replicates that converge to culture, for simulations with parameterisations that corresponds to the "co-ordinates" in which the pixel is located. Blue regions indicate that the explorative advantage of imitation errors is not great enough for disengagement to occur early on during the optimisation process. The lower region shows cases where social learners invade and fix in the population. In all these cases the values for the selection exponents are set to $x_d = 1$ and $x_r = 10$ creating a selective environment with strong survival selection and relatively weak reproductive selection. For μ values larger than 0.005 the trends observed at the positive edge of this heat map extrapolate until $\mu_p = 0.015$ (Higher values of μ_p were not explored). The constrained range has been selected to increase resolution on contrasting regions. Black arrows indicate the points in parameter-space where the single runs shown in Figure 3.6 (i.e., A, B and C) are located.

Further exploration of pairs of selection coefficients for survival and reproductive selection strengths is presented in the array of heat maps in Figure 3.9; each having the format explained in the previous paragraph. Combinations of selection coefficient values of 1, 5 and 10 for x_d and x_r were selected for their contrasting results. Figure 3.9 can be interpreted as a four-dimensional parameter-space representation showing the distribution of probabilities for the fixation of social learners. For relatively strong survival selection, and relatively high imitation error rates, social learners fix following pattern B in Figure 3.6. For mutation and imitation error rates that are low and similar, social learners fix following pattern C in Figure 3.6. Given that our model does not feature decision-making attributes, and that both error rates and the survival value of the trait being evolved are not functions of previously acquired traits, we consider these examples of unbiased social learning originating by means of its adaptive value alone.

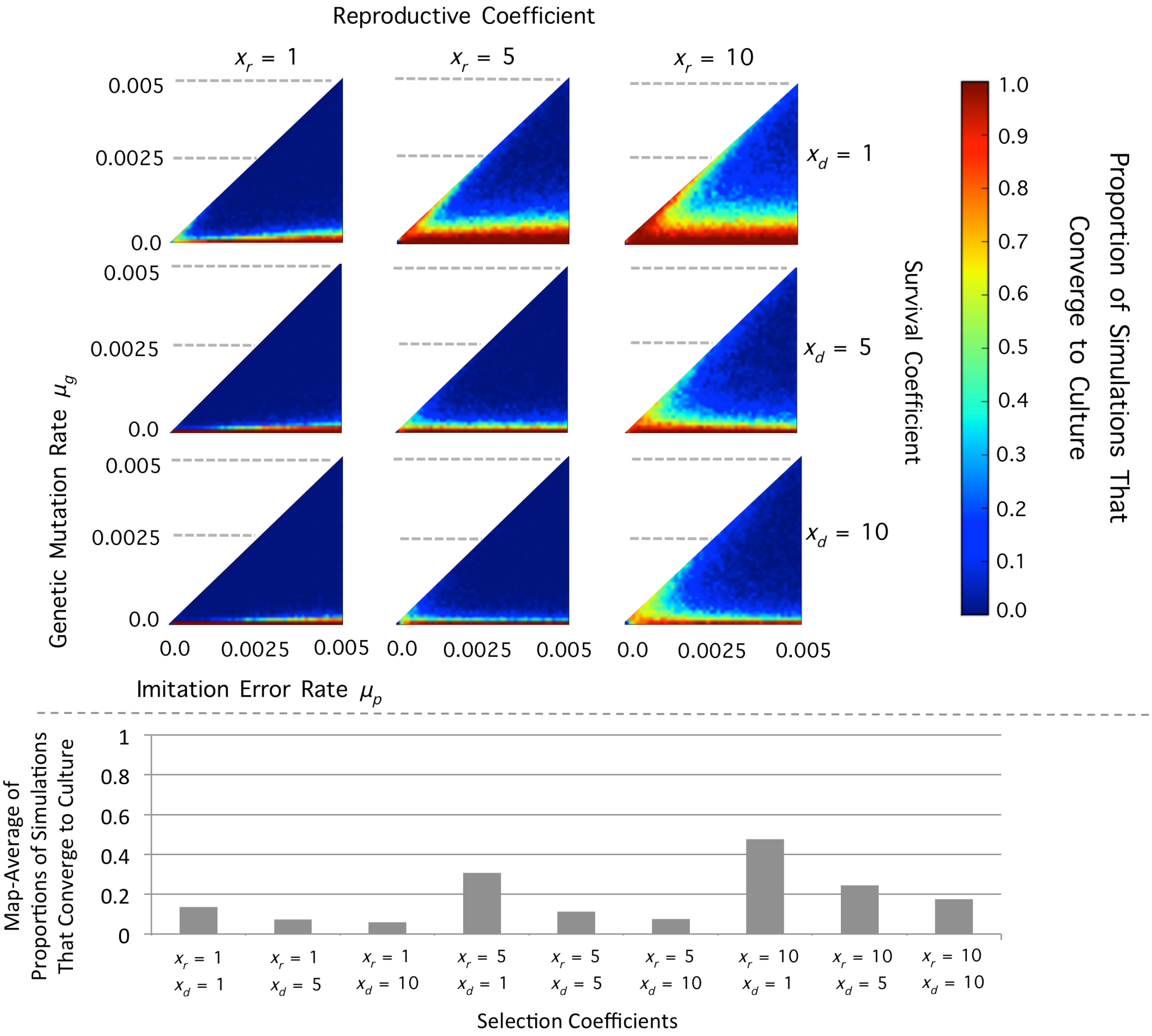

FIGURE 3.9: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where social learners fix in a high proportion of replicates through genotype-phenotype disengagement. From left to right, columns of maps have x_r values of 1, 5 and 10. From top to bottom, rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. The length of the bars represents the average colour value/temperature of the corresponding plot (i.e., the average proportion of simulations that converge to culture for the entire map). Simulations with parameterisations in which (1) survival selection is strong (i.e., simulations with relatively low values of x_d), and (2) reproductive selection is weak (i.e., simulation with high values of x_r), are more likely to converge to culture.

To discuss the mechanisms that underpin the origin and evolution of culture in our individual-based simulation model, it is important to firstly describe what we refer to as *culture* in the context of our results. Here, culture is the aggregate of phenotypes of all individuals in the population *after* the disengagement of phenotype and genotype has occurred. Therefore, in populations where culture has evolved, the behavioural phenotype of individuals is different from the behavioural phenotype they would exhibit if they were expressing their genetic information. In cases B and C of Figure 3.6, the green line after genotype-phenotype disengagement represents the average fitness of the phenotype in the population and can be said to represent culture's average fitness as well. These cases (i.e., cases B and C in Figure 3.6) also show that the emergence of culture occurs at the same time as social learners fix. Note that genotype-phenotype disengagement (i.e., the process by which the number of 1s in the phenotype strings of the population diverges form the number of 1s in the genotype strings), is an outcome of our simulation model and *not* an implemented or externally enforced process.

Our model highlights two important requirements for the emergence of culture. First, a strong selection pressure on survival relative to selection pressure on reproduction, and, second, an error rate for imitation that is greater than that for genetic inheritance. The culture-enhancing effects of these conditions have been discussed in existing literature (Avital and Jablonka, 2005; Richerson and Boyd, 2006; McElreath and Strimling, 2008; Mesoudi, 2011). However, their sufficiency for the origin of adaptive culture has not been discussed.

The existence of a strategy space and a sustained competition between genes and social learners are key features of our individual-based approach that are not present in classical models (Enquist et al., 2008; Mesoudi, 2011; Boyd and Richerson, 1988b; Rendell et al., 2010b; Rogers, 1988). Understanding the effects of these features demands analysis of individual simulation runs. In the next subsection, we explain the underlying principles that make social learners more likely to invade a population of instinctive actuators when survival selection is stronger than reproductive selection. The following two subsections describe two specific mechanisms for the fixation of social learners and the emergence of a non-genetic evolutionary system in our model; these correspond to scenarios B and C in Figure 3.6. Both mechanisms co-occur in all simulations where social learners fix. However, as will be explained, the mechanism that drives scenario B is dominant for simulations where $\mu_p \gg \mu_g$ and the one that drives scenario C is dominant for simulations where $\mu_p \simeq \mu_g$.

3.2.1 Under Relatively Strong Survival Selection Social Learners and Instinctive Actuators Have Equal Selective Pressure

In the results shown in Figure 3.9, we can appreciate a clear tendency for simulations with strong survival selection (i.e., low x_d values) and weak reproductive selection (i.e., high x_r values) to converge towards social learning in larger areas of the parameter space formed by mutation and imitation error rates (i.e., μ_g and μ_p, respectively). Before explaining why the areas of high convergence are distributed in such a way within each heat map, we first here explain why relatively strong survival selection benefits social learners.

Let's first consider that when reproductive selection is dominant, the individuals that reproduce more frequently are, by definition, those that have relatively fitter phenotypes. In a population of instinctive actuators, this means that the phenotype of a newly added offspring is, on average, more fit than the current mean of the population. In contrast, in a population of unbiased social learners, the offspring of the fittest individuals (i.e., those that have imitated fit behaviours) will imitate phenotypes chosen at random from the current population; hence the phenotypes introduced by newly added offspring from social learners will have, on average, the same value as the current mean. For this reason, it is generally held that unbiased social learning causes a regression to the mean compared to genetic reproduction (Richerson and Boyd, 2006; Cavalli-Sforza and Feldman, 1983). Moreover if we also assume that imitation is error-free, it follows that unbiased social learning will reduce the phenotypic variance of the population, a result that is expected if we substitute a randomly chosen phenotype in the population for the copy of another.

When reproductive selection is weak or more precisely, absent, the effect of the regression to the mean caused by unbiased social learning is no longer a disadvantage in comparison to instinctive actuation. In this scenario, the offspring of the latter is also introducing phenotypes that are chosen at random, therefore regressing to the mean to the same extent as social learners do. Moreover, if we assume that genetic replication is error-free, the introduction of the offspring of an instinctive actuator also reduces the variance of the population to the same extent as a newly added social learner with perfect imitation.

When the regression to the mean caused by reproductive selection is equal for unbiased social learners and instinctive actuators (i.e., when the strength of such form of selection is none), the genetic and non-genetic evolutionary systems compete under equal selective pressures (i.e., the selective pressure produced by survival selection) to find fitter solutions in the strategy space. Under these conditions, the relationship between the genetic mutation rate and the imitation error rate becomes a decisive factor in the fixation of either type of agent, as these rates control the overall variance of phenotypes for each evolutionary system (i.e., genetic and cultural). During early stages of optimisation, a high error rate is advantageous (exploration stage); in later stages, where phenotypes are closer to the optimum, a high error rate becomes deleterious (exploitation stage).

We understand that this explanation may be counterintuitive, and therefore unpack our rationale with the following considerations:

1. In a population with no reproductive selection, individuals are selected at random to reproduce, regardless of whether they are instinctive actuators or imitators.
2. A randomly selected individual will have, on average, a survival fitness value equal to the current mean of the population.
3. If the newly introduced offspring is an instinctive actuator (and genotype-phenotype disengagement has not occurred), this offspring will have, on average, a phenotype with mean survival fitness (± mutation).
4. If the newly introduced offspring is an imitator, it will pick a random phenotype to imitate and therefore it will also have a survival fitness value that is, on average, equal to the current population mean (± imitation error rates).
5. Survival selection selects individuals based exclusively on their phenotype regardless of whether they are instinctive actuators or imitators. Fitter individuals survive for longer, and therefore can be unbiasedly selected for reproduction or as models-to-imitate in more iterations.

The logical conclusion of points 1-5 is that the eventual evolution and fixation of social learners depends on whether or not they find fitter solutions than instinctive actuators through imitation errors. For certain combinations of μ_p and μ_g, the exploration advantage of social learners can even offset the relative disadvantage induced by low levels of reproductive selection (i.e., a selective pressure which only improves the evolution rate of instinctive actuators). Therefore, social learners can invade populations of instinctive actuators, even for parameterisations where reproductive selection is not completely absent. As we will see in the next subsection, this needs to occur during the early stages of optimisation, as higher error rates become disadvantageous once phenotypes have genetically evolved above certain quality threshold.

3.2.2 When Imitation Error is Much Higher Than Genetic Mutation, Social Learners Fix During Early Stages of Optimisation

To explain how this process works, we need to define two periods occurring in every simulation, each period having distinctive properties where error rates and survival enhancing traits lead to different outcomes. Plot A in Figure 3.10 shows these periods during the first ten thousand iterations of a simulation run where social learners fix as per scenario B of Figure 3.6. In Figure 3.10, the value of α, which marks the transition between the two periods, has been estimated using the method discussed in Section 3.3. The value of α corresponds to the phenotype sum-of-1s value above which the high imitation error rate of social learners, in balance with selection, produces on average relatively less fit variants than the variants produced by a genetic evolutionary system.

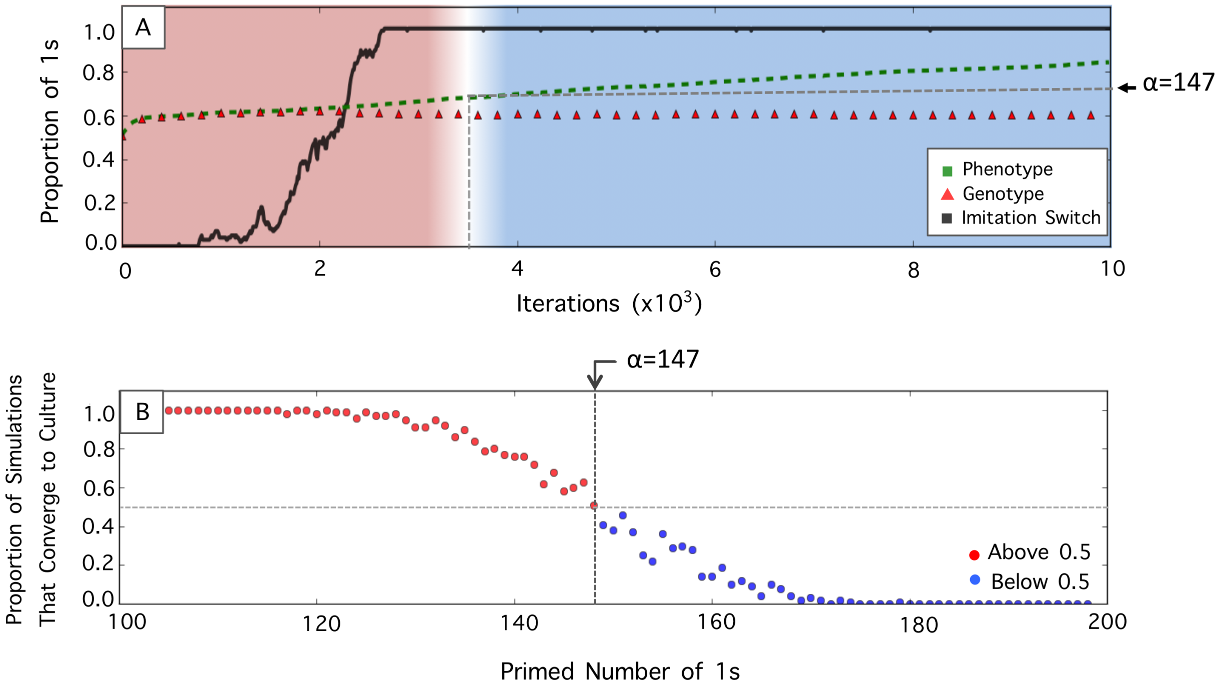

FIGURE 3.10: Plot A shows the first ten thousand iterations of a single run, with two distinct periods highlighted. The green line represents the average proportion of 1s in the phenotypes, the red triangles do the same for the genotypes and the black line represents the proportion of imitators in the population. The first period is here shaded in red and extends from the beginning of the simulation to around iteration 3700 when the second period begins (shaded in blue). During the first period, the exploration advantage of imitators must offset the penalty of their lower selective pressure (i.e., the lower heritability of their imitated behaviour) or imitation will not evolve. If imitators do not fix during the first period, genetic evolution will continue to improve phenotypes towards sum-of-1s values where the high imitation error rate becomes even more disadvantageous. The limit α in the vertical axis corresponds to the phenotype sum-of-1s value, above which, the high imitation error rate of social learners, in balance with selection, produces more deleterious than beneficial errors (the method for the estimation of α is shown in Section 3.3). Consequently, if the initial genetic evolutionary system evolves the phenotype to this value before social learners fix and cause genotype-phenotype disengagement, the non-genetic evolutionary system is not likely to emerge. In Plot B, we show further evidence to support this claim. For each point in this plot, we measure the proportion of 100 replicate runs that converge to culture (i.e., the proportion of simulations that maintain a frequency of social learners above 0.95, after 5×10^4 iterations). For each point, replicates were initialised with a primed number of 1s in their phenotype and genotype strings (horizontal axis). In contrast to our standard simulation setup, populations in plot B start with 50% of their individuals as social learners. This adjustment guarantees that the genetic and cultural evolutionary systems are fully instantiated from the beginning of the simulation run. Note that in our standard simulation setup, social learners need to slowly increase in frequency, and form chains of intimations, in order for the cultural evolutionary system to be instantiated. In plot B, we measure the dominance between cultural and genetic systems as a function of the primed number of 1s in the phenotype strings. We define α as the sum-of-1s value, above which, the genetic evolutionary system, with its higher transmission fidelity, is more likely to dominate over the cultural evolutionary system. Note that the estimated value of $\alpha = 147$ coincides with the transition point where simulations have a 50% chance of converging to culture. For higher sum-of-1s values, *less than* 50% of simulations converge to culture (blue dots). For lower sum-of-1s values, *more than* 50% of simulations converge to culture (red dots). These results are consistent with our calculations as shown in Section 3.3. (Parameter Set: $x_d = 1$, $x_r = 10$, $\mu_g = 0.0005$, $\mu_p = 0.005$).

The first period runs from the start of the simulation when the population comprises only non-imitative, instinctive actuators that map their own genotype into their phenotype and reproduce by replicating their genotype with a mutation rate μ_g. Individuals are selected to reproduce on the basis of their phenotype, and also selected to die on the basis of the same bit string. With a probability μ_c, a new agent can turn its genetically inherited imitation switch from 0 to 1, and hence obtain its phenotype by imitation with error rate μ_p per bit. During this period, three processes occur. First, a marginal frequency of imitators emerges by effect of mutation drift alone. Second, these imitators start exploring the strategy space with error rate μ_p, picking on strategies from imitators and instinctive actuators alike. Third, incipient chains of imitation (i.e., imitators imitating other imitators) arise.

In the first iterations of this period, instinctive actuators are the large majority and their phenotype improves with selective pressure from both reproductive and survival selection. They pass their phenotype to their offspring via vertical genetic inheritance with high fidelity ($\mu_g < \mu_p$). By contrast, imitator phenotypes improve only in response to selective pressure for survival selection. This is because phenotypes expressed by imitators cannot increase their copy numbers by means of reproductive selection as their offspring do not inherit the parental phenotype, but instead imitate any model at random (i.e., there is no assumption of discriminatory, guided, or adaptive imitation in this model). By contrast, survival selection still affects imitators, as a better phenotype reduces the chance of being selected by the death function, and therefore leads to a longer life. Longer-living agents have a higher chance of being used as another imitator's model.

At this stage, imitation has an exploratory advantage, since imitation error is higher than genetic mutation (Figure 3.11). But social learners can only fix if this exploratory advantage overcomes the selection handicap that imitators suffer from. This is more likely to occur if the trait evolved has little effect on fertility but a large positive effect on longevity (i.e., $x_r \gg x_d$).

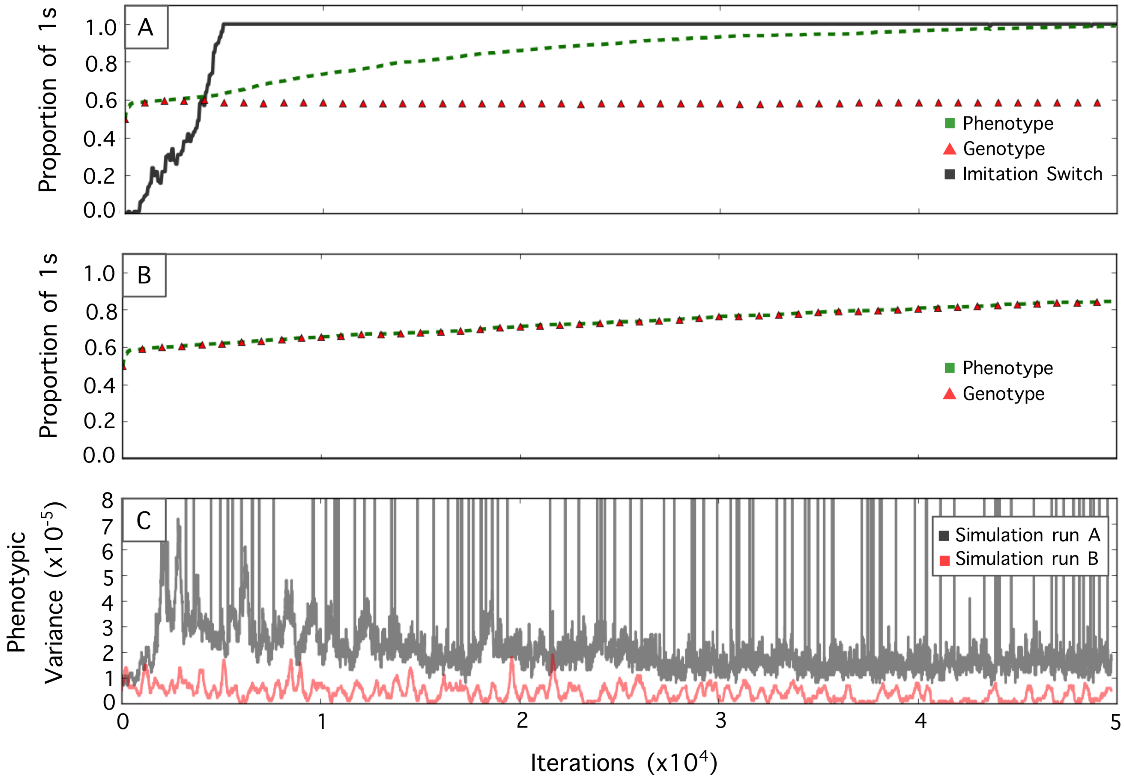

FIGURE 3.11: Comparison between a simulation run where social learning is evolvable (plot A), and a simulation run where social learning is not evolvable (plot B). In plots A and B, the green dotted line represents the proportion of 1s in the phenotypes of the population and the red triangles the proportion of 1s in the genotypes. The black line in plot A represents the frequency of social learners. For both runs, $\mu_g = 0.0005$ and for run A $\mu_p = 0.005$. Selection strength coefficients are set to $x_d = 1$ and $x_r = 10$. Plot C shows the variance of the sum-of-1s values for the phenotypes in the populations of each run. Throughout iterations, run A has a consistent higher phenotypic variance compared to run B. This is explained by the fixation of social learners in run A, and the consequent emergence of the non-genetic evolutionary system in which phenotypes are imitated with an error rate higher than genetic mutation. The grey line (i.e., variance in simulation run A) shows large peaks of variance that extend over the upper limit of the plot. These are caused by mutation events where the offspring of a social learner turns into an instinctive actuator and expresses its genotype (which has been disengaged and hence has a lower sum-of-1s value than the population average) as its phenotype. These mutants are selected against and therefore the variance returns rapidly to the baseline range of fluctuation.

For social learning to evolve and genotype-phenotype disengagement to take place as per scenario B in Figure 3.6, the imitator minority has to find better solutions at a rate that overcomes its disadvantage from having only one source of selective pressure: survival. For all cases where $\mu_p \gg \mu_g$, this needs to occur quickly. Once genetic evolution, guided by instinctive actuators, improves the phenotype beyond a certain number of 1s (i.e., α value in Figure 3.10), the likelihood of imitators taking over the population drops to zero. At this point, the average phenotype solution starts benefiting from lower mutation rates due to its closeness to the optimum and the exploration advantage of social learners turns into a disadvantage. In Figures 3.8 and 3.9, social learners do not fix in simulations where μ_g is above certain values, precisely because the population evolves phenotypes beyond this threshold of number of 1s before social learners can rise in frequency and establish imitation chains. As a result of this process, the area with a high proportion of simulations converging to social learners is limited to a horizontal band at the bottom of the plot.

The second period, shaded in blue in Figure 3.10, starts from the point at which the higher rate of imitation error turns into a disadvantage compared to the mutation rate of genetic replication. If fixation of imitators has already been achieved during the first period, then the genotype has been disengaged and is no longer subject to selection. In this case, we say that a non-genetic evolutionary system has emerged, and only the phenotype will continue to improve with imitation error as its source of variation and survival selection pushing it towards the optimum. This non-genetic evolutionary system manifests as a "pool" of variants maintained by unbiased imitation alone.

An alternative way to understand this mechanism and the importance of the relative difference between imitation error rate and genetic mutation rate is by imagining two hypothetical populations under the same degree of survival selection and in the absence of reproductive selection: one population of unbiased imitators ($P_{imitate}$) and one population of instinctive actuators ($P_{inherit}$).

When both populations are initialised with random genotypes/phenotypes (i.e., populated on average with 50% 1s and 50% 0s), the accumulation of fit alleles in $P_{inherit}$ is limited by the mutation rate, while the accumulation of fit alleles in $P_{imitate}$ is limited by the imitation error rate. If these two rates are equal, while imitators in $P_{imitate}$ can copy any strategy currently in the population, this does not give them an advantage or disadvantage compared with $P_{inherit}$, since the strategies available to be copied in $P_{imitate}$ are no more diverse than the ones that are available to reproduce in $P_{inherit}$.

If the imitation error rate in $P_{imitate}$ were higher than the genetic mutation rate in $P_{inherit}$, then $P_{imitate}$ would initially have an advantage and would accumulate fit alleles at a higher rate than $P_{inherit}$. If, conversely, the mutation rate in $P_{inherit}$ were higher than the error rate in $P_{imitate}$ (an unrealistic assumption. See paragraph 3.2), then $P_{inherit}$ would have an advantage and would accumulate fit alleles faster than $P_{imitate}$.

However, when $P_{imitate}$ and $P_{inherit}$ are both initialised with very fit solutions, the population suffering the highest error rate would be disadvantaged since errors would tend to degrade the fit solutions more often. At this point, if the imitation error rate were lower than the mutation rate, it would always pay $P_{inherit}$ to swap to the higher fidelity mechanism. But for $P_{imitate}$, even if the mutation rate were much lower than the imitation error rate, it would only benefit them to become instinctive actuators if their genotypes were as fit as their phenotypes. But during the time that members of $P_{imitate}$ have been imitators, their genotypes have been under no selection pressure, and will therefore have drifted towards the mean, which means that $P_{imitate}$ is unlikely to give up imitation even when its high error rate is counter-productive. This asymmetry creates a "ratchet" in the system, which ensures that even when the conditions that allowed social learners to successfully invade instinctive actuators have changed, instinctive actuators cannot recover.

In Figure 3.12, we test this regime-change explanation by comparing the results in heat map A, showing our standard simulation results, against heat map B, where simulations are initialised with all phenotypes at a much higher sum-of-1s value (i.e., exactly 175 out of 200 instead of, on average, 100 out of 200). As expected, the bottom band becomes thinner in B, and cases that maintain a large proportion of simulations converging to social learning are only those where: (1) genetic mutation is close to 0, or (2) $\mu_g \simeq \mu_p$ and μ_p is low (i.e., less than 0.0025). The initial sum-of-1s value of 175 was selected for simulations in plot B; larger values do not have a further effect in the distribution of areas where social learners fix.

When genetic mutation is equal to 0, social learning is the only way to optimise phenotypes. Therefore, it is expected that in these marginal cases (i.e., lower red strip in map B of Figure 3.12), social learners will fix even when their high imitation error rate is likely to produce more deleterious than advantageous variants. Areas where $\mu_g \simeq \mu_p$ and μ_p is low (i.e., lower left corner) retain a high proportion of simulations converging to social learning (in plot B of Figure 3.12) due to a secondary mechanism that will be discussed in the next subsection.

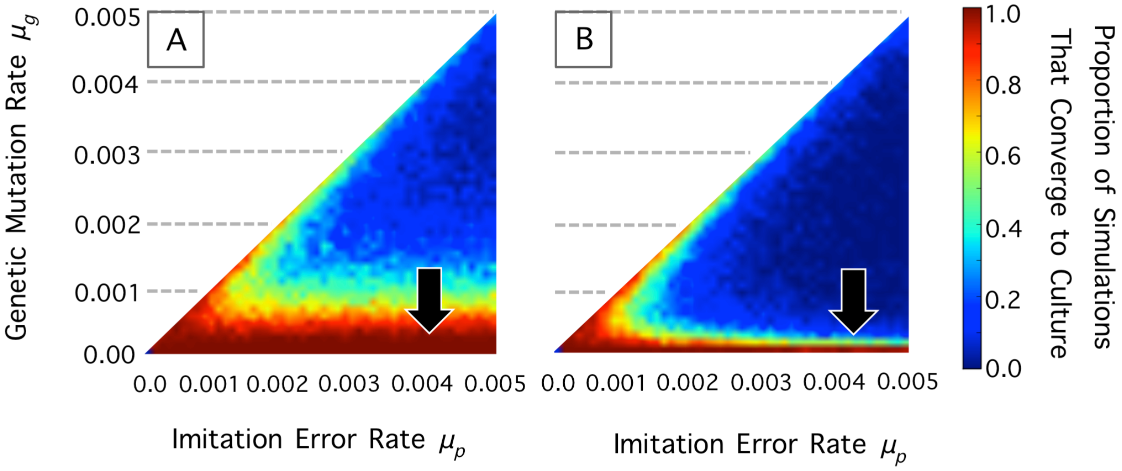

FIGURE 3.12: Comparison between heat map A, where the initial population of instinctive actuators is initialised with an average phenotype proportion of 1s equal to 0.5 (i.e., 50% of bits in the phenotypes are in position 1); and heat map B, where the population is initialised with an exact phenotype proportion of 1s of 0.875. Red areas indicate a high proportion of simulations converging to social learning (i.e., where social learners fix). Simulations with low genetic mutation (but not 0) and relatively high imitation error rates (i.e., areas pointed to by black arrows) converge to social learning in plot A and not in plot B. Simulations in plot B being initialised with phenotypic sum-of-1s values above the critical threshold α (where α is the sum-of-1s value above which the large imitation error rate of imitators produces, on average, less fit variants than the genetic system would), explain this difference. In both maps, the selection exponents are set to $x_d = 1$ and $x_r = 10$.

As further support for our regime-change rationale and to explain the horizontal (i.e., independent of μ_p) transition at $\mu_g = 0.001$ in our original heat map results (see Figure 3.8), we present Figure 3.13, where our standard simulation results in map A are compared against results for simulations where the population is initialised with a 0.5 proportion of imitators (i.e., 50% of the initial population have their imitation switch in position 1) and each individual phenotype is initialised with exactly half of its bits in position 1. These modifications produce two complementary effects that eliminate the relative advantage of instinctive actuators and the genetic evolutionary system over imitators and the non-genetic evolutionary system at the start of the simulation run according to our explanation (i.e., it levels the field between both evolutionary systems, at initial conditions): (1) they increase the probability for social learners to form chains of imitation (i.e., consecutive events of imitators imitating other imitators) from the beginning of the simulation (i.e., half of the population is already a social learner); and (2) they decrease the probability that instinctive actuators evolve phenotypes with sum-of-1s values above the discussed α threshold in the first couple of hundred iterations (i.e., as they would do by evolving the fittest individual phenotype from the initial high-variance population for simulations in map A: where each phenotype bit is initialised with equal chance of being 1 or 0).

In accordance with our explanation (i.e., where we described how initial chains of imitation are required in order to instantiate a cultural evolutionary system), large regions of the parameter space converge to social learning in map B, including areas where μ_g values are above 0.001 (i.e., for simulations that do not converge to social learning in map A). Moreover, as we use larger μ_p values (i.e., large imitation error rates), social learners are able to fix in simulations with increasing values of μ_g (notice the inclined boundary between red and blue regions in map B in contrast to the horizontal boundary at $\mu_g = 0.001$ in map A). Such a pattern is the result of the non-genetic evolutionary system having an equal competitive advantage at initial conditions, and therefore imitators being more likely to find and fix fitter behavioural variants in the early stages of optimisation. As imitation error is set to higher values, the non-genetic evolutionary system maintained by social learners can outcompete genetic systems with higher mutation rates (but never higher than imitation error rates) as per scenario B in Figure 3.6. Once social learners fix, and the non-genetic evolutionary system is instantiated, the process of genotype-phenotype disengagement hampers latter reinvasions by instinctive actuators.

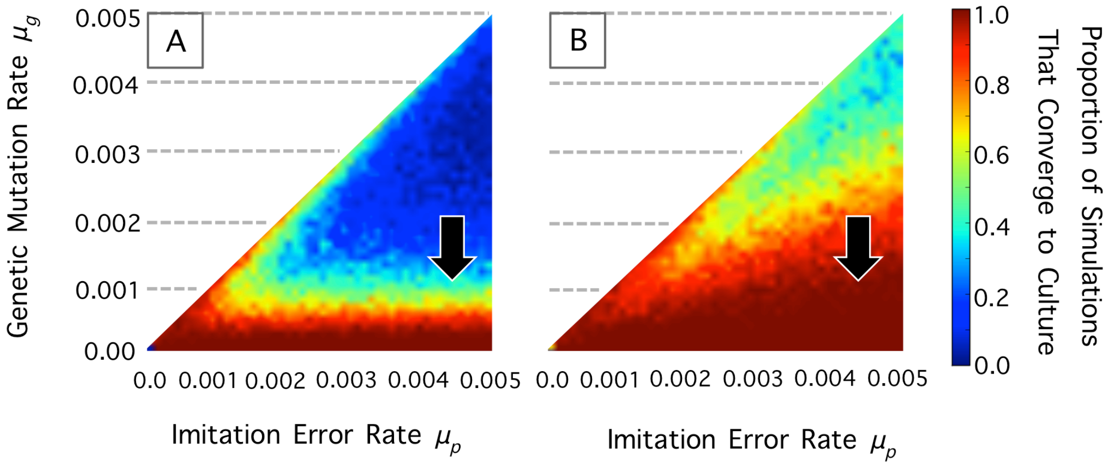

FIGURE 3.13: Comparison between map A, showing our standard simulation results where populations are initialised with all individuals as instinctive actuators and phenotype string bits are set to 1 or 0 with equal probability, and map B, where simulations are initialised with half the population as social learners and all individual phenotypes are initialised with exactly half of their bits as 1s. In map B, social learners fix in areas where the imitation error rate is high, and the mutation rate is above 0.001. For increasing imitation error rate values (μ_p), social learners fix for higher values of mutation rate (μ_g) in map B. Contrasting areas are indicated with black arrows. In both maps, the selection exponents are set to $x_d = 1$ and $x_r = 10$.

3.2.3 Social Learners Can Also Fix Due to Drift and the Irreversibility of Genotype-Phenotype Disengagement

In our model, when social learners fix, this fixation is irreversible. Once imitators have invaded the population and caused genotype-phenotype disengagement, genotypes cannot re-engage in genetic evolutionary adaptation because the information contained in these bit strings drifts at random and is no longer subject to selective pressure (i.e., since it is not expressed in phenotypes). Therefore, once the population starts evolving phenotypes via the non-genetic evolutionary system, it is unlikely to go back to genetic evolution. In scenario C in Figure 3.6, this simple process alone (i.e., with no significant assistance from advantageous high error rates) can fix social learners in the long term.

Scenario C in Figure 3.6 occurs under conditions where both genetic mutation (μ_g) and imitation error (μ_p) are relatively low (i.e., below 0.0025). In these simulations, tens of iterations can pass between mutation (or imitation error) events that introduce fitter variants in the population. Consequently, not only do social learners have a large window of opportunity to fix (as per the mechanism explained in the previous subsection), but also the number of imitators can slowly increase as a result of an asymmetry in the effect of imitation-switch mutations.

To understand this mechanism, we must remember that since social learner genotypes are not expressed, they are shielded from selection and therefore tend to accumulate deleterious mutations. Consequently, the offspring of a long line of imitators is likely to be less fit if it mutates into an instinctive actuator, whereas the same is not true if the offspring of an instinctive actuator mutates into a social learner. Due to the asymmetry between these two mutation events, social learners can fix without any significant explorative advantage.

In Figure 3.14, we prove our rationale by comparing three different heat maps for: (A) our original simulation model, (B) simulations where phenotypes are initialised with a large proportion of 1s (i.e., 0.875 rather than 0.5), and (C) Lamarckian simulations in which the genotype of a social learner is updated to be a perfect copy of the phenotype that it achieved through social learning. Comparisons between map (A) and (B) (i.e., same plots from Figure 3.12) serve to highlight areas where a large proportion of simulations still converge to social learning despite phenotypes being initialised past the critical threshold (α) discussed in the previous section. The remaining red areas in map (B) correspond to cases where: (1) scenario C in Figure 3.6 occurs ($\mu_g \simeq \mu_p$), and (2) genetic mutation is non-existent ($\mu_g = 0$).

Forcing social learners to keep their genotype string as a copy of their phenotype prevents genotype-phenotype disengagement in map C of Figure 3.14. Under such a condition, social learners do not fix for any area in the map. In these simulations, genotype and phenotype effectively become a single string. Social learners are however still distinct from instinctive actuators, as the former obtains its phenotype through horizontal transmission (i.e., non-parental imitation), whereas the latter does it through vertical genetic inheritance.

In plot F of Figure 3.14, we show a single representative run where the prevention of genotype-phenotype disengagement (i.e., by keeping phenotypes and genotypes identical within individuals) makes the invasion of social learners reversible (compare to plot D in Figure 3.14). In such a situation, it is unlikely that individual simulations would converge to the fixation of social learners (i.e., would have a final frequency of social learners that is higher than 0.95 after 5×10^4 iterations). Instead, we observe a permanent "drifting" pattern where social learners fluctuate in frequency throughout the simulation run. In the simulation run presented in plot F, this "drifting" pattern continues indefinitely (i.e., it continues after 5×10^4 iterations).

The results shown in maps B and C, and plots E and F of Figure 3.14 serve to illustrate our rationale with regards to the fixation of social learners through the slow drifting mechanism explained in this subsection. This mechanism depends on a fundamental asymmetry between social learners and instinctive actuators, where the former is likely to irreversibly fix in a population even when having no significant advantage in comparison to the latter.

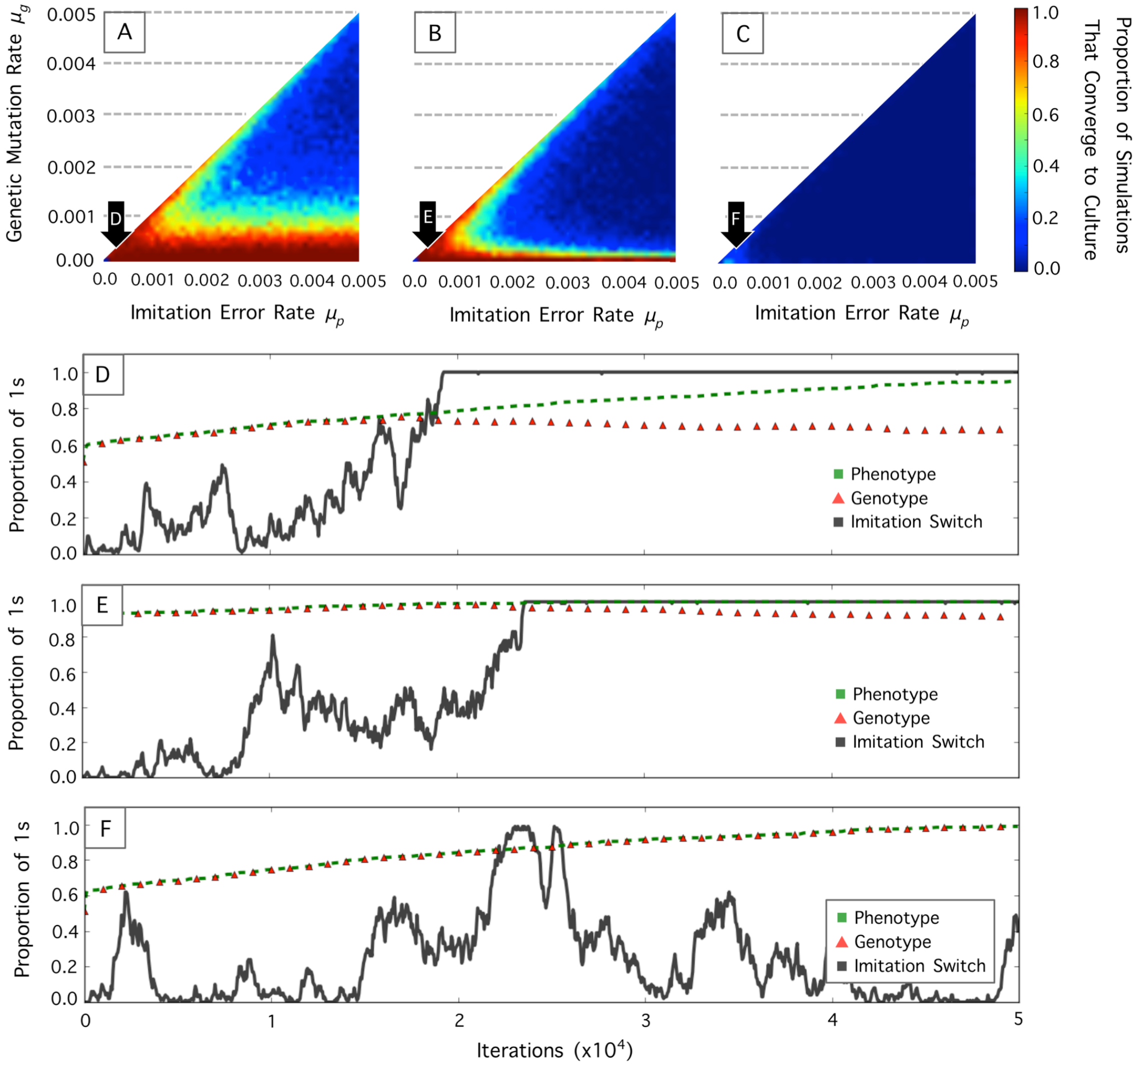

Figure 3.14: Comparison of results from three different simulation set-ups demonstrates the effect of genotype-phenotype disengagement on the fixation of social learners. In maps A, B and C, red colours indicate combinations of μ_g and μ_p for which a large proportion of simulations converge to social learning. Plots D, E and F show individual simulation runs where $\mu_g = 0.0002$ and $\mu_p = 0.0002$ for respective maps A, B and C (parameter-space position pointed by black arrows). Map A shows results for our standard simulation model where the initial population of instinctive actuators is initialised with an average phenotype proportion of 1s equal to 0.5 (i.e., 50% of bits in the phenotypes are in position 1) and where genotype-phenotype disengagement is not hampered. In map B, disengagement is also not hampered, but the population is initialised with an exact phenotype proportion of 1s of 0.875. In this map, areas with low (but not zero) values of μ_g and relatively high values of μ_p do not converge to social learning as per scenario B in Figure 3.6. However, initialising phenotypes with a large proportion of 1s does not affect areas with low and similar values of μ_g and μ_p. In map C, simulations are initialised with an average phenotype proportion of 1s equal to 0.5 but genotype-phenotype disengagement is hampered (i.e., the genotype of social learners is forced to be a copy of their imitated phenotype). This alteration makes the invasion of imitators reversible; and therefore the fixation of social learners through the "drifting" mechanism explained in this subsection is hampered (as shown in plot F). For simulations shown in plots D and E, the genotype strings slowly converge towards the neutral average (i.e., sum-of-1s = 0.5), in the long term. In all maps and plots, the selection exponents are set to $x_d = 1$ and $x_r = 10$.

3.3 Analytical Estimations: Testing Our Mechanistic Explanation

In Figure 3.10, we showed how simulation runs are separated into two different periods represented by the red and blue shaded areas in the plot's background. However, the point that separates these two regions (i.e., the white shaded strip) has only been described in qualitative terms so far. In this section, an analytical estimation of this point is derived and compared with our experimental results.

Let us define α as the average sum-of-1s value for the phenotype strings in a population at the point that separates the red and blue areas in Figure 3.10. According to our mechanistic description, a population that genetically evolves its phenotype beyond this point (i.e., a population that has not experienced genotype-phenotype disengagement when reaching $L_1 = \alpha$) is a population that we would not expect to develop culture. In other words, α is the average sum-of-1s value where the offspring of a social learner would have, on average, the same fitness as the offspring of an instinctive actuator.

In our simulations, agents start with a sum-of-1s value of half the length of their bit strings (i.e., $L_1 = 100$). However, not all parameterisations where $\alpha > 100$ are expected to converge to the fixation of social learners. The value of α should be understood as a boundary rather than a predictor for culture (i.e., if social learners fix in a population, we expect them to do so before the population has reached α).

In the following paragraphs, we derive an expression for the expected value of α that at first glance will seem to be in disagreement with our experimental results. However, further analysis shows how initial conditions and particular stochastic processes in our simulations can explain the areas where this disagreement is apparent. We conclude this section by showing that once these factors have been taken into account, our expected α values work as a boundary for the evolution of culture according to the mechanism explained in the previous section.

To estimate α, we need to calculate the expected fitness of instinctive actuators and social learners considering the strength of selection and their respective mutation and imitation error rates. We will start by calculating the expected frequencies for all values of sum-of-1s created by the variation functions on each system (i.e., mutation and imitation error). After this first step, each frequency is multiplied by its respective fitness value, as calculated considering reproduction and survival selection. The α value is then the average sum-of-1s value in a population for which the expected fitness of cultural and genetic systems are equal.

Let $p(i|L, \mu)$ be the probability of i bits changing their value in a string of length L with a bit-flip rate of μ. In other words, $p(i|L, \mu)$ is the probability a bit-string of length L

moves a *hamming distance* of magnitude i after a bit-flip mutation of rate μ (Equation 3.3).

$$p(i|L,\mu) = \frac{L!}{i!(L-i)!}\mu^{i}(1-\mu)^{L-i} \tag{3.3}$$

In our model, the loss of fitness caused by losing one unit in the sum-of-1s value is not equal in magnitude to the gain in fitness caused by gaining one (i.e., our fitness functions for survival and reproduction are non-linear functions of the sum-of-1s value). This means that calculating the number of changes occurring during mutation is not sufficient for our estimation. We must separate the probabilities of obtaining strings with sum-of-1s values above and below the current value. To do this we differentiate the two types of flip that may occur in a bit string (i.e., flips from 0 to 1 vs. flips from 1 to 0).

Let $q(b|a,\mu)$ be the probability that a string with sum-of-1s value a turns into a string with sum-of-1s value b under bit-flip mutation μ, where $b = a + (k - i)$ with k being the number of 0s that turned into 1s and i the number of 1s that turned into 0s. Using Equation 3.3, $q(a + (k - i)|a, \mu)$ can be expressed as the product of the probability of existing 1s turning into 0s (i.e., $p(i|a,\mu)$) multiplied by the probability of existing 0s turning into 1s (i.e., $p(k|L-a,\mu)$) (Equation 3.4).

$$q(a+(k-i)|a,\mu) = \frac{a!}{i!(a-i)!}\mu^{i}(1-\mu)^{a-i}\frac{(L-a!)}{k!(L-a-k)!}\mu^{k}(1-\mu)^{L-a-k} \tag{3.4}$$

To estimate α, we need to compare the expected next-generation fitness value for the phenotype distribution produced by genetic mutation, against the value produced by the distribution created by imitation errors. These expected fitness values are calculated by adding the products of the probabilities of producing strings with sum-of-1s values b multiplied by the corresponding fitness values for $L_1 = b$, for all possible values of b.

As we explained in the previous section, the genetic system is affected by both survival and reproductive selection. With this in mind, we estimate the fitness value of an individual with sum-of-1s value b living in a population where the average fitness value is equal to that produced by a sum-of-1s value a, as $\omega_g(b,a)$ (Equation 3.5).

$$\omega_g(b,a) = \frac{e^{\frac{b}{x_d}}}{e^{\frac{a}{x_d}}}\frac{e^{\frac{b}{x_r}}}{e^{\frac{a}{x_r}}} \tag{3.5}$$

Conversely in a cultural system, that is an evolutionary system where fitness is only affected by survival selection, the expected fitness value of a string with sum-of-1s b is $\omega_p(b,a)$ (Equation 3.6).

$$\omega_p(b,a) = \frac{e^{\frac{b}{x_d}}}{e^{\frac{a}{x_d}}} \tag{3.6}$$

In our equations, we simplify the estimation of α by assuming that all individuals have the same sum-of-1s value, to calculate the next-generation distributions for the genetic and cultural systems. This simplification saves us the complicated steps involved in the calculation of the exact form for the moving distribution of sum-of-1s during the optimisation process. As we will see when comparing analytical estimations and experimental results in Figure 3.18, this simplification does not introduce large errors.

The expression for the expected next-generation average fitness value of an individual with sum-of-1s value a in a genetic system (i.e., $\omega_{gw}(a)$) is then estimated as the sum of the product for all probabilities of obtaining strings with sum-of-1s values $b = a + (k - i)$ multiplied by their respective expected fitness.

$$\omega_{gw}(a) = \sum_{k=0}^{L-a} \sum_{i=0}^{a} q(a + (k-i)|a, \mu_g)\omega_g(a + (k-i), a) \tag{3.7}$$

The equivalent expression for a cultural system (i.e., where only survival selection affects fitness) is $\omega_{pw}(a)$ in Equation 3.8.

$$\omega_{pw}(a) = \sum_{k=0}^{L-a} \sum_{i=0}^{a} q(a + (k-i)|a, \mu_p)\omega_p(a(k-i), a) \tag{3.8}$$

In our simulations, the length of all bit-strings is 200 ($L = 200$) and, on average, agents start a simulation run with a bit sequence where 50% of their bits are 1s (i.e., each bit is initialised as 1 or 0, with equal chance. See Section 3.1). This means that α values of interest (i.e., α values we can compare with our results) are those between 100 and 200. Furthermore, for culture to evolve by the mechanism we described in the previous section, $\omega_{pw}(a)$ must be greater than $\omega_{gw}(a)$ for $a = 100$. In other words, the cultural system must have a fitness advantage over its genetic counterpart at the beginning of every simulation run. Using Equations 3.7 and 3.8, we can numerically find α by minimizing the difference between $\omega_{pw}(a)$ and $\omega_{gw}(a)$ in the range $100 < a < 200$.

Figure 3.15 shows the value for $\omega_{pw}(a)$ and $\omega_{gw}(a)$ for $100 < a < 200$, where $\mu_g = 0.0001$, $\mu_p = 0.0002$, $x_d = 1$ and $x_r = 10$. The point where the green and red lines cross corresponds to α. We predict that a simulation with the pattern shown in this figure can only evolve culture if social learners fix *before* the average sum-of-1s in the population has reached α.

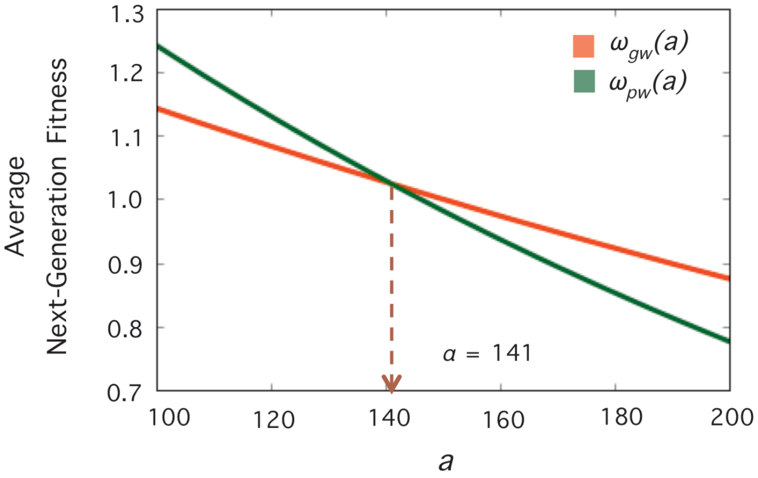

FIGURE 3.15: Values of the expected average fitness in a population of instinctive actuators as a function of a (i.e., $\omega_{gw}(a)$) are shown in red. The corresponding expected fitness for an alternative system formed by social learners (i.e., $\omega_{pw}(a)$) is shown in green. The intersection between these two lines represents α. In simulations with these parameters, we expect social learners to fix, if they do at all, *before* the genetic system moves the phenotype sum-of-1s beyond this point. Notice that $\omega_{gw}(100) < \omega_{pw}(100)$ is a requirement for culture to evolve according to our mechanism. Parameters used: $x_d = 1$, $x_r = 10$, $\mu_g = 0.0001$, $\mu_p = 0.0002$.

For cases where $\omega_{gw}(100) > \omega_{pw}(100)$, culture does not evolve according to our mechanism. Cases where $\omega_{gw}(100) > \omega_{pw}(100)$ and where $\omega_{gw}(a)$ and $\omega_{pw}(a)$ "cross" in the range $100 < a < 200$ occur for the parameter region where $\mu_g > \mu_p$, but as discussed above, this is a condition that we assume to be unrealistic (see Paragraph 3.2).

In Figure 3.16, the estimated α values for combinations of μ_g and μ_p drawn from the set $[1 \times 10^{-4}, 2 \times 10^{-4}, \ldots, 50 \times 10^{-4}]$ and where $\mu_p > \mu_g$ are shown side-by-side with our simulation results for $x_r = 10$ and $x_d = 1$. The blue regions in the right-hand figure represent cases where $\omega_{gw}(100) > \omega_{pw}(100)$ and the evolution of culture is therefore not anticipated. In these regions, α is undefined, but is set to 0 for plotting purposes.

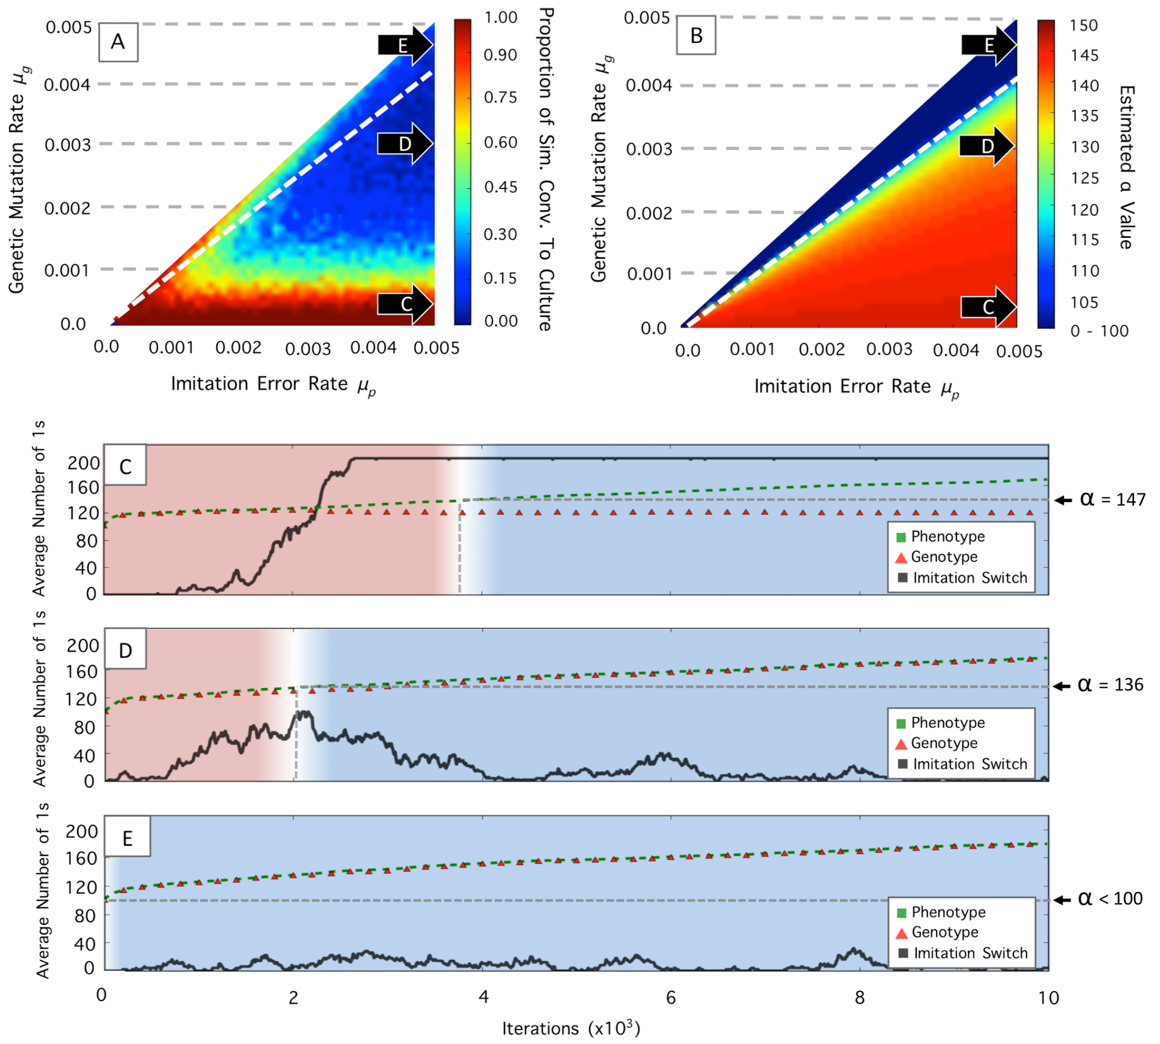

FIGURE 3.16: Map A shows the simulation results described in Figure 3.8; red indicates a large proportion of simulations converging to culture. Map B shows the respective α values as calculated in our estimations. For all parameterisations where $\omega_{gw}(100) > \omega_{pw}(100)$, the value of α is undefined in our range of interest and hence set to 0 for plotting purposes. The white dotted lines indicate the point at which α transitions to the region $100 < \alpha < 200$. Comparisons between these two plots show a large area of mismatch between regions where $\alpha > 100$ (on map B) and combinations of μ_g and μ_p with a high proportion of simulations converging to culture (on map A). Black arrows with labels C, D and E point to parameterisations for which respective plots C, D and E are representative simulation runs. On each of these plots the red region represents the period of optimisation for which an instantiated cultural evolutionary system would improve phenotypes faster than its genetic counterpart (i.e., while the average number of 1s in the phenotypes is lower than α). Conversely, the blue region indicates the period in which the genetic system produces phenotypes superior to those generated by social learning. Notice that a higher value of α increases the length of the red region and therefore the chance for a cultural evolutionary system to emerge before phenotypes genetically evolve beyond α. However, for a large region of the parameter space in which $\alpha > 100$, culture does not emerge. Plot D shows a representative simulation run where this is the case. In these parameterisations, the genetic evolutionary system evolves phenotypes too rapidly for the initial minority of social learners to increase in frequency and consolidate a competitive cultural evolutionary system before phenotypes reach α. Notice that in plot D, genotype-phenotype disengagement started to occur around iteration 2×10^3, but then stopped and reversed. According to our explanation, this is the result of instinctive actuators in the population reaching α before the complete fixation of social learners occurred. Plot E shows a representative simulation run for cases in which the emergence of culture is not expected to occur. For these cases, $\alpha < 100$. In all our simulations, the average sum-of-1s value in the phenotype strings is initialised at 100. Therefore, we do not expect culture to emerge in case E, as phenotypes already start with sum-of-1s values greater than α (i.e., within the blue-shaded period). Parameters used: $x_d = 1$ and $x_r = 10$.

In Figure 3.16, the white dashed line derived in map B does not do a good job of separating runs that result in culture from runs that do not. Figure 3.17 shows similar areas of mismatch for other combinations of x_r and x_d (compare to Figure 3.9).

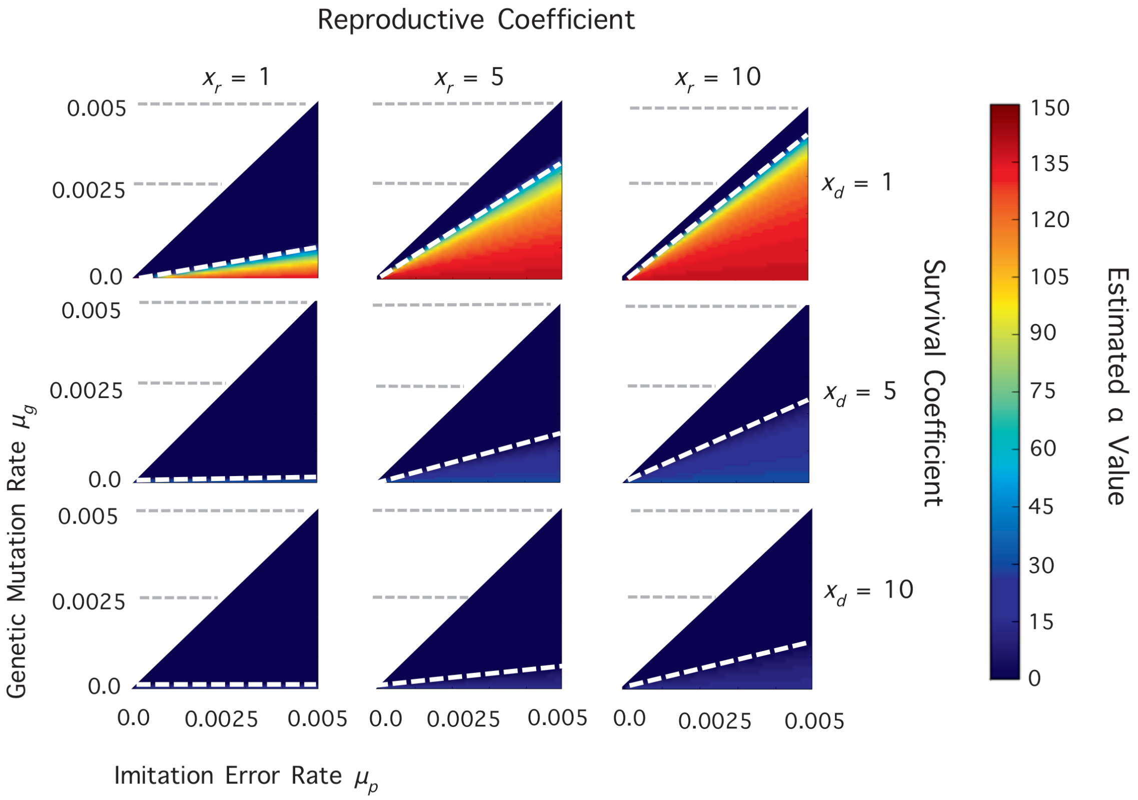

FIGURE 3.17: Heat map array showing estimated α values for combinations of genetic mutation rates, imitation error rates, reproductive coefficients and survival coefficients. From left to right, columns of maps have x_r values of 1, 5 and 10. From top to bottom, rows of maps have x_d values of 1, 5 and 10. In each sub-plot, a range of combinations for μ_g and μ_p is explored for the corresponding x_r and x_d values. For parameterisations where $\omega_{gw}(100) > \omega_{pw}(100)$, the value of α is set to 0. The white dotted lines indicate transitions between $\alpha = 0$ and $\alpha > 100$ areas. When comparing this figure with our simulation results (Figure 3.9), large areas of mismatch between cases where $\alpha > 100$ and cases where a large proportion of simulations converge to culture are noticeable.

We must remember that Equations 3.7 and 3.8 are approximations of the next-generation average fitness in stand-alone genetic and cultural systems (i.e., as opposed to an hybrid system), a simplification that does not capture the full range of stochastic processes and interactions that the hybrid system produces in our simulations. Therefore, α must be understood as a bound defining the point, during individual runs, after which we do not expect social learners to fix (if they have not already). Values of α larger than 100 are then not a predictor for the emergence of culture.

In the following paragraphs, we will explain the reasons for the mismatch observed between our analytical and simulation results. We further show that for simulation runs where culture evolves (i.e., cases where social learners fix in the population), the point of fixation of social learners (i.e., the sum-of-1s value of phenotypes at the point where social learners invade) is bounded by our analytically estimated α values. We conclude this section by showing that the simulation results are a rough subset of the analytical ones.

The discrepancies found between regions where $\alpha > 100$ and regions where culture evolves (i.e., discrepancies of red areas between maps in Figure 3.16) are caused by the effect of relatively large genetic mutation rates. In these cases, the population reaches the critical value of α through genetic evolution *before* the initial minority of social learners reaches the critical mass required to establish a competitive cultural system. In cases where a critical number of social learners already exist at the beginning of the simulation, or where the target sequence is continuously changed, this discrepancy is reduced (see Figure 3.18). In other words, for some cases where $\alpha > 100$, there is a window for social learners to exploit, but parameters and initial conditions ensure that these agents cannot move fast enough to make use of such a window.

In support of this explanation, Figure 3.18 shows a pair of maps similar to Figure 3.16, but in this case the experimental results are obtained from simulations where the optimal bit string sequence (i.e., the sequence we initially established as a string with all bits in position 1) is changed with a given rate. This is a change that will be addressed extensively in Chapter 5. In a model with this extension, the fitness values are recalculated as a function of the match between an agent's phenotype and the changing optimal sequence. In these simulations, it takes longer for the genetic system to reach α because the target is moving around. Therefore, social learners are able to increase to a minimal density and establish an incipient cultural system before $L_1 = \alpha$. The area of discrepancy between analytical estimations and experimental results is thus expected to be smaller.

Note that a rate of change for the optimal sequence can be understood as equivalent to a rate of environmental change. In this chapter, we use the example in Figure 3.18 to explain why some areas in Figure 3.17 where $\alpha > 100$ do not evolve culture. However, in Chapter 5, we will further explain and discuss the effects of environmental change on the evolution of culture by analysing the effects of this model extension.

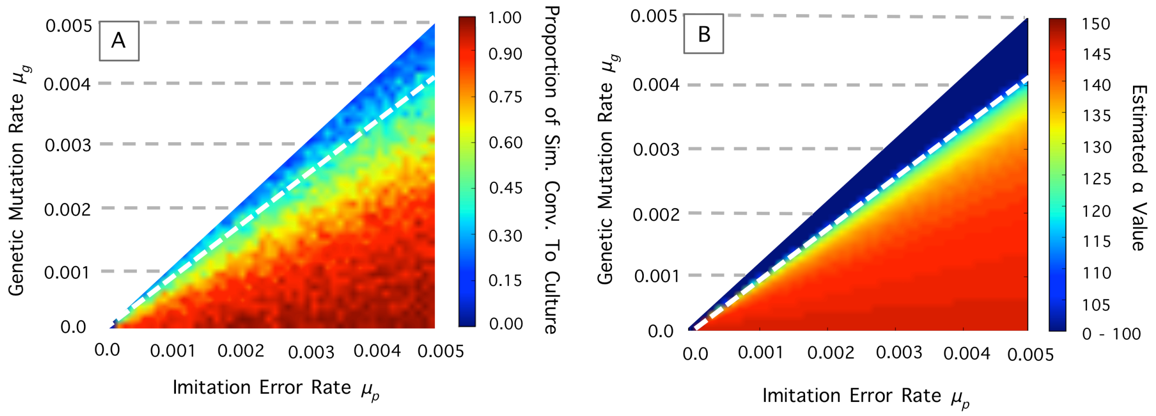

FIGURE 3.18: The left plot shows the proportion of simulations that converge to culture for specific μ_g and μ_p values in the range $[0.0, 0.005]$. The simulations used to construct this plot include a rate of change for the optimal bit-string sequence against which each agent's phenotype is match-compared when calculating their fitness value (i.e., the equivalent to an unchanging string of all 1s in our original model). The plot on the right shows the respective α values as calculated in our estimations. For all parameterisations where $\omega_{gw}(100) > \omega_{pw}(100)$, the value of α is set to 0. The white dotted lines indicate the transition between $\alpha = 0$ and $\alpha > 100$ areas. In contrast with Figure 3.16, the area of mismatch between cases with high convergence to culture in our simulation results and cases where $\alpha > 100$ in our estimations is reduced. Parameters used: $x_d = 1$ and $x_r = 10$.

To further test the validity of the mechanism presented in Subsection 3.2.2 (i.e., where social learning evolves due to an exploratory advantage during the early stages of optimisation), we now need to compare the average sum-of-1s in the population phenotypes when social learners invade (i.e., when genotype-phenotype disengagement occurs) with our estimated α values. According to our mechanism, we should expect this average sum-of-1s to be less than the corresponding bound: α.

Figure 3.19 shows the surface created by our estimated α values for $x_r = 10$ and $x_d = 1$ (3D-map A) and the same surface plotted together with a_c (3D-map B). We define a_c as the average sum-of-1s value for the phenotype strings at the point where the frequency of social learners rises above 50% during a simulation run (i.e., the phenotype sum-of-1s value when genotype-phenotype disengagement occurs).

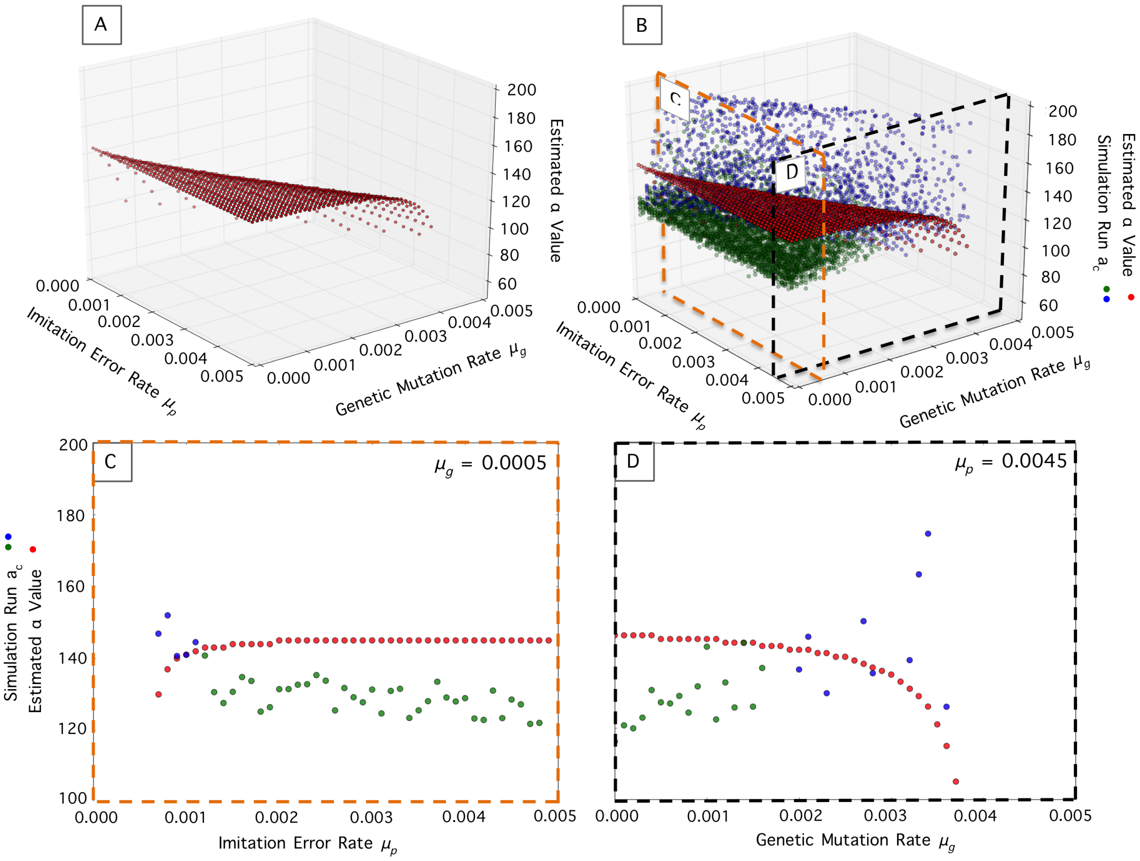

FIGURE 3.19: Plots showing a perspective of the 3D space formed by mutation rates μ_g, imitation error rates μ_p and α values. Map A shows the estimated value of α for combinations of μ_g and μ_p in the range $[0.0, 0.005]$. When projected on the $\mu_g\mu_p$-plane, the surface formed by these red points corresponds to the red region shown in map B of Figures 3.16 and 3.18, where $\alpha > 100$. Plot B shows the same α estimates along with the a_c values registered on simulation runs with corresponding values of μ_g and μ_p. We define a_c as the average sum-of-1s value for the phenotype string at the point where the frequency of social learners rises above 50% during a simulation run (i.e., the phenotype sum-of-1s value at genotype-phenotype disengagement). Each green or blue point represents a simulation run where culture emerged. In most simulation runs, genotype-phenotype disengagement occurs before the average phenotypic sum-of-1s value in the population reaches α. Blue points represent simulations that are in close proximity to the diagonal where $\mu_g = \mu_p$. In these simulations, culture evolves through the "drifting" mechanism explained in Subsection 3.2.3; therefore, we do not expect a_c to be upper-bounded by α. Plots C and D show two-dimensional slices of the 3D space in plot B. In plot C, we fixed $\mu_g = 0.0005$ and showed a_c and α values for a range of μ_p values. In plot D, we fixed $\mu_p = 0.0045$ and also showed a_c and α values, but for a range of μ_g values. These plots, which are in agreement with our explanations, show that most simulations converge to culture at a point of the optimisation process in which the sum-of-1s value of the phenotype strings is still below the estimated α value. Parameters used: $x_d = 1$ and $x_r = 10$.

The red "surface" in Figure 3.19 corresponds to the triangular region of cases where $\alpha > 100$ shown in Figure 3.16 but seen from a lateral perspective. In this figure, instead of using a colour gradient to identify α values, we directly use positions in the vertical axis.

Figure 3.19 B shows the simulation results for cases where $\mu_g < \mu_p$, each green and blue point representing a single simulation run. The vertical axis projection of these points is the population average sum-of-1s value for the phenotype bit-string as calculated when the frequency of social learners rose above 50% for the first time (a_c) (i.e., the average sum-of-1s value when culture emerged).

Figure 3.19 supports our mechanistic explanation of how culture emerges in the simulation model. The presence of a large aggregation of green points located *below* the surface formed by α estimations is a clear indication that culture does indeed emerge due to its initial advantage as an evolutionary system compared to its genetic counterpart. This cloud corresponds to the red band at the bottom of the simulation results plot in Figure 3.16. For culture to evolve in these cases, social learners must have fixed before the genetic system alone has reached α. Once social learners fix in the population, the discussed process of genotype-phenotype disengagement maintains culture. The pattern observed in Figure 3.19 is also present for other combinations of x_r and x_d.

The relatively few cases where culture evolved above our estimated "α surface" are all in close proximity to the diagonal where $\mu_g = \mu_p$, these cases are represented by blue dots in Figure 3.19. The evolution of culture in regions close to this diagonal can be attributed to chance (i.e., the stochastic nature of our simulation) rather than to a specific mechanism, since both the genetic and cultural system have very similar parameters. Once social learners fix by chance, genotype-phenotype disengagement makes instinctive actuators unlikely to re-invade the population (see Subsection 3.2.3 for a detailed explanation). As in cases where $\mu_g > \mu_p$, we consider parameterisations close to this diagonal somewhat unrealistic, as they imply similar values for imitation error rates and genetic mutation rates.

3.4 Limitations

As explained in the introduction of this thesis, our model intends to show how an evolutionary system driven by social learning can evolve to displace its genetic counterpart without relying on decision-making traits. However, in pursuit of simplicity, we have intentionally overlooked some known processes and characteristics of natural populations, including: complex fitness landscapes and individual learning. We believe the conclusions derived from our model hold even when taking into account these factors, but nevertheless find it important to discuss their potential effects. These discussions frame our model as a theoretical contribution rather than a predictive tool in the context of natural populations where a variety of processes might simultaneously affect evolutionary dynamics.

The topology of the fitness landscape (i.e., the phenotype-fitness map) is an important factor to take into consideration when evaluating evolutionary algorithms (DeJong, 2002). In this chapter, we have decided to limit our analysis to a monotonic and static fitness landscape where the optimal fitness value is produced by a string of all 1s. The problem represented by this topology is known as the *onemax* problem (DeJong, 2002). In Chapter 5, we extend our model by considering a dynamic landscape where the optimal string sequence changes throughout a simulation run, but this change does not fundamentally affect the topology of the fitness landscape. In this dynamic landscape, at any given point in time, agents are still obtaining a monotonic increase in fitness as a function of how many bits in their phenotype match the optimal sequence.

Limiting our analysis to the onemax problem throughout this thesis leads to a better understanding of the processes at play in our model. However, we cannot claim our results are extendable to *all* possible fitness landscapes. Nevertheless, we hypothesize that the advantage of cultural systems with relatively high imitation error rates, over genetic evolutionary systems, can be said to hold true even when a certain degree of "roughness" is introduced. We assume this is due to social learners having a greater chance of jumping fitness valleys with their relatively large imitation error rates compared to instinctive actuators.

Previous work in the field of "administrative science" has tried to analyse the effect of social learning on evolutionary dynamics when applied to rough fitness landscapes (Lazer and Friedman, 2007). Their conclusions however suggest that populations of social learners have less chances of overcoming fitness valleys compared to instinctive actuators. These types of work argue that in systems where individuals imitate one another, the overall phenotype variance diminishes, a process that keeps the population stuck at local optima. We believe our model reaches a different conclusion because we have included an explicit imitation error rate that gives social learners the ability to jump across fitness valleys, whereas other authors assume imitation occurs with perfect fidelity.

As mentioned in Chapter 2, and in contrast to other models in the literature, we have not included individual learning as a competing strategy. In the context of our model, individual learners are agents that acquire better-than-average behaviours by paying a fixed cost in fitness, this cost representing the process of trial-and-error involved in acquiring the behaviour itself. We have omitted this type of learning for two reasons. First and foremost, we wanted to avoid introducing any external advantage for social learners compared to instinctive actuators; the presence of individual learners indirectly gives social learners an advantage over their genetic counterpart as the former has access, through imitation, to the better-than-average solutions individual learners find. Second, including an individual learning strategy adds an extra set of variables to our model, which introduces unnecessary complications when answering the research questions of this thesis. The claims of our work revolve around the merits of social learning as a stand-alone evolutionary system able to find solutions faster than genes. Including individual learning introduces confounding effects between the two learning strategies, which are not conducive to our objective. Furthermore, individual learning allows agents to improve on instinctive or acquired behaviours, an ability that can be understood as a general form of guided variation, a decision-making capability. Introducing individual learning would therefore force us to consider social learners, in the hypothetical species being modelled, as able to have some form of guided variation (i.e., we assume that the neurological machinery that supports one capability is likely to support the other).

In natural populations, several species that exhibit culture also have some degree of individual learning (Richerson and Boyd, 2006; Jablonka and Lamb, 2006; Mesoudi, 2011). Although we warrant our results might change when considering the presence of this learning strategy, they still provide a theoretical insight into a previously unidentified advantage of social learners over instinctive actuators. Our model is admittedly more applicable, although still largely simplified, to scenarios where individual learning is an unviable strategy. These scenarios include cases where the cost of an error in the trial-and-error process of individual learning is large enough to offset the net benefit of applying such a learning strategy in the first place (e.g., a mountain goat learning which path to choose on a steep rock face where any misstep can lead to death) and cases where the cognitive requirements to solve a particular problem are far greater than those present in any individual in the population.

3.5 Discussion

We understand that our model is an extreme simplification of a natural process. It captures, however, a set of sufficient conditions with far less assumptions than previously thought for cultural systems to emerge. Our results suggest that populations encountering a new environmental challenge can evolve culture, provided that potential solutions in the strategy space enhance survivability more than fertility, and that their variation during social transmission tends to be higher than that of genetic inheritance, consequently, affording social learners an exploration advantage. In our results, disengagement should not be interpreted as the complete substitution of entire genomes by cultural information. After all, culture cannot survive without its biological vehicles. We rather consider the bit strings in our model to represent solutions to specific challenges that can be solved behaviourally through either imitated traits or genetically encoded traits.

In nature, the map between genetic information and the expressed phenotype is incredibly complex (Pigliucci, 2010). We are yet to untangle the intricate relations within and between genes and their developmental environment. The potential exploration capabilities of gene networks might be greater than currently thought. Nevertheless, we assume complex behaviours are harder to find genetically than by social learning. The argument is that solutions might be out of reach for genomes due to further developmental constraints and contingencies. While our model does not account for any of these complexities explicitly, it assumes that they would add more constraints on strategy-space exploration for genes than for memes, and therefore including them would more likely increase rather than decrease the chances of culture evolving.

Our simulation model shows that a high imitation error rate can be detrimental to the evolution of social learners when behaviours are close to the optimum. Some analytical models explicitly include a fixed fitness cost for social learners as a way to represent this imitation error rate (i.e., the "cost" of social learning) (Mesoudi, 2011). Here, we show that this high error rate can be both an enabler and an inhibitor of the fixation of social learners, depending on the optimisation problem and the value of the imitation error rate compared to genetic mutation. In our model, no further costs associated with social learning are introduced. All costs related to the development and maintenance of the physiological hardware that allows imitation are considered negligible compared to the fitness value produced by the imitated behaviour itself. All agents compete for the same pool of resources and confront the same challenges when doing so. When population structure, multiple pools of resources, and optimal foraging theory are considered, the fixation of unbiased social learners might not happen in the way described here.

In all of our simulations, the value of μ_c has been set to 0.01. This means, on average, once in one hundred iterations, an instinctive actuator turns into an imitator or vice versa. The argument for fixing this rate independently from μ_g is one of simplicity, not artefact. Analyses of our model show that a critical density of imitators is required to start chains of imitation (imitators copying imitators), as sequential events of high-error transmission are required to start a non-genetic evolutionary system that is competitive compared to the genetic one (see Appendix C). The lower the chance of these chains forming, the less likely it is for social learners to fix, even under otherwise favourable conditions. Fixing the mutation rate for the imitation switch ensures that a constant "background" density of imitators is created by mutation alone, independent of the fidelity of both genetic and social transmission, and that this density is resistant to stochastic fluctuations throughout the simulation. This simplification does not decrease the validity of our analyses, as the critical initial density of imitators is still small (i.e., 1 mutation in 100 births). It is also not the only way of increasing the tendency for imitative chains to form. We hypothesise the same effect can be achieved with $\mu_c = \mu_g$ by imposing a simple spatial population structure and restraining imitation to only take place locally. However, as mention in the introduction, this caveat/requirement should be taken into consideration when interpreting our results in comparison with those of other models in the literature. In these models, the invasion of social learners can start with a initial differential proportion of social learners (see Section 3.1).

Our conclusions not only have implications for the specific case of horizontal imitation, but also for the evolution of horizontal information transfer in general, compared to vertical imitation.

For example, a mechanistically analogous scenario to the one described in this chapter could occur in a population of hosts and symbionts. In this analogous scenario, the spread of symbiotic bacteria from one host to another would be equivalent to the spread of variants in cultural systems. We could imagine members of the host population having a "switch" with two states: one in which they obtain a useful enzyme by inheriting an enzyme-producing bacterial strain with a low mutation rate from their vertically inheriting parents, and another in which they rely on the horizontal acquisition of this bacterial symbiont in order to produce the enzyme for them. This acquisition process could take many forms, from any close-range interactions to specific types of transmission, such as sharing food or consuming faecal matter. We assume that horizontally transmitted bacteria have higher mutation rates and shorter replication times compared to the parentally inherited bacterial strains, which ensures an exploratory advantage for horizontal transmission analogous to the relatively high error rate associated with imitation in social learners.

In this scenario, our model predicts that horizontal contagion by symbiotic bacteria, and the associated behaviour that enhances the chances of contagion, can evolve in the host population if the enzyme extends the lifespan of the hosts (hence the chances of infecting others) more than it increases their reproductive output. The evolution of horizontal contagion also depends on the relative mutation rates of horizontally infecting strains and vertically inherited ones, with the former needing to be large enough such that the exploration advantage of horizontal transmission is able to displace vertical transmission during early stages of the optimisation process (i.e., during the "first period" described in Figure 3.10).

Finally, it is important to mention that environmental change is not explored in this chapter. However, in Chapter 5 we show that adding environmental fluctuations to our model, by frequently changing the target phenotype (e.g., from all 1s to some other string), increases the evolvability of culture in the four-dimensional parameter space and culture's adaptive value compared to genetic evolution. This effect contrasts with some existing models of the evolution of culture (i.e., Rogers' model) that consider environmental change to be a hindrance to the fixation of social learners.

3.6 Conclusions

Our model demonstrates that a basic cultural system can emerge in a population when selection pressure for survivability is stronger than selection pressure for reproduction. An extended strategy space (where exploration over a large sequence space is required) distinguishes our approach from previous models and is essential for understanding the exploration advantage of culture versus genes. Analysis of our results leads to a consistent explanation of cultural emergence where genotype and phenotype disengage, with the former evolving exclusively by (non-critical/non-discriminatory/unguided/unbiased) social learning. During this process of disengagement, the imitator minority must be able to offset its lower selection pressure with its higher exploration capability, a condition that is facilitated when the survivability of the evolved strategy is more important than its reproductive fecundity. Our simulation model offers a very simple framework for the emergence of cultural systems and serves as a basis for future research extensions. This simplicity, in particular removing the need for decision-making cognitive abilities such as those that result in transmission bias and/or guided variation, lowers the cognitive behavioural threshold that a species must overcome in order to evolve cultural information transmission, and therefore makes it relevant for a broader range of organisms.

In contrast to previous models, our work shows that the emergence of a cultural evolutionary system can be caused by the invasion and eventual fixation of unbiased social learners without the need for individual learners. In other models, the invasion of unbiased social learners and the consequent displacement of instinctive actuators only occur with the presence of individual learners. Therefore, existing models of the evolution of culture converge to mixed populations of social and individual learners instead of a monomorphic population of social learners; as shown in our results (Mesoudi, 2011; Richerson and Boyd, 2006; Rogers, 1988; Henrich and Gil-White, 2001; Enquist et al., 2008; McElreath and Henrich, 2007). This key difference in our model further supports our claim: the cognitive machinery required for individuals to directionally improve upon, or differentiate between behavioural variants (i.e., the presence of individual learning, guided variation or biased imitation), is not a requisite for the evolution of social learning and by extension the emergence of cultural evolutionary systems.

Chapter 4

The Role of Horizontal Imitation in the Emergence of Cultural Evolutionary Systems

In the last chapter, we showed that social learners (i.e., individuals that imitate randomly selected members of the population) could fix under strong survival selection when they have an exploratory advantage compared to instinctive actuators (i.e., when imitation error is larger than genetic mutation). In our original model, each agent is represented by two binary sequences of length $L = 200$ that respectively code for genotype and phenotype, and a binary imitation switch that controls whether an agent is an instinctive actuator (i.e., its phenotype sequence is a perfect copy of its own genotype) or a social learner (i.e., its phenotype is imitated at birth from a randomly selected member of the population). In that model, the fixation of social learners is irreversible after *genotype-phenotype disengagement* has occurred, because, by this point, the imitated phenotypes expressed in the population are fitter than the phenotypes that would result from the expression of social learner genotypes, since the genotypes of social learners are shielded from selection (see Chapter 3 for a comprehensive explanation).

As explained in the previous chapter, genotype-phenotype disengagement is more likely to occur when social learners imitate with an error rate larger than genetic mutation, and when survival selection is stronger than reproductive selection. However, to produce disengagement, imitated phenotypes do not need to be horizontally transmitted (i.e., imitated from randomly selected non-parental individuals). As we will explain in this chapter, exclusive imitation of parental phenotypes (i.e., vertical imitation) can also produce similar results. In a system where social transmission (i.e., non-genetic transmission) occurs only from parent to offspring, the decomposition of selection into survival and reproduction is not necessary to explain the fixation of social learners and the evolution of culture. This is because, in contrast to horizontal imitators, agents that copy their parents benefit from reproductive selection to the same extent as instinctive actuators do.

In order to explain the value of horizontal imitation, we must then explain its evolution in the presence of the alternative: vertical imitation. Horizontal imitation is precisely what makes social learning algorithmically different from genetic evolution. Social learning can occur through horizontal and/or vertical transmission, but genetic inheritance only occurs through vertical lines. For this reason, explaining the advantage (if any) of horizontal transmission is important if we are to understand social learning as something more than an alternative substrate for genetic evolution. Furthermore, the existence of dominant horizontal transmission is key to explain important processes of established cultural evolutionary systems: processes that occur at much faster time scales than a generation (e.g., rapid propagation of new ideas) (Richerson and Boyd, 2006; Perreault, 2012; Harris, 1999).

Most existing models discussed so far do not consider both kinds of imitators and explain the evolution of only the horizontal kind by: (1) arguing that horizontal imitators can iterate through many models in a lifetime, (2) introducing decision-making attributes for social learners (e.g., biased imitation towards fitter models), and (3) introducing individual learners (i.e., agents that pay a fixed trial-and-error cost but have a fitter phenotype behaviour as a result of it) in the population (Richerson and Boyd, 2006; Mesoudi, 2011). When iterating through many models in a lifetime and/or biasing imitation towards fitter models, horizontal imitators experience an additional selective force compared to instinctive actuators. When introducing individual learners in the population, horizontal imitators displace instinctive actuators through an information-parasitic process; they can copy fitter strategies from individual learners without paying the corresponding trial-and-error fitness cost. In both cases, the evolution of horizontal imitation is explained with the assumption that individuals can discern between strategies, or can improve and/or learn strategies on their own (i.e., a non-Darwinian process).

In existing work, the evolution of horizontal imitation is therefore not explained by *only* the advantage of randomly imitating non-parental models (i.e., a Darwinian process), but instead by the added benefits of: (1) expressing and evaluating many behavioural variants in a single lifetime, (2) decision-making attributes, and (3) the presence of individual learners (Richerson and Boyd, 2006; Mesoudi, 2011; Cavalli-Sforza and Feldman, 1981; Takahashi, 1999; McElreath and Strimling, 2008).

Considering this, we frame the work in this chapter as a theoretical effort to assess the value of horizontal imitation, even above and beyond that of having an imitation error rate that is higher than genetic mutation (see Chapter 3). We do so by explaining and discussing an extended version of our original model. In this extension, social learners can imitate both vertically (i.e., from their parents) and horizontally (i.e., from a randomly chosen non-parental member of the population). Similar to our original model, and in contrast to other models in the literature, imitation occurs only once in an agent's lifetime (i.e., at birth), without imitation bias, and without individual learners being present in the population.

The results of this chapter are divided into three sections. In the first section, we show that adding the option of vertical imitation increases the volume of the model's parameter space associated with a convergence to culture, i.e., combinations of survival selection strength, reproductive selection strength, genetic mutation rate and imitation error rate that did not originally allow convergence to culture do allow this after the introduction of vertical imitation. In the second section, we focus on parameterisations where the strength of survival selection is far greater than that of reproductive selection: a condition where both vertical and horizontal imitators are under the same selective pressure. In that section, we show that *when imitation error rates are equal for both horizontal and vertical imitation*, horizontal imitators dominate over the vertical kind (i.e., are more likely to fix in the population). This is the most important result in this chapter, as it shows that horizontal imitation has value above and beyond that of a high imitation error rate (i.e., compared to genetic mutation). As we will explain in our results, horizontal imitators maintain their dominance over vertical imitators (and instinctive actuators) because they are more likely to acquire trending behavioural variants (i.e., behavioural variants that are increasing in frequency because they are fitter than average). In the discussion section we further explain how horizontal imitation can rapidly increase the frequency of fitter variants without relying on decision-making attributes and/or the presence of individual learners.

In the third section, we show that *when the error rates for each type of imitation are different* (i.e., high for horizontal imitation and low for vertical imitation), a sequential invasion of horizontal imitators (displacing the initial population of instinctive actuators) and then vertical imitators (displacing the horizontal imitators that invaded) takes place. This sequential pattern corresponds to two periods in the evolutionary process, an initial period where horizontal imitators with a relatively high error rate have an exploration advantage, and a second period where vertical imitators with a relatively low error rate have an inheritance fidelity advantage.

The chapter concludes with a discussion of the role of horizontal imitation in the evolution of social learning and the emergence of basic cultural evolutionary systems in human and animal populations.

4.1 Model Extension with Vertical and Horizontal Imitation

By modifying our basic simulation model, we can explore the effects of vertical imitation on the evolution of social learning. Here, we are mainly interested in whether vertical imitators will (1) displace horizontal imitators, (2) coexist alongside horizontal imitators, or (3) be displaced by horizontal imitators.

In this version of the model, we have included an extra state for the imitation switch. Here, individuals can either be instinctive actuators with a switch value of 0, horizontal imitators with a switch value of 1, or vertical imitators with a switch value of 2. As in the original model, an instinctive actuator expresses the information contained in their own genome as their phenotype (without error), and a horizontal imitator selects a random member of the population (including its parents) from whom to copy their phenotype with a mutation rate μ_h. A vertical imitator copies the phenotype of its parent with an imitation error rate μ_v. All other features of the original model from Chapter 3 remain unchanged in this extension. A diagram of the algorithm for the extended model is explained in Figure 4.1.

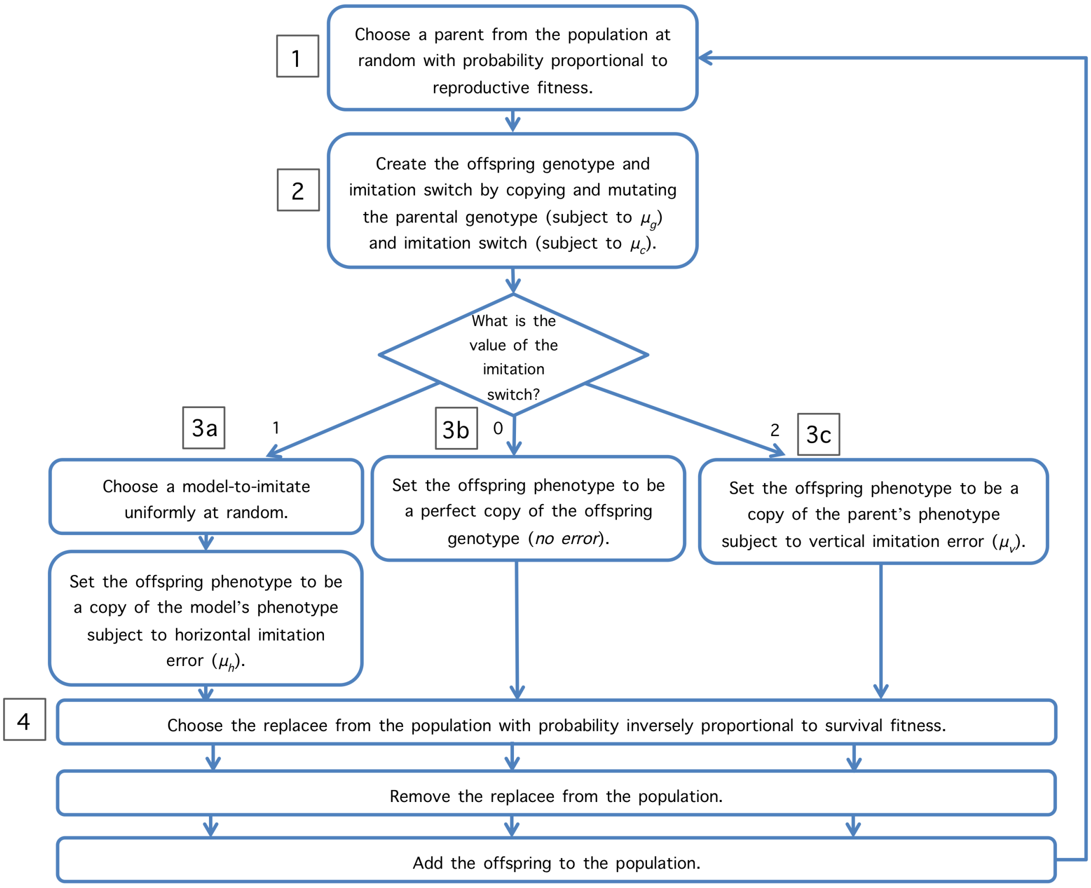

FIGURE 4.1: An algorithmic representation of the extended model's logic. (1) The selection function picks an individual from the population. (2) Its genotype string and imitation switch are copied (mutations occur at a rate of μ_g per genotype bit and μ_c for the imitation switch). (3a) If the imitation switch has value 1, a random individual from the population is selected and its phenotype will be copied (imitation errors occur at a rate of μ_h per bit). (3b) If the imitation switch has value 0, the phenotype will be a perfect copy (i.e., without errors) of the individual's own genotype. (3c) If the imitation switch has value 2 the parental phenotype will be copied (imitation errors occur at a rate of μ_v per bit). (4) The resulting combination of phenotype and genotype will replace an individual selected by the death function.

As in the original model, instinctive actuators mutate into imitators at a rate of μ_c. If an instinctive actuator mutates in this way, it has an equal chance of becoming a vertical imitator (imitation switch $= 2$) or a horizontal imitator (imitation switch $= 1$). Conversely, imitators of both types (vertical and horizontal) can mutate with rate μ_c into either instinctive actuators or the other kind of imitator, also with equal chance. Every imitator performs only one type of imitation in their lifetime and only imitates once. The type of imitation and the imitated phenotype are both set at birth. For all results in this chapter, the mutation rate for the imitation switch is fixed at $\mu_c = 0.01$. This value is equal to the one used in the original model and is fixed for the same reason; to maintain a minimal frequency of all types of agents (see Chapter 3 for a comprehensive explanation).

For all results in this chapter, population size is set to one hundred individuals ($N = 100$), the length of the phenotype and genotype strings are set to two hundred bits ($L = 200$), and both types of selection are controlled by Boltzmann-weighted functions of the sum-of-1s in the phenotype strings (see Equation 3.1 and Equation 3.2). A string of all-1s (i.e., $L_1 = L$) represents the optimal solution.

Simulations were initialised with all agents as instinctive actuators (i.e., every imitation switch set to 0) and each bit of their genotype and phenotype bit strings set to 1 or 0 with uniform random probability per bit. At the start of the simulation, each agent's phenotype and genotype strings are perfect copies of one another. These initial conditions prevent mutation biases from imposing directed drift on the genotypes or phenotypes, even in the absence of selective pressure.

4.1.1 Simulation Steps

In this subsection, steps presented in Figure 4.1, that were not present in Figure 3.1, are explained in detail. In a similar manner to Subsection 3.1.1, in this subsection, we put emphasis on clarifying the internal logic of each function and how it affects the bit strings that define individual agents in the simulation. This information can also be found in Appendix B, along with the description of the original simulation model (see Chapter 3) and the details of the extension explained in Chapter 5. To assess the internal logic of our model, before generating our analytical results, all individual functions here described were "unit tested". That is, for each function, the expected outputs for a range of possible inputs were evaluated.

4.1.1.1 Three-Way Mutation Implementation

In the model extension explained in this chapter we include a third type of agent: a vertical imitator. Therefore, the logic behind the mutation function for the imitation switch must allow for mutations between three values, instead of only two, as in the original model. Notice that in this model extension, when the imitation switch of the new individual, created in step 2, has a value of 0, the individual is considered an *instinctive actuator* and will obtain its phenotype through step 3b of Figure 4.1, using the same process described in Subsection 3.1.1.5 (i.e., by copying their own genotype without error). If the imitation switch of the new individual has a value of 1, the individual is considered an *horizontal imitator* (referred to as a "social learner" in Chapter 3) and will obtain its phenotype through step 3a of Figure 4.1, using the same process described in Subsection 3.1.1.4 (i.e., by imitating the phenotype string of a randomly selected individual from the population, subject to mutation rate μ_h). If the imitation switch of the new individual has a value of 2, the individual is considered a *vertical imitator* and will obtain its phenotype through step 3c of Figure 4.1 by a process that will be described in detail in the next subsection.

The mutation function can change the imitation switch from one value to any of the other two. In each iteration, when producing a new individual/offspring, the imitation switch will change from its current value with probability $\mu_c = 0.01$. If the imitation switch is mutated, one of the other two values will be selected with equal probability (i.e., 0.5). For example, if an individual has a switch value of 0, and it engages in a mutation event (i.e., with probability μ_c), then the value of the switch has a 50% chance of turning into 1 and a 50% chance of turning into 2. Figure 4.2, shows a diagram representing the logic behind this mutation function.

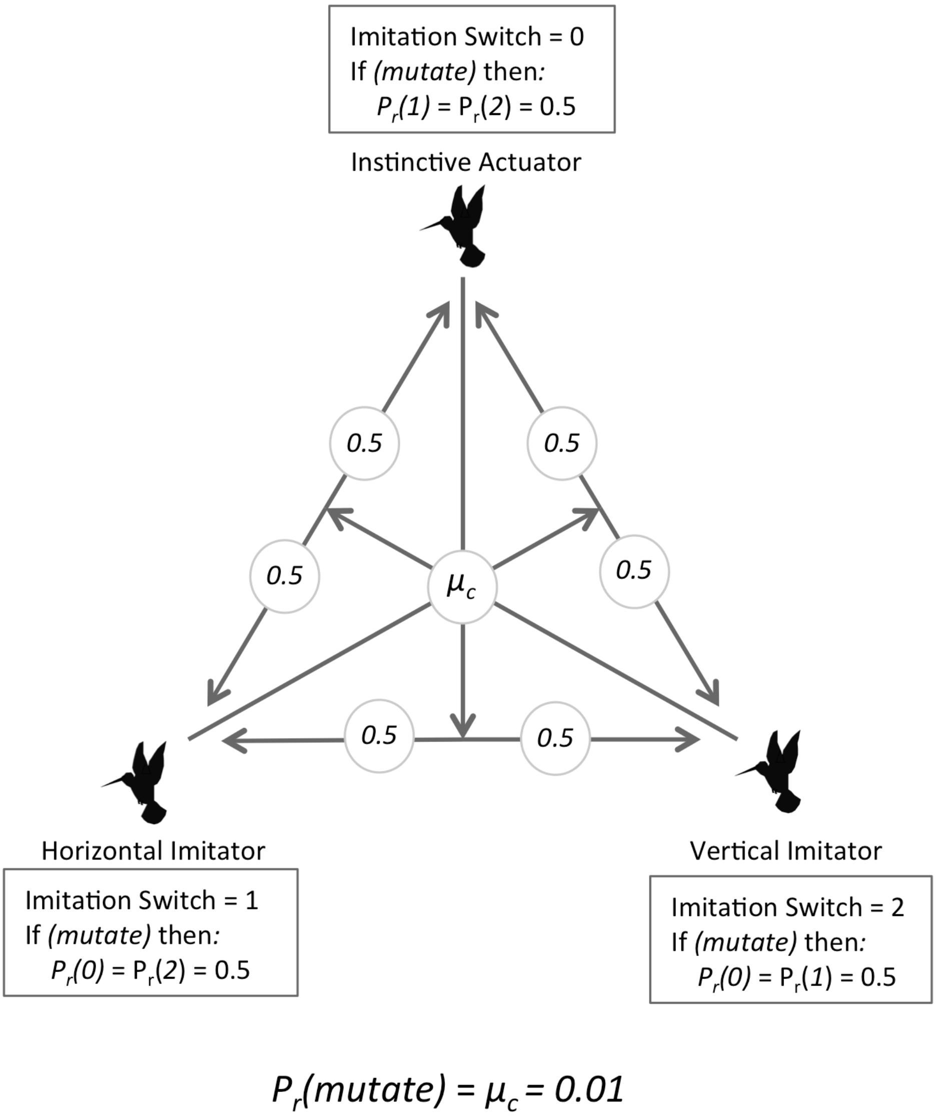

Figure 4.2: Diagram representing the logic of our three-way mutation function. The new individual/offspring, created in step 2 of Figure 4.1, can mutate its imitation switch with probability μ_c. If a mutation occurs, one of the two other values for the imitation switch will be selected, with equal probabilities. For example, if the value of the imitation switch is 2 (i.e., the individual is a vertical imitator), and a mutation event occurs (i.e., with probability μ_c), the switch will mutate into 1 (i.e., the individual will become a horizontal imitator) with a 50% chance, and into 0 (i.e., the individual will become an instinctive actuator) with a 50% chance.

4.1.1.2 Step 3c: If the Offspring is a Vertical Imitator

In step 2 of Figure 4.1, we produced a new individual/offspring with a genotype string and an imitation switch, but without a phenotype string. In this step, we will create a phenotype for this individual assuming that its imitation switch is 2 (i.e., the individual is a vertical imitator). When the imitation switch is set to 2, the phenotype of the new individual is a copy of the phenotype string of its own parent (i.e., the individual selected in step 1). This process of phenotype imitation occurs once at birth. The phenotype string is copied with an error rate of μ_v per bit. This means, each bit copied into the phenotype of the new individual will change its value (i.e., changing from 0 to 1, or from 1 to 0) with probability μ_v. Notice that by imitating the parental phenotype, vertical imitators can still produce genotype-phenotype disengagement (see Figure 4.5), as the substrate of evolution, i.e., the phenotype (i.e., the information that is ultimately optimised), can be inherited independently from the genotype. In this chapter, we refer to vertical and horizontal imitators as *social learners*. Figure 4.3 shows a diagram explaining this step.

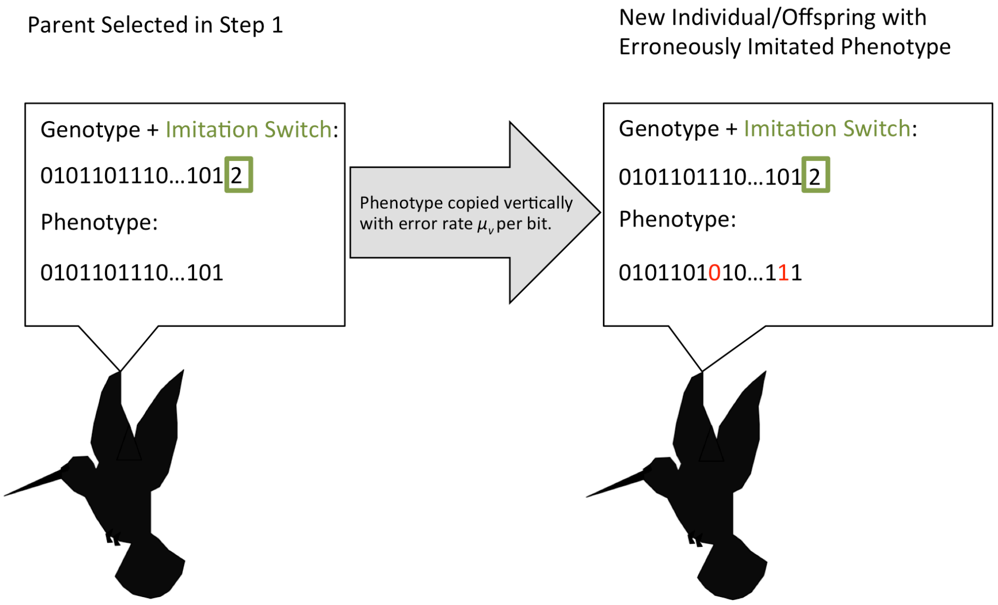

FIGURE 4.3: Diagram representing the third step of our simulation model (if the new individual is a vertical imitator). If the value of the imitation switch of the new individual is equal to 2 (i.e., the individual is a vertical imitator), we create its phenotype by copying the entire phenotype string of its own parent (i.e., the individual selected in step 1). Each phenotype bit is copied subject to a mutation rate of μ_v. Notice that the imitation switch is not part of the phenotype. Mutated bits are shown in red.

4.2 Results for the Model Extension with Vertical Imitation

The results of this chapter are organised into three distinct subsections. In the first subsection, we explore a range of selection strength coefficients, genetic mutation rates and imitation error rates. In that subsection, we show that social learners can fix under strong reproductive selection when they imitate their parents (Figure 4.4).

In the second and third subsections, we carry out an in-depth analysis of simulations where (1) survival selection is stronger than reproductive selection, and (2) at least one of the two types of imitation has a relatively large error rate when compared to the rate of genetic mutation. In these simulations, social learners can easily fix and displace instinctive actuators. However, the type of imitator that dominates the population (i.e., that is more frequent in the long term) is different depending on the imitation error rates for vertical (μ_v) and horizontal (μ_h) imitation.

4.2.1 Social Learners Can Fix Under Strong Reproductive Selection When They Imitate Their Parents

The effect on the proportion of simulations that converge to culture of adding vertical imitation to the model is shown in Figure 4.4. Each map in the set corresponds to a combination of survival and reproductive selection strength coefficients: x_d and x_r, respectively. Points inside each map correspond to combinations of μ_v and μ_g. In these plots, the error rate of horizontal imitation is equal to that of vertical imitation ($\mu_h = \mu_v$). Red colours indicate a large proportion of simulations where social learners irreversibly fix after 5×10^4 iterations (i.e., simulations where culture emerges).

For each map, one hundred replicates were carried out for each combination of mutation rate, μ_g, and imitation error rate, μ_v, from the set $[1 \times 10^{-4}, 2 \times 10^{-4}, \ldots, 50 \times 10^{-4}]$, and where $\mu_v > \mu_g$, with fixed values of survival and reproductive selection coefficients (x_d, x_r). The proportion of replicates that converge to culture (i.e., where the proportion of imitators of any type is above 0.95 by the end of the simulation run) is presented in dark red. These maps only show combinations below the diagonal $\mu_v = \mu_g$, where imitation error is greater than genetic mutation.

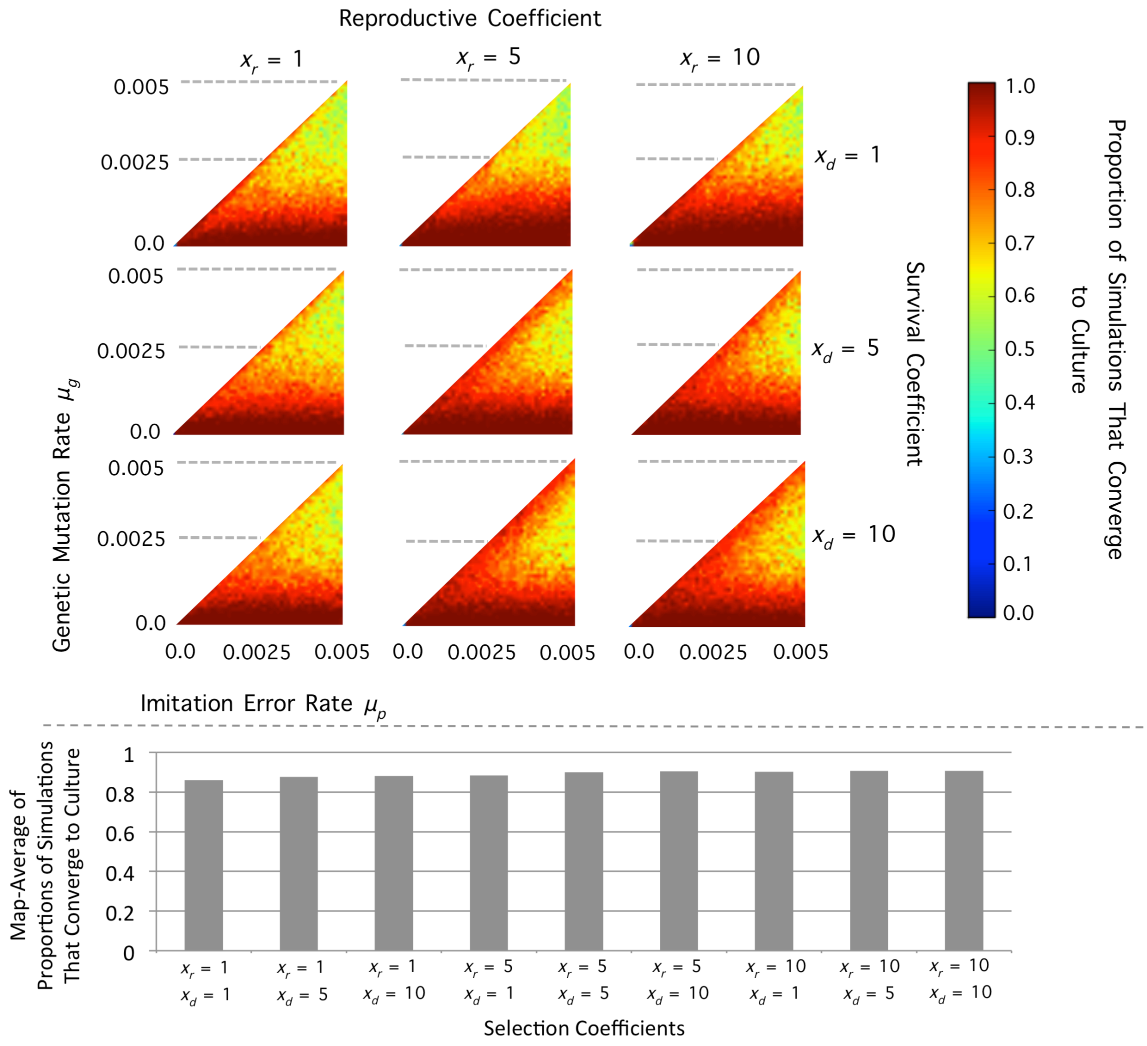

FIGURE 4.4: Heat map array depicting how the tendency for social learners to emerge in a system of instinctive actuators, horizontal imitators and vertical imitators varies with their reproductive and survival coefficients (where higher coefficients indicate weaker selection), genetic mutation rate, and imitation error rate. Within each plot, red data points indicate that a high proportion of 100 replicates evolved culture through genotype-phenotype disengagement. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. The length of the bars represents the average colour value/temperature of the corresponding plot (i.e., the average proportion of simulations that converge to culture for the entire map). For all parameterisations, the proportion of simulations that converge to culture increases in comparison with the results in Figure 3.9, where social learners can only imitate horizontally. These results show that the presence of vertical imitators increases the chances of social learners fixing. In simulations that converge to culture, a population of vertical imitators, a population of horizontal imitators, or a mixed population of vertical and horizontal imitators is found at the end of the run (no instinctive actuators).

In Figure 4.4, we show how social learners can fix under a large variety of parameterisations when strictly vertical imitation is an evolvable trait (see Figure 3.9 for comparison).

Cases with strong reproductive selection ($x_r = 1$) are of special interest; in contrast to the results of our original model, parameterisations with strong reproductive selection show a high proportion of simulations converging to culture. In Figure 4.5, we compare the proportion of simulations that converge to culture in our original model (map A) against the equivalent plot in the extended version including vertical imitation (map B). The simulations used to produce these plots were run under strong reproductive selection ($x_r = 1$).

In plots C and D, we show individual simulation runs where $\mu_g = 0.001$ and $\mu_p = 0.005$, for the original and the extended model. In maps A and B black arrows highlight the area where these single-runs are represented. For this parameterisation, there is a low proportion of simulations converging to social learning in our original model. However, when vertical imitation is available, most of these simulations converge on culture (i.e., a non-genetic evolutionary system). In plot C, the phenotype evolves genetically, and genotype-phenotype disengagement does not occur. In plot D, genotype-phenotype disengagement occurs early in the simulation, although this is the result of an invasion by vertical rather than horizontal imitators. In our results, we count all events of genotype-phenotype disengagement (accompanied by the fixation of any kind of imitator) as the emergence of culture (i.e., a non-genetic evolutionary system). However, we consider that when disengagement occurs due to the fixation of vertical imitators, this non-genetic evolutionary system is algorithmically equivalent, in terms of inheritance direction, to a genetic evolutionary system (i.e., to a population of instinctive actuators with the same mutation rate as the error rate of vertical imitators).

The fact that vertical imitators are affected by both survival and reproductive selection to the same extent as instinctive actuators explains the difference between these plots. In our original model, imitators (i.e., horizontal imitators) are not affected by reproductive selection (i.e., their offspring do not inherit their phenotype) and therefore can only displace instinctive actuators if their exploration advantage compensates for their weaker selective pressure. This is more likely to occur when survival selection is stronger than reproductive selection. In contrast, vertical imitators do not need to compensate for weaker selection compared to instinctive actuators. This means vertical imitators can produce genotype-phenotype disengagement even when reproductive selection is strong, by taking advantage of the increased exploratory capabilities that result from their relatively high imitation error (i.e., higher than the genetic mutation rate of instinctive actuators). However, in order for vertical imitators to fix, phenotype fitness disengagement still needs to occur before this high imitation error rate becomes destructive (i.e., before the α sum-of-1s value is reached (see Chapter 3)).

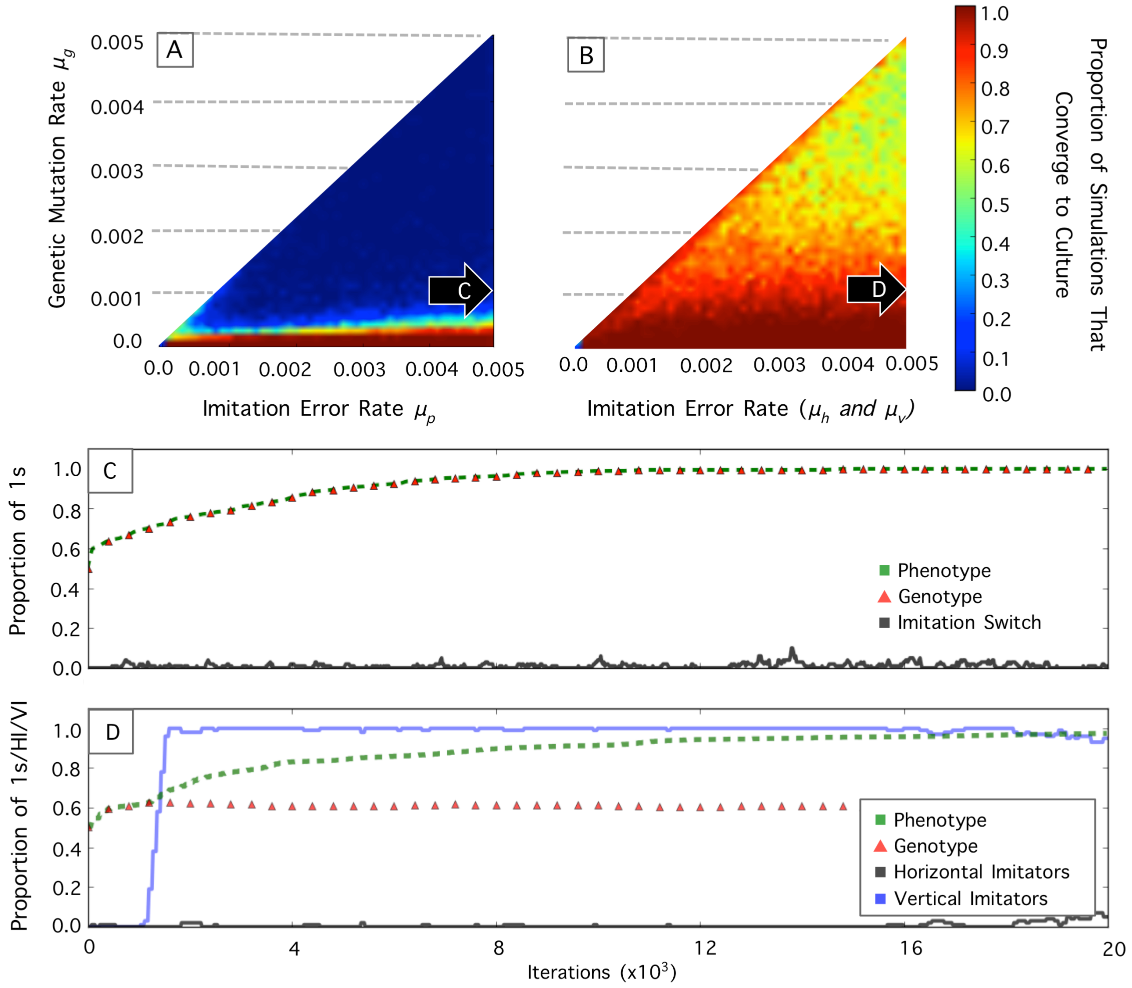

Figure 4.5: Comparison between the behaviour of the original model and the extended version including vertical imitators. Heat maps A and B show the proportion of simulations converging to social learning for a range of imitation error rates and mutation rates, μ_v (equal to μ_h) and μ_g. In both maps survival and reproduction selection coefficients are set to $x_r = 1$ and $x_d = 1$. The black arrows in plots A and B indicate points where $\mu_v = \mu_h = 0.005$ and $\mu_g = 0.001$. Representative single runs with these values are shown in plots C and D, corresponding to arrows C and D, respectively. In plot C, the phenotype (green dotted line) is optimised along with the genotype (red triangles) through genetic evolution; horizontal imitators (black line) do not fix. Plot D shows genotype-phenotype disengagement produced by vertical imitators (blue line). The presence of vertical imitators leads to the fixation of social learners (albeit of the vertical kind) even when reproductive selection is strong.

4.2.2 Horizontal Imitation Dominates Over Vertical Imitation Under Weak Reproductive Selection

In the previous subsection, we showed that vertical imitators could fix and produce genotype-phenotype disengagement for parameterizations where horizontal imitators cannot (i.e., when reproductive selection is strong). However, when selective pressure for horizontal and vertical imitators is similar (i.e., when reproductive selection is non-existent), it is not clear which kind of imitator will cause genotype-phenotype disengagement (if any), or if one type of imitator will dominate over the other (i.e., one type of imitator will be present at higher frequency in the long term). To clarify this, we will focus on areas of the parameter space (shown in Figure 4.4) where both vertical and horizontal imitators have equal selective pressure and where a large proportion of the simulations converge to social learning. Like in the previous subsection, we will maintain the same error rate for both types of imitation ($\mu_h = \mu_v$), so that any advantage of one type over the other can be attributed entirely to the sourcing of imitated strategies (vertical or horizontal) and not to an exploratory advantage produced by differences in error rates.

We will focus our attention on the case where $\mu_h = \mu_v = 0.01$, $\mu_g = 0.001$, $x_r = 1 \times 10^3$ and $x_d = 1$. In this case, horizontal and vertical imitation error rates are equal to each other and higher than the genetic mutation rate, while reproductive selection is very weak compared to survival selection. Qualitatively similar results are achieved for similar parameterisations. Notice that the reproductive selection coefficient for this parameterisation is set to a higher value compared to other simulations ($x_r = 1 \times 10^3$). Setting this coefficient with such a high value effectively eliminates reproductive selection. We have done this to guarantee that both vertical and horizontal imitators are subject to the same selective pressure. Note that, as discussed in Section 3.2, the cultural evolutionary system maintained by horizontal imitators (equivalent to "social learners" in Chapter 3), is only affected by survival selection (see Appendix C).

Figure 4.6 shows data for the first 5×10^5 iterations of three different simulation runs. The green dotted line represents the average phenotype in the population, the red triangles the average genotype, and the blue and black lines represent the proportion of vertical and horizontal imitators, respectively. In each simulation run different patterns are observed; in some cases, genotype-phenotype disengagement occurs through an invasion of vertical imitators (e.g., replicate 1 and 2), and in others through an invasion of the horizontal type (e.g., replicate 3). In most replicates, alternating invasions of vertical and horizontal imitators occur after disengagement (e.g., replicates 1, 2 and 3). Purely stochastic processes account for the differences between these simulation runs.

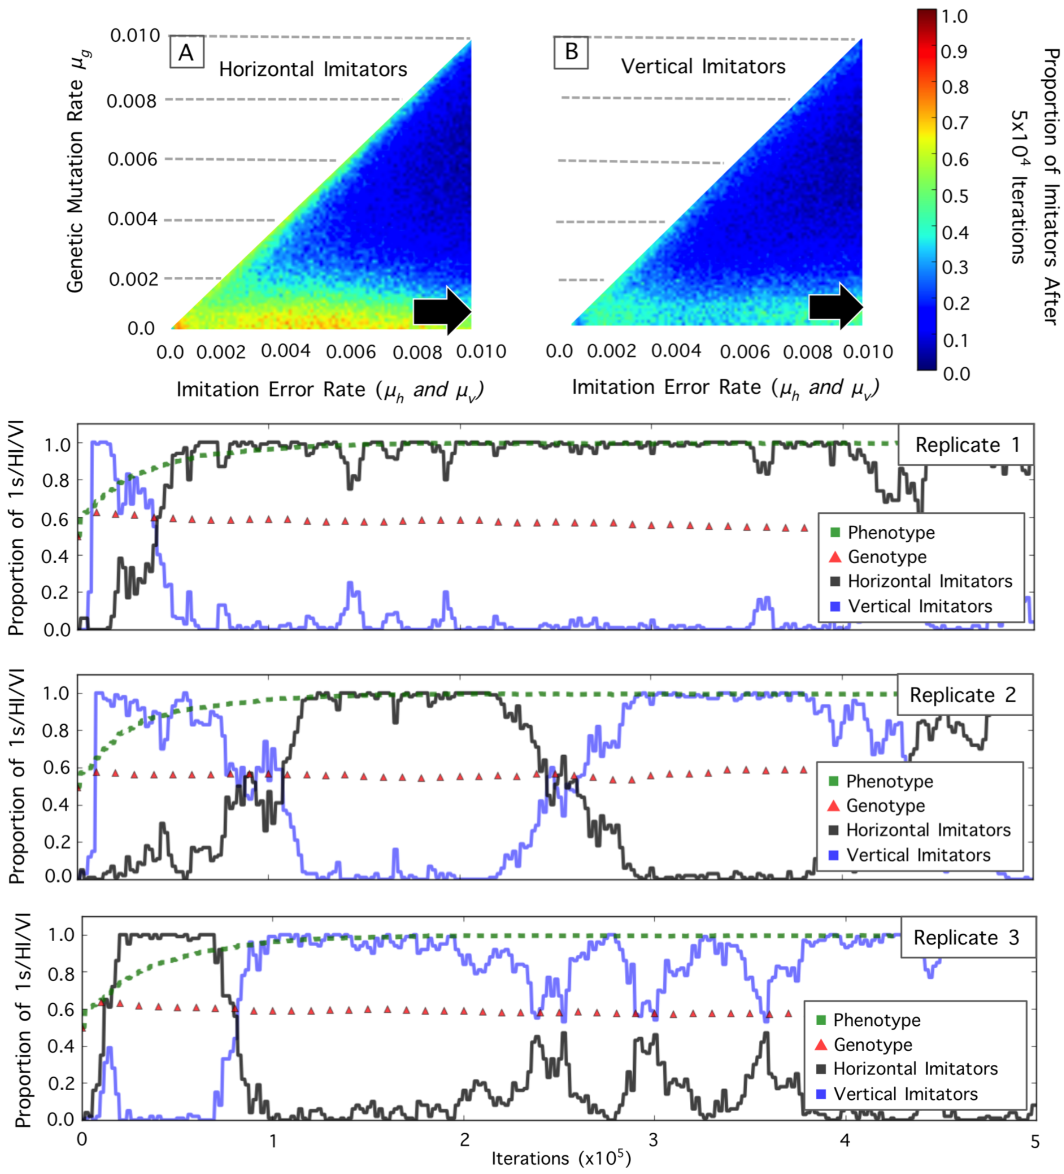

FIGURE 4.6: Heat maps showing the average proportion of horizontal imitators (map A) and vertical imitators (map B) after 5×10^4 iterations, for simulations where reproductive selection is effectively absent (i.e., where only survival selection applies). The selection coefficients for these maps are $x_d = 1$ and $x_r = 1 \times 10^3$. The black arrows in maps A and B indicate the parameterisation point for the three replicate runs shown below the maps. For these replicates $\mu_h = \mu_v = 0.01$ and $\mu_g = 0.001$. These runs are representative of the results in the lower horizontal band in maps A and B (i.e., regions with hotter colours), where a high proportion of simulations converge to either horizontal or vertical imitation. The green dotted line indicates the proportion of 1s in the phenotypes of the population, the red triangles indicate the proportion of 1s in the genotype, the black line indicates the proportion of 1s in the imitation switch (i.e., proportion of horizontal imitators) and the blue line indicates the proportion of 2s in the imitation switch (i.e., proportion of vertical imitators). In the first replicate, vertical imitators invade during genotype-phenotype disengagement and are then reinvaded by horizontal imitators. In the second replicate, long-term alternating invasions of horizontal and vertical imitators occur after genotype-phenotype disengagement. In the third replicate, an invasion of horizontal imitators takes place during genotype-phenotype disengagement, but vertical imitators later displace them.

Initially we hypothesised that both the initial invasion of social learners and the long-term fluctuations (i.e., the sequential re-invasions after genotype-phenotype disengagement) would have occurred with the same probability for each type of imitator (i.e., neither type of imitator would be dominant). This could mean: (1) in half of the iterations (sampled across many simulation replicates), the populations are mostly formed by one type of imitator; or (2) in all iterations there are mixed populations of vertical and horizontal imitators where both types appear in similar proportion. The second hypothesis would imply that neither horizontal imitation nor vertical imitation dominates over the other.

To test this, we ran 1×10^4 replicate simulations and averaged the proportion of 1s for the genotype, phenotype, imitation switch and the proportion of vertical imitators (i.e., the proportion of 2s in the imitation switch) for each iteration across the replicates (e.g., we average the proportion of 1s in the genotype for the first iteration of all 1×10^4 replicates to produce the first iteration-average point for the proportion of 1s in the genotype in Figure 4.7).

To our surprise, the distinctive pattern in Figure 4.7 emerged. It is important to clarify that no single simulation run exhibits the results shown in this figure. This is particularly true for the averages of horizontal and vertical imitators. While the average phenotype and genotype data have very small standard deviations (i.e., less than 1%), averages for the proportion of horizontal and vertical imitators have very large standard deviation values (i.e., more than 50%). Also, the data used to calculate each average proportion (i.e., the average on each iteration) has a clear bimodal distribution with most values either close to 1.0 or 0.0. This indicates that in individual simulation runs, like the ones in Figure 4.6, the invasion of one type of imitator is fast and almost complete (i.e., they displace other types of agents).

In Figure 4.7, the average phenotype (green dotted line) and genotype (red triangles) disengage in the first iterations of the simulation runs. For most runs, this co-occurs with the fixation of vertical imitators (blue diamonds) rather than horizontal imitators (black diamonds). However, after disengagement has occurred, horizontal imitators become dominant over the vertical kind. This does not imply that vertical imitators cannot invade once horizontal imitators have fixed (see replicate 2 in Figure 4.6); it rather means that horizontal imitators reinvade more frequently and remain at higher frequencies for longer, compared to vertical imitators.

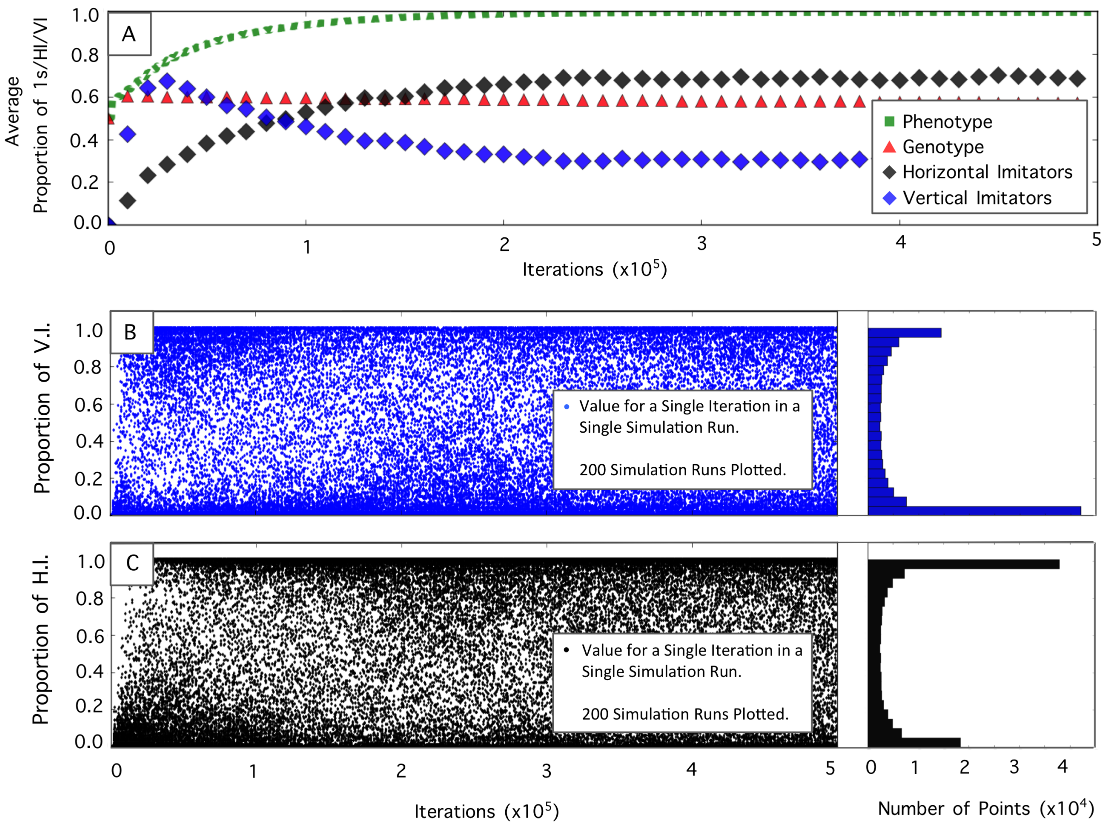

FIGURE 4.7: In plot (A), the average trajectory for 1×10^4 individual simulation replicates is shown for the extended model with vertical and horizontal imitators. The green dotted line represents the average proportion of 1s in the phenotypes and the red triangles the average proportion of 1s in the genotypes. Disengagement between these two lines occurs early in the simulations, while the average proportion of vertical imitators rises steeply (blue diamonds). After disengagement takes place, non-genetic evolution is temporarily dominated by vertical imitation until its horizontal counterpart takes over and dominates in the long term. This plot is intended as a representation of average statistical tendencies in the model. It should be noted that no single run of the numerical simulation generates these patterns. More specifically, longstanding mixed populations of horizontal and vertical imitators are rare in individual runs. In around 60% to 70% of the replicates, this model will achieve a population mostly composed of horizontal imitators. Plots B and C, show the distribution of trajectory points for 200 replicate simulation runs. Each point in these plots represents the proportion of vertical imitators (plot B), or horizontal imitators (plot C), for one iteration, for one of 200 simulation runs. As shown in the bar plots on the right side of plots B and C, the distribution of these points is bimodal. Selection strength coefficients are fixed at $x_d = 1$ and $x_r = 1 \times 10^3$. Imitation error and genetic mutation were set at $\mu_h = \mu_v = 0.01$ and $\mu_g = 0.001$

The pattern observed in Figure 4.7 can be explained in two different stages, one occurring early in the simulation before horizontal imitators establish chains of imitation (i.e., sequential imitation of imitators), and another starting after horizontal imitators engage in these chains.

At the beginning of every simulation run, both vertical and horizontal imitators are absent from the population. When the first imitators start to appear (i.e., when the first mutations from instinctive actuators to imitators occur), the vertical type has a clear advantage over the horizontal type, as the latter is more likely to imitate instinctive actuators (i.e., agents with poor exploration capability) rather than other imitators. Imitating instinctive actuators in the early stages of optimisation is disadvantageous because whichever phenotype they have is a perfect copy of their genotype, which itself has been evolving with a low mutation rate compared to the imitation error rate. In the early stages of optimisation, a higher error rate is advantageous (see Chapter 3).

In contrast, vertical imitators readily start accumulating advantageous changes by imitating with a high error rate in their vertical inheritance lines; every vertical imitator will imitate an imitator (i.e., their parent). Because every imitation event comes with a high error rate, and vertical imitators sequentially undergo more of these events (i.e., by imitating their parents) compared to horizontal imitators, vertical imitators are more likely to *err into* (i.e., erroneously imitate into) the fittest strategies in the early stages of optimisation.

To test this claim, we have produced plot A of Figure 4.8 which shows the iteration averages for the first 5×10^5 iterations of 1×10^4 replicates for simulations starting with *equal* proportions of vertical and horizontal imitators (50:50), rather than a monomorphic population of instinctive actuators. Under this condition, horizontal learners are shown to be dominant during and after disengagement, as they can readily form imitation chains from the start of the simulation run (i.e., they are very likely to imitate other imitators). This result does not explain why horizontal imitation is more frequent in the long term than the vertical kind (see Figure 4.7), but rather why vertical imitators are the ones involved in genotype-phenotype disengagement in the early stages of optimisation in Figure 4.7.

To explain the long-term dominance of horizontal over vertical imitators we must discuss how each type accesses trending variants. Here, we have described a *trending variant* as those phenotypes that have a higher-than-average fitness and are increasing in frequency in the population.

Horizontal imitators can easily acquire these variants when vertical imitators err into them, and reproduce them. In contrast, vertical imitators have a restricted access to trending variants that horizontal imitators have erred into. To understand this asymmetry, let us compare the cases where: (1) a vertical imitator adopts a new variant that is fitter than any other in the population, and (2) a horizontal imitator adopts a new variant that is fitter than any other in the population.

1. When a vertical imitator adopts a new fittest variant:

   When a vertical imitator errs into the fittest current variant (i.e., by imitating its parent with an error), it lives longer than the rest of the population and will produce more offspring than other individuals (i.e., will be randomly selected for reproduction more often as a result of living longer). The fittest variant spreads in the population along with the vertical imitator lineage that carries it, and the frequency of horizontal imitators is reduced. In co-occurrence with this process, the likelihood of one of the remaining horizontal imitators copying the trending variant increases (i.e., horizontal imitators select models to imitate at random from the population; the higher the frequency of a variant, the more likely that it will be copied). Once a horizontal imitator copies the fitter variant, the horizontal imitator will tend to survive (and reproduce), on average, for as long as the agent it copied.

   Therefore, it is very unlikely that in the event of an invasion by vertical imitators, horizontal imitators would completely disappear. In a population dominated by vertical imitators, the horizontal imitator minority is still imitating with a high error rate and can err into the *new* fittest variant in the population by error-prone imitation of any type of agent.

2. When a horizontal imitator adopts a new fittest variant:

   When a horizontal imitator errs into the fittest variant in the population (i.e., by imitating a randomly selected individual with an error), it survives longer than other agents and therefore its phenotype becomes more likely to be imitated by its offspring and the offspring of other horizontal imitators. The horizontal imitator with the fittest variant in the population does not have a guaranteed transference of their phenotype to their offspring (i.e., like vertical imitators do). However, it compensates for this by benefitting from the offspring of other horizontal imitators copying its phenotype more frequently, a by-product of living longer and thus being randomly selected for imitation in more iterations. The invasion of horizontal imitators accelerates and eventually most vertical imitators are displaced. In contrast with horizontal imitators, the vertical kind cannot access the trending variants found and spread by horizontal imitators as a function of their increasing frequency. For a vertical imitator to imitate a variant found by a horizontal imitator, a rare mutation of the imitation switch needs to occur (i.e., the offspring of a horizontal imitator needs to mutate into a vertical imitator). However, even when this rare mutation occurs, the vertical imitator is unlikely to invade the population of horizontal imitators and displace them. If, by chance, vertical imitators start increasing in the population (e.g., they err into an even fitter variant), the offspring of some horizontal imitators would imitate phenotypes from the fitter vertical imitator lineage and rapidly catch on.

This means that there is an asymmetry between the two types of imitators in terms of their access to imitation models: *horizontal imitators can easily access the vertical kind, but vertical imitators can rarely access the horizontal kind*. This asymmetry ensures that horizontal imitators will be more likely to adopt the fittest strategies in the long term whether by erring into it or, alternatively, by imitating fitter strategies discovered by vertical imitators. A diagram showing the asymmetry of information transmission between different type of imitators, at different stages of the simulation, is presented in Figure 4.9.

To test this idea, plot B of Figure 4.8 shows the averages per iteration for the first 5×10^5 iterations of 1×10^4 replicates for simulations where horizontal imitators can only imitate other horizontal imitators (i.e., a horizontal imitator imitates another randomly chosen horizontal imitator). Like the simulations used for plot A in the same figure, vertical and horizontal imitators start with equal proportions (50:50). In plot B, the sourcing of information for horizontal imitators is limited to a sub-group, and hence the advantage of horizontal imitators (i.e., being able to copy any type of agent in the population, and quickly acquire trending variants) is hampered. This condition leads to simulations where neither type of imitator is dominant over the other.

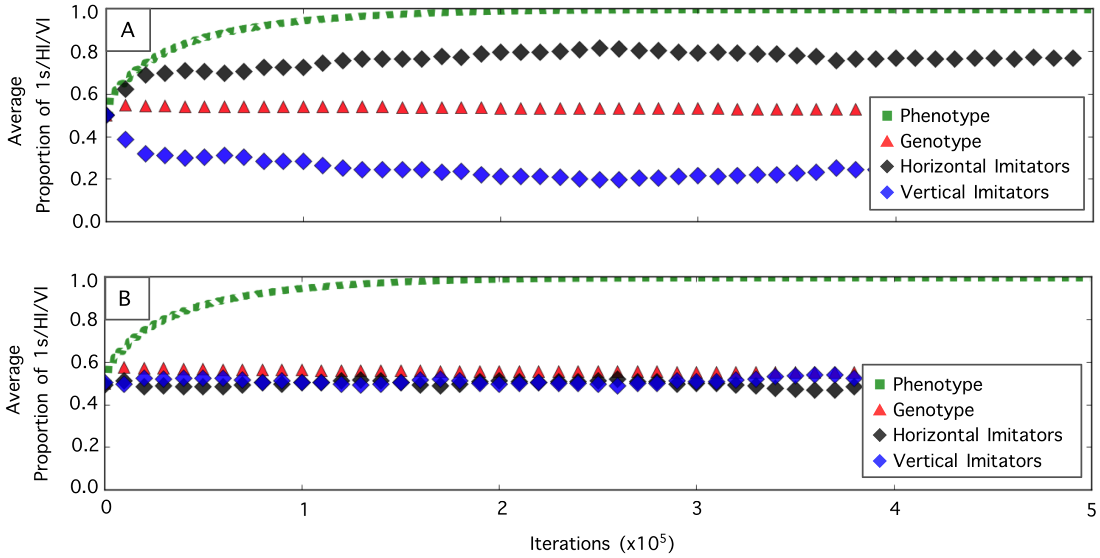

FIGURE 4.8: The average trajectories for 1×10^4 replicates is shown for two modified models with vertical and horizontal imitators. The green dotted line shows the average proportion of 1s in the phenotypes, the red triangles the average proportion of 1s in the genotype, the black diamonds the proportion of 1s in the imitation switch (i.e., proportion of horizontal imitators) and the blue diamonds the proportion of 2s in the imitation switch (i.e., proportion of vertical imitators). These plot are intended as a representation of overall statistical tendencies in the model. It should be noted that no single run of the numerical simulation generates these patterns. In plot A, the simulations start with half of the population as vertical imitators and the other half as horizontal imitators. On average, under these conditions, horizontal imitators dominate. Simulations used to generate plot B, apart from starting with half the population as vertical imitators and the other half as horizontal imitators, also restrict horizontal imitators such that they may only imitate the phenotype of other horizontal imitators. Under these conditions, neither vertical nor horizontal imitators dominate. These plots are shown as evidence of processes that explain results in Figure 4.7 (see text). Selection strength coefficients are fixed at $x_d = 1$ and $x_r = 1 \times 10^3$. Imitation error and genetic mutation were set at $\mu_h = \mu_v = 0.01$ and $\mu_g = 0.001$

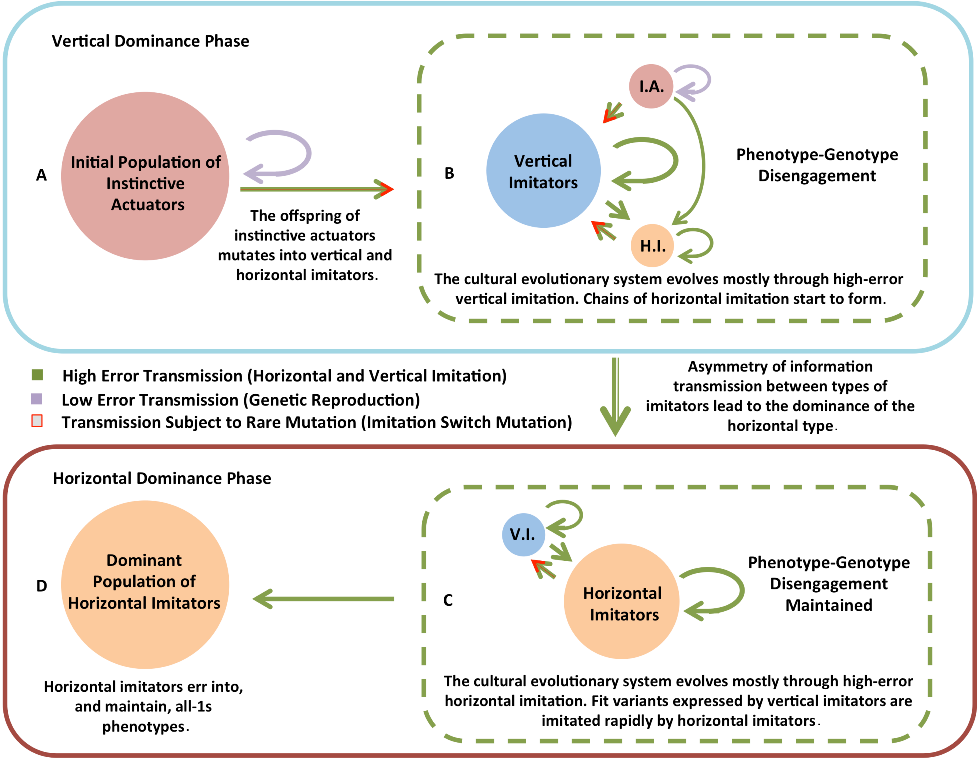

FIGURE 4.9: The diagram shows the dynamic of a system with horizontal and vertical imitation, where both types of imitation have equal error rates ($\mu_h = \mu_v$). (A) The first step takes place when mutant horizontal and vertical imitators copy information that was initially expressed in the phenotype of an instinctive actuator. (B) Once the information gets into a line of vertical imitators (V.I.), genotype-phenotype disengagement occurs and the non-genetic evolution towards fitter variants is then driven by high-error vertical imitation. In this stage, horizontal imitators (H.I.) start to form small chains of imitation (i.e., imitators imitating other imitators). However, compared to vertical imitators, horizontal imitators suffer the disadvantage of sometimes imitating phenotypes that are evolving with a low mutation rate, i.e., the phenotypes of instinctive actuators (I.A.). (C) The proportion of imitators increases rapidly and instinctive actuators practically disappear. In this stage, the disadvantage of horizontal imitators is no longer present, and their ability to imitate trending variants from any type of individual in the population now gives them an advantage. If, by chance, any instinctive actuator or vertical imitator errs into a fitter variant during this period, horizontal imitators are likely to incorporate it, and further evolve it, when it starts increasing in frequency. In contrast, vertical imitators do not have access to fitter variants erred into by horizontal imitators, except for the rare case where the offspring of a horizontal imitator mutates into a vertical imitator. (D) Horizontal imitators dominate the non-genetic evolutionary system due to the asymmetry in information transmission between vertical and horizontal imitators. Rare temporary reinvasions by vertical imitators might occur.

4.2.3 Low Vertical Imitation Error and High Horizontal Imitation Error Produces Sequential Invasions

In the previous section, great care was taken to ensure that both vertical and horizontal imitators enjoyed an equal exploration advantage in order to test the value of horizontal imitation above and beyond that of its higher imitation error rate. In this section, however, we set the vertical imitation error rate to a lower value than that of horizontal imitation in order to explore a different question: *Can cultural evolutionary systems adaptively adjust their exploration and exploitation capabilities by switching between types of imitation?* Lowering this error rate is a sensible assumption, as in natural populations the fidelity of vertical imitation can be strengthened by constant interaction and spatio-temporal correlation between parents and offspring, compared with horizontal imitation (i.e., non-parental imitation) (Richerson and Boyd, 2006; Mesoudi, 2011; Findlay et al., 1989; Aunger, 2000; Harris, 1999).

For this modified version of the model, we will focus on a region of the parameter space within which the imitation error rate for vertical imitators is $\mu_v = 0.001$ (i.e., the same value as μ_g) and the horizontal imitation error rate is $\mu_h = 0.005$. All other parameters are set to the same values described in the previous section (for comparison purposes). Similar results can be obtained with qualitatively similar parameterisations (i.e., strong survival selection, weak reproductive selection, low vertical imitation error and high horizontal imitation error).

Our results show a distinctive pattern in which horizontal imitators produce genotype-phenotype disengagement in the early stages of the simulation, but are displaced subsequently by vertical imitators. Vertical imitators then dominate non-genetic evolution until the optimum is reached. Plot A in Figure 4.10 shows the typical pattern of an individual simulation with two distinctive periods; an initial period where horizontal imitators produce genotype-phenotype disengagement, and a subsequent period where vertical imitators invade once the phenotypes are close to the optimum. Plot B shows the average trajectories for 1×10^4 replicates. Similar to the plots in Figures 4.7 and 4.8, plot B does not represent a specific simulation run.

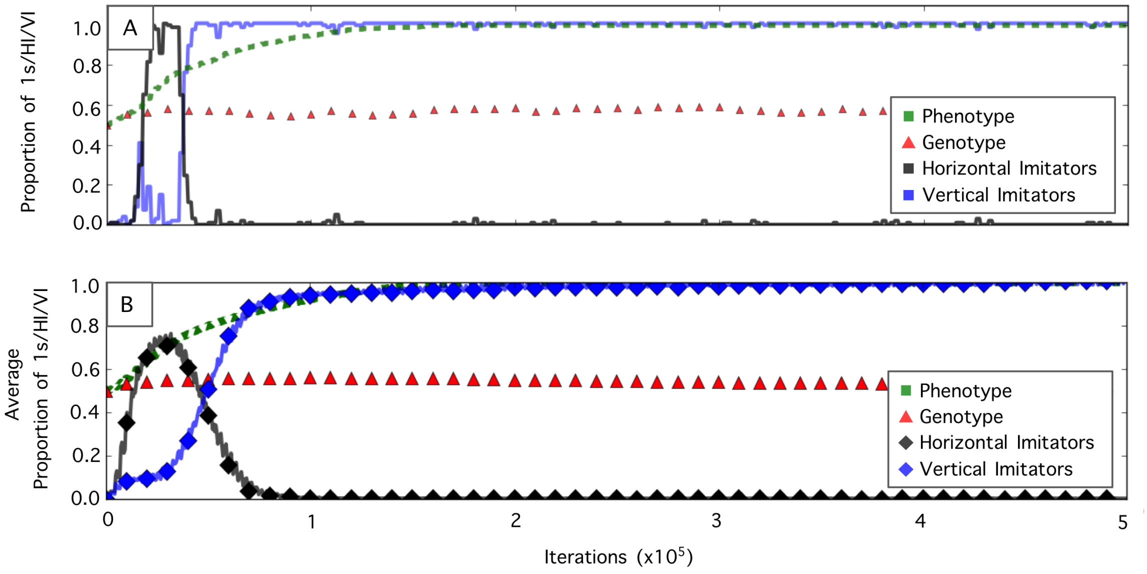

FIGURE 4.10: Plot A shows a typical individual simulation run for a model with different imitation error rates for horizontal and vertical imitators. Early genotype-phenotype disengagement occurs due to the invasion of horizontal imitators; vertical imitators invade after disengagement (and displace horizontal imitators) when solutions are close to the optimum. In plot B, the average trajectories for 1×10^4 replicates are shown for simulation replicates with both vertical and horizontal imitators. The green dotted line shows the average number of 1s in the phenotypes and the red triangles the average number of 1s in the genotypes. Disengagement between these two lines occurs early on while the average frequency of horizontal imitators rises steeply (black line/diamonds). After disengagement takes place, horizontal imitators momentarily drive the non-genetic system until vertical imitators take over (blue line/diamonds). For both these plots, vertical imitation occurs with high inheritance fidelity (i.e., low imitation error rate) and horizontal imitation occurs with a high error rate, $\mu_v = 0.001$ and $\mu_h = 0.01$. Boltzmann selection coefficients are fixed at $x_d = 1$, $x_r = 1 \times 10^3$. Genetic mutation is fixed at $\mu_g = 0.001$

This result is not surprising, but it is interesting from a biological and algorithmic point of view. Early in the simulation, when half of the phenotypic bits are zero, high imitation error rates create, on average, fitter variants than low imitation error rates. During this phase the high error rates of horizontal imitators can be said to be “constructive” (i.e., exploration stage). This process drives genotype-phenotype disengagement. Horizontal imitation, with its high error rate, broadens the gap between the proportion of 1s in the phenotype and the genotype until phenotypes are optimised to a point where high imitation error rates create, on average, less fit variants than low imitation rates (i.e., exploitation phase).

When entering this phase, the previously discussed model with no vertical imitation (see Chapter 3) did not allow reinvasion of instinctive actuators because the disengaged (and hence randomly drifting) genotypes could not compete with the imitated phenotypes with much higher fitness values. In this extended version, however, vertical imitators (i.e., parental imitators with low error rates) are able to imitate the fit phenotypes evolved through horizontal imitation if they are born as the mutated offspring of a horizontal imitator (i.e., an offspring that mutated its imitation switch value from 1 to 2). If this mutation happens after the average phenotype in the population has evolved beyond the critical sum-of-1s value α (i.e., the value where high mutation rates start being destructive), the high fidelity of vertical imitators provides them with an advantage over the relatively low fidelity of horizontal imitators.

In the simulations used in this subsection, the difference between error rates for vertical and horizontal imitation is large, hence the sourcing advantage of horizontal imitation (i.e., rapid access and spread of trending variants) (see previous subsection) has little effect compared to the high inheritance fidelity advantage of vertical imitation. Under these conditions, in contrast to the results of the previous subsection, vertical imitators will persist in the long term. Figure 4.11 explains this process in a diagram form.

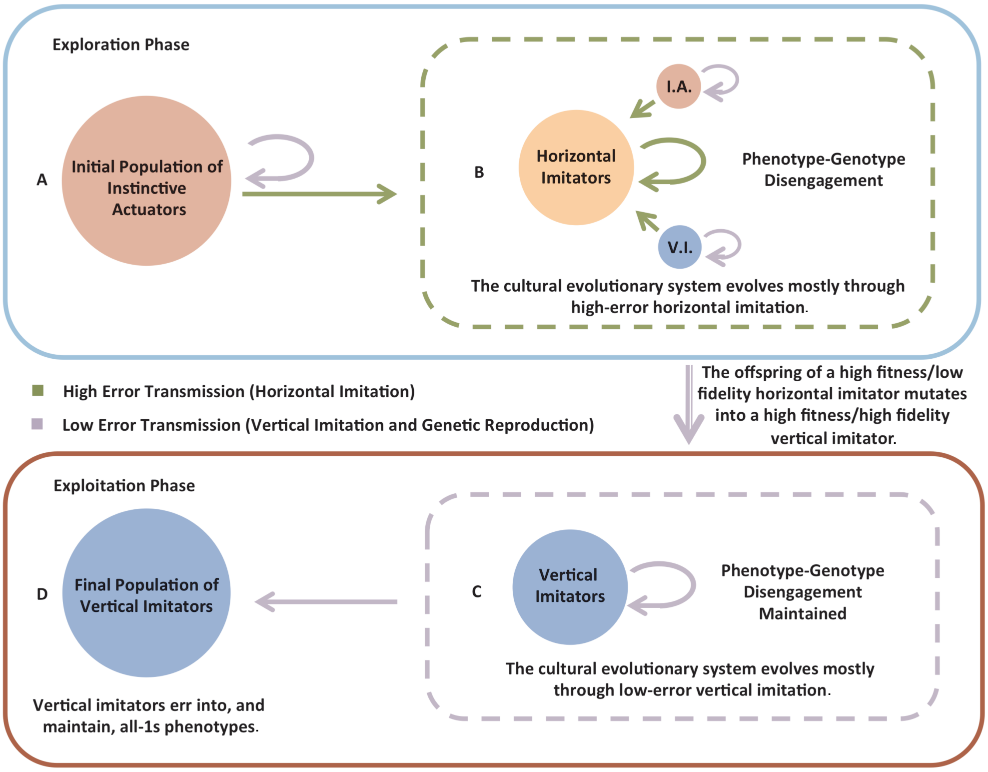

Figure 4.11: The diagram shows the two-phase dynamic of a system with high-error horizontal imitation and low-error vertical imitation. (A) The first step takes place when mutant horizontal imitators copy information that was initially expressed in the phenotype of an instinctive actuator. (B) Once the information gets into a chain of horizontal imitators, genotype-phenotype disengagement occurs and the non-genetic evolution towards fitter variants is then driven by high-error imitation. If, by chance, any instinctive actuator or vertical imitator errs into a fitter variant during this period, horizontal imitators are likely to incorporate it (and further evolve it) when it starts increasing in frequency. Due to the poor average performance of the variants existing in the population, high error rates tend to be constructive in this phase (i.e., exploration phase). (C) Phenotypes are optimised (mostly by horizontal imitators) to a point where high imitation error rates generate deleterious alleles faster than it finds beneficial ones. At this point, a population of horizontal imitators enters a destructive phase. Vertical imitation then becomes advantageous and takes over the non-genetic evolutionary system, further evolving phenotypes with a low-error rate during transmission. Some vertical imitators obtain the fit phenotypes evolved by the previous invasion of horizontal imitators when the offspring of horizontal imitators mutates their imitation switch from 1 to 2. (D) Vertical imitators then drive the non-genetic evolutionary system and eventually find the optimal solution.

This model extension is interesting from a biological and an algorithmic perspective. Biologically, it shows how basic properties of information transmission in social learning systems (i.e., differences in imitation error rates for vertical and horizontal imitation) can drastically affect evolutionary dynamics. Algorithmically, the two-phase pattern in our results is highly desirable in some optimisation methods, where the balance between exploration and exploitation is adaptively shifted through an optimisation run (DeJong, 2002).

In natural populations, each process of information inheritance generally has its own associated mutation/error rate. The model described here suggests that both exploitative and explorative attributes can exist in minimalistic social learning systems featuring two partially isolated types of imitators (i.e., vertical and horizontal imitators) that can occasionally exchange information (i.e., when one type copies the other, or when one type mutates into the other). In populations of human and non-human primates, a similar process of sequential invasions can occur within social networks, where early innovations are rapidly modified and transmitted with a high error rate; but once a sufficiently good solution is found, further variation becomes more conservative (Kadushin, 2011).

4.3 Discussion

In this chapter, we have explored a model extension that includes vertical imitation as an evolvable trait in the population. The resulting model is very simple, but still has the same limitations discussed in our original version (see Chapter 3). As in the original model, we have not included individual learners or decision-making attributes (e.g., biased imitation, or individual improvement of imitated behaviours) in order to test that social learners (using either vertical or horizontal imitation) can fix in a population through an entirely Darwinian process (i.e., without being parasitic on individual learners or depending on additional processes of selection). We have framed the work in this chapter as a theoretical effort to understand the advantages and disadvantages of horizontal information transmission compared to vertical information transmission, beyond the exploration benefit introduced by high imitation error rates (i.e., compared to genetic mutation).

4.3.1 Horizontal Imitation is Key to Social Learning

Most of the processes that make social learning interesting are directly dependent on horizontal transmission (Mesoudi, 2011; Richerson and Boyd, 2006). In other words, the non-genetic evolutionary system driven by social learning is algorithmically different from genetic evolution precisely because information is inherited through different pathways, both vertical and horizontal. For this reason, explaining the advantage (if any) of horizontal transmission is important if we are to understand social learning as something more than an alternative substrate for genetic evolution.

In human and animal populations, the effects and relative dominance of vertical and horizontal imitation are unclear (Shennan and Steele, 1999; Laland et al., 2000; Harris, 1999). Animal species with social learning tend to obtain a large component of their behavioural repertoire from vertical imitation (Findlay et al., 1989; Aunger, 2000; Hewlett and Cavalli-Sforza, 1986; Shea, 2009). Humans are also thought to imitate most of their early-life behavioural variants from their parents (Aunger, 2000; Hewlett and Cavalli-Sforza, 1986; Mulder et al., 2006; Shennan and Steele, 1999; Laland et al., 2000). However, many of these studies have been criticized for overestimating the contribution of vertical imitation (Harris, 1999). According to Harris (1999), behavioural correlation between parent and offspring is almost non-existent when correcting for genetic inheritance (Harris, 1999; McElreath and Strimling, 2008). Therefore, it has been suggested that the dominant form of transmission in social learning is horizontal (Harris, 1999). Supporters of this idea argue that dominant horizontal transmission is required to explain important processes of cultural evolution, processes that occur at much faster time scales than a generation (e.g., rapid propagation of new ideas) (Richerson and Boyd, 2006; Perreault, 2012; Harris, 1999).

4.3.2 Existing Theoretical Work Makes Additional Assumptions to Explain the Evolution of Horizontal Imitation

As discussed by McElreath and Strimling, theoretical work that focuses on the evolution of horizontal imitation in the presence of vertical imitation is rare (McElreath and Strimling, 2008). Existing contributions rely on analytical models with no extended strategy space or imitation error rates (Cavalli-Sforza and Feldman, 1981; Takahashi, 1999; McElreath and Strimling, 2008) and only explain the evolution of horizontal imitators in the context of additional decision-making attributes for social learners or the presence of individual learners in the population. For example, in Takahashi (1999), vertical imitators are favoured (i.e., dominant) for a large range of conditions, including strong survival selection (Takahashi, 1999). However, horizontal imitators can only invade when imitation is biased towards fitter-than-average variants due to the aforementioned issue of phenotype-fitness de-correlation (see Subsection 1.2.1) (Takahashi, 1999; McElreath and Strimling, 2008).

More recently, McElreath and Strimling extended Rogers' (1988) basic analytical model to include vertical imitators (McElreath and Strimling, 2008). In this work, the authors compare the effect of survival selection (selection for viability) and reproductive selection (selection for fecundity) on the dominance of each type of social learner (McElreath and Strimling, 2008). Some of their findings are in agreement with our results (e.g., when survival selection is strong, vertical imitators dominate). However in their model, individual learners are present at low frequencies at equilibrium and imitators do not fix. Furthermore, for horizontal imitators to dominate, they not only require the absence of reproductive selection (similar to our model) and the presence of individual learners (unlike our model), but also a moderate amount of environmental fluctuation. Similar to Rogers' (1988), environmental fluctuation is represented by the rate of change of a two-state variable (i.e., an extended strategy space is not considered) (McElreath and Strimling, 2008).

As discussed in Subsection 4.2.2, despite the lack of both environmental fluctuations and individual learners, our model can explain the long-term dominance of horizontal imitators. However, McElreath and Strimling's conclusions on how horizontal imitators dominate under moderate environmental fluctuations rely on a process with some similarities to that described in our simulation results. According to these authors, horizontal imitators outcompete vertical imitators because they acquire variants created by individual learners more easily (i.e., variants that are guaranteed to be adapted to the current environmental state) (McElreath and Strimling, 2008). Similarly, we have shown that when selective pressure and imitation error rates are equal in vertical and horizontal imitators, the advantage of the horizontal kind lies in their ability to incorporate trending variants (i.e., variants that are fitter than average) faster than vertical imitators (see subsection 4.2.2). In both models, every horizontal imitator has a chance to copy fitter-than-average variants (i.e., trending variants in our model; individually learned variants in McElreath and Strimling's model) from any other individual in the population. In contrast, also for both models, vertical imitators rarely acquire variants from fitter horizontal imitators (our model), or individual learners (McElreath and Strimling's model) (McElreath and Strimling, 2008).

Unlike previous work, our model explains the dominance of horizontal imitation over vertical imitation without making unnecessary assumptions, such as: horizontal imitators having decision-making capabilities, individual learners being present in the population, and environmental fluctuations occurring during the evolution of social learning. Showing that horizontal imitation can evolve and dominate over vertical imitation without these requirements broadens the range of species and evolutionary scenarios where its presence can be explained.

4.3.3 Horizontal Imitation as a Special Case of Horizontal Transmission

Our findings can be extended to other non-genetic evolutionary systems besides social learning. As a general case, our model shows how indiscriminate horizontal transmission can be dominant over its vertical counterpart for traits that increase viability (i.e., survivability) rather than fecundity (i.e., reproduction), even when the substrates for each kind of transmission are not isolated. In the context of our alternative example, where a symbiotic bacteria (i.e., a bacteria that benefits its host) can be horizontally or vertically transmitted by its host (e.g., through faecal matter consumption), results suggest that a bias towards horizontal transmission (e.g., consuming the faecal matter of non-parental individuals) will evolve even if mutation rates are equal for vertical and horizontal channels (see Chapter 3 for context). Furthermore, when vertical and horizontal error rates differ (i.e., when vertical transmission has a lower error rate), horizontal transmission can initially explore novel environmental challenges (i.e., with high error rates and rapid fixation of trending variants), and vertical transmission can later exploit and maintain solutions (i.e., through low error rates) when phenotypes are closer to the optimum.

This sequential process of exploration and then exploitation attenuates the shortcomings of social learners compared to instinctive actuators in the original model in Chapter 3. In particular, it allows social learners to reach a selection-error balance with fitter phenotypes (i.e., with a larger average sum-of-1s value in their phenotype), similar to the phenotypes that a population of monomorphic instinctive actuators would find. In this model extension, we have not, however, shown that the fixation of social learners is adaptive at the population level (i.e., that it increases the long-term average fitness of the population) compared to a monomorphic population of individual learners. In the next chapter, we will prove that this is possible when a moderate amount of environmental change is introduced.

4.4 Conclusions

In this chapter, the model introduced in the previous chapter was extended in order to explore the effects of horizontal imitation. The three main findings of this chapter are:

1. In a system featuring the availability of both vertical and horizontal imitation under strong reproductive selection, the emergence of non-genetic cultural evolutionary systems is brought about by an invasion of vertical rather than horizontal imitators. From an initial population of instinctive actuators, vertical imitators fix (and cause genotype-phenotype disengagement) due to the combination of their high imitation error rate (i.e., higher than genetic mutation) and strong selective pressure (i.e., stronger than horizontal imitators).
2. Everything else being equal, horizontal imitation is more likely to dominate the long-term evolutionary process in a population with strong survival selection and weak (or non-existent) reproductive selection. The process by which horizontal imitators gather information from the whole population (i.e., by selecting a random agent to imitate) is advantageous when compared to vertical imitation. Horizontal imitation allows individuals to acquire trending variants when they start rising

in frequency, regardless of which individuals express these trending variants (i.e., horizontal imitators, vertical imitators or instinctive actuators). In contrast, vertical imitators are constrained to vertical lines of inheritance; if, by chance, an adaptive variant is expressed by another type of agent (i.e., horizontal imitators or instinctive actuators), vertical imitators do not have access to this variant and hence are more likely to be displaced by it (see subsection 4.2.2).

3. When horizontal imitation has a high-error rate compared to vertical imitation, a two-phase dynamic unfolds during the optimisation process. In the first phase, the non-genetic evolutionary system emerges with the invasion of horizontal imitators (i.e., genotype-phenotype disengagement is brought about by horizontal imitators). In the second phase, vertical imitators displace horizontal imitators when the high error rate of horizontal imitation becomes destructive (i.e., when it creates, on average, less fit variants than the alternative; vertical imitation). This two-phase dynamic is interesting from a biological and algorithmic perspective as it allows the population to adapt the exploration-exploitation trade-off at different stages of the evolutionary process.

In contrast to existing theoretical models our simulation results explain the evolution of horizontal imitation without assuming that social learners have decision-making attributes or individual learners are present in the population. As in the model introduced in Chapter 3, we have achieved this by including explicit mutation and imitation error rates, an extended strategy space, and by separating selection into survival and reproductive components. In our simulation, the process of genotype-phenotype disengagement leads to the fixation of culture in the same way it does in the original model in Chapter 3. Once the cultural system emerges, the dominant type of imitation varies depending on the relative strength of each kind of selection (i.e., reproductive and survival) and the relative imitation error rates of vertical and horizontal imitators.

Chapter 5

Unbiased Social Learning is Adaptive Under Strong Survival Selection and Moderate Environmental Change

In Chapter 3, we showed that the most unassuming form of social learning, where imitation is unbiased and guided variation is not present, could invade and fix in a population of instinctive actuators without requiring the presence of individual learners (see Chapter 3). We also showed that this fixation is not caused only by the higher imitation error of social learners, compared to genetic mutation, but also by the advantage of horizontal transmission compared to the vertical kind. However, in the previous chapters, we did not show that the fixation of social learners increases the overall mean fitness of the population, as a monomorphic population of instinctive actuators can also reach the optimum when social learners do not fix.

In our model, social learners fix because high imitation error rates and horizontal transmission (i.e., unbiased non-parental transmission) are advantageous during the early stages of optimisation and cause genotype-phenotype disengagement. Once disengagement occurs, it is highly unlikely that instinctive actuators reinvade the population, even when their low error inheritance would be advantageous in the later stages of optimisation (i.e., when low error rates better conserve the optimised variants).

In existing literature, the long-term adaptive value of social learning, and by extension of cultural evolutionary systems, has been the focus of many theoretical models (Rogers, 1988; Henrich and Gil-White, 2001; Castro and Toro, 2004; Enquist et al., 2008; McElreath and Henrich, 2007). Since Rogers' (1988), the long-term adaptive value of unbiased social learners has been questioned. As we explained in Chapter 2, the fitness of social learners at equilibrium in Rogers (1988) is equal to the fitness of individual learners and hence their presence does not increase the mean fitness of the population (i.e., the population does not perform on average fitter behaviours when social learners invade).

After Rogers (1988), the question of how social learning could increase the overall mean fitness of a population was explored in numerous theoretical models (Richerson and Boyd, 2006; Mesoudi, 2011; Henrich and Gil-White, 2001; Castro and Toro, 2004; Enquist et al., 2008; McElreath and Henrich, 2007). Most of these models have already been discussed in previous chapters; all of them include decision-making attributes (i.e., fitness-biased imitation or individual improvement of imitated variants) and/or the presence of individual learners in the population (Richerson and Boyd, 2006; Mesoudi, 2011; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007; Henrich and Gil-White, 2001; Castro and Toro, 2004; Enquist et al., 2008; McElreath and Henrich, 2007; Borg and Channon, 2012).

So far, our results have explained how the fixation of social learners can occur without any of these assumptions (see Chapter 3). In this chapter, we further demonstrate that the fixation of unbiased social learners performing horizontal imitation can increase the overall mean fitness of a population compared to a monomorphic population of instinctive actuators. In the process, this work answers our third and fourth research questions: *Does environmental change increase the chance of social learners fixing in a population without individual decision-making capabilities? Does the fixation of social learners increase the overall mean fitness of a population under moderate environmental change?* To achieve this, we will add an extension to our original model in Chapter 3: a rate of environmental change. This feature is generally included in most theoretical models as a binary change of the environmental state (Rogers, 1988; Henrich and Gil-White, 2001; Castro and Toro, 2004; Enquist et al., 2008; McElreath and Henrich, 2007; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007), but here we will include it as a rate of change of the optimal sequence (i.e., the sequence we have considered to be all-1s in Chapter 3 and Chapter 4).

Before explaining in detail how our model extension works, we will discuss the existing models that explore the effect of environmental change on the evolution of social learning. The general consensus in the literature is that moderate environmental change has a positive effect on the evolution of social learning (i.e., social learners are more likely to invade under moderate environmental change) (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007). However, to reach this conclusion existing models must include individual learners (i.e., they only explain the evolution of social learners as part of a mixed population) (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007). As we will see, our results agree with the conclusions of this consensus, but we reach those conclusions with a different mechanism and without requiring individual learners.

Existing theoretical approaches explore the effect of environmental change using dynamical system models with expressions for the frequencies of three kinds of agents: instinctive actuators, individual learners, and social learners. As explained before, instinctive actuators inherit their behaviour from their parents, individual learners pay a trial-and-error cost in fitness to always perform a behaviour that matches their environment, and social learners imitate other agents in the population (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007). Environmental change is represented as a probability function for match (or mismatch) between a behaviour obtained by a particular type of agent and the current environmental state; a better matching rate translates into higher fitness for that type. Selection only affects fecundity (i.e., there is no survival selection) and the explored strategy space is binary (i.e., there are only two fitness levels: environment-matching or environment mismatching). Under this set-up, all models reach similar conclusions regarding the dominance of each type of agent for low, moderate and extreme rates of environmental change (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

In existing models, low rates of environmental change lead to the evolution of monomorphic populations of instinctive actuators. Under these conditions individuals experience the same environmental state as their parents, and therefore the genetically evolved instincts of instinctive actuators are likely to match the current environmental state. In contrast, individual learners pay a fixed trial-and-error cost for the same environmental match. According to these models, social learners do not have an advantage over instinctive actuators in this scenario, as unbiased imitation does not yield a reproductively fitter behaviour than the average of instinctive actuators. The authors of these models conclude that an invasion of social learners is therefore unlikely for low rates of environmental change (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

When environmental change is moderate, the inherited behaviours of instinctive actuators are less likely to match the current environmental state, and hence their fitness becomes lower than other agents. On the other hand, individual learners are guaranteed to match their current environmental state and therefore never pay the cost of environmental mismatch (only the trial-and-error cost). As long as the risk-cost of environmental mismatch for instinctive actuators is larger than the trial-and-error cost of individual learners, the latter can displace the former. Under this condition, the advantage of social learning increases compared to instinctive actuators; and at the same time, their fitness becomes dependent on the frequency of individual learners (i.e., social learners imitate the environment-matching behaviours learnt by individual learners with a rate proportional to their frequency, but do not pay the trial-and-error cost to learn them) (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

As the rate of environmental change increases in magnitude, the mixed population of social and individual learners at equilibrium starts tending towards a majority of individual learners. If the rate of environmental change is high enough, the risk of environmental mismatch incurred by social learners becomes a disadvantage. Under this condition, individual learning is the fittest learning strategy for all frequencies of social and individual learners, and therefore the existing models converge towards a monomorphic population of individual learners (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007). Adding decision-making abilities for social learners (e.g., biased imitation) diminishes their risk of environmental mismatch even for high rates of environmental change and therefore increases their frequency at equilibrium.

In our simulation, we do not include individual learners but we still conclude that a moderate rate of environmental change increases the chances of social learners fixing in the population. Furthermore, we show that the evolution and fixation of social learners under a moderate rate of environmental change increases the overall mean fitness of the population compared to a population of instinctive actuators. This is the first theoretical effort to have achieved this result without invoking decision-making attributes or the presence of individual learners.

In the following two sections, we introduce the model extension and analyse its results in three subsections. The first subsection shows how a moderate rate of environmental change increases the area of the parameter space where culture emerges. In the second subsection, we focus on an area of the parameter space where the proportion of simulations that converge to culture drastically increases with moderate rates of environmental change. In this subsection, we also show how our results agree with the bounds estimations from Chapter 3 (see Section 3.3). Finally, the last subsection shows that for certain levels of environmental change, a fixed population of social learners increases the overall mean fitness of the population without requiring biased imitation or the existence of individual learners. A final section with discussions and general conclusions closes this chapter.

5.1 Model Extension with Environmental Change

As in the original model, each agent is represented by two sequences of bits of length 200 (i.e., $L = 200$) corresponding to phenotype and genotype, and a single extra bit corresponding to the imitation switch. When individuals reproduce (i.e., when asexual reproduction occurs), both genotype and imitation switch are inherited from parent to offspring with a mutation rate per bit of μ_g for the genotype and μ_c for the imitation switch.

Agents born with an imitation switch in position 0 are instinctive actuators and express their genotype into their phenotype with a 1-to-1 map (i.e., their phenotype is always an exact copy of their genotype). When the agent is born with an imitation switch in position 1, it is a social learner, and therefore its phenotype will be an imitated copy of the phenotype of a randomly selected individual in the population. During imitation, errors occur with an error rate of μ_p per bit. In contrast with the model extension in Chapter 4 and similar to the original model in Chapter 3, the imitation switch in this model only has two positions: horizontal imitation (i.e., unbiased social learning) and instinctive actuation. Each individual only expresses one phenotype throughout its lifetime regardless of whether this phenotype is genetically inherited (i.e., instinctively actuated) or imitated from a random individual at birth (i.e., socially learned).

As in the original model, selection is introduced by a death function and a reproduction function. On each iteration, an individual is selected for reproduction with a probability proportional to the fitness of their phenotype, and another individual is selected for death (i.e., to be replaced by the reproducing agent's offspring) with a probability inversely proportional to the fitness of their phenotype. Fitness is calculated using the sum-of-environment-matching-bits (L_m) instead of the sum-of-1s value (L_1) used in the original model in Chapter 3. In the extension explained in this chapter, the optimal sequence (which represents the environment-to-be-matched) changes with a given probability per iteration; each time the sequence changes, all agents in the population have their sum-of-environment-matching-bits (L_m) recalculated.

The strength of these selection processes is controlled by the reproductive selection coefficient x_r and the survival selection coefficient x_d in Boltzmann-weighted functions (Equations 5.1 and 5.2). These equations calculate the probability for a given phenotype i or j, to be selected for reproduction P_{ri} or death P_{dj} on each iteration, respectively.

$$P_{ri} = \frac{e^{\frac{L_{mi}}{x_r}}}{\sum_{k=1}^{N} e^{\frac{L_{mk}}{x_r}}} \tag{5.1}$$

$$P_{dj} = \frac{e^{\frac{L-L_{mj}}{x_d}}}{\sum_{k=1}^{N} e^{\frac{L-L_{mk}}{x_d}}} \tag{5.2}$$

The optimal sequence (i.e., the sequence used to calculate fitness by counting the number of phenotype bits that match to it) is initialised as a random bit string. All phenotypes and genotypes in the initial population are also initialised at random. All agents start as instinctive actuators (i.e., their genotype and phenotype sequences are identical and their imitation switch value is 0). The optimal sequence changes through the simulation as a function of the parameter δ_e, which has a fixed value for each simulation run. The parameter δ_e represents the probability that *a randomly selected bit* in the optimal sequence, flips from 1 to 0, or from 0 to 1, in an iteration (see Subsection 4.1.1). For all results shown in this chapter, we will specify the rate of environmental change using δ_e, and additionally, we will provide the equivalent μ_e value for the given δ_e. The value of μ_e represents the estimated probability for *each bit* in the optimal sequence to change from 1 to 0, or from 0 to 1 (i.e., $\mu_e = \delta_e/L$). The value of μ_e is comparable against mutation rates and imitation error rates (i.e., μ_g and μ_p), which are defined as bit-flip probabilities for *each bit* in a bit string. Equations 5.1 and 5.2 are similar to Equations 3.1 and 3.2 from Chapter 3, but apply to any optimal sequence (i.e., not only a sequence of all-1s).

The environmental change function is implemented as an extra step in our simulation model's logic (see Figure 5.1). In this step, a randomly selected bit in the current optimal sequence is changed with probability δ_e. When the optimal sequence changes, all fitness values are recalculated immediately afterwards.

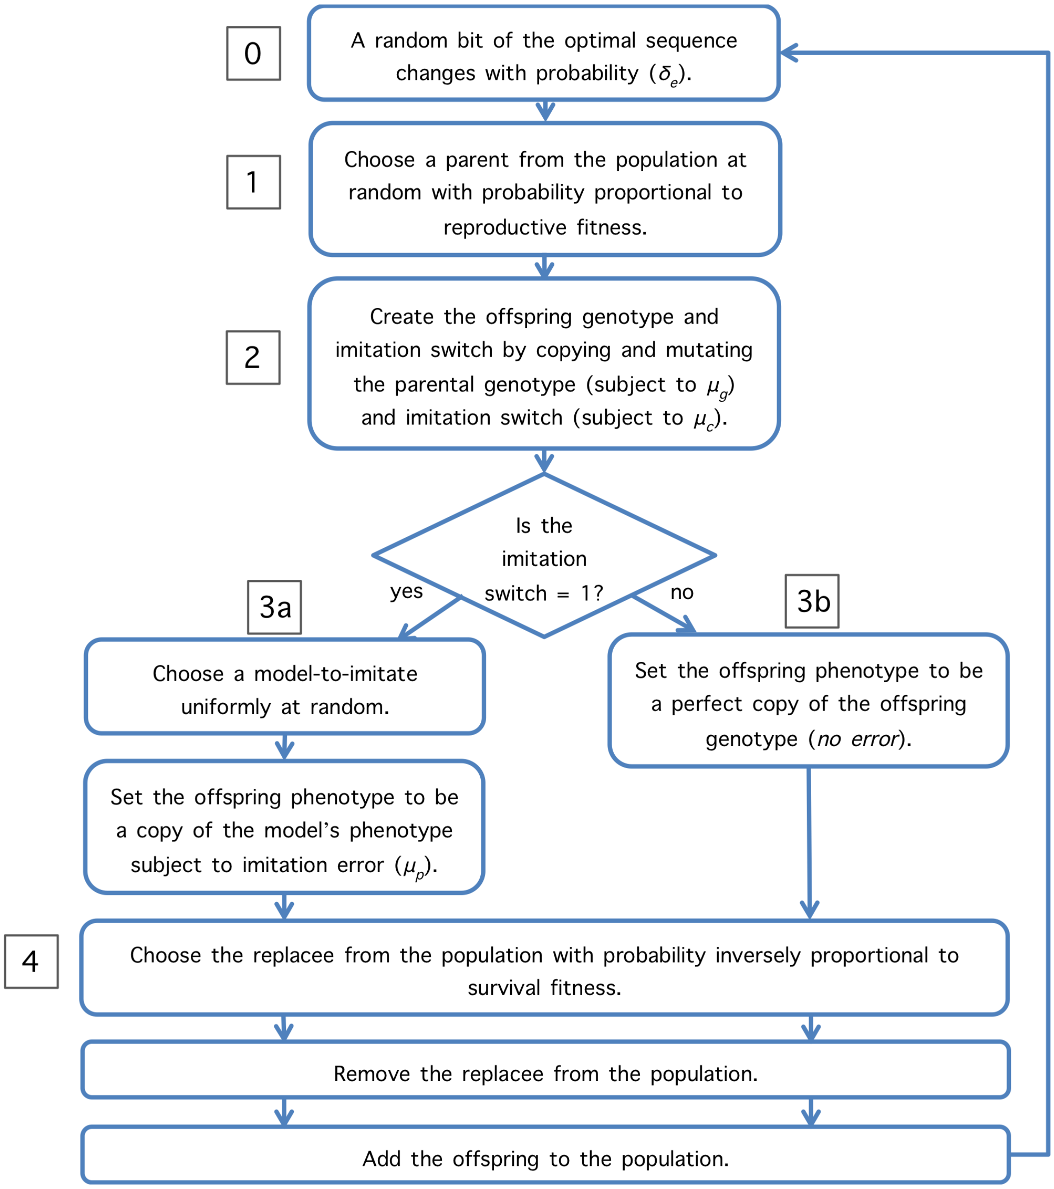

FIGURE 5.1: An algorithmic representation of the extended model's logic. (0) With probability δ_e a randomly selected bit in the optimal sequence is flipped (i.e., from 1 to 0, or from 0 to 1). If the optimal sequence changes, all fitness values are recalculated. (1) The selection function picks an individual from the population. (2) Its genotype string is copied including the imitation switch (mutations occur at a rate of μ_g per genotype bit and μ_c for the imitation switch). (3a) If the imitation switch has value 1, a random individual from the population is selected and its phenotype will be copied (imitation errors occur at a rate of μ_p per bit). (3b) If the imitation switch has value 0, the phenotype will be a perfect copy of the individual's genotype. (4) The resulting combination of phenotype and genotype will replace an individual selected by the death function.

It is important to point out that the introduction of a rate of environmental change does not alter the character of the optimisation problem the population is "solving". In Chapter 3, we referred to this problem as the "onemax" problem, where fitness was simply a function of the sum-of-1s value in the phenotype sequence and the optimal sequence was all-1s. In this chapter, we do not intend to change the character of the problem, which remains the same when the optimal string changes (i.e., fitness is still calculated as the sum of matching bits between phenotype and the optimal sequence). The distinction between environmental change and change of phenotype-fitness landscape (i.e., changes in the number of optima/ruggedness of the fitness landscape) is important for reasons that will be discussed later in this chapter, when we explain how the latter can also make social learning adaptive, but through a different mechanism than the former.

Compared to existing theoretical models that explore the effect of environmental change on the evolution of social learning, our model is distinctive in that it: (1) separates natural selection into survival and reproduction; and (2) includes an extended strategy space with an explicit optimisation problem. As we will explain in the next section, these features allow us to reach conclusions similar to the ones reached by previous models, but for different reasons. In contrast with other models, our work shows that a moderate rate of environmental change makes unbiased social learning increase the overall mean fitness of the population without requiring the presence of individual learners. This is, to our knowledge, the first theoretical model to reach such a conclusion (Richerson and Boyd, 2006; Mesoudi, 2011; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

5.1.1 Simulation Steps

In this subsection, steps presented in Figure 5.1, that were not present in Figure 3.1, are explained in detail. In a similar manner to Subsections 3.1.1 and 4.1.1, in this subsection, we put emphasis on clarifying the internal logic of each function and how it affects the bit strings that define individual agents in the simulation. This information can also be found in Appendix B, along with the description of the original simulation model (see Chapter 3) and the details of the extension explained in Chapter 4. Once again, we want to mention that the internal logic of our model, before generating our analytical results, was "unit tested". That is, for each function, the expected output for a range of possible inputs was evaluated.

5.1.1.1 Step 0: Implementation of the Environmental Change Function

In the model extension explained in this chapter, we included a rate of environmental change (δ_e). The parameter δ_e represents the probability that a *randomly selected bit* in the optimal sequence (i.e., the sequence that encodes the optimal solution) changes from 1 to 0, or from 0 to 1, in a single iteration. Therefore, the implementation of the environmental change function occurs in two steps: (1) a single bit in the optimal sequence is selected with uniform random probability, and (2) the selected bit changes from 1 to 0, or from 0 to 1, with probability δ_e. Notice that the environmental change function is therefore different from the previously described mutation and imitation error functions. When a genotype string is mutated with rate μ_g, each bit in the string changes from 1 to 0, or from 0 to 1, with probability μ_g. Similarly, when a phenotype string is erroneously imitated with error rate μ_p, each bit in the string changes from 1 to 0, or from 0 to 1, with probability μ_p.

In this chapter, we have selected the 2-step implementation of the environmental change function due to its superior computational efficiency, compared to an implementation with an independent bit-flip change probability for each bit. Nevertheless, in order to make environmental change rates comparable to our mutation and imitation error rates, we have included the equivalent μ_e rate for every δ_e used in our results. The value of μ_e represents the estimated probability for any given bit in the optimal sequence to change from 1 to 0, or from 0 to 1. This value is calculated as the product of: (1) the probability that a given bit would be selected during uniform random selection (i.e., $1/L$), where L is the length of the optimal sequence, and (2) the probability that the selected bit changes its value (i.e., δ_e). Therefore, $\mu_e = \delta_e/L$. Figure 5.2 shows a diagram explaining the implementation of the environmental change function.

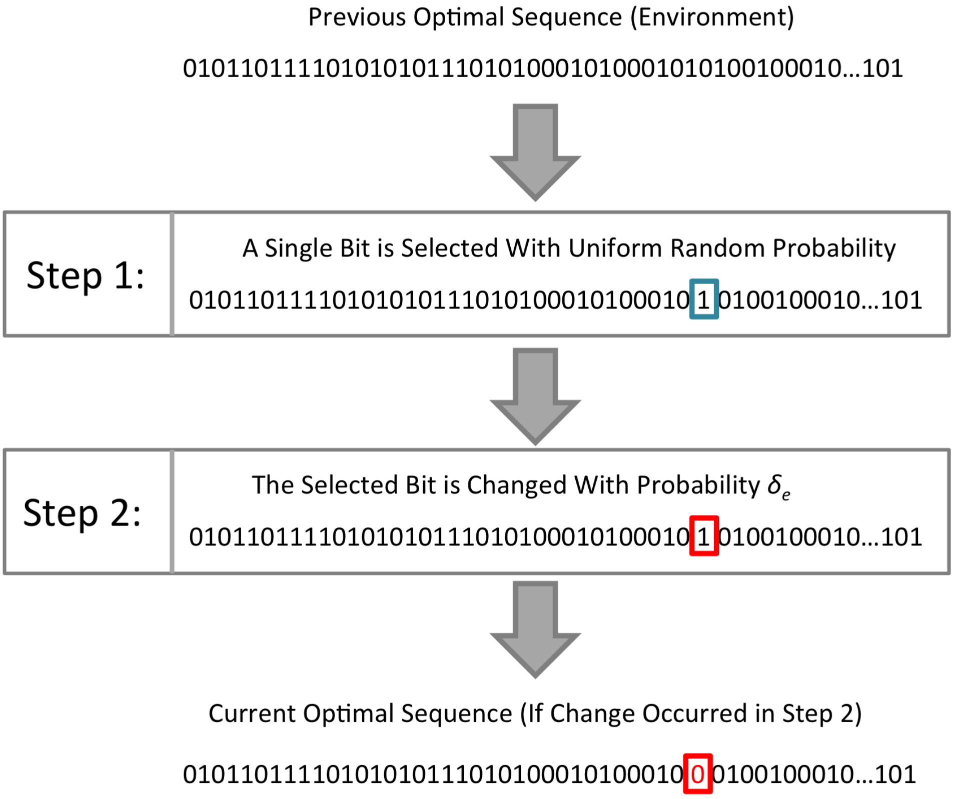

Figure 5.2: Implementation of the environmental change function. Step 1: with uniform random probability, a single bit in the optimal sequence is selected. Step 2: the bit selected in step 1, changes from 1 to 0, or from 0 to 1, with probability δ_e. For each iteration, the resulting sequence is used to calculate reproduction and survival fitness.

5.1.1.2 Fitness Calculations in a Changing Environment

The incorporation of an environmental change function in our model, forces us to frequently recalculate *reproductive selection scores* and *death selection scores*. Every time a bit changes in the optimal sequence, the scores of the entire population have to be recalculated in order to represent the appropriate fitness values for the current environmental state. In the original model, where the optimal sequence was fixed as a string of all-1s, reproductive selection scores are calculated using the sum-of-1s values of the phenotype strings (L_1) and the death selection scores are calculated using the complement of the sum-of-1s values ($L - L_1$) (see Equations 3.1 and 3.2). In this extension, the reproductive selection scores are calculated using the number of phenotypic bits that match their respective positions in the optimal sequence (L_m). Conversely, the death selection scores are calculated using the number of phenotypic bits that do not match their respective positions in the optimal sequence ($L - L_m$) (see Equations 5.1 and 5.2).

5.2 Results of the Model Extension with Environmental Change

The results of this chapter are separated into three subsections: In the first subsection, we show that moderate environmental change increases the area of the parameter space defined by imitation error rate, genetic mutation rate, survival selection strength and reproductive selection strength that converges to culture (i.e., where social learners fix). In this subsection, we also show that this increase agrees with the boundary conditions discussed in Chapter 3 (i.e., culture is still less likely to emerge when genetic mutation rates are advantageous compared to imitation error rates (see Chapter 3)).

In the second subsection, we explain why social learners are favoured by moderate environmental change. We put this explanation to the test by analysing individual simulation runs of our extended model (i.e., the model explained in the previous section) and comparing them to a modified version of our simulation, for which our theory predicts social learners will not fix.

In the last subsection, we show that the fixation of social learners is adaptive in the long term, as it increases the overall mean fitness of the population compared to a monomorphic population of instinctive actuators. In this subsection, we prove that for moderate rates of environmental change, the optimisation process gets stuck in a "red queen" state (Van Valen, 1973). In this state, the rate of adaptation (i.e., the rate at which evolution fixes environment-matching bits) and the rate of environmental change (i.e., the rate at which phenotypes un-match the environment) maintain an equal pace (i.e., the population will no longer optimise its average fitness value after reaching this state). In this situation, the higher imitation error of social learners combined with the advantage of horizontal transmission (see Chapter 4) allows them to track the environment faster than instinctive actuators, and therefore produce fitter phenotype sequences at this *equal-pace* state.

Every simulation run reported here uses a population of one hundred individuals ($N = 100$) and an imitation switch mutation rate of $\mu_c = 0.01$. These values are equal to those used in Chapter 3 and Chapter 4 and are here maintained for comparative purposes.

5.2.1 Social Learners are More Likely to Fix Under Moderate Environmental Change

To explore the effect of environmental change on the fixation of social learners we use an approach similar to that of Chapter 3 and Chapter 4 (i.e., an array of heat maps that explore the multidimensional parameter space of our model); Figure 5.3. Each heat map in the array contains simulation results for fixed values of survival selection strength (i.e., x_d) and reproductive selection strength (i.e., x_r), for all combinations of imitation error rates (μ_p) and genetic mutation rates (μ_g) in the set $[1 \times 10^{-4}, 2 \times 10^{-4}, \ldots, 50 \times 10^{-4}]$. The colour of each coordinate represents the proportion of simulations, from a set of 100 replicates, that converges to culture (i.e., where the proportion of social learners is above 0.95 by the end of 5×10^4 iterations). Figure 5.3 shows a set of nine heat maps for different combinations of survival selection strength and reproductive selection strength; results are shown only for cases where $\mu_p > \mu_g$. Simulations used in this figure have an environmental change rate of $\delta_e = 0$ (i.e., equivalent to those depicted in Figure 3.9).

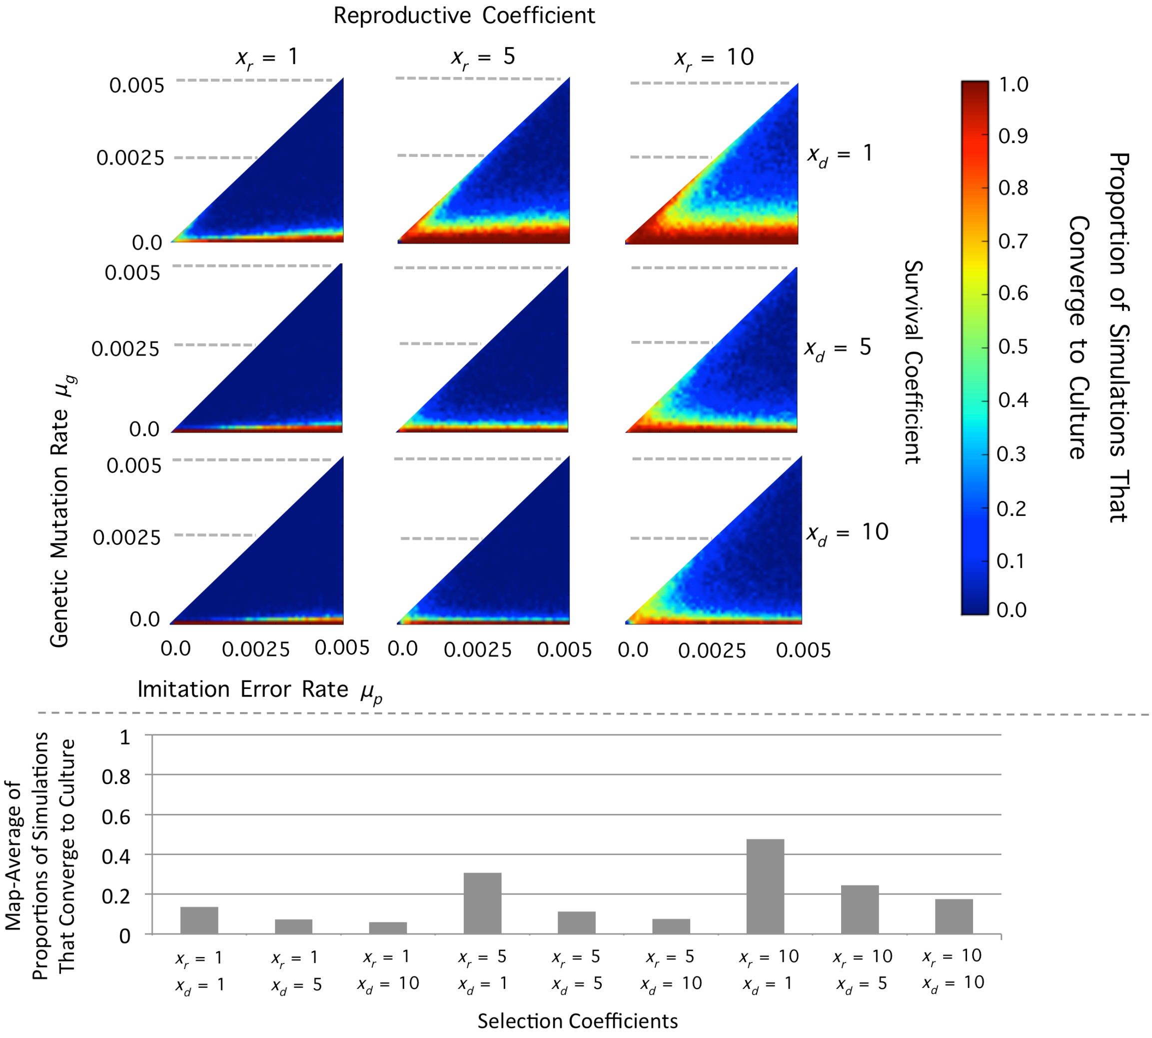

FIGURE 5.3: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right, columns of maps have x_r values of 1, 5 and 10. From top to bottom, rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. In all runs, the rate of environmental change is set to $\delta_e = 0.0$ (i.e., $\mu_e = 0.0$). A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array.This figure is equivalent to Figure 3.9.

The proportion of simulations that converge to culture increases when we change the average rate of environmental change to $\delta_e = 0.04$ (i.e., $\mu_e = 2 \times 10^{-4}$)(see Figure 5.4); and $\delta_e = 0.10$ (i.e., $\mu_e = 5 \times 10^{-4}$) (see Figure 5.5). This is particularly apparent for regions of the parameter space where survival selection is relatively strong compared to reproductive selection (i.e., $x_d = 1$, $x_r = 10$). Notice how areas that previously had no replicates converging to culture now show a high proportion of convergence, and some areas that already had a high proportion of simulations converging to culture increase their proportions even higher. We have included more heat map arrays with various of rates environmental change in the Appendix D.

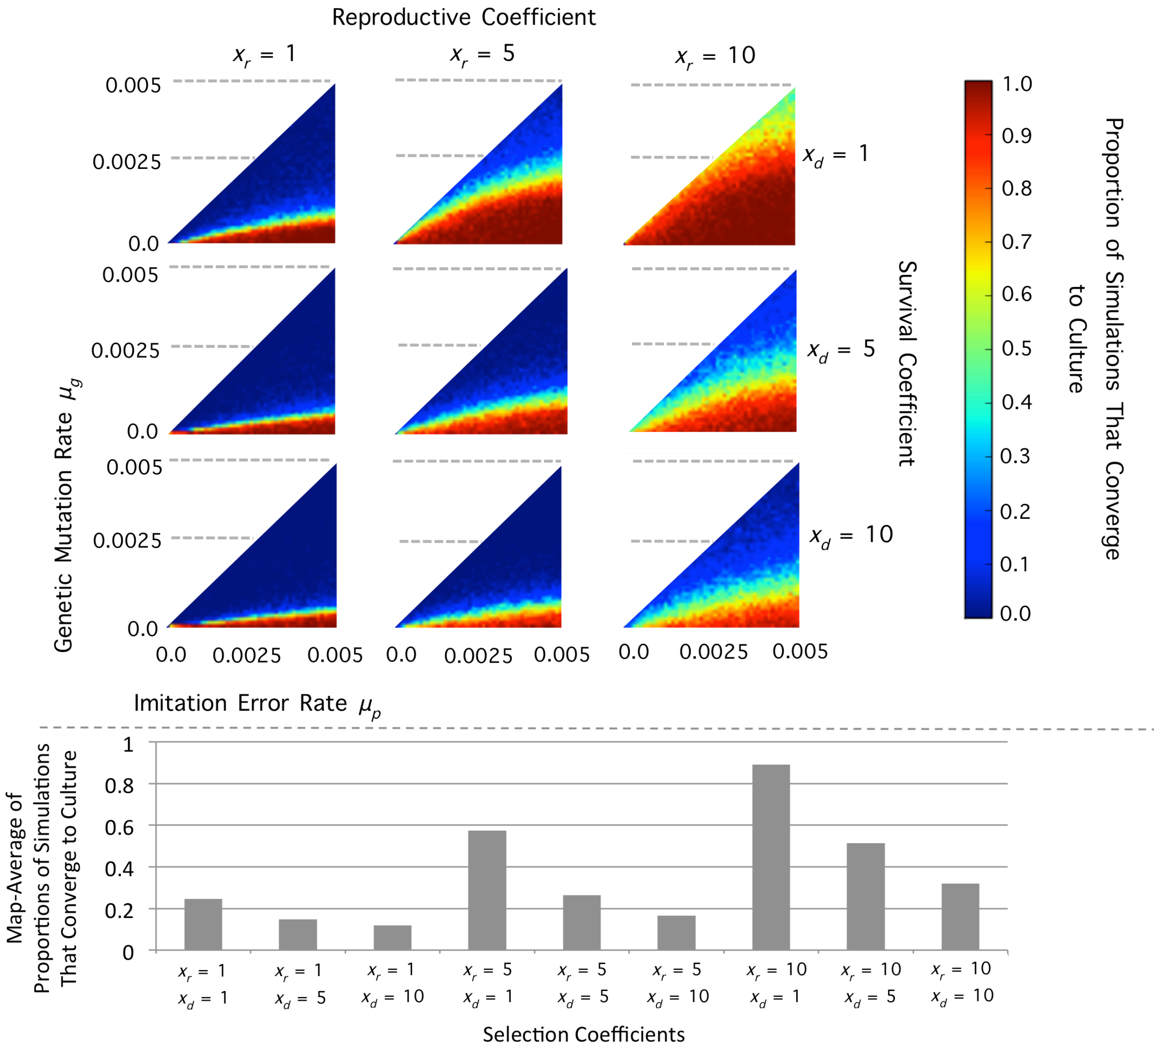

FIGURE 5.4: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right, columns of maps have x_r values of 1, 5 and 10. From top to bottom, rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In contrast to Figure 5.3, in all runs in this figure, the rate of environmental change is set to $\delta_e = 0.04$ (i.e., $\mu_e = 2 \times 10^{-4}$). Compared to Figure 5.3, this heat map array has a larger area of the parameter space converging towards culture.

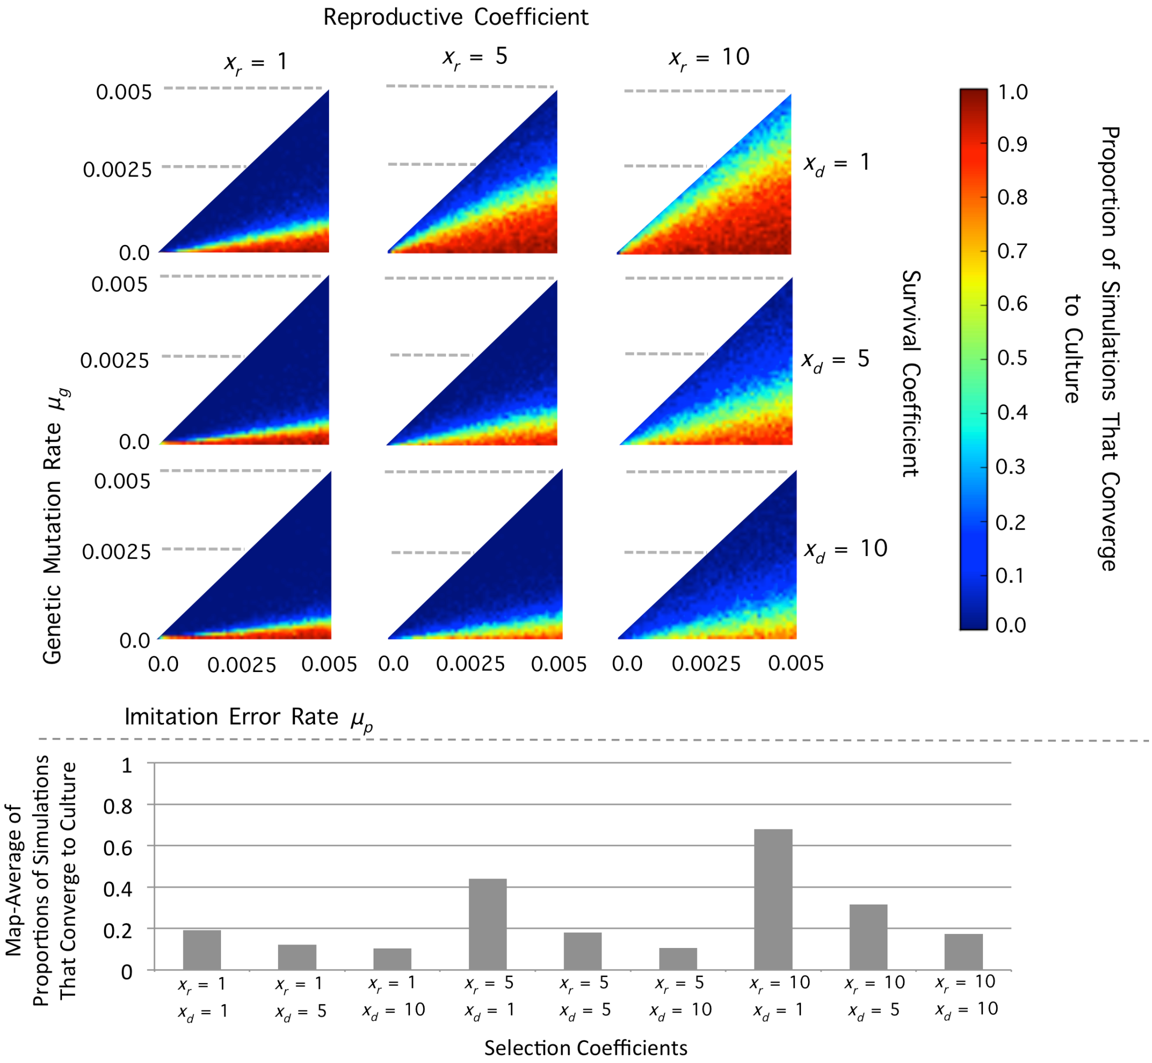

FIGURE 5.5: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right, columns of maps have x_r values of 1, 5 and 10. From top to bottom, rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs, the rate of environmental change is set to $\delta_e = 0.1$ (i.e., $\mu_e = 5 \times 10^{-4}$). Compared to Figure 5.3, this heat map array has a larger area of the parameter space converging towards culture. However, compared to Figure 5.4, some of the areas of high convergence have decreased their proportion of simulations converging to culture (e.g., red colours are less dark).

For all combinations of survival (x_d) and reproductive selection (x_r), the effect of moderate environmental change on the evolution of social learning is positive (i.e., slight increases of environmental change lead to more simulations in which culture emerges).

However, for combinations where survival selection is stronger than reproductive selection, this effect is most apparent (see Figures 5.3, 5.4 and 5.5). For comparison purposes, let us then focus on the heat maps where $x_d = 1$ and $x_r = 10$ (i.e., where survival selection is stronger than reproductive selection) for the rest of our discussions in this subsection.

In Chapter 3, we concluded that strong survival selection (compared to reproductive selection) was necessary for the evolution of social learning, as only this type of selection affects the non-genetic evolutionary system. For scenarios with both kinds of selection (i.e., survival and reproduction), the genetic evolutionary system tends to find solutions faster, as it is pressured by both types of selection. We also concluded that a relatively high imitation error rate compared to genetic mutation favours the emergence of culture; with a higher error rate, social learners are able to find fitter variants faster and, in consequence, be more likely to produce genotype-phenotype disengagement (i.e., irreversibly fix in the population).

In Figure 5.3 (equal to Figure 3.9), we note the existence of a threshold μ_g value, above which social learners do not fix. In these heat maps, parameterisations that converge to culture are restricted to horizontal bands at the bottom of each map (i.e., where genetic mutation is low).

As previously explained, for high values of genetic mutation, instinctive actuators will evolve their phenotypes beyond a critical number of environment-matching bits (i.e., α) before the cultural system can emerge. In the context of this chapter, α is the number of environment-matching phenotype bits for which a social learner (whose phenotype is only affected by survival selection and gets imitated with error rate μ_p), living in a population of instinctive actuators, produces offspring with the same average fitness as that of an instinctive actuator (whose phenotype is affected by both types of selection and gets replicated with mutation rate μ_g). When the optimisation process has reached α, the low mutation rate of genetic reproduction is advantageous compared to the high imitation error rate (which produces deleterious mutations, i.e., environmentally mismatched bits, more frequently than beneficial ones). Social learners are therefore not expected to fix for simulations that genetically evolve their phenotypes beyond the α value.

To complete this explanation, we must remember that simulations start with a monomorphic population of instinctive actuators, and therefore the critical mass of social learners must appear by mutation for culture to emerge; this process takes a certain number of iterations. In the meantime, genetic evolution may be able to optimise the phenotype of the population beyond α before social learners reach the critical frequency to cause genotype-phenotype disengagement (regardless of the μ_p value). For high mutation rates (i.e., high values of μ_g), this is more likely to occur; such a process explains the discrepancies between our estimated boundary (i.e., the heat map line below which we should expect social learners to fix) and the experimental results (see heat map A in Figure 5.6; also, see Section 3.3 for calculations of the boundary line).

In Figure 5.6, we present four heat maps that differ only in the rate of environmental change. The maps show results for simulations with δ_e values of 0.0 (i.e., no environmental change), 0.04, 0.1 and 0.4. The map with no environmental change presents a large area with a low proportion of simulations converging to culture, under the boundary line. However, when a moderate rate of environmental change is included, this area shows a larger proportion of simulations converging to culture.

In simulations with a moderate rate of environmental change, the optimal sequence is constantly altered and therefore instinctive actuators take a larger number of iterations to reach the aforementioned α value. When instinctive actuators take a longer time to reach the α value, social learners have a larger chance of producing genotype-phenotype disengagement, as they are more likely to increase in frequency and start forming chains of imitation (i.e., imitators imitating other imitators through several generations) before it is too late (i.e., before phenotypes are evolved genetically beyond the α value).

Regardless of the rate of environmental change however, the evolution of social learning is still bounded to parameterisations where the non-genetic evolutionary system (i.e., a system with high imitation error rates, horizontal transmission and exclusively under survival selection) can produce, on average, better behavioural variants than a genetic system (i.e., a system with low mutation rates, vertical transmission and under the combined effect of reproductive and survival selection) (see Section 3.3).

As our simulation is initiated with an average 50% match between phenotypes and the optimal sequence (i.e., environment); only parameterisations where $L/2 < \alpha < L$ are expected to converge to culture. The white line in all maps in Figure 5.6 represents the boundary for this condition. All points in this line have $\alpha = L/2$; simulations above the line ($\alpha < L/2$) are not expected to converge to culture (i.e., all proportions of simulations converging to culture should be equal or lower than 0.5 above the line, as the likelihood of these simulations actively evolving and fixating social learners tends to 0 and only stochastic processes apply). The noticeable exception being areas where μ_g and μ_p are low and similar in maps with low rates of environmental change (i.e., lower left corners of maps A and B in Figure 5.6). In these areas, social learners fix through the drifting mechanism explained in the results subsection titled “Social learners can also fix due to drift and the irreversibility of genotype-phenotype disengagement” (Subsection 3.2.3). When environmental change is relatively high (i.e., in maps C and D of Figure 5.6), this mechanism no longer fixes social learners, as the rate of optimisation for both evolutionary systems (i.e., cultural and genetic) is too slow to find solutions compared with the rate at which environmental change renders those solutions obsolete. In such a scenario, genotype-phenotype disengagement does not occur, as neither string of bits (i.e., phenotype or genotype) can evolve better solutions than the neutral average where 50% of their bits match the current environmental optimum.

The boundary line in Figure 5.6 is produced by numerically estimating the μ_p and μ_g values for which $\alpha = L/2 = 100$. A detailed explanation of this estimation can be found in Section 3.3.

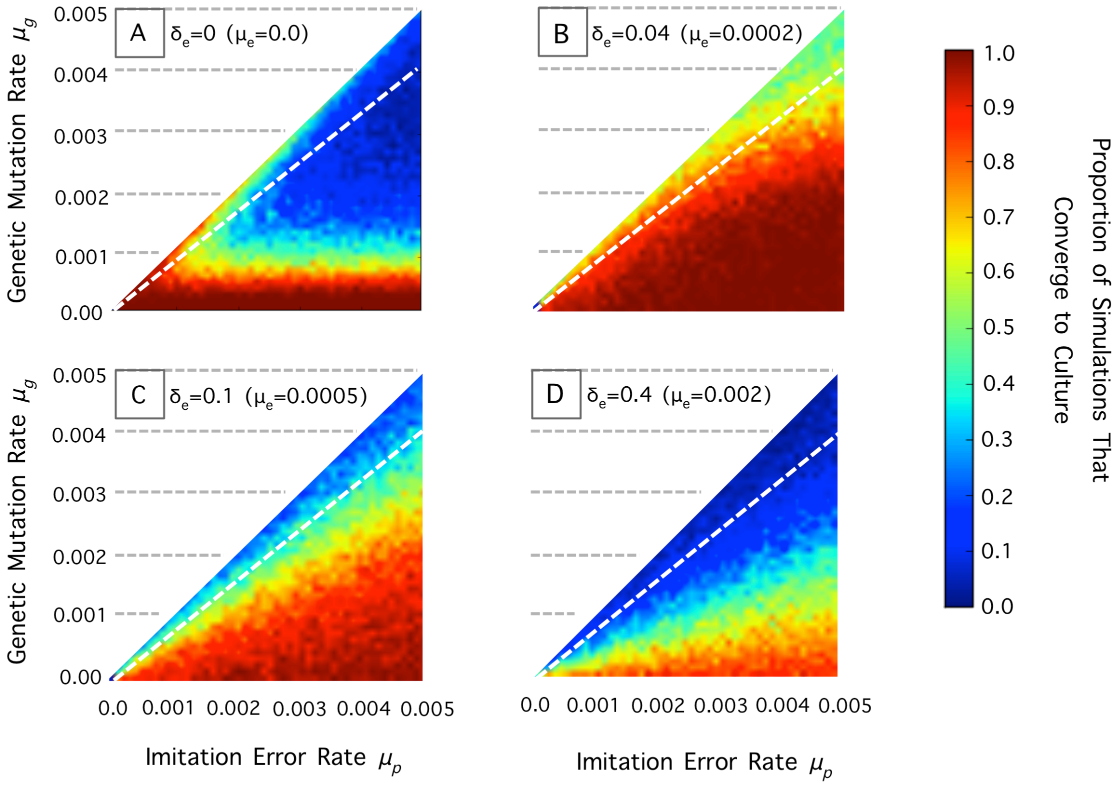

FIGURE 5.6: Four heat maps with different rates of environmental change (i.e., δ_e). Within each map, red areas represent parameter sets where a high proportion of replicates evolve culture through genotype-phenotype disengagement. For moderate rates of environmental change, large areas of the heat map converge to culture. The white dotted line in the maps represents the estimated boundary line (α) above which we do not expect to find a high proportion of simulations converging to culture, regardless of the rate of environmental change. The noticeable exception being areas where μ_g and μ_p are low and similar in maps with low rates of environmental change. Parameterisations at any point in this line have an estimated α value of $L/2$, which means that even at the start of the run, the imitation error rate is not advantageous to social learners. All points above the line have estimated α values lower than $L/2$. All maps show the same parameter region, where mutation rates are between 0.0 and 0.005, and $\mu_p > \mu_g$. For all maps, the strength of survival and reproductive selection is fixed at $x_d = 1$ and $x_r = 10$.

For extreme rates of environmental change, neither evolutionary system (i.e., cultural or genetic) is able to keep track of the changing optimal sequence (i.e., optimisation is hampered in simulations with extreme environmental change). In this scenario, genotype-phenotype disengagement is less likely to occur, as social learners cannot evolve phenotypes to the point where a reinvasion of instinctive actuators would be impossible (i.e., social learners cannot optimise phenotypes towards fitness values higher than those of randomly drifting genotypes). For this reason, the area under the dotted line in Figure 5.6 has a lower proportion of simulations converging to culture for the map where $\delta_e = 0.4$ compared to the map where $\delta_e = 0.1$. However, as we will explain in the next subsection, parameterisations with low genetic mutation rates can still be invaded by social learners, albeit temporarily, even under high rates of environmental change.

5.2.2 In a Changing Environment Social Learners Have a Broader Time-Window to Produce Genotype-Phenotype Disengagement

As explained in Chapter 3 and recapitulated in the previous subsection, it is crucial for parameterisations where social learners have the potential to fix, that they do so before the α value is reached through genetic evolution. Adding environmental change delays this process and therefore increases the chance for social learners to achieve genotype-phenotype disengagement. To further explain and test this idea, we will analyse: (1) heat map results where the population of instinctive actuators is initialised above the α value, and (2) individual simulation runs in a point of the parameter space where the proportion of simulations converging to culture is strongly influenced by the rate of environmental change.

In Figure 5.7, we show that starting simulations with an average sum of environmentally matching bits (L_m) higher than α does not affect the distribution of areas in the parameter space that converge to culture, when the population is subjected to moderate environmental change. Maps A and B show results for standard simulation runs where phenotypic bits are initialised at random, with a 0.5 chance of matching their corresponding bit in the optimal sequence (i.e., the sequence that represents the initial environmental state). In map A, there is no environmental change. In map B, there is a moderate amount of environmental change (i.e., $\delta_e = 0.04$). Map A clearly shows a threshold value of μ_g above which simulations do not converge towards culture, even for large values of μ_p; in map B, this threshold does not restrict the area where simulations converge to culture. In simulations for map B, the low mutation rate of instinctive actuators cannot evolve phenotypes beyond a certain number of environmentally matching bits, as the rate of environmental change tends to depress their fitness values to the neutral average (i.e., 50% of matching bits). In contrast, the relatively high rate of imitation error can effectively evolve those phenotypes above this limit, (i.e., the limit imposed by an equal-pace state between the rate of genetic evolution and the rate of environmental change). For this reason, the area of the parameter space that converges to culture is drastically extended in map B compared to map A.

In contrast to maps A and B, maps C and D show simulation results for runs where all phenotypes are initialised with a proportion of environmentally matching bits equal to 0.875 (i.e., $L_m = 175$ for $L = 200$). This proportion is well above the estimated α value for all points in the parameter space. In map C, we observe the expected outcome: a reduced area where simulations converge to culture in the lower region of the map (i.e., area indicated by a black arrow). In map D, however, this change has no significant effect on the areas that converge to culture. This lack of effect proves that: (1) a constant change of the optimal environmental sequence rapidly erases the initial advantage of genetic evolution created by initialising phenotypes with a high proportion of environmentally matching bits, and (2) the sum of environmentally matching bits that the genetic evolutionary system can maintain in an equal-pace state with the rate of environmental change is lower than the sum achieved by the non-genetic evolutionary system (i.e., the system maintained by high error imitation). More on this in the next subsection.

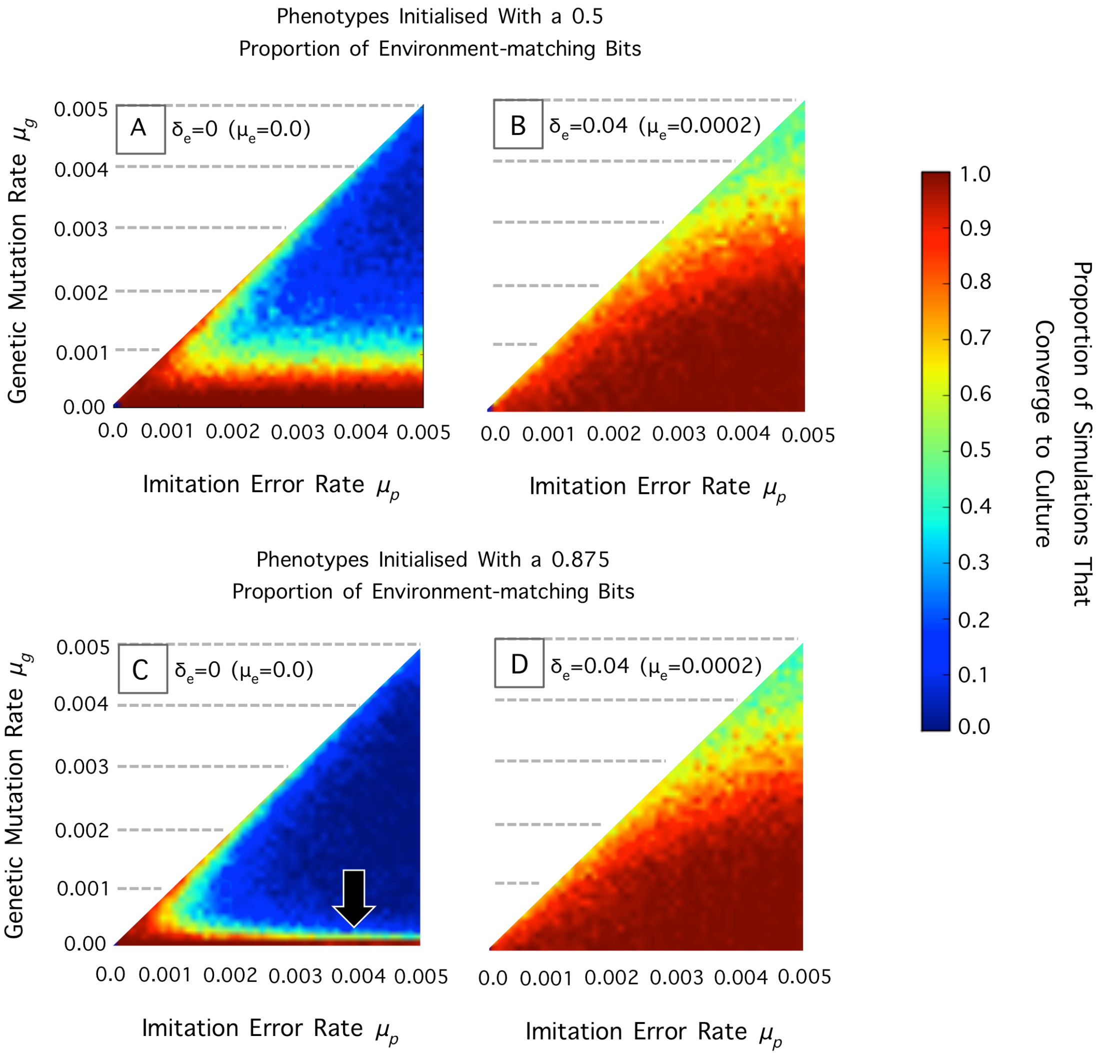

Figure 5.7: Comparative heat maps showing the effect of initialising phenotypes with a high proportion of environment-matching bits for cases with and without environmental change. Maps A and B show the proportion of simulations converging to culture for runs where phenotypes are initialised with an average of 50% of their bits matching the initial optimal sequence. In map A, there is no environmental change. In map B, the rate of environmental change (i.e., the rate at which the optimal sequence is changed) is set to $\delta_e = 0.04$. Maps C and D show the proportion of simulations converging to culture for runs where all phenotypes are initialised with exactly 87.5% of their bits matching the initial optimal sequence (i.e., $L_m = 175$ for $L = 200$). In map C, there is no environmental change. In map D, the rate of environmental change is set to $\delta_e = 0.04$. Initialising phenotypes with a high proportion of environment-matching bits drastically decreases the area where simulations converge to culture for cases with no environmental change (black arrow in map C). However, once a moderate rate of environmental change is introduced, this effect disappears (map D). For all maps, the strength of survival and reproductive selection is fixed at $x_d = 1$ and $x_r = 10$.

To further understand how environmental change affects our simulation model, we focus our attention on the dynamics of single runs located in areas of the parameter space where environmental change has a drastic effect on the convergence to culture. In Figure 5.8, we show three different heat map results for simulations with no environmental change (map A), moderate environmental change (map B) and extreme environmental change (map C). For each of these simulations, we analyse a typical run in the area that is most affected by environmental change (indicated by black arrows). This particular point was selected arbitrarily; similar conclusions can be derived from any other point in the surrounding area.

Plots D, E and F show representative simulation runs in areas indicated by black arrows in maps A, B and C, respectively. In plot D, phenotypes evolve genetically to the optimum; the relatively high mutation error (i.e., compared to lower regions in the plot) optimises the sum of environmentally matching bits before social learners can originate a non-genetic evolutionary system. When the environment does not change, this process occurs rapidly, as the evolved phenotypes are not rendered obsolete by changes in the optimal target sequence. In contrast, plot E shows how genetic evolution cannot evolve phenotypes beyond a certain fitness value ($L_m \simeq 120$) when the optimal sequence is changed with a moderate rate. Under this condition, social learners have an extended window of time to emerge and eventually fix in the population. The higher rate of imitation (i.e., compared to genetic mutation) allows the cultural evolutionary system (i.e., the system maintained by social learners) to evolve phenotypes beyond the restricted value maintained by the genetic system during the first 15×10^3 iterations. Once social learners invade and genotype phenotype disengagement occurs, the population evolves culturally until reaching a new equal-pace state.

When the rate of environmental change (i.e., the rate of change for the optimal sequence) is increased even further, neither the genetic nor the cultural system can keep pace. Under this condition, genotype-phenotype disengagement cannot occur, as potentially disengaged genotypes will still have average values similar to the phenotypes of social learners. This scenario is clearly shown in plot F of Figure 5.8.

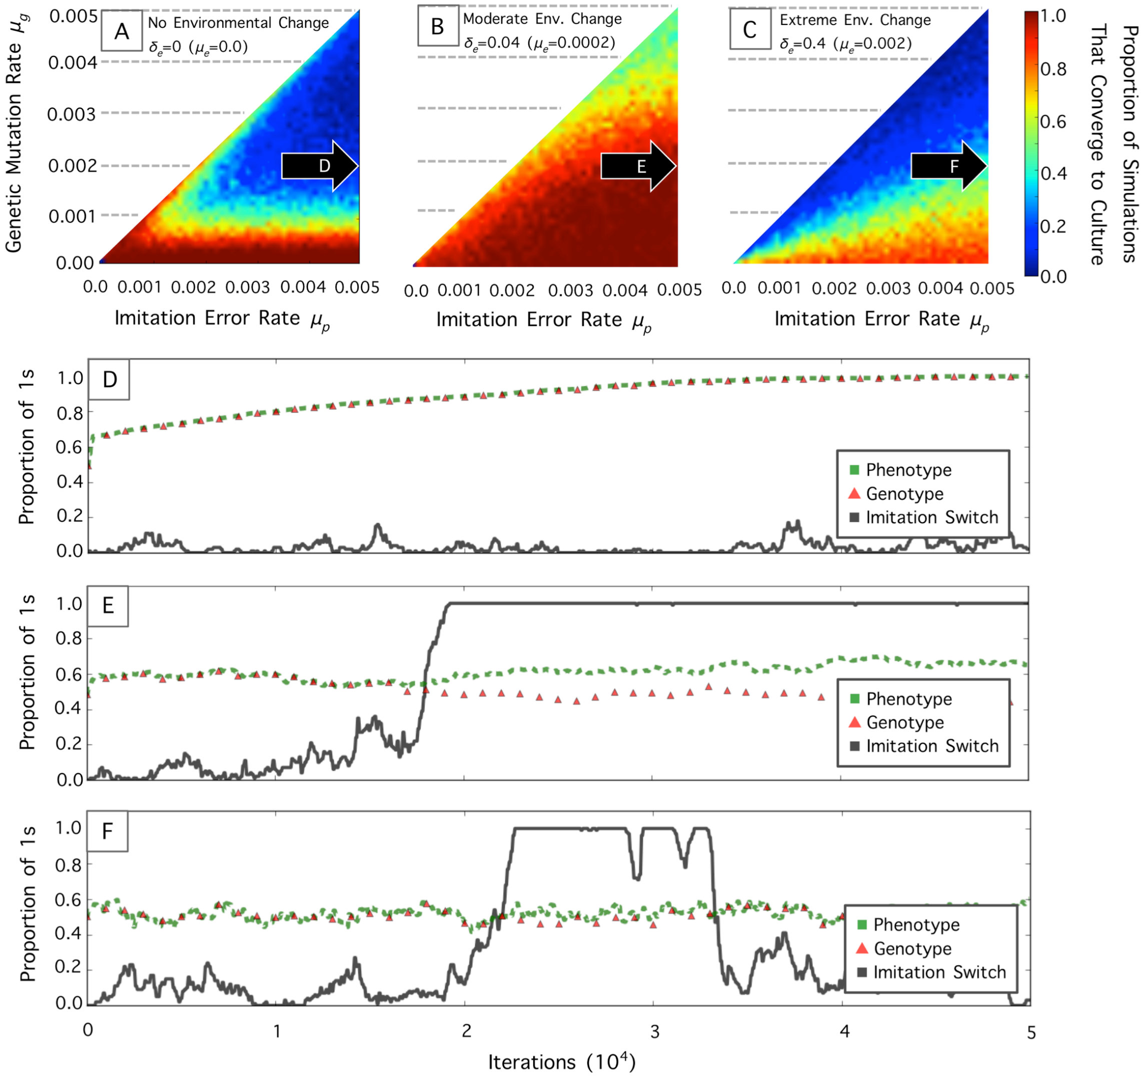

FIGURE 5.8: Comparative heat maps showing the effect of different rates of environmental change on the proportion of simulations that converge to culture (maps A, B and C), and comparative results for representative simulation runs for regions of the parameter space where the rate of environmental change has a drastic effect on the proportion of simulations that converge to culture (plots D, E and F). In Map A, when the rate of environmental change is equal to zero, areas with a high proportion of simulations converging to culture are restricted to a horizontal band at the bottom of the map (i.e., where mutation is low). In map B (where the rate of environmental change is moderate), this area increases in size; for higher rates of imitation error, simulations with higher rates of genetic mutation still converge to culture. When the rate of environmental change further increases, the areas of high cultural convergence are reduced (map C). Parameterisations that do not converge to culture in map A, but do in map B, genetically evolve their phenotypes beyond α before social learners can instantiate a cultural system, under no environmental change (plot D). However, under moderate environmental change, the genetic system cannot optimise phenotypes beyond α and therefore social learners are more likely to produce genotype-phenotype disengagement and fix in the population (plot E). For extreme rates of environmental change, neither evolutionary system (i.e., cultural or genetic) can optimise phenotypes far from the neutral average proportion of 0.5. Therefore, irreversible genotype-phenotype disengagement cannot occur. Plots D, E and F are representative runs for $\mu_g = 0.002$ and $\mu_p = 0.005$ in respective maps A, B and C (positions are indicated by black arrows). For all maps and plots shown, the strength of survival and reproductive selection is fixed at $x_d = 1$ and $x_r = 10$.

5.2.3 Unbiased Social Learning Can Increase the Overall Mean Fitness of a Population Under Strong Survival Selection and Moderate Environmental Change

Taking into consideration previous results, it's now easy to conceive that for some areas of the parameter space, the relatively high error rate of social learners will reach an equal-pace state with the rate of environmental change where phenotypes have a higher fitness value than they would in a genetic evolutionary system. To confirm and further support this claim, we show two complementary figures in this subsection, Figures 5.9 and 5.10. In Figure 5.9, we compare the results between simulations where social learning is evolvable (left-side column) and simulations where social learning is not evolvable (right-side column). In the later case, we simply set the mutation rate of the imitation switch to zero $\mu_c = 0$.

In each map in Figure 5.9, we show the proportion of environment-matching bits for all phenotypes of the population after 5×10^4 iterations averaged over one hundred replicates; for all combinations of μ_g and μ_p in the set $[1 \times 10^{-4}, 2 \times 10^{-4}, \ldots, 50 \times 10^{-4}]$, where $\mu_p > \mu_g$. For the first row, simulations are run with no environmental change. Under this condition phenotypes can reach values close to the optimum for all combinations of mutation and imitation error rates except for the trivial cases where $\mu_g = 0$ and $\mu_p = 0$. [1]

For the second row of maps, simulations are under a moderate rate of environmental change ($\delta_e = 0.04$). In the left map (i.e., results for simulations where social learning is evolvable), there is a clear horizontal gradient in which parameterisations that have a high imitation error rate also have a high number of environment-matching bits in their phenotypes. In contrast, the right map (i.e., results for simulations where social learning is not evolvable) has a vertical gradient where higher rates of genetic mutation yield a higher proportion of phenotypic environment-matching bits. These results are easily explained by the increased capacity of evolutionary systems with higher imitation error (or mutation) rates to track environmental changes; with higher mutation and imitation error rates comes a higher proportion of environment-matching bits when reaching an

---

[1] Notice that the chromatic scale used to represent different proportions of environment-matching phenotype bits is constrained between 0.75 and 0.5, with values above and below this range being represented by the respective boundary colours. We selected this constrained scale to better display contrasting regions in the second and third row maps. However, we think it is important to point out that using the chromatic scale for the full range of proportions (i.e., from 0.0 to 1.0) does not affect the qualitative results of maps in the first row (i.e., the differences between average phenotype values within the red coloured regions in maps of the first row are too small to create a perceivable chromatic difference (also see map A of Figure 5.10)). Moreover, when these maps are plotted with a constrained chromatic range between 0.98 and 1.0 only small differences are observed; with regions that converge to culture having only a slightly inferior average proportion of environment-matching bits in their phenotypes. This pattern is easily explained by the relatively large imitation error rate, compared to genetic mutation. This large error rate diminishes the fidelity at which the optimal sequence is maintained in the population.

equal-pace state with the rate of environmental change (i.e., when the established evolutionary system optimises phenotypes at the same rate as environmental change turns them obsolete).

In the previous subsections, we have shown that simulations in which social learning is evolvable can reach an equal-pace state by either the genetic or the cultural evolutionary system. For parameterisations where this state is reached by the genetic system, we don't expect a difference in the proportion of environment-matching phenotype bits between the two maps (i.e., left and right maps in Figure 5.9). For parameterisations where this equal-pace state is reached by the cultural system, we expect some regions to optimise their phenotypes above the fitness value that a genetic evolutionary system could achieve.

In the third row of maps in Figure 5.9, we show results for simulations under extreme environmental change (i.e., where $\delta_e = 0.4$). As mentioned in the previous subsections, neither the genetic nor the cultural evolutionary system can optimise phenotypes under this rate of environmental change. Therefore, as expected, the proportion of environment-matching bits in both maps remains quite similar and very close to the neutral average of 0.5. However, there are still barely noticeable gradient patterns in both plots, similar to those observed where the rate of environmental change is moderate. This gradient is also produced by the ability of evolutionary systems with higher rates of imitation error and genetic mutation to better track environmental change. In contrast to simulations in the maps of the second row, social learners do not irreversibly fix under extreme environmental change and therefore the measured proportion of environment-matching bits after 5×10^4 iterations comes from a fluctuating mixed population of social learners and instinctive actuators (see plot F of Figure 5.8).

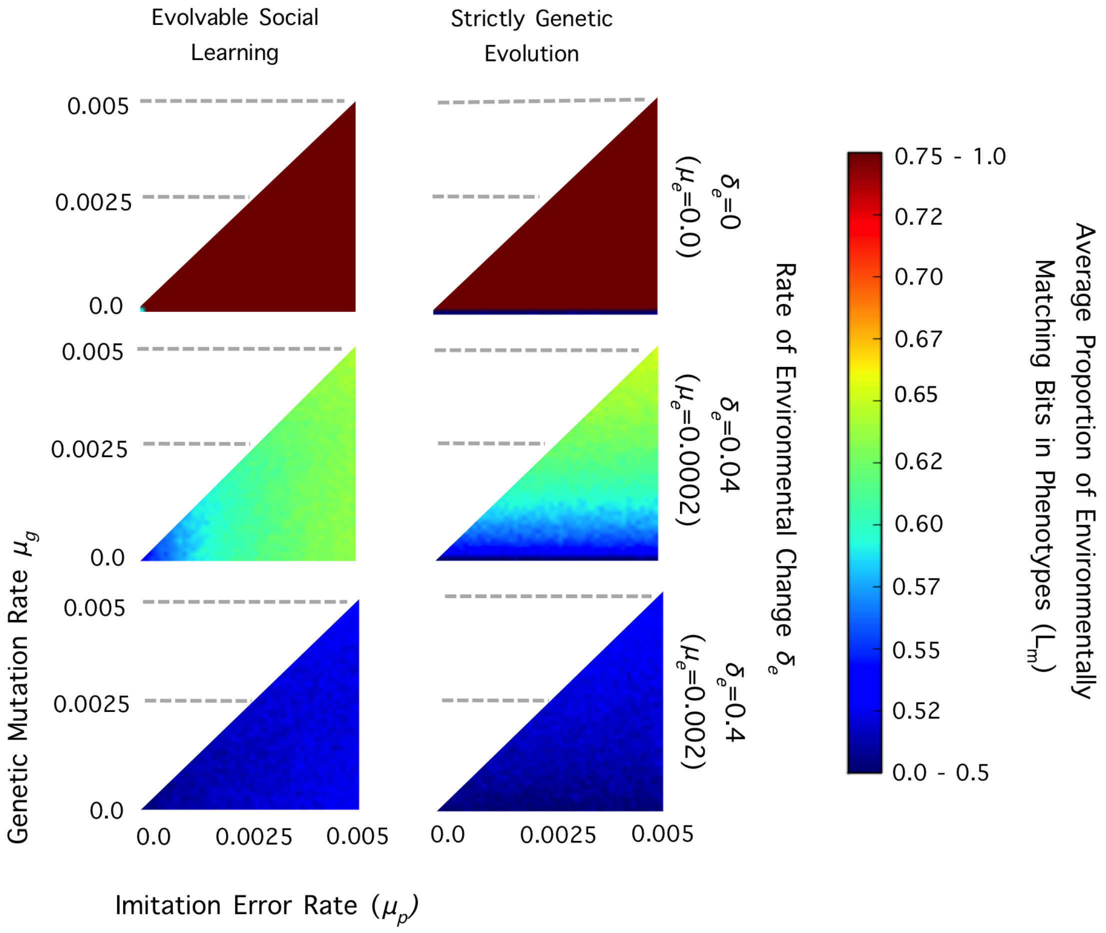

FIGURE 5.9: Comparative heat maps showing the average proportion of environmentally matching bits in phenotypes for simulations where social learning is evolvable (left column) and simulations where social learning is not evolvable (right column). Within each map, colours closer to red indicate combinations of μ_g and μ_p where the average phenotypes after 5×10^4 iterations (for 100 replicates) are evolved to a higher number of environment-matching bits. The pair of maps in the first row shows results for simulations with no environmental change (i.e., where the optimal sequence is unchanged through simulation runs). The second row shows results for simulations under moderate environmental change (i.e., where, on average, 0.04 bits of the optimal sequence are changed every iteration). The last row shows results for simulations under extreme environmental change (i.e., where, on average, 0.4 bits of the optimal sequence are changed every iteration). Under moderate environmental change, a distinctive horizontal gradient appears for the map on the left where social learning is evolvable, and a vertical gradient for the map on the right where social learning is not evolvable. These gradients are explained by the ability of evolutionary systems (either genetic or cultural) to better track environmental changes when their mutation/imitation error rates are relatively high. For all maps, the strength of survival and reproductive selection is fixed at $x_d = 1$ and $x_r = 10$.

Figure 5.10 is complementary to Figure 5.9 and shows the areas of the parameter space where the evolution of social learning increases the overall mean fitness of the population (i.e., where the average proportion of environment-matching bits its greater for simulations in which social learning is evolvable than for simulations where it is not). In maps A, B and C, we show three levels of environmental change that respectively correspond to each row in Figure 5.9 (i.e., where the rate of environmental change is: none, moderate and extreme). Red coloured areas indicate combinations of μ_g and μ_p for which the evolutionary option of social learning increases the proportion of phenotypic bits that match the current environment. Blue regions correspond to areas where the evolutionary option of social learning decreases this proportion.

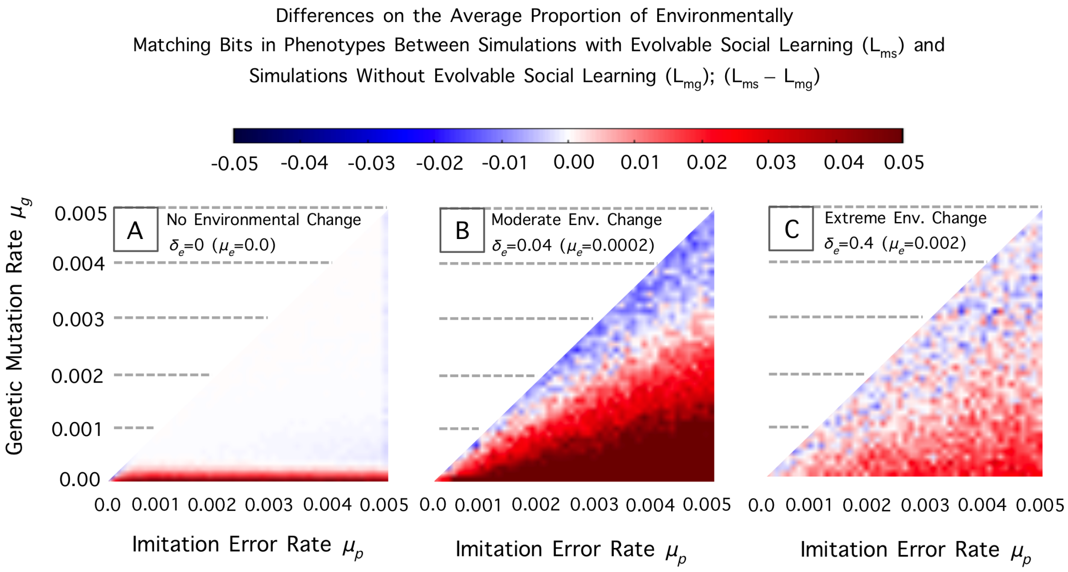

Figure 5.10: Three maps showing the differences in the average proportion of environment-matching bits (L_m) after 5×10^4 iterations, for 100 replicates, between simulations where social learning is evolvable (L_{ms}) and simulations where social learning is not evolvable (L_{mg}). In map A, the environment does not change throughout the simulation. Simulations of map B are under a moderate rate of environmental change (i.e., on average, 0.04 bits are flipped in the optimal sequence every iteration). Simulations of map C are under an extreme rate of environmental change (i.e., on average 0.4 bits in the optimal sequence are flipped every iteration). In map A, no significant differences are appreciated except for the trivial cases where $\mu_g = 0$. Map B clearly shows how simulations with large values of $\mu_p = 0$ are more likely to evolve social learning systems that are adaptive (i.e., social learning systems that increase the overall mean fitness of the population) when under a moderate rate of environmental change. Map C shows a pattern similar to that of map B. However, in map C, this pattern is less pronounced, as the adaptive effect of social learning is not as marked when the high rate of environmental change maintains phenotypes relatively close to the neutral average (i.e., where only 50% of bits in the phenotypes match the current environmental state). For all maps, the strength of survival and reproductive selection is fixed at $x_d = 1$ and $x_r = 10$.

Without environmental change, the differences between the proportion of environment-matching bits in phenotypes for simulations where social learning is evolvable and simulations where it is not evolvable are almost non-existent (i.e., most of map A in Figure 5.10 is coloured white). The only noticeable exceptions being the trivial cases where $\mu_g = 0$. In these exceptional cases, genetic evolution cannot explore the strategy space; hence only the cultural evolutionary system is able to optimise phenotypes. It therefore follows that phenotypes of simulations where social learning is evolvable will outperform those of simulations where it is not.

For moderate rates of environmental change, there is a large set of combinations of μ_g and μ_p for which the evolutionary option of social learning increases the overall mean fitness of the population (notice the red triangular area at the bottom of map B in Figure 5.10). These parameterisations overlap nicely with regions where a large proportion of simulation replicates converge to culture (see comparative maps A and B in Figure 5.11). In those scenarios, the rate of optimisation for the genetic evolutionary system keeps pace with the rate of environmental change at a phenotypic value (i.e., number of environment-matching bits) lower than that for which a population of social learners would keep pace with the same rate of environmental change. Therefore, moderate environmental change not only increases the chances for social learners to fix (i.e., by preventing genetic optimisation before phenotypes reach the α value), but also makes the fixation of social learners adaptive (i.e., makes a population of social learners more fit than a monomorphic population of instinctive actuators) (see plot E of Figure 5.11).

For some simulations that converge to culture under moderate environmental change, the overall mean fitness of the population is not larger than that achieved by a population of instinctive actuators (see white and blue areas in Figure 5.11). To explain these cases, let us recapitulate the two mechanisms that lead to the fixation of social learners: (1) social learners consistently find fitter solutions than instinctive actuators during early stages of optimisation, and (2) the frequency of social learners "drifts" to fixation in simulations where mutation and imitation error rates are low and similar to one another (see Subsection 3.2.2 and Subsection 3.2.3 for a detailed explanation). In the previous subsections of this chapter, we explained that mechanism 1 is more likely to occur under moderate environmental change because the rate at which phenotypes improve during the early stages of optimisation (i.e., where evolution is driven by instinctive actuators) is either delayed (i.e., it takes more iterations to reach the critical α value) or suppressed (i.e., the equal-pace state between the rate of environmental change and the rate of genetic evolution maintains phenotypes under the critical α value). For cases where genetic evolution reaches an equal-pace state below the α value, the cultural evolutionary system is bounded to produce average phenotypes that are fitter than those produced by a population of instinctive actuators (i.e., a genetic evolutionary system). In contrast, if social learners evolve only because of a momentary advantage in their optimisation rate (i.e., compared to genetic evolution), a monomorphic population of instinctive actuators (e.g., like the one in simulations where social learning is not evolvable) would eventually evolve phenotypes to equal (or higher) fitness values than those produced by the cultural evolutionary system.

Areas of the parameter space where social learning evolves only due to a momentary advantage in the early stages of optimisation (i.e., rather than a constantly present advantage due to genetically evolved phenotypes being maintained below the α value by the rate of environmental change) are located in the upper region of the diagonal that separates areas with high convergence rates from areas with low convergence rates (i.e., region C, indicated by black arrows in maps A and B in Figure 5.11). In these simulations, the rate of genetic mutation is high enough to maintain a large number of phenotypic environment-matching bits when keeping pace with the rate of environmental change. However, the even larger rate of imitation error is relatively advantageous during the early stages of optimisation (i.e., compared to genetic mutation) and therefore allows social learners to invade and produce genotype-phenotype disengagement. Once social learners irreversibly fix, the optimisation of the phenotypes continues until reaching equal pace with the rate of environmental change. For these regions in particular, the proportion of phenotypic environment-matching bits reached at this higher-than-α equal-pace state by a population of social learners is equal to or lower than the proportion that would be achieved by a monomorphic population of instinctive actuators (i.e., where social learning is not evolvable).

Cases in which mechanism 2 is dominant (i.e., where the frequency of social learners "drifts" to fixation) occur in regions of the parameter space where μ_g and μ_p are small, and $\mu_g \simeq \mu_p$ (i.e., region D, indicated by black arrows in maps A and B in Figure 5.11). In these regions, both cultural and genetic systems are competitively similar when finding new variants in the optimisation process. However, as explained in Chapter 3, the combination of a slow rate of innovation (i.e., a slow rate at which new fitter variants are introduced) and the similarity between mutation and imitation error rates increases the probability that social learners could perform the fittest (current) variant in the population for an extended number of continuous iterations, even when social learning is not statistically expected to be the most advantageous strategy in the long term. When this occurs, an otherwise partial invasion of social learners leads to genotype-phenotype disengagement, which prevents the reinvasion of instinctive actuators and leads to the irreversible instalment of a cultural evolutionary system (see Chapter 3). In consequence, simulations in this area are likely to evolve culture even when the rate of genetic mutation is slightly more advantageous than the imitation error rate in the long term (i.e., even when the genetic system would have reached a higher, or similar, proportion of environment-matching bits when keeping pace with the rate of environmental change).

When the rate of environmental change is too high for either evolutionary system to optimise phenotypes, genotype-phenotype disengagement does not occur. In these cases, the frequency of social learners fluctuates throughout the simulation run, while the average values of phenotypes and genotypes remain close to the neutral average (i.e., where only 50% of the bits in the strings are matching the current environmental state). This process produces the results observed in map C of Figure 5.10, where for all combinations of μ_g and μ_p, the differences between the average number of environment-matching bits in the phenotypes of simulations where social learning is evolvable and the phenotypes of simulations where social learning is not evolvable is minimal (i.e., notice the less pronounced pattern in map C compared to map B in the same figure). However, in the less pronounced pattern in map C of Figure 5.10, we can still see a similar colour distribution as in map B. Such pattern is explained by the discussed relative differences between the ability of mutation rates and imitation error rates to track environmental changes.

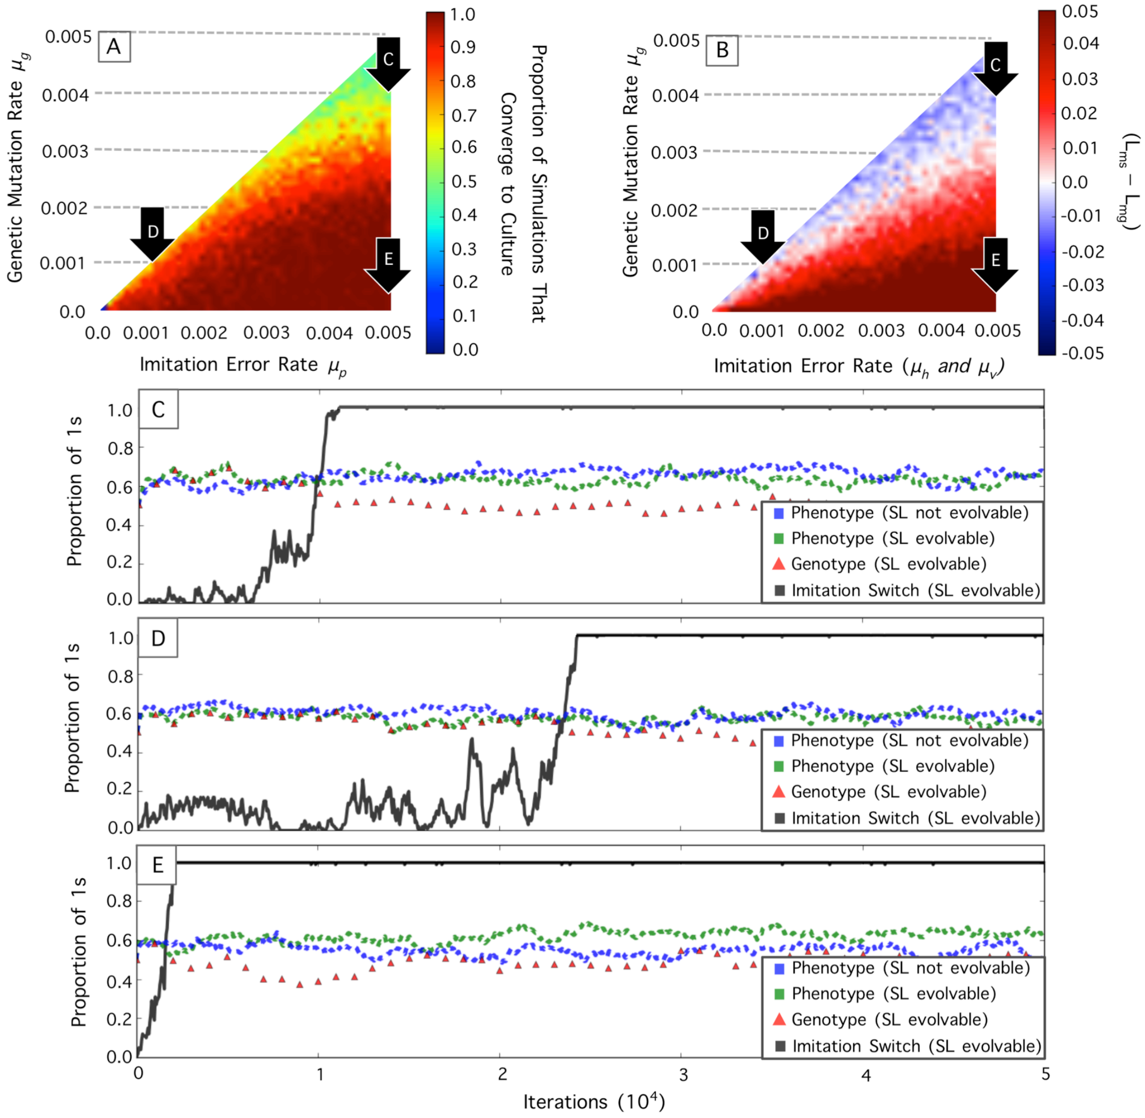

FIGURE 5.11: Comparison between three distinctive regions of the parameter space for parameterisations under moderate environmental change (i.e., $\delta_e = 0.04$, equivalent to $\mu_e = 2 \times 10^{-4}$). Map A shows the proportion of simulations, out of 100 replicates, that converge to culture after 5×10^4 iterations. Map B shows the difference in the proportion of environment-matching bits between simulations where social learning is evolvable and simulations where it is not. Combinations of μ_g and μ_p for which the evolution of social learning is adaptive are coloured in red on map B. Black arrows in maps A and B correspond to parameterisation points where the evolution of social learning has different effects on the long-term average fitness value of phenotypes. Plots C, D and E correspond to representative runs for the areas indicated by marked arrows. In these plots, the blue line shows the proportion of environment-matching bits in the phenotypes of a population where social learning cannot evolve; the green line and red triangles correspond to the proportion of environment-matching bits on the phenotype and genotype strings for a run where social learning is evolvable. Finally, the black line corresponds to the proportion of social learners. In plot C (where $\mu_g = 0.04$ and $\mu_p = 0.05$), social learning evolves due to a transient advantage during the early stages of optimisation; for this simulation, the emergence of culture is not adaptive. In plot D (where $\mu_g = 0.01$ and $\mu_p = 0.01$), social learning evolves due to a "drifting" mechanism; for this area, the emergence of culture is also not adaptive. In plot E (where $\mu_g = 0.005$ and $\mu_p = 0.05$), social learning evolves due to a constant advantage over the genetic evolutionary system; for this area of the parameter space, the evolution of social learning is adaptive, as it improves the proportion of phenotypic environment-matching bits compared to a population without social learning. For all maps and plots in this figure, $x_d = 1$ and $x_r = 10$.

5.3 Discussion

The model extension introduced in this chapter supports the claim that unbiased social learning can be adaptive without the inclusion of decision-making attributes or individual learners. In this brief discussion, section we will focus on the main difference of our approach compared to previous models, and further describe a simple alternative scenario where unbiased social learning is adaptive compared to a monomorphic population of instinctive actuators. A more general discussion regarding the results of this chapter, Chapter 3 and Chapter 4 can be found in the next chapter.

5.3.1 Comparing Evolutionary Systems to Solve Rogers' Paradox

In this chapter, we have found conditions under which social learning is adaptive without the presence of individual learners or decision-making attributes like biased imitation or guided variation; specifically: when survival selection is stronger than reproductive selection, the imitation error rate is higher than genetic mutation, and the rate of environmental change is moderate. However, to frame our claims within the context of the existing literature, we must now recapitulate and discuss our results in contrast to existing analytical models.

In this thesis, we have clearly taken a simulation modelling approach, rather than the more common analytical approach; the differences between these can make it difficult to compare our work directly with existing seminal (Rogers, 1988; Cavalli-Sforza and Feldman, 1983) and derivative models (Henrich and Gil-White, 2001; Castro and Toro, 2004; Enquist et al., 2008; McElreath and Henrich, 2007). Nevertheless, we believe there is a contrasting consideration that has been overlooked in existing literature and that is key to our approach. In this thesis, we broach the question: "How can unbiased social learning increase the overall mean fitness of the population?" (i.e., Rogers' Paradox) by focusing on the properties of the cultural evolutionary system maintained by social learners, rather than on the individual traits that social learners have (or should have) in order to outcompete other strategies.

In contrast to Rogers (1988), Cavalli-Sforza and Feldman (1983) and all the subsequent theoretical models in the field (Henrich and Gil-White, 2001; Castro and Toro, 2004; Enquist et al., 2008; McElreath and Henrich, 2007; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007), our focus is oriented at explaining how the most simple form of a cultural evolutionary system could outcompete its genetic counterpart, with the invasion and eventual fixation of social learners being just a consequence of that process. In our simulation, the extended strategy space is a distinctive element that allows both evolutionary systems (i.e., genetic and cultural) to compete against each other in terms of finding fitter solutions to a given environmental challenge. This competition is fundamentally different from the individual level competition (between instinctive actuators and social learners) that is investigated in existing literature (Henrich and Gil-White, 2001; Castro and Toro, 2004; Enquist et al., 2008; McElreath and Henrich, 2007; Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Aoki et al., 2005; Borenstein et al., 2007).

5.3.2 Alternative Mechanism for the Fixation of Adaptive Social Learning Systems

In a general sense, the addition of a moderate rate of environmental change produces a scenario where the genetic system is less likely to find and/or maintain phenotypes that are as fit as those found and/or maintained by the cultural evolutionary system. Considering this, it is easy to imagine at least one other mechanism that could lead to adaptive social learning in our model: the inclusion of a rugged strategy space where fitness valleys (i.e., unavoidable low-fitness regions that interrupt an otherwise monotonic gradient) are more likely to be crossed by social learners (i.e., with relatively high imitation error rates) than instinctive actuators (i.e., with relatively low mutation rates). When comparing simulations where social learning is evolvable with simulations where social learning is not evolvable, there will be combinations of μ_g and μ_p for which the evolution of social learning is advantageous (i.e., where a high μ_p value allows social learners to jump fitness valleys that a monomorphic population of instinctive actuators is extremely unlikely to cross). In these cases, we can claim that the social learning system is also adaptive without the requirement of individual learners or decision-making attributes, as it is able to reach fitter sequences in the strategy space than its genetic counterpart.

In the next chapter, an extended "Future Work" section describes this alternative mechanism and two more in detail (see sections 6.9, 6.10 and 6.11).

5.4 Conclusions

In this chapter, the effects of adding a rate of environmental change (i.e., a rate of change for the optimal sequence) were explored using an extension of our original simulation model. The three main findings of this chapter are:

1. Environmental change increases the likelihood of social learners fixing in a population of instinctive actuators for a large area of the parameter space defined by μ_g, μ_p, x_d and x_r. This effect is especially strong when $x_r > x_d$ (i.e., when survival selection is stronger than reproductive selection) and $\mu_p > \mu_g$ (i.e., when imitation error rates are higher than genetic mutation).

2. For extreme rates of environmental change, constructive evolution is unviable, as rapidly changing environments are impossible to track by either evolutionary system (i.e., cultural or genetic). Under this condition, the population cannot produce consistent improvements of the phenotypes (i.e., adaptive evolution does not occur) and hence genotype-phenotype disengagement is hampered. In this scenario, the average proportion of environment-matching bits in all phenotypes and genotypes in the population remains close to the neutral average, where only 50% of their bits match the current environmental state.

3. A moderate rate of environmental change makes unbiased social learning adaptive at the population level (i.e., social learning increases the overall mean fitness of the population compared to a monomorphic population of instinctive actuators) when $\mu_p \gg \mu_g$ and survival selection is stronger than reproductive selection. This effect is not conditioned to the presence of decision-making attributes in social learners or the presence of individual learners in the population. This is an important finding, as it challenges the existing assumptions regarding the adaptive value of unbiased imitation (Richerson and Boyd, 2006; Rogers, 1988; Cavalli-Sforza and Feldman, 1983).

Chapter 6

General Discussions, Future Work and Conclusions

In the last three chapters, I have investigated all research questions of this thesis (see Section 1.5). For each question, an in-depth analysis of a different variant of my individual-based simulation provided support for my claims and conclusions. In this closing chapter, I comment on general points of discussion that relate to all previous chapters, describe potential future work derived from my results, and state the final conclusions of this thesis.

6.1 Cultural Evolutionary Systems Can Emerge Without Requiring Alternative Sources of Selection

In this section I discuss the results of my investigation into research question 1: *What are the conditions of variation and selection in evolutionary systems that lead to the invasion of social learners and the emergence of culture without invoking individual decision-making capabilities?*

An evolutionary system, whether genetic or non-genetic (i.e., cultural), must instantiate the fundamental elements of reproduction, inheritance, mutation and selection. In this chapter, I reflect on the mechanisms that provide these elements in my model, and the shortcomings of prior models at doing so. My model approach explicitly focuses on the conditions for culture to take over the genetic evolutionary system, rather than the traditional focus of previous models (i.e., explaining the evolution of social learning as an individual trait). This approach led me to build a model that included all the elements of an evolutionary system, and then find conditions through simulations where this system could displace an established genetic one.

The process of social learning implicitly introduces the first three elements when information is *imperfectly* transmitted from one individual to the next (Richerson and Boyd, 2006). As shown, the balance between mutation and inheritance is as fundamental in these systems as it is in any evolutionary process. On the one hand, too much mutation drives the population into an error catastrophe, unable to improve on existing variants. On the other, excessive fidelity retards evolution when little variation exists for selection to act upon (DeJong, 2002). As discussed in Chapter 4 and Chapter 5, the direction of inheritance also plays an important role in cultural systems, as horizontal transmission drastically changes the cultural evolutionary system compared to its vertical counterpart.

In most models, directionality (for improvement) in the cultural system is introduced by decision-making processes such as *transmission bias* and/or *guided variation.* Transmission bias introduces an explicit form of selection by equipping social learners with a direct or indirect bias towards imitating variants with higher reproductive fitness (Richerson and Boyd, 2006; Enquist et al., 2008). Models using this approach include frequency dependent bias, also referred to as conformity bias (Henrich and Boyd, 1998) and social status bias (Richerson and Boyd, 2006). It is important to point out, however, that once a cultural system is established, these two types of biases can evolve without being proxies for fitness bias (Richerson and Boyd, 2006), but these biases lead to the evolution of social learning only if they act as indirect forms of fitness-biased imitation. Guided variation introduces an internalised form of selection within every imitator. Models including guided variation assume individuals select amongst the variants that they have copied and improve them, so that their expressed behaviours are fitter than the ones that they originally imitated (Richerson and Boyd, 2006; Mesoudi, 2011; Aoki et al., 2005; Borenstein et al., 2007).

In contrast to existing models, my results show that when survival selection is stronger than reproductive selection, the strategy-space exploration advantage of social learners can effectively offset the reduced selective pressure on their emerging non-genetic system. The separation of survival and reproductive components of selection is a key feature of my models and contrasts with the classical assumption that survivability is merely a proxy for reproductive fitness, a shortcoming of most existing models in the literature. This feature was included in one previous analytical model by McElreath and Strimling (2008), where an extension of Rogers' basic (1988) model evaluates conditions under which vertical imitation is favoured over horizontal imitation (McElreath and Strimling, 2008). The authors showed that vertical imitation (i.e., imitating one's own parents) is favoured over horizontal imitation (i.e., imitating non-parental agents) when reproductive selection is stronger than survival selection (McElreath and Strimling, 2008). However, social learners in McElreath and Strimling's model (i.e., both vertical and horizontal imitators) rely on the presence of individual learners in order to invade, as they do in the original Rogers (1988) model (i.e., by imitating fit behaviours from individual learners without themselves having to pay the trial-and-error cost of learning). In both models, therefore, social learners do not fix as they are only able to invade until they reach an equilibrium frequency in a mixed population (McElreath and Strimling, 2008; Rogers, 1988). The models proposed by McElreath and Strimling (2008), and Rogers (1988), do not include an extended strategy space nor genetic mutation or imitation error rates, as mine does. Therefore, their models do not consider genes and culture as evolutionary systems in competition to achieve an optimal solution. In my work, I have shown that only when (1) I differentiate between reproductive and survival selection types and (2) I consider these mutation/error rates in an extended strategy space (i.e., not a simple two-strategy model) can unbiased social learners irreversibly fix in a monomorphic population of instinctive actuators (i.e., can culture evolve). In existing models, social learners parasitically rely on the directionality introduced by transmission bias, guided variation and individual learning in order to partially invade a population of instinctive actuators. In this thesis, I have shown that social learners can fix without these sources of directionality, only relying on natural selection.

6.2 Horizontal Imitation is Advantageous When Survival Selection is Strong

In this section I discuss the results of my investigation into research question 2: *Does horizontal transmission enhance the competitive ability of cultural evolution relative to genetic evolution and vertical cultural inheritance?*

In a system where social learners can be either vertical or horizontal imitators and are pressured by both kinds of selection (i.e., survival and reproductive), the emergence of a cultural evolutionary system is produced by the invasion of the vertical rather than the horizontal kind. This occurs because both types of selection (i.e., survival and reproductive) affect vertical imitators (and the cultural evolutionary system they maintain), whereas horizontal imitators are only affected by survival selection. Therefore, during the first stages of optimisation, vertical imitators fix (and cause genotype-phenotype disengagement) due to the combination of high imitation error rates (i.e., higher than genetic mutation, but equal to horizontal imitation) and stronger selective pressure (i.e., stronger than horizontal imitators).

However, when survival selection is much stronger than reproductive selection, horizontal imitators are more likely to dominate the evolutionary process, as both types of imitators perceive a similar selective pressure but the horizontal type acquires trending variants faster, regardless of which type of individual adopted them first in the population (i.e., horizontal imitator, vertical imitator or instinctive actuator). In contrast, vertical imitators are constrained to vertical lines of inheritance, so if by chance an adaptive variant is adopted first by the other type of imitator (i.e., an horizontal imitator), vertical imitators will not have access to this variant when it starts increasing in frequency and are therefore likely to be displaced by the invasion of the horizontal type, which imitates and propagates the trending variant.

This asymmetry makes vertical imitators more likely to be displaced by horizontal imitators during the evolutionary process than horizontal imitators are to be displaced by vertical imitators. Horizontal imitation is therefore the dominant kind of imitation in simulations where both types are evolvable and where survival selection is much stronger than reproductive selection.

As a more general conclusion, I can say that horizontal transmission is more advantageous than vertical transmission when survival selection is strong. This means that the exploratory advantage of social learners (i.e., horizontal transmitters) over instinctive actuators (i.e., vertical transmitters) in my model, is not only a product of their higher imitation error rate but also their ability to copy phenotypes horizontally.

6.3 Environmental Change Increases the Long-Term Overall Mean Fitness of the Population

In this section I discuss the results of my investigation into research questions 3 and 4: *Does environmental change increase the chance of social learners fixing in a population without individual decision-making capabilities? Does the fixation of social learners increase the overall mean fitness of a population under moderate environmental change?*

The model presented in Chapter 3 explores minimal conditions for the evolution of unbiased social learning using an extended strategy space defined by the "onemax" problem on a string of length $L = 200$. For all parameterisations where social learners fix, the process of genotype-phenotype disengagement prevents instinctive actuators from re-invading the population, therefore maintaining the long-term stability of the non-genetic evolutionary system over its genetic counterpart. However, the genetic system on its own (i.e., a population with no social learners) can also find the optimal solution in this space eventually. It is not an explicit claim of my original model in Chapter 3 that unbiased social learning can increase the long-term overall mean fitness of the population compared to instinctive actuators. Instead, I merely show in that chapter that by increasing the overall mean fitness *only* during the early stages of optimisation (i.e., red region in Figure 3.10) social learners can irreversibly fix (see Chapter 3).

Significant effort in evolutionary theory has been focused on explaining the long term advantage of social learning (i.e., how social learners can increase the overall mean fitness of a population) (Rogers, 1988; Richerson and Boyd, 2006). In Chapter 5, I showed a way in which this could be achieved, by including a rate of environmental change.

In my original model (i.e., in Chapter 3) my fitness function used the sum-of-1s value as an input and considered a string of all-1s as the optimal sequence (i.e., fitness varies inversely with the hamming distance between phenotype and the string of all-1s). In Chapter 5, however, I considered an equivalent function where fitness varies inversely with the hamming distance to some arbitrary bit string (i.e., this string is considered to be the optimal solution). Given this equivalence, altering the optimal solution with a given rate throughout the optimisation process and recalculating the corresponding fitness values after each such alteration represents a rate of environmental change.

Under a moderate rate of this type of environmental change, I showed that neither of the evolutionary systems (i.e., genetic or non-genetic) would be able to drive the average phenotype of the population to the optimum, as environmental change would constantly move this target. I further showed, that if the rate of environmental change keeps the average phenotype value below the critical value α (i.e., where social learners find solutions faster than instinctive actuators), unbiased social learners (with their higher imitation error rate and horizontal imitation) could increase the overall mean fitness of the population. Therefore, under this scenario, the evolution and fixation of social learners improves adaptiveness at a population level when compared to a monomorphic population of instinctive actuators subjected to the same rate of environmental change. The idea that moderate rates of environmental change can favour the evolution of social learning is in agreement with conclusions derived from existing models (Kameda and Nakanishi, 2002; Wakano and Aoki, 2006; Rendell et al., 2010a; Boyd and Richerson, 1995, 1996). However, in contrast to existing models, I achieve this without the presence of individual learners or decision-making attributes of any kind.

6.4 The Effect of Individual Learners

When considering the question of how culture evolved, most evolutionary theorists have focused their research on finding conditions under which social learners can invade a population of individual learners (Boyd and Richerson, 1988a; Blute, 1987; Rogers, 1988; Richerson and Boyd, 2006; Enquist et al., 2008; Mesoudi, 2011; Aoki et al., 2005). In these models, individual learners perform a fitter-than-average behaviour for the current environmental state, but pay a fitness cost in learning this behaviour. This cost is generally associated with the process of trial-and-error (Richerson and Boyd, 2006; Mesoudi, 2011). In these models, the environment changes with a given frequency, but individual learners always match their behaviour to the current environmental state (Boyd and Richerson, 1988a; Blute, 1987; Rogers, 1988; Richerson and Boyd, 2006). In a population of individual learners, a mutant social learner with unbiased imitation can easily increase in frequency because the social learner is spared the trial-and-error cost of individual learning, but can acquire, through imitation, the fit behaviour produced by it (Rogers, 1988; Boyd and Richerson, 1988a; Blute, 1987). However, as the frequency of social learners increases so does the chance of imitating a behaviour that does not match the current environmental state (e.g., a population of only social learners cannot keep track of environmental changes), therefore the frequency of social learners will only increase to an equilibrium point where both individual and social learners have the same fitness value (i.e., a mixed population) (Richerson and Boyd, 2006; Rogers, 1988). The fitness of individual learners is not density dependent. Hence, at equilibrium, the mixed population has an overall fitness value equal to that of a monomorphic population of individual learners (Rogers, 1988). According to these models, decision-making attributes are thus required to explain the complete fixation of social learners (i.e., to entirely displace individual learners) and increase the overall mean fitness of the population compared to a monomorphic population of individual learners (Richerson and Boyd, 2006; Rogers, 1988; Enquist et al., 2008; Mesoudi, 2011; Aoki et al., 2005).

In Rogers (1988) and most existing models, social learners can invade (but not fix) because the presence of individual learners introduces a density-dependent benefit to social learners compared to individual learners and instinctive actuators; social learners have costless access to the fitter-than-average behaviours found by individual learners, but instinctive actuators do not. Therefore, the underlying source of fitter-than-average variants that allow social learners to invade when individual learners are present is similar to that found in models where social learners have decision-making abilities that improve on imitated behaviours (i.e., guided variation). In both cases, the valuable information found by social learning is ultimately derived from some form of individual learning (i.e., the individual ability to optimise behavioural fitness). In my model, social learners fix due to a purely Darwinian process where behavioural variants (generated without bias) are the units of selection. In contrast, we can view previous models as Lamarckian adaptation processes, necessitating mechanisms of directed improvement (i.e., directionality) acting on innate or acquired variants, in this case via individual learning or guided variation.

Individual learners exist in current natural populations alongside social learners and may have been present prior to the evolution of social learners (Avital and Jablonka, 2005). But, as a theoretical exercise, I exclude individual learning from my models as it was necessary to test the hypothesis that it is not required to explain the fixation of social learning and the emergence of a corresponding cultural evolutionary system. Accordingly, I am able to investigate the possibility of a generic mechanism for the emergence of other kinds of simpler alternative evolutionary systems through the evolution of unbiased horizontal information transfer where individual learning may not be relevant (e.g., the example of bacterial transfer explained in Section 3.5).

6.5 A Comment on the Baldwin Effect

In theory, processes like *the Baldwin Effect*, where evolution gradually converts learned behaviours into instincts, could cause the reinvasion of instinctive behaviours and the displacement of their socially learned equivalents in the long term (Simpson, 1953). However, this process is unlikely to occur in my models (with the exception of cases featuring extreme environmental change) as the rate of optimisation of the non-genetic evolutionary system (i.e., the system driven by social learners) keeps socially learned behaviours far fitter than any instinctive behaviour. After disengagement, the large fitness difference between socially learned behaviours and instincts creates a scenario where no genetic lineage of instinctive actuators could be maintained for long enough to reinvade. Genotypes in the population of social learners accumulate a large number of random mutations. Therefore whenever one of these agents' offspring mutates into an instinctive actuator, their expressed behaviour is sure to be outcompeted by the behaviours of social learners. The relatively high rate at which genotypes accumulate mutations, compared to the rate at which they are exposed to any selective pressure, makes it impossible for instinctive actuators to reinvade after genotype-phenotype disengagement has occurred.

6.6 A Comment on the Cost of Social Learning

My simulation models show that a high imitation error rate can be detrimental to the evolution of social learners when behaviours are close to the optimum. Some analytical models explicitly include a fixed fitness cost for social learners as a way to represent this imitation error rate (i.e., the "cost" of social learning) (Mesoudi, 2011). Here I show this high error rate can be both an enabler and an inhibitor of the fixation of social learners, depending on the optimisation problem and the value of the imitation error rate compared to genetic mutation.

In my models, no further costs associated with social learning are introduced. All costs related to the development and maintenance of the physiological hardware that allows imitation are considered negligible compared to the fitness value produced by the imitated behaviour itself. All agents compete for the same pool of resources and confront the same challenges when doing so. When population structure, multiple pools of resources, and optimal foraging theory are considered, the fixation of unbiased social learners might not happen in the way described in my simulation models.

6.7 A Comment on the Two-State Imitation Switch in My Simulation Model

In my model, individuals have a genotypic bit which acts as a switch determining whether or not the individual engages in social learning. As I explained in Chapter 3, if the value of the switch is 0, the individual's phenotype is simply a copy of its own genotype. However, if the value of this switch is 1, the individual will express an "imitated" phenotype, randomly selected from the population.

It is important to note that preliminary analyses of simulations with continuous "imitation probabilities" were performed (see Appendix C). In these simulations, the imitation "switch" can take values between 0.0 and 1.0. For each individual, this value represents the probability of imitating a behaviour selected at random from the population, instead of expressing their own genotype. For example, if the imitation probability of an individual is 0.4, that individual will obtain its behaviour by imitating the behaviour of a randomly selected member of the population with a probability of 0.4, and will express the behaviour encoded in its genotype with a probability of 0.6. A comparison between the results of the original "two-state switch" model and the results of the "continuous probability" model did not show any noticeable differences. Therefore, for simplicity, I made the decision of developing all models in this thesis using the "two-state" switch version.

6.8 A Comment on Evolutionary Stability in My Simulation Model

From a dynamical systems perspective, the fixation of social learners in Chapters 3, 4 and 5 can be interpreted as a transient state. In theory, there is a highly unlikely, but still present, possibility that an instinctive actuator expresses the optimal sequence after genotype-phenotype disengagement. Note that after disengagement, the genotype strings of the entire population are expected to drift close to the neutral average (i.e., close to a 50% mix of 1s and 0s). However, with non-zero probability, some of these strings can encode the optimal sequence. In a population where social learners have fixed, it is technically possible that an individual with the optimal genotype string also happens to be an individual that mutates into an instinctive actuator. If this occurs, genotype-phenotype disengagement would reverse and instinctive actuators would reinvade the population. Note that this process occurs because instinctive actuators inherit their genetic strings (which are ultimately expressed as their phenotype) with higher fidelity than social learners imitate the phenotype of other agents. When the population has reached the optimal sequence, this higher fidelity is advantageous.

This caveat of my model should be taken into consideration when comparing my simulation results against the results of mathematical models in the literature, especially when claiming that social learning in my simulation model is an *evolutionarily stable strategy* (i.e., a strategy that cannot be invaded by any other type of strategy). For all results and discussions in this thesis, unless stated otherwise, I consider the probability that genotype-phenotype disengagement could be reversed as negligible. Note that in natural populations, optimisation problems are generally open-ended (i.e., they are always open to further optimisation), therefore, the likelihood of reversing disengagement decreases rapidly as cultural evolution drives phenotypes further away from the drifting, genetically encoded solutions.

In my model, the feature that makes the reinvasion of instinctive actuators highly unlikely is the large size of the strategy space. Note that when the strategy space is small, and all strategies in that space are sampled with significant probability (e.g., through random drift), the fixation of social learners is likely to be reversed as there are other evolutionary accessible strategies that could invade. In contrast, when the strategy space is large, the reinvasion of instinctive actuators becomes highly unlikely (See Section C.4 in Appendix C for more information). In this work, I have implicitly assumed that the fixation of social learners in a model with a large strategy space could be considered a type of evolutionary pseudo-stability. However, as mentioned in the previous paragraph, this pseudo-stability is different from the standard evolutionary stability in mathematical models.

In principle, it would be possible to devise a model where, in addition to the probability of reinvasion being extremely low, a variation/increase in the frequency of instinctive actuators would be selected against. However, such model would require some form of cost, or frequency dependence, that is not included in my simulation.

6.9 Future Work: In a Rugged Fitness Landscape Unbiased Social Learning Could Increases the Overall Mean Fitness of the Population

With the intention of keeping my fitness function as simple as possible, I have focused only on the “onemax” problem in this thesis. However, fitness functions can get far more complicated than this (DeJong, 2002). In particular, there can be functions where no monotonically increasing paths between local and global optima exist (DeJong, 2002; Kauffman, 1993). In these cases, populations can get stuck in local optima, unable to “jump” the gap formed by sequences with lower-than-average fitness (e.g., by mutation) or reach a gradient leading to sequences with higher fitness values.

Previous authors have suggested imitation can aggravate this situation by reducing the overall variation in the population (Lazer and Friedman, 2007), a claim that is only true if social learners are assumed to imitate with perfect fidelity. However, in my model, I explicitly include an imitation error rate that is larger than genetic mutation. Therefore, I hypothesise the non-genetic system maintained by unbiased social learning would have a greater chance of crossing low-fitness gaps between optima in a "rugged" landscape (DeJong, 2002; Weissman et al., 2009). This ability increases the long-term mean fitness of a population of social learners compared to a population of instinctive actuators, as the former could reach optima inaccessible to the latter. A short extension of my model, where a rugged landscape is introduced, could show that this is a plausible mechanism for the emergence of adaptive unbiased social learning.

Such extension, however, could also show the opposite. To understand why, we should keep in mind that a rugged fitness landscape is in essence an *epistatic* fitness landscape. In an epistatic landscape, fitness is the result of specific and sometimes intricate non-linear interactions between alleles. High-fitness phenotypes can therefore turn into catastrophically low-fitness ones with a single mutation event. In such a scenario, an evolutionary system with a low mutation rate (i.e., a genetic evolutionary system) might produce fitter individuals than one with a high imitation error rate (i.e., a cultural evolutionary system). We therefore cannot conclude at first glance that a rugged fitness landscape will always increase the likelihood of the emergence of cultural systems, or increase their adaptive value over their genetic counterpart.

The results of the model extension suggested in this section would offer insights into the effect of fitness landscape ruggedness on the evolution of social learning systems and their relative adaptive value.

As a side note on this section, I would like to mention that the combination of two adaptive processes, operating at two different timescales (e.g., genetic and cultural), can produce adaptations in rugged landscapes that are not possible with only one timescale at any mutation rate (Watson and Szathmary, 2016). Therefore, it is possible for an extended version of my model to explain the evolution of culture in rugged landscapes by giving social learning a different timescale than that of genetic inheritance (e.g., by allowing social learners to imitate more than once during their lifetime).

6.10 Future Work: The Effect of Socio-Spatial Structure on the Evolution of Social Learning

The original simulation model and extensions used in this thesis have only considered a "well-mixed" scenario where: (1) fitness normalisation occurs across the entire population (i.e., the fitness of each individual is calculated in relation to all other members of the population), (2) social learners select models to imitate from the whole population, and (3) the selective environment (i.e., the optimum string against which all phenotypes are compared in order to calculate fitness) is the same for the whole population.

I used this scenario in order to keep my model as simple as possible and to illustrate what I believe to be the most unassuming mechanism to explain the evolution of unbiased social learning. However, changing these assumptions could lead to an alternative explanation for the evolution of adaptive unbiased social learning (i.e., one that does not rely on differences between mutation rates and imitation error rates).

The rationale that leads me to this hypothesis is based on the idea that after migrating to a new group, individuals are likely to benefit from imitating behaviours rather than following their own instincts, which might have evolved in a different environment (Richerson and Boyd, 2006; Mesoudi, 2011). Therefore, when a population is fragmented into groups (i.e., becomes a metapopulation) and: (1) there is a non-zero rate of migration between these groups, and (2) there are differences in the environment each group lives in, social learners will outcompete instinctive actuators due to their ability to imitate phenotypes from their host groups. In contrast, when migrating from one group to another, instinctive actuators express their vertically inherited genotype (i.e., genotypes which have evolved in a different group/environment) as their phenotypes and therefore are likely to mismatch the current environment they live in.

My model framework can easily be extended to control: (1) the fragmentation of the original population (i.e., how many sub-groups exist), (2) the migration rate between groups (i.e., the likelihood of an agent moving from one group to another) and (3) the variance of the optimum sequences between groups (i.e., how different are the environments each group experiences). When considering these three variables, I hypothesize, for certain parameterisations (i.e., for certain values/degrees of 1, 2, and 3), social learners could maintain an overall mean fitness that is higher than that of a population of instinctive actuators.

6.11 Future Work: Unbiased Social Learning can Prove Advantageous When Imitators Construct Their Phenotypes by Recombining Those of Several Individuals

In the fields of evolutionary biology and evolutionary computation, it is generally agreed that one of the key advantages of sexual reproduction is the overall increase in phenotypic variance (DeJong, 2002; Futuyma, 2009; Hartfield and Keightley, 2012). In this thesis, I have established that a higher phenotypic variance in an evolutionary system (i.e., either cultural or genetic) can be advantageous at certain stages of optimisation and under certain parameterisations. However, the way in which sexual recombination increases this variance is fundamentally different from the way mere mutation or imitation errors do. During recombination, the genotypes from two different individuals (i.e., parents) are recombined to produce a new one; in some cases, the result is equal or less fit than the parental genotypes; in other cases, the result is more fit than either of them (Futuyma, 2009). In a population that is evolving solutions to match a current environmental sequence (i.e., as in my simulation model), individuals might have similar fitness values as a consequence of their phenotypes matching two different subsets of the optimal string. In this scenario, an asexual population will need to undergo several generations to optimise phenotypes, with mutation rates introducing small errors in each generation. In contrast, a sexual population could reach a perfect match in as little as a single generation, provided that the recombination of two existing phenotypes leads to the environment matching sequence.

For simplicity, my model only includes asexual reproduction for instinctive actuators and the equivalent for social leaners (i.e., imitators only copy a single individual in the population, once in their lifetime). However, I hypothesize that extending the model to include sexual and cultural recombination could show that there are alternative conditions under which unbiased social learners can outcompete instinctive actuators during the early stages of optimisation (i.e., without requiring different rates of mutation/imitation errors).

When constructing this extension, adding sexual recombination for genetic reproduction would simply imply using the reproductive selection function to pick two (i.e., rather than one) parental individuals and then recombine their genotypes to produce their offspring's genotype. The equivalent recombination process for the cultural system would randomly pick two or more phenotypes and recombine them. The hypothesis that social learners could outcompete instinctive actuators in this extension relies on the idea that social learners are not limited to a two-parent recombination process. If multi-parent cultural recombination produces fitter solutions than sexual reproduction, for any stage of the optimisation process, we expect social learners to invade the population and produce genotype-phenotype disengagement. In support of this idea, existing literature in evolutionary computation has already shown that multi-parent recombination can indeed outcompete two-parent recombination for many classical optimisation problems (Eiben et al., 1994; Tsutsui and Jain, 1998).

6.12 Final Conclusions

My model demonstrates that a basic non-genetic evolutionary system can emerge in a population when selection pressure for survivability is stronger than selection pressure for reproduction. An extended strategy space (where exploration over a large sequence-space is required) distinguishes my approach from previous models and is essential for understanding the exploration advantage of social learning versus genetic inheritance. Analysis of my results leads to a consistent explanation of the emergence of a cultural evolutionary system where phenotype and genotype disengage, with the former evolving exclusively by social learning that is unbiased (i.e., non-critical/non-discriminatory/unguided). During this process of disengagement, the imitator minority must be able to offset its lower selection pressure with its exploratory advantage, a condition that is facilitated when the survivability of the evolved strategy is more important than its reproductive fecundity. This exploratory advantage is not only a product of higher imitation error rates (i.e., compared to genetic mutation), but also of the ability to imitate non-parental models (i.e., horizontal imitation).

Under a moderate rate of environmental change, social learners are more likely to fix in the population, and moreover increase its overall mean fitness compared to a monomorphic population of social learners. To achieve this conclusion, my model does not require decision-making attributes or the presence of individual learners in the population. Therefore, to my knowledge, this work offers the first explanation of the evolution of adaptive culture (i.e., culture that increases the overall mean fitness of the population) that does not involve forms of directional change other than purely Darwinian natural selection (i.e., that does not involve alternative optimisation drivers like biased imitation, guided variation or individual learning).

My original simulation model offers a very simple framework for the emergence of non-genetic systems and serves as a tool for future research extensions. This simplicity, in particular removing the need for decision-making abilities such as those that result in transmission bias and/or guided variation, lowers the minimal number of individual capabilities that a species must have in order to evolve social learning and therefore broadens its potential applications.

Appendix A

Glossary of Terms

Individual Learning / Individual Learners:

A type of learning strategy in which individuals learn new behavioral variants without relying on imitative processes (i.e., without relying on social learning). In existing models in the literature, individual learners pay a learning cost (i.e., a trial-and-error cost), ultimately expressed in fitness, and always express environment-matching behaviors (i.e., behaviors that are beneficial in the current environmental state). In our simulation model, we do not include individual learners.

Social Learning / Social Learners:

A type of learning strategy in which individuals acquire behavioral variants from others, through any process that results in the imitation of the behaviours of one individual (i.e., the model to imitate) by another individual (i.e., the social learner). In our simulation models in Chapter 3 and 5, social learners imitate, at birth, the phenotype of a randomly selected individual in the population. In our simulation model in Chapter 5, social learners can be horizontal imitators (i.e., they imitate a randomly selected individual from the population) or vertical imitators (i.e., they imitate their parents).

Culture / Cultural Evolutionary System:

A set of behavioural traits that are not the direct result of genetic expression but the product of an evolving pool of variants (i.e., behaviours and ideas) stored and imperfectly transmitted within and/or between overlapping generations by means of social learning. A cultural evolutionary system, is the adaptive system that maintains and evolves culture. In this system, behaviors are inherited and replicated through social learning, random variations are introduced by imitation errors, and different forms of natural selection and transmission biases act as selection functions. In a cultural evolutionary system, the evolution of the phenotype occurs independent from the evolution of the genotype (see *genotype-phenotype disengagement*).

Genetic Evolutionary System:

A genetic evolutionary system is an adaptive system maintained by genetic replication and reproduction. In this system, random variations are introduced by genetic mutation and natural selection acts as a selection function. In a genetic evolutionary system, genotype and phenotype are evolved in parallel, with the phenotype being the result of the expression of the genotype.

Instinctive Actuators / Innate Agents:

Individuals whose behaviour is entirely the product of genetic expression (i.e., individuals whose behaviour is instinctive). In our simulation model, the phenotype of instinctive actuators is an exact copy (i.e., a copy without errors) of their genotype.

Fixation:

The complete invasion of one type of agent over all others, in a population. Logically, this process leads to a monomorphic population of the invading type (e.g., the fixation of social learners in a population that was initially formed by instinctive actuators leads to a monomorphic population of social learners).

Phenotype-Fitness De-correlation:

De-correlation between the phenotype of social learners and their reproductive output. This process occurs when social learners imitate any member of the population at random (unbiased imitation) and are under natural selection (i.e., where fitness is a function of reproductive output alone). Under these conditions, the pool of strategies that social learners imitate from (i.e., the population) has, logically, an average fitness value. Since every social learner's offspring obtains a strategy to imitate from this pool (at random), their behaviour is not defined by their genetic endowment and therefore the link between an increased reproductive output and the increase in frequency of the strategy that produced that output is broken. This breakage is Phenotype-Fitness De-correlation.

Social Learning Stagnation:

Decrease of phenotypic variance in populations of social learners, as a consequence of error-free imitation. This process is present in existing models in the literature, but not in our simulation model. In our simulation model, random variation is introduced by imitation error. In existing models, social learning stagnation occurs when social learners continuously imitate one another, perfectly, and therefore, as a population, loose track of environmental changes.

Critical Social Learning / Biased Imitation:

A form of social learning in which individuals bias their imitation towards behaviours/-models with relatively high fitness values.

Unbiased Imitation / Random Imitation:

A form of social learning in which individuals imitate behaviours selected at random from the entire population, i.e., unbiased social learners do not bias imitation towards high-fitness behaviours/models.

Guided Variation:

Process defined by social learners improving upon imitated behaviours. In existing models in the literature, guided variation is introduced as a form of individual learning. This process can be considered an additional form of selection, one that occurs within each individual, when poor-performing variations of the originally imitated behaviours are discarded in favour of fitter, individually enhanced ones.

Behavioural Variant:

A type of behavioral solution for a given behavioural challenge/trait. Example: a mating dance with circular patterns of movement is a *behavioural variant* for the behavioural challenge of "attracting a suitable mate". A mating dance with triangular patterns of movement is another behavioural variant, for the same behavioural challenge. In our simulation model, each specific sequence of the entire phenotypic string of bits, is a behavioural variant.

Cultural Variant:

A behavioural variant that is transmitted and reproduced through social learning in a cultural evolutionary system.

Genetic Variant / Genetic Allele:

A behavioural variant (or any other type of phenotypic variant) that is transmitted and reproduced through genetic replication in a genetic evolutionary system.

Decision-Making Abilities:

Set of cognitive abilities that allow individuals to differentiate between alternative behavioral variants based on the fitness values of such variants. Critical social learning and guided variation are two examples of decision-making abilities.

Horizontal Imitation / Horizontal Imitators:

A from of social learning in which individuals imitate members of the population with whom they have no genetic relation. Some models in the literature distinguish between *horizontal imitation* and *oblique imitation.* In these models, horizontal imitation refers to imitation between members of the same generation, and oblique imitation refers to imitation of non-parental members of the parental generation. In our simulation model, we define horizontal imitators as individuals who imitate, at birth, the phenotype of an individual selected at random from the population (parents included). Our simulation model is a steady-state model without discrete generations; therefore, we do not distinguish between horizontal imitation and oblique imitation.

Vertical Imitation / Vertical Imitators:

A form of social learning in which individuals imitate their parents. In our simulation model in Chapter 4, vertical imitators copy the phenotype of their parents.

Genotype-Phenotype Disengagement:

In our simulation, genotype-phenotype disengagement is defined as the de-correlation of genotype and phenotype information, accompanied by the fixation of social learners. This process is a result of genetic information not being expressed in the phenotype of social learners and, in consequence, selection not acting on genotypes. When selection does not act on genotypes, genetic drift moves the information encoded in these strings of bits towards poor-performing solutions. In turn, this situation makes it even less likely for an innate mutant (i.e., a social learners' offspring that mutates into an instinctive actuator) to reinvade. Once disengagement starts, behavioural variants keep evolving by cultural evolution, while the average fitness of the genotype strings drifts towards the neutral average (i.e., a string of randomly selected bits). The reversal of this process is highly unlikely.

Survival Selection / Viability Selection:

Survival selection is defined as a form of selection that reduces life-span, but not reproductive output. In our model, this form of selection is implemented using a Boltzmann function (see Equations 3.2 and 5.2). The strength of this form of selection is controlled by the exponent x_d in the aforementioned equations, i.e., the effect that the number of environment-matching bits in the phenotype has on the relative probability of an individual being picked for death, is controlled by the exponent x_d. High values of x_d reduce the strength of survival selection. When $x_d = \infty$, the death function effectively converges to a random selection function, where each individual in the population has the same probability of being selected for death.

Reproductive Selection / Fecundity Selection:

Reproductive selection is defined as a from of selection that reduces reproductive output. In most models in evolutionary biology, this from of selection is equated to natural selection. In our simulation, we implement survival selection using a Boltzmann function (see Equations 3.1 and 5.1). The strength of this form of selection is controlled by the exponent x_r in the aforementioned equations, i.e., the effect that the number of environment-matching bits in the phenotype has, on the relative probability of an individual being picked for reproduction, is controlled by the exponent x_r. High values of x_r reduce the strength of reproductive selection. When $x_r = \infty$, the reproductive function effectively converges to a random selection function, where each individual in the population has the same probability of being selected for reproduction.

Imitation Error Rate (μ_p):

Rate at which social learners introduce errors when imitating the behavior of other individuals in the population. In our simulation model, the imitation error rate μ_p represents the probability that each bit in the phenotype string of a social learner is changed from 1 to 0, or from 0 to 1, during the imitation process.

Genetic Mutation Rate (μ_g):

Rate at which errors in genetic replication occur during reproduction. In our simulation model, the genetic mutation rate μ_g represents the probability that each bit in the genotype string of an individual is changed from 1 to 0, or from 0 to 1, during genetic replication.

Imitation Switch:

In our simulation model, the imitation switch is a genetically transmitted bit that determines whether an individual is an instinctive actuator (switch value equal to 0), a social learner / horizontal imitator (switch value equal to 1), or a vertical imitator (switch value equal to 2). During genetic reproduction, this bit is genetically inherited from parents, with probability μ_c of changing from 1 to 0, or from 0 to 1.

Fidelity / Transmission Fidelity:

In our simulation model, fidelity is the degree of exactness with which phenotype and genotype information is transmitted during phenotypic imitation and genetic replication, respectively.

Trial-and-Error Cost:

Cost that represents the time and resources that are consumed in the process of learning a new behaviour / skill. In existing models in the literature, trial-and-error is represented as a fix cost in fitness (i.e., a cost in relative reproductive output). In a fluctuating environment, this cost is paid by individual learners in order to “learn” the specific behaviour that matches the current environmental state. In our simulation model, we do not include individual learners; therefore, we do not include trial-and-error costs.

Appendix B

Simulation Details

In this appendix we describe our basic individual-based simulation model and the extensions that derive from it. For these descriptions, we have used text already provided in Chapters 3, 4 and 5. However, here, we present this text in a continuous format, and explain how the model extensions in Chapters 4 and 5 derive from our basic model (i.e., explained in Chapter 3). The text contains detailed descriptions of simulation steps with particular emphasis on the internal logic of each step. Additionally, this appendix contains functional definitions of processes and types of agents within the context of our model (e.g., *social learner, instinctive actuator, horizontal imitator, vertical imitator, environmental change, genetic mutation, imitation error, selection coefficients, survival selection, reproductive selection.*).

B.1 Chapter 3: The Basic Model

Our individual-based simulation model represents each individual by two strings of bits. The first string represents an individual's phenotype and the second represents its genotype. The fitness of an individual is determined by considering only the phenotypic information. Individuals also have a single additional genotypic bit which acts as a switch determining whether or not the individual engages in social learning. If the value of the switch is 0, the individual's phenotype is simply a copy of its own genotype (for the purposes of our research questions, a one-to-one identity mapping from genotype to phenotype is sufficient). However, if the value of this switch is 1, then rather than express a phenotype derived from its own genotype, an individual will express an "imitated" phenotype randomly selected from the population. These social learners will each imitate a randomly selected phenotype at birth and their phenotypes will then remain unchanged for the remainder of their lifetimes.

All the agents in the initial population have an imitation switch value of 0 and each bit of their genotype bit string set to 1 or 0 with uniform random probability per bit.

Each initial agent's phenotype is set to be a perfect copy of its genotype. These initial conditions prevent mutation biases from imposing directed drift on the genotypes or phenotypes, even in the absence of selective pressure. For all the results shown in this thesis, unless stated otherwise, the length of the phenotype and genotype bit strings is 200 bits ($L = 200$), and the size of the population is 100 individuals ($N = 100$). We have selected these values considering, (1) the need for short computational times per simulation run (i.e., small populations and short bit strings require less computational time), and (2) the need for minimal bit string length and population sizes in order to show key processes described in this thesis (i.e., if bit strings and/or population sizes are too small some processes described in this thesis will not occur). Selection is defined by a Boltzmann-weighted sum-of-1s in the string of phenotypic bits. A phenotypic string of all-1s represents the optimal solution.

Qualitatively similar results can be reproduced with larger populations and longer bit strings. Smaller populations produce variation between individual simulation runs, but on average they also agree qualitatively with the results described in this thesis. Shorter bit-strings can produce different results, as they represent simpler search problems, and a sustained search and optimisation period of several iterations is required in order to produce the results reported here. For example, a single-bit (two-strategy) model will not reproduce our results.

Selection is established by the joint action of a reproduction function and a death function, which impose *selection for reproduction* and *selection for survival*, respectively. Reproduction selects an individual i from the population with a probability P_{ri} using a Boltzmann-weighted function of its sum of phenotypic 1s (i.e., L_{1i}) normalised across the population. The Boltzmann exponent x_r in the function, controls the strength of selection for reproduction. High values of x_r reduce the effect of reproductive selection (Equation B.1).

$$P_{ri} = \frac{e^{\frac{L_{1i}}{x_r}}}{\sum_{k=1}^{N} e^{\frac{L_{1k}}{x_r}}} \tag{B.1}$$

The death function selects an individual j from the population with a probability P_{dj} using a Boltzmann-weighted function of the complement of its sum of phenotypic 1s (i.e., $L - L_{1j}$). Thus, P_{dj} is the relative probability of dying, or anti-fitness. The Boltzmann exponent x_d in the function, represents the strength of selection for survival. High values of x_d reduce the effect of survival selection (Equation B.2).

$$P_{dj} = \frac{e^{\frac{L-L_{1j}}{x_d}}}{\sum_{k=1}^{N} e^{\frac{L-L_{1k}}{x_d}}} \tag{B.2}$$

A key feature of our model is the ability to change the relative balance between reproductive selection and survival selection. Higher values for exponents x_r and x_d reduce the effect of the sum of phenotypic 1s on probabilities of reproduction and death, respectively.

The simulation is a steady-state model, which means that, in each iteration, a single individual is selected for reproduction by the reproduction function and a single individual is selected for death by the death function (i.e., the model produces overlapping "generations"; therefore, agents created at different iterations coexist at any given point of a simulation run). The genotype and imitation switch of the reproducing individual are mutated, and they then replace the genotype and imitation switch of the dying individual. Genotype mutation occurs as a uniform bit-flip probability of μ_g per bit, except for the imitation switch which has a uniform bit-flip probability of μ_c per bit. If the imitation bit of the new copy has value 0, the new phenotype bit string is a one-to-one copy of the new genotype (not including the imitation switch). If the switch has value 1, a random phenotype from the population is copied (unbiased horizontal imitation) with a bit-flip error rate of μ_p per bit. The imitation switch can only be passed genetically and is not part of any phenotype.

For simplicity, individuals reproduce asexually with no "crossover" during genetic replication or phenotypic imitation and there is no environmental change altering the optimum bit-string (i.e., the string where $L_1 = L$), although an extension on the effect of the latter is presented in Chapter 5.

Changing the values of μ_g and μ_p controls the rates of mutation and imitation error, respectively. We fix $\mu_c = 0.01$ for all results reported here and vary μ_g and μ_p relative to this value. Figure B.1 illustrates an algorithmic implementation of the model.

As discussed in Section 3.5, we chose to fix $\mu_c = 0.01$ to incentivise the formation of *chains of imitation* (i.e., sequential imitation among imitators) in the early stages of our simulation runs. For lower values of μ_c, the density of social learners produced by mutation alone (i.e., before any form of selection acts on the population), is too low for imitation among imitators to occur with high enough frequency as to kick-start an evolutionary system (i.e., culture). Values of μ_c that are larger than those minimally required to start chains of imitation, produce a large initial density of social learners without the intervention of any evolutionary process. For obvious reasons, a high initial density of social learners would make our model unsuitable to describe the emergence of culture, as this process is precisely the outcome of the evolution of social learning. The value of $\mu_c = 0.01$ was selected as a compromise between these two constraints. Please see Section 3.5 for more details.

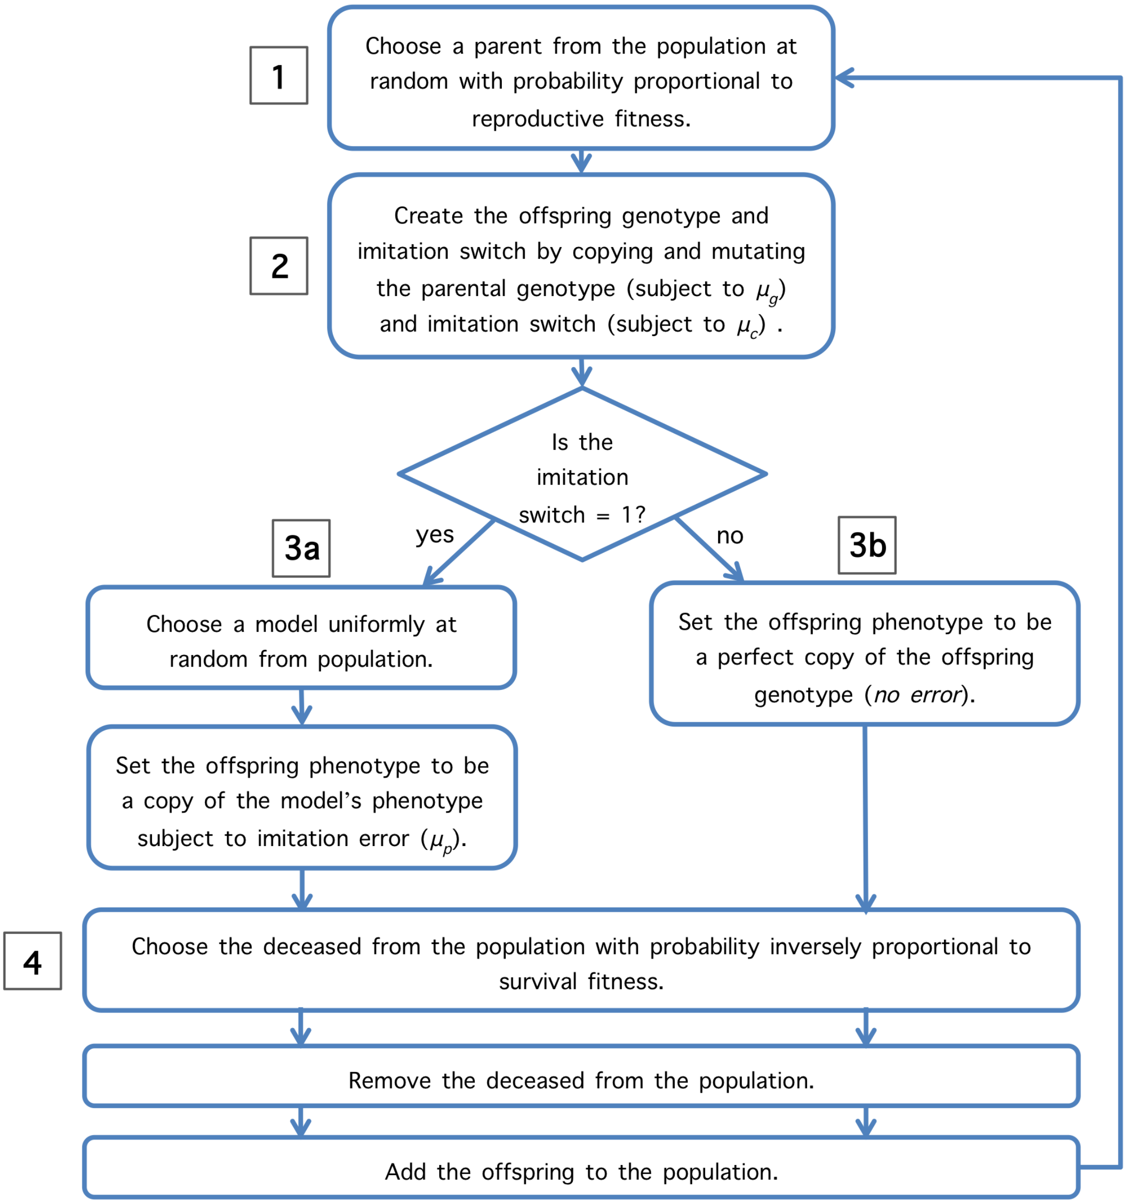

FIGURE B.1: An algorithmic representation of the model's logic. (1) The selection function picks an individual from the population. (2) Its genotype string is copied including the imitation switch (mutations occur at a rate of μ_g per genotype bit and μ_c for the imitation switch). (3a) If the imitation switch has value 1, a random individual from the population is selected and its phenotype will be copied (imitation errors occur at a rate of μ_p per bit). (3b) If the imitation switch has value 0, the phenotype will be a perfect copy of the individual's genotype. (4) The resulting combination of phenotype and genotype will replace an individual selected by the death function.

B.1.1 Simulation Steps

In this subsection, the model presented in Figure B.1 is explained in detail, with particular emphasis on clarifying the internal logic of each function and how it affects the bit strings that define individual agents in the simulation. To assess the internal logic of our model, before generating our analytical results, all individual functions described here were “unit tested”. That is, for each function, the expected outputs for a range of possible inputs were evaluated.

B.1.1.1 Population Setup

To setup a simulation run, we initialise a population of 100 agents (i.e., $N = 100$) by allocating two bit strings of 200 bits per individual (i.e., $L = 200$). One string of 200 bits represents the phenotype, and the other string of 200 bits represents the genotype. For all simulations in this thesis, unless stated otherwise, the value of each bit is assigned at random during the setup (i.e., each bit has equal probability of being 1 or 0). Phenotype and genotype strings represent encoded solutions for a behavioural challenge (e.g., how to cross a river safely, build a nest or sing a song).

In addition to these two strings, each individual also has an extra “genotype” bit. This bit acts as a switch. If the switch has value 1, the individual will substitute, at birth, their entire 200-bit phenotype string for the phenotype string of a randomly selected individual from the population. If the switch has value 0, the individual will express an exact copy of its 200-bit genotype string as their phenotype (without including the imitation switch). Unless stated otherwise, all individuals in the population are initialised with their imitation switch in position 0. This means all individuals are *instinctive actuators* (i.e., their genotype string and phenotype string are exact copies). Figure B.2 shows a diagram explaining this setup.

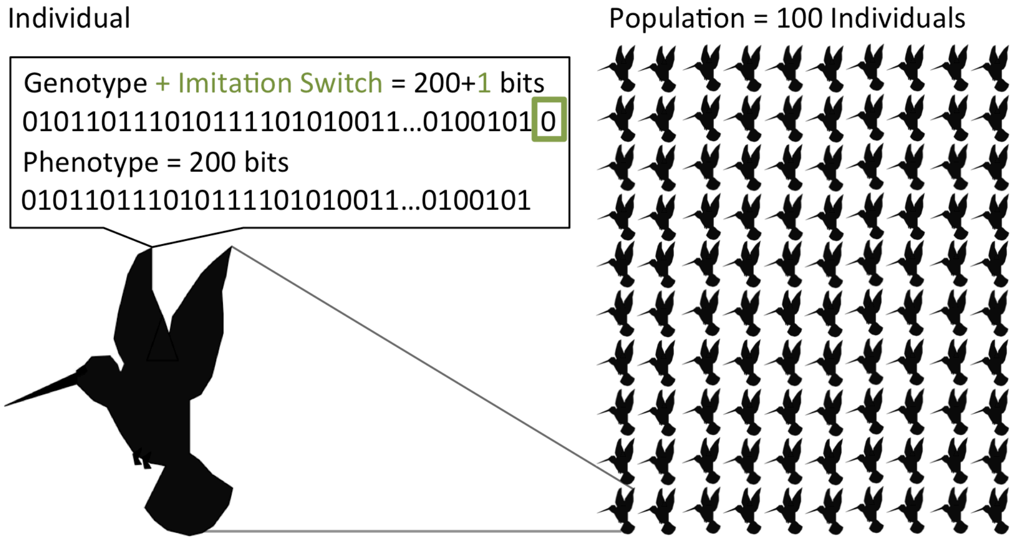

FIGURE B.2: Simulation setup of a population with 100 individuals. Each individual is represented by two strings of 200 bits. One string represents the phenotype, and the other string represents the genotype. In addition to these strings, each agent has an imitation switch bit (framed in green) which is inherited during genetic reproduction but not expressed in the phenotype. All agents are initialised with their imitation switch set to 0.

For each individual we calculate *reproductive selection scores* as $e^{\frac{L_1}{x_r}}$, where L_1 is the sum-of-1s in the phenotype string and x_r is the Boltzmann coefficient that controls the strength of reproductive selection in the simulation run. In a similar manner, we calculate *death selection scores* for every individual as $e^{\frac{L-L_1}{x_d}}$, where $L-L_1$ is the sum-of-0s in the phenotype string, and x_d is the Boltzmann coefficient that controls the strength of survival selection in the simulation run.

Note that x_r and x_d are global parameters in the simulation (i.e., we use the same x_r and x_d to calculate the reproductive selection scores and death selection scores of all individuals, respectively). To avoid repeating computationally expensive operations, we store these selection scores and only recalculate them when necessary.

B.1.1.2 Step 1: Select a Parent With Probability Proportional to Reproductive Fitness

In the first step of our algorithm, we select an individual using the reproductive selection scores calculated during the setup. An individual i will be selected for reproduction with probability P_{ri} as described in Equation B.1. This means, the probability of selecting any given individual for reproduction is equal to the reproductive selection score of that individual divided by the sum of the reproductive selection scores of all individuals in the population (i.e., the scores are normalised).

B.1.1.3 Step 2: Create the Offspring Genotype and Imitation Switch

In the second step of our algorithm, we create the genotype and the imitation switch of a new individual. This new individual will be introduced into the population in step 4, and represents the offspring of the parent selected in step 1. To produce its genotype, we copy the genotype string of the parent with an error rate of μ_g per bit. This means, each bit copied into the genotype of the new individual will change its value (i.e., changing from 0 to 1, or from 1 to 0) with probability μ_g. The imitation switch of the parent is also copied into the offspring. However, this switch will change its value with probability μ_c, instead of μ_g. Figure B.3 shows a diagram explaining this step.

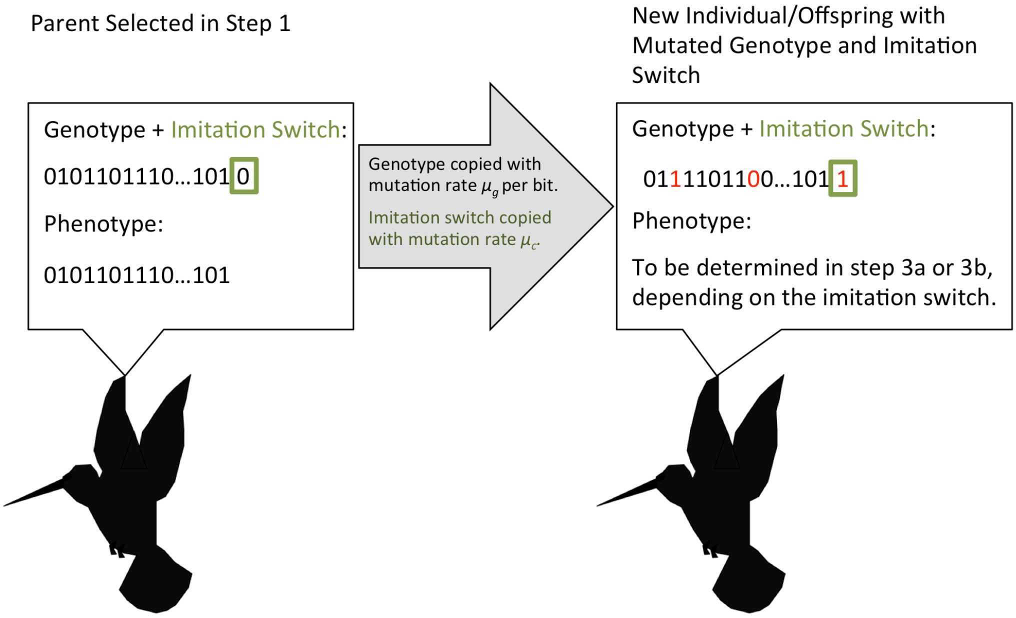

Figure B.3: Diagram representing the second step of our simulation model. We create a new individual using the individual selected in the first step (i.e., parent). The genotype of the new individual is a copy of the parental genotype. Each genotype bit is copied subject to a mutation rate of μ_g. The imitation switch of the new individual is a copy of the parental imitation switch. This switch is copied subject to a mutation rate of μ_c. Mutated bits are shown in red.

B.1.1.4 Step 3a: If the Offspring is a Social Learner

In the previous step, we produced a new individual/offspring with a genotype string and an imitation switch, but without a phenotype string. In this step, we will create a phenotype for this individual assuming that its imitation switch is 1 (i.e., the individual is a social learner). When the imitation switch is set to 1, the phenotype of the new individual is a copy of the phenotype string of another individual, selected at random from the entire population (parent included). This process of phenotype imitation occurs once at birth. The phenotype string is copied with an error rate of μ_p per bit. This means, each bit copied into the phenotype of the new individual will change its value (i.e., changing from 0 to 1, or from 1 to 0) with probability μ_p. Notice that by imitating the phenotype of an individual selected at random from the population, social learners are effectively engaging in *unbiased imitation*. Figure B.4 shows a diagram explaining this step.

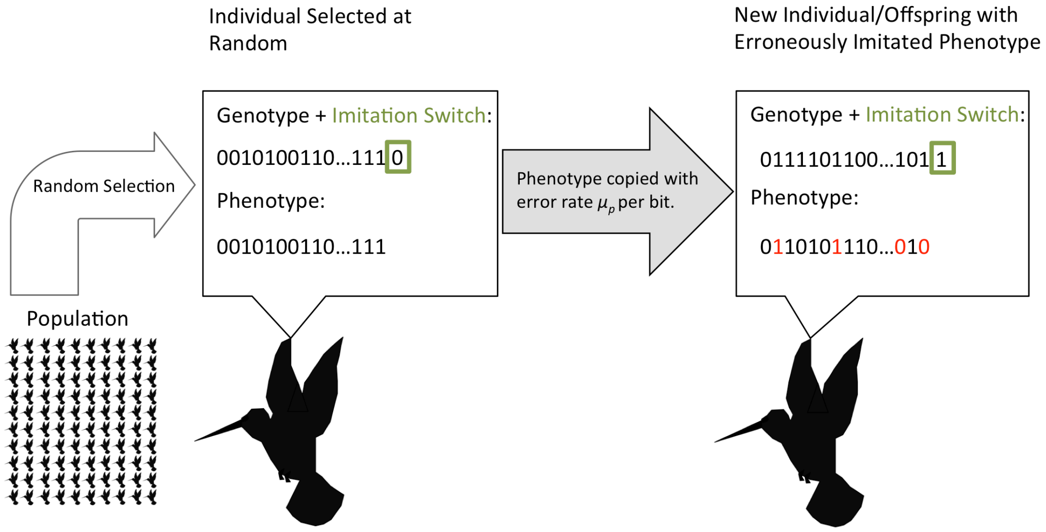

FIGURE B.4: Diagram representing the third step of our simulation model (if the new individual is a social leaner). If the value of the imitation switch of the new individual is equal to 1 (i.e., the individual is a social learner), we create its phenotype by copying the entire phenotype string of an individual selected at random from the population. Each phenotype bit is copied subject to a mutation rate of μ_p. Notice that the imitation switch is not part of the phenotype. Mutated bits are shown in red.

B.1.1.5 Step 3b: If the Offspring is an Instinctive Actuator

In step 3a we created the phenotype of the new individual/offspring by copying the entire phenotype of a randomly selected agent. We used this method because the imitation switch of the new individual had a value of 1. In contrast, in this alternative step, which we labelled as step 3b, we will describe the method for creating the phenotype of the new individual when its imitation switch is set to 0. It is important to note, that new individuals will create their phenotypes using step 3a or step 3b, but not both.

If the imitation switch of the new individual is set to 0, an exact copy of its genotype string will be used to create its phenotype. This copying process represents genetic expression, and occurs without error. For the focus of our research questions, a 1-to-1 map between phenotype and genotype will suffice. The imitation switch is not part of the phenotype, and therefore is not copied. Figure B.5 shows a diagram explaining this step.

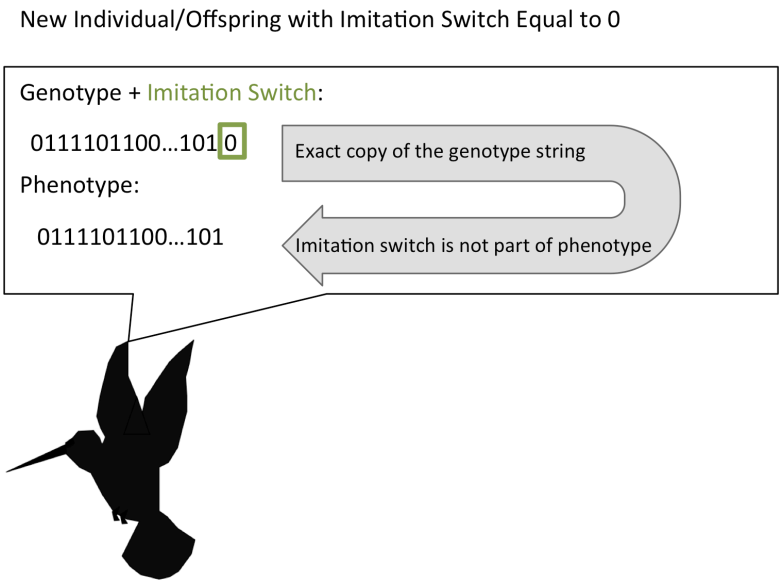

Figure B.5: Diagram representing the third step of our simulation model (if the new individual is an instinctive actuator). If the value of the imitation switch of the new individual is equal to 0 (i.e., the individual is an instinctive actuator), its phenotype is an exact copy of its own 200-bit genotype string. Each bit is copied without error. This step represents genetic expression (i.e., the phenotypic expression of the behavioural variant encoded in the genotype). Notice that the imitation switch is not part of the phenotype; therefore it is not "expressed".

B.1.1.6 Step 4: Select a Replacee With Probability Inversely Proportional to Survival Fitness

In this step, the new individual created through step 3a or 3b is introduced in the population. However, in order to maintain a constant population size, another individual needs to be removed/killed/replaced. The replacee is selected using the death selection scores mentioned in our setup. An individual j will be selected for replacement/death with probability P_{dj} as described in Equation B.2. This means, the probability of selecting any given individual for death is equal to the death selection score of that individual divided by the sum of the death selection scores of all individuals in the population (i.e., the scores are normalised). Note that the death selection scores use the sum-of-0s instead of the sum-of-1s. Therefore, the probability of being selected for death is inversely proportional to the total number of 1s. Also, note that the parent selected for reproduction in step 1 can also be selected as the replacee in this step.

After a replacee as been selected, the new individual takes its place in the population. The reproduction selection scores and death selection scores of the new individual are calculated and stored. These scores will be used in the next iteration of the model, starting on step 1.

B.1.2 System Details

As it can be appreciated, the core simulation of our work is very simple. During the research and development phase of this thesis, our model was independently implemented in three different computer languages: NetLogo 4.0, Python 2.7 and C. Results from each implementation were crosschecked to guarantee consistency and reproducibility. Single-run results, shown in Section 3.2, were produce using the implementation in python, on an Apple-MacBook Pro laptop (2.3 Intel Core i5 - 8 GB of RAM) and an Apple-IMac desktop (3.4 GHz Intel Core i7 - 8 GB of RAM). However, implementations in NetLogo and Python were deemed too slow to produce results that required a large number of replicated runs, in a sensible amount of time (e.g., Figures 3.8 and 3.9 in Section 3.2).

For results that required a large number of replicates, a trivially parallelised version of our implementation in C was executed using OpenMP (Open Multi-Processing) on the IRIDIS4 supercomputing cluster, at the University of Southampton. The IRIDIS4 supercomputing cluster is one of the largest computational facilities in the United Kingdom; it is professionally managed, and available to the university's entire research community. IRIDIS4 delivers a reliable batch service for distributed memory parallel jobs. For this thesis, a single node in the cluster was used. A node has 16 Intel Sandy Bridge CPUs at 2.6 GHz with 64 GB of RAM. All plotting and visualisations were produced using matplotlib 4: a python visualisation library.

B.2 Chapter 4: Extension with Vertical and Horizontal Imitation

By modifying our basic simulation model, we can explore the effects of vertical imitation on the evolution of social learning. Here, we are mainly interested in whether vertical imitators will (1) displace horizontal imitators, (2) coexist alongside horizontal imitators, or (3) be displaced by horizontal imitators.

In this version of the model, we have included an extra state for the imitation switch. Here, individuals can either be instinctive actuators with a switch value of 0, horizontal imitators with a switch value of 1, or vertical imitators with a switch value of 2. As in the original model, an instinctive actuator expresses the information contained in their own genome as their phenotype (without error), and a horizontal imitator selects a random member of the population (including its parents) from whom to copy their phenotype with a mutation rate μ_h. A vertical imitator copies the phenotype of its parent with an imitation error rate μ_v. All other features of the original model from Chapter 3 remain unchanged in this extension. A diagram of the algorithm for the extended model is explained in Figure B.6.

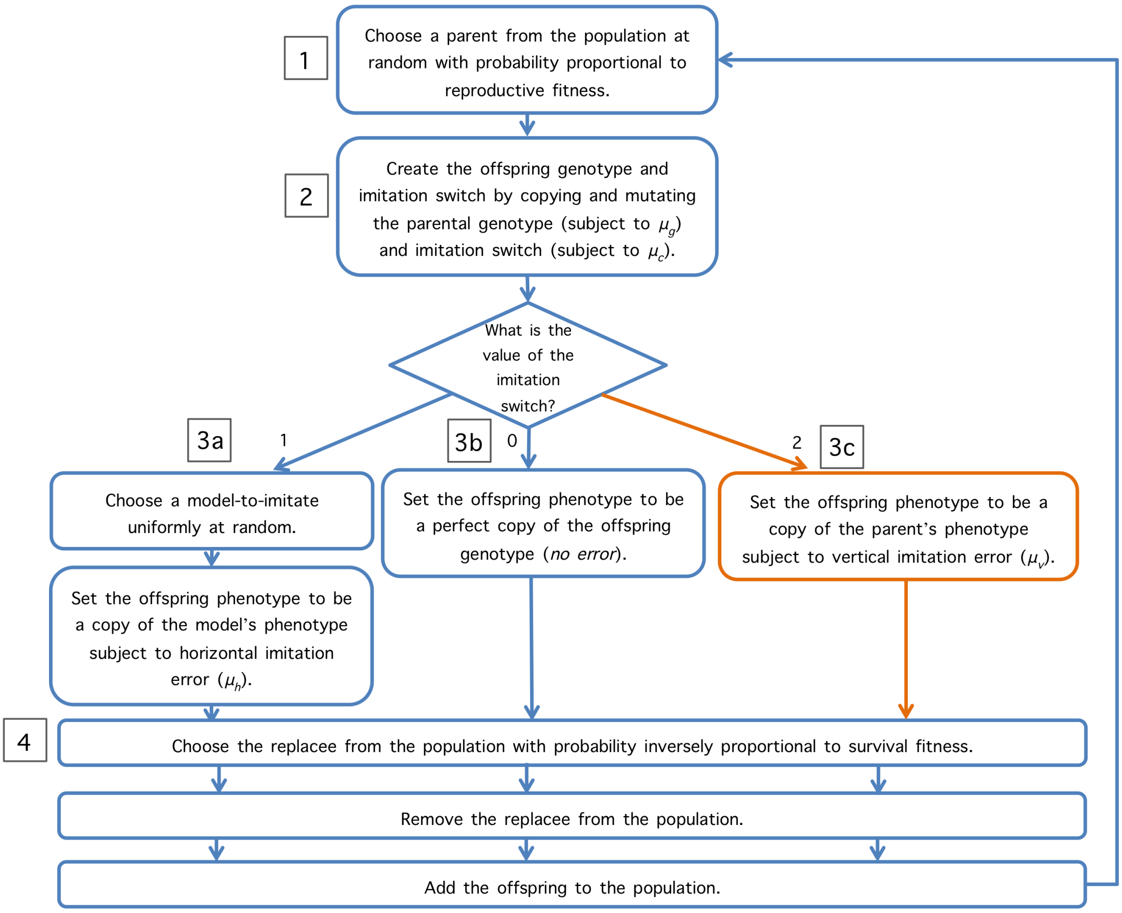

FIGURE B.6: An algorithmic representation of the extended model's logic. (1) The selection function picks an individual from the population. (2) Its genotype string and imitation switch are copied (mutations occur at a rate of μ_g per genotype bit and μ_c for the imitation switch). (3a) If the imitation switch has value 1, a random individual from the population is selected and its phenotype will be copied (imitation errors occur at a rate of μ_h per bit). (3b) If the imitation switch has value 0, the phenotype will be a perfect copy (i.e., without errors) of the individual's own genotype. (3c) If the imitation switch has value 2 the parental phenotype will be copied (imitation errors occur at a rate of μ_v per bit). (4) The resulting combination of phenotype and genotype will replace an individual selected by the death function. Steps that have been added in this extension, compared to the basic model in Figure B.1, are coloured in orange.

As in the original model, instinctive actuators mutate into imitators at a rate of μ_c. If an instinctive actuator mutates in this way, it has an equal chance of becoming a vertical imitator (imitation switch $= 2$) or a horizontal imitator (imitation switch $= 1$). Conversely, imitators of both types (vertical and horizontal) can mutate with rate μ_c into either instinctive actuators or the other kind of imitator, also with equal chance. Every imitator performs only one type of imitation in their lifetime and only imitates once. The type of imitation and the imitated phenotype are both set at birth. For all results in this chapter, the mutation rate for the imitation switch is fixed at $\mu_c = 0.01$. This value is equal to the one used in the original model and is fixed for the same reason; to maintain a minimal frequency of all types of agents (see Chapter 3 for a comprehensive explanation).

In this extension, population size is set to one hundred individuals ($N = 100$), the length of the phenotype and genotype strings are set to two hundred bits ($L = 200$), and both types of selection are controlled by Boltzmann-weighted functions of the sum-of-1s in the phenotype strings (see Equation B.1 and Equation B.2). A string of all-1s (i.e., $L_1 = L$) represents the optimal solution.

Simulations were initialised with all agents as instinctive actuators (i.e., every imitation switch set to 0) and each bit of their genotype and phenotype bit strings set to 1 or 0 with uniform random probability per bit. At the start of the simulation, each agent's phenotype and genotype strings are perfect copies of one another. These initial conditions prevent mutation biases from imposing directed drift on the genotypes or phenotypes, even in the absence of selective pressure.

B.2.1 Simulation Steps

In this subsection, steps presented in Figure B.6, that were not present in Figure B.1, are explained in detail. To assess the internal logic of our model, before generating our analytical results, all individual functions here described were "unit tested". That is, for each function, the expected outputs for a range of possible inputs were evaluated.

B.2.1.1 Three-Way Mutation Implementation

In this model extension we include a third type of agent: a vertical imitator. Therefore, the logic behind the mutation function for the imitation switch must allow for mutations between three values, instead of only two, as in the original model. Notice that in this model extension, when the imitation switch of the new individual, created in step 2, has a value of 0, the individual is considered an *instinctive actuator* and will obtain its phenotype through step 3b of Figure B.6, using the same process described in Subsection B.1.1.5 (i.e., by copying their own genotype without error). If the imitation switch of the new individual has a value of 1, the individual is considered an *horizontal imitator* (referred to as a "social learner" in Chapter 3) and will obtain its phenotype through step 3a of Figure B.6, using the same process described in Subsection B.1.1.4 (i.e., by imitating the phenotype string of a randomly selected individual from the population, subject to mutation rate μ_h). If the imitation switch of the new individual has a value of 2, the individual is considered a *vertical imitator* and will obtain its phenotype through step 3c of Figure B.6 by a process that will be described in detail in the next subsection.

The mutation function can change the imitation switch from one value to any of the other two. In each iteration, when producing a new individual/offspring, the imitation switch will change from its current value with probability $\mu_c = 0.01$. If the imitation switch is mutated, one of the other two values will be selected with equal probability (i.e., 0.5). For example, if an individual has a switch value of 0, and it engages in a mutation event (i.e., with probability μ_c), then the value of the switch has a 50% chance of turning into 1 and a 50% chance of turning into 2. Figure B.7, shows a diagram representing the logic behind this mutation function.

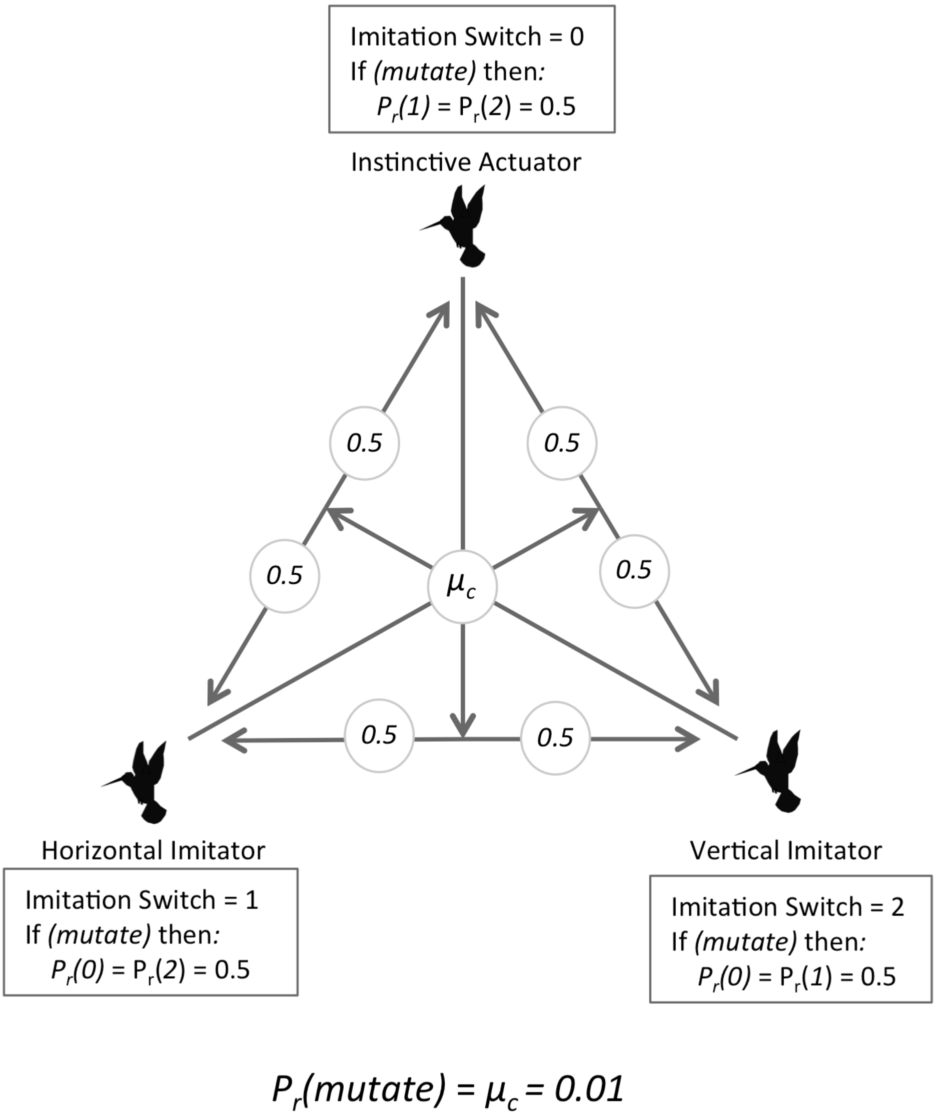

FIGURE B.7: Diagram representing the logic of our three-way mutation function. The new individual/offspring, created in step 2 of Figure B.6, can mutate its imitation switch with probability μ_c. If a mutation occurs, one of the two other values for the imitation switch will be selected, with equal probabilities. For example, if the value of the imitation switch is 2 (i.e., the individual is a vertical imitator), and a mutation event occurs (i.e., with probability μ_c), the switch will mutate into 1 (i.e., the individual will become a horizontal imitator) with a 50% chance, and into 0 (i.e., the individual will become an instinctive actuator) with a 50% chance.

B.2.1.2 Step 3c: If the Offspring is a Vertical Imitator

In step 2 of Figure B.6, we produced a new individual/offspring with a genotype string and an imitation switch, but without a phenotype string. In this step, we will create a phenotype for this individual assuming that its imitation switch is 2 (i.e., the individual is a vertical imitator). When the imitation switch is set to 2, the phenotype of the new individual is a copy of the phenotype string of its own parent (i.e., the individual selected in step 1). This process of phenotype imitation occurs once at birth. The phenotype string is copied with an error rate of μ_v per bit. This means, each bit copied into the phenotype of the new individual will change its value (i.e., changing from 0 to 1, or from 1 to 0) with probability μ_v. Notice that by imitating the parental phenotype, vertical imitators can still produce genotype-phenotype disengagement (see Figure 4.5), as the substrate of evolution, i.e., the phenotype (i.e., the information that is ultimately optimised), can be inherited independently from the genotype. In this extension, we refer to vertical and horizontal imitators as *social learners*. Figure B.8 shows a diagram explaining this step.

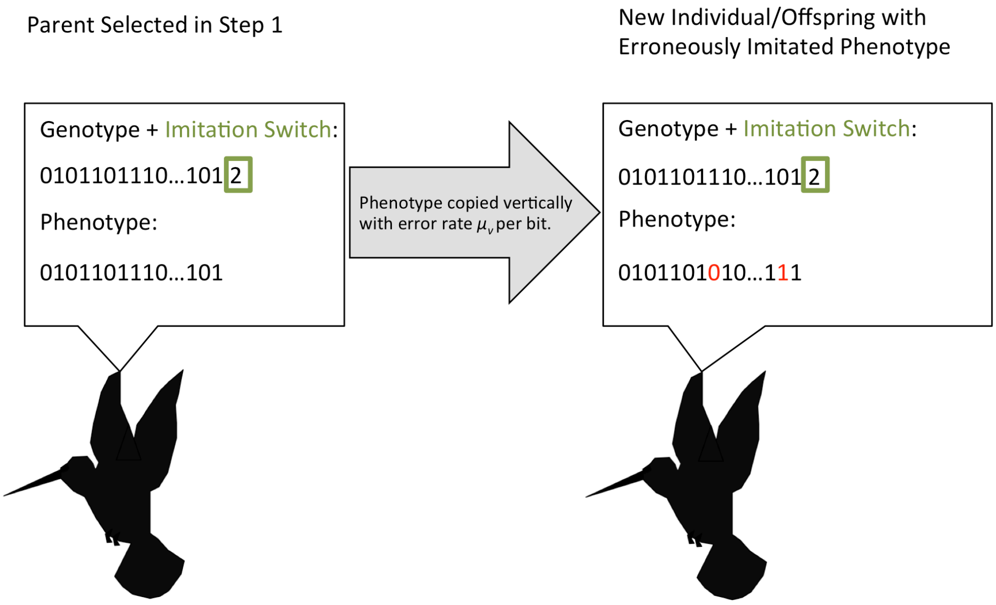

FIGURE B.8: Diagram representing the third step of our simulation model (if the new individual is a vertical imitator). If the value of the imitation switch of the new individual is equal to 2 (i.e., the individual is a vertical imitator), we create its phenotype by copying the entire phenotype string of its own parent (i.e., the individual selected in step 1). Each phenotype bit is copied subject to a mutation rate of μ_v. Notice that the imitation switch is not part of the phenotype. Mutated bits are shown in red.

B.3 Chapter 5: Extension with Environmental Change

As in the original model, each agent is represented by two sequences of bits of length 200 (i.e., $L = 200$) corresponding to phenotype and genotype, and a single extra bit corresponding to the imitation switch. When individuals reproduce (i.e., when asexual reproduction occurs), both genotype and imitation switch are inherited from parent to offspring with a mutation rate per bit of μ_g for the genotype and μ_c for the imitation switch.

Agents born with an imitation switch in position 0 are instinctive actuators and express their genotype into their phenotype with a 1-to-1 map (i.e., their phenotype is always an exact copy of their genotype). When the agent is born with an imitation switch in position 1, it is a social learner, and therefore its phenotype will be an imitated copy of the phenotype of a randomly selected individual in the population. During imitation, errors occur with an error rate of μ_p per bit. In contrast with the model extension in Chapter 4 and similar to the original model in Chapter 3, the imitation switch in this model only has two positions: horizontal imitation (i.e., unbiased social learning) and instinctive actuation. Each individual only expresses one phenotype throughout its lifetime regardless of whether this phenotype is genetically inherited (i.e., instinctively actuated) or imitated from a random individual at birth (i.e., socially learned).

As in the original model, selection is introduced by a death function and a reproduction function. On each iteration, an individual is selected for reproduction with a probability proportional to the fitness of their phenotype, and another individual is selected for death (i.e., to be replaced by the reproducing agent's offspring) with a probability inversely proportional to the fitness of their phenotype. Fitness is calculated using the sum-of-environment-matching-bits (L_m) instead of the sum-of-1s value (L_1) used in the original model in Chapter 3. In the extension explained in this chapter, the optimal sequence (which represents the environment-to-be-matched) changes with a given probability per iteration; each time the sequence changes, all agents in the population have their sum-of-environment-matching-bits (L_m) recalculated.

The strength of these selection processes is controlled by the reproductive selection coefficient x_r and the survival selection coefficient x_d in Boltzmann-weighted functions (Equations B.3 and B.4). These equations calculate the probability for a given phenotype i or j, to be selected for reproduction P_{ri} or death P_{dj} on each iteration, respectively.

$$P_{ri} = \frac{e^{\frac{L_{mi}}{x_r}}}{\sum_{k=1}^{N} e^{\frac{L_{mk}}{x_r}}} \tag{B.3}$$

$$P_{dj} = \frac{e^{\frac{L-L_{mj}}{x_d}}}{\sum_{k=1}^{N} e^{\frac{L-L_{mk}}{x_d}}} \tag{B.4}$$

The optimal sequence (i.e., the sequence used to calculate fitness by counting the number of phenotype bits that match to it) is initialised as a random bit string. All phenotypes and genotypes in the initial population are also initialised at random. All agents start as instinctive actuators (i.e., their genotype and phenotype sequences are identical and their imitation switch value is 0). The optimal sequence changes through the simulation as a function of the parameter δ_e, which has a fixed value for each simulation run. The parameter δ_e represents the probability that *a randomly selected bit* in the optimal sequence, flips from 1 to 0, or from 0 to 1, in an iteration (see Subsection B.2.1). For all results shown in this chapter, we will specify the rate of environmental change using δ_e, and additionally, we will provide the equivalent μ_e value for the given δ_e. The value of μ_e represents the estimated probability for *each bit* in the optimal sequence to change from 1 to 0, or from 0 to 1 (i.e., $\mu_e = \delta_e/L$). The value of μ_e is comparable against mutation rates and imitation error rates (i.e., μ_g and μ_p), which are defined as bit-flip probabilities for *each bit* in a bit string. Equations B.3 and B.4 are similar to Equations B.1 and B.2 from Chapter 3, but apply to any optimal sequence (i.e., not only a sequence of all-1s).

The environmental change function is implemented as an extra step in our simulation model's logic (see Figure B.9). In this step, a randomly selected bit in the current optimal sequence is changed with probability δ_e. When the optimal sequence changes, all fitness values are recalculated immediately afterwards.

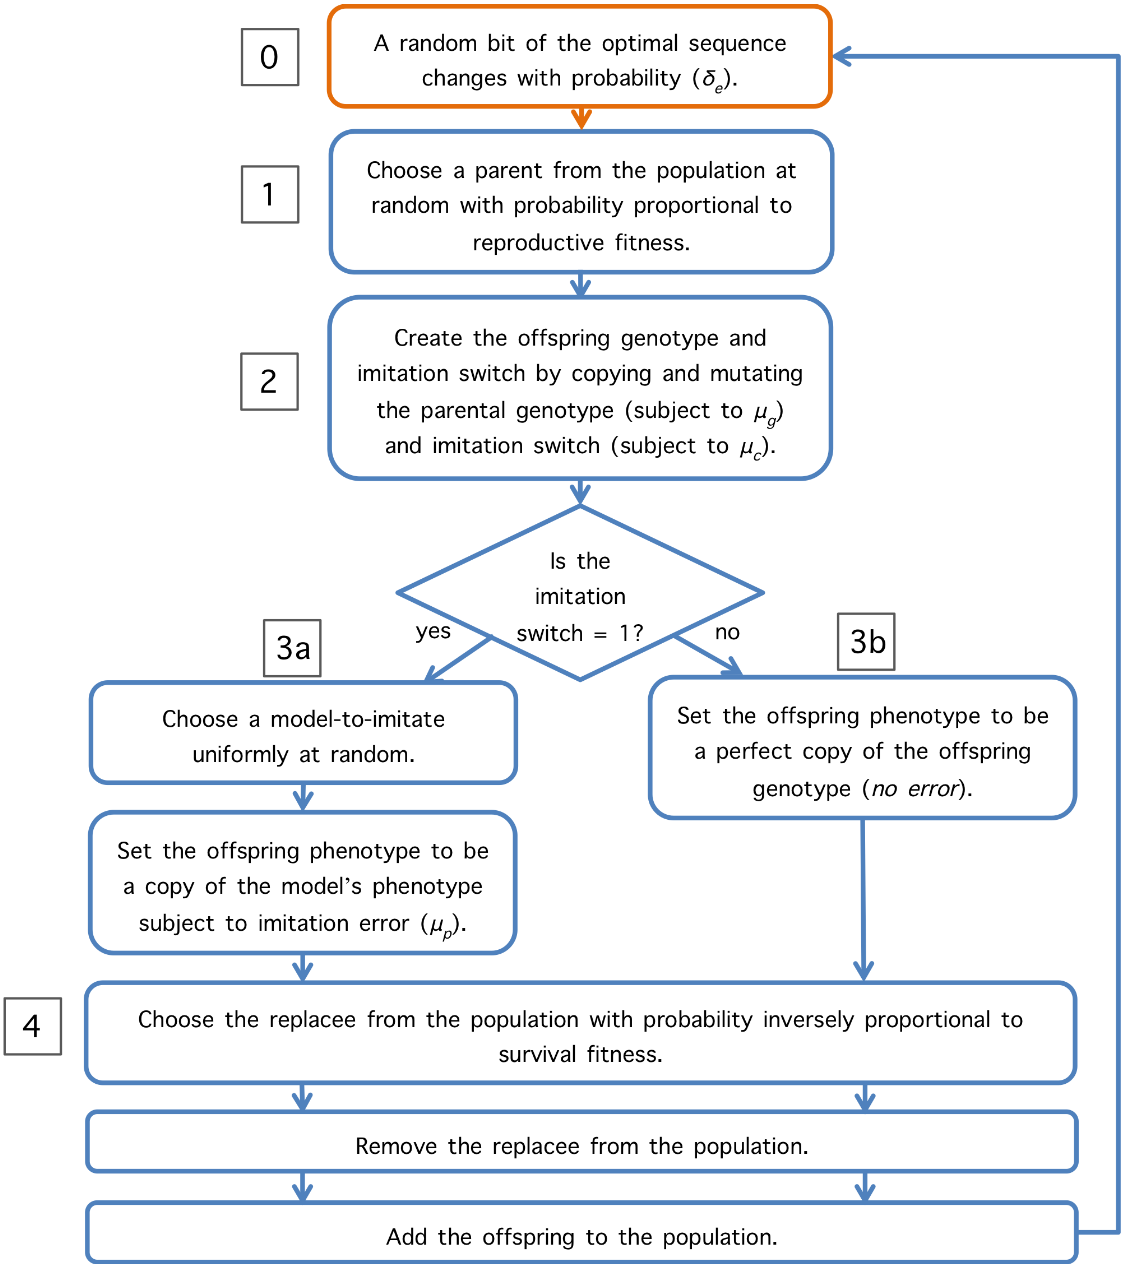

FIGURE B.9: An algorithmic representation of the extended model's logic. (0) With probability δ_e a randomly selected bit in the optimal sequence is flipped (i.e., from 1 to 0, or from 0 to 1). If the optimal sequence changes, all fitness values are recalculated. (1) The selection function picks an individual from the population. (2) Its genotype string is copied including the imitation switch (mutations occur at a rate of μ_g per genotype bit and μ_c for the imitation switch). (3a) If the imitation switch has value 1, a random individual from the population is selected and its phenotype will be copied (imitation errors occur at a rate of μ_p per bit). (3b) If the imitation switch has value 0, the phenotype will be a perfect copy of the individual's genotype. (4) The resulting combination of phenotype and genotype will replace an individual selected by the death function. Steps that have been added in this extension, compared to the basic model in Figure B.1, are coloured in orange.

B.3.1 Simulation Steps

In this subsection, steps presented in Figure B.9, that were not present in Figure B.1, are explained in detail. Once again, we want to mention that the internal logic of our model, before generating our analytical results, was “unit tested”. That is, for each function, the expected output for a range of possible inputs was evaluated.

B.3.1.1 Step 0: Implementation of the Environmental Change Function

In the model extension explained in Chapter 5, we included a rate of environmental change (δ_e). The parameter δ_e represents the probability that a *randomly selected bit* in the optimal sequence (i.e., the sequence that encodes the optimal solution) changes from 1 to 0, or from 0 to 1, in a single iteration. Therefore, the implementation of the environmental change function occurs in two steps: (1) a single bit in the optimal sequence is selected with uniform random probability, and (2) the selected bit changes from 1 to 0, or from 0 to 1, with probability δ_e. Notice that the environmental change function is therefore different from the previously described mutation and imitation error functions. When a genotype string is mutated with rate μ_g, each bit in the string changes from 1 to 0, or from 0 to 1, with probability μ_g. Similarly, when a phenotype string is erroneously imitated with error rate μ_p, each bit in the string changes from 1 to 0, or from 0 to 1, with probability μ_p.

In this work, we have selected the 2-step implementation of the environmental change function due to its superior computational efficiency, compared to an implementation with an independent bit-flip change probability for each bit. Nevertheless, in order to make environmental change rates comparable to our mutation and imitation error rates, we have included the equivalent μ_e rate for every δ_e used in our results. The value of μ_e represents the estimated probability for any given bit in the optimal sequence to change from 1 to 0, or from 0 to 1. This value is calculated as the product of: (1) the probability that a given bit would be selected during uniform random selection (i.e., $1/L$), where L is the length of the optimal sequence, and (2) the probability that the selected bit changes its value (i.e., δ_e). Therefore, $\mu_e = \delta_e/L$. Figure B.10 shows a diagram explaining the implementation of the environmental change function.

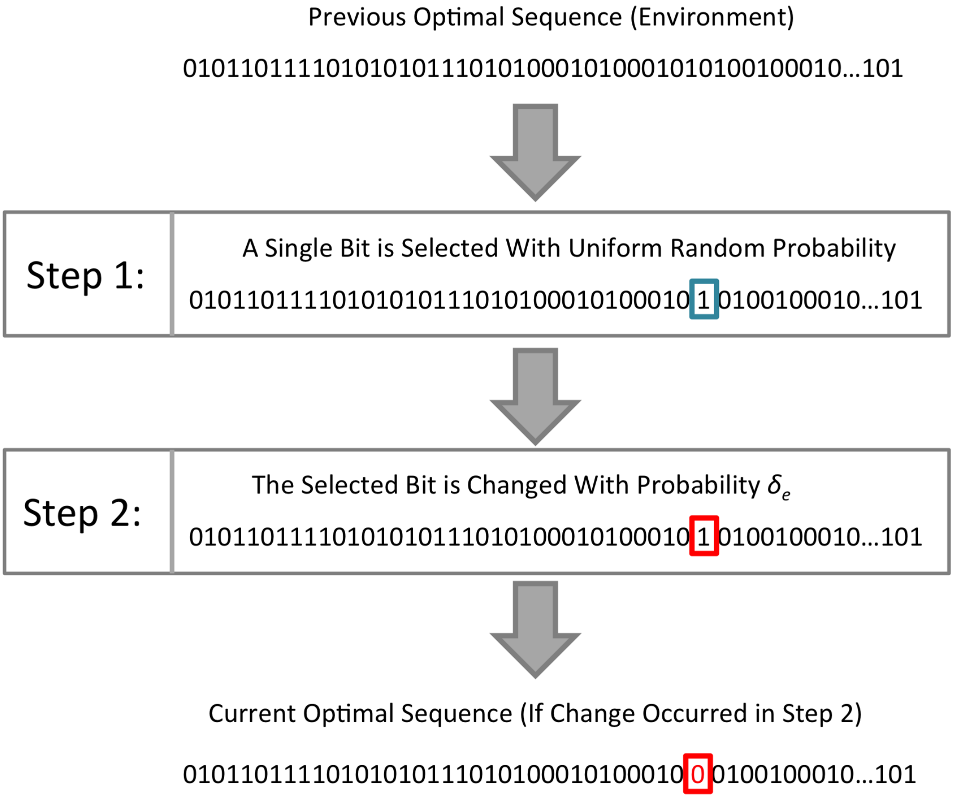

Figure B.10: Implementation of the environmental change function. Step 1: with uniform random probability, a single bit in the optimal sequence is selected. Step 2: the bit selected in step 1, changes from 1 to 0, or from 0 to 1, with probability δ_e. For each iteration, the resulting sequence is used to calculate reproduction and survival fitness.

B.3.1.2 Fitness Calculations in a Changing Environment

The incorporation of an environmental change function in our model, forces us to frequently recalculate *reproductive selection scores* and *death selection scores*. Every time a bit changes in the optimal sequence, the scores of the entire population have to be recalculated in order to represent the appropriate fitness values for the current environmental state. In the original model, where the optimal sequence was fixed as a string of all-1s, reproductive selection scores are calculated using the sum-of-1s values of the phenotype strings (L_1) and the death selection scores are calculated using the complement of the sum-of-1s values ($L - L_1$) (see Equations B.1 and B.2). In this extension, the reproductive selection scores are calculated using the number of phenotypic bits that match their respective positions in the optimal sequence (L_m). Conversely, the death selection scores are calculated using the number of phenotypic bits that do not match their respective positions in the optimal sequence ($L - L_m$) (see Equations B.3 and B.4).

Appendix C

Preliminary Analyses, Simulation Environment and Parameter Setup

In this appendix, we describe a series of preliminary analyses that were used to setup our simulation environment, test basic assumptions of our model, and justify the selection of parameter values used in the results sections of Chapter 3, 4 and 5. Each section in this appendix describes an specific analysis. For consistency, we have used fragments of the text already contained in the body of the thesis, to describe the context in which each analysis is relevant. Note that all the parameter setups used in this thesis were selected, ultimately, *with a focus on the scope of our research questions.* These questions are concerned, mostly, with the *possibility* of culture emerging in a evolutionary system without decision-making processes. Therefore, the specific parameter values/ranges in this work (e.g., μ, x_d and x_r values), were selected because they produce results with direct relevance to our questions.

C.1 Selecting the Range of Mutation and Imitation Error Rates (μ_g, μ_p, μ_h and μ_v)

In our simulation models, two evolutionary systems, cultural and genetic, compete in the process of optimising a string of bits. In both system, the rate of mutation is a key parameter that can determine how rapidly solutions are found, and with how much fidelity they are maintained. During a preliminary assessment, we considered a wide range of variation rates (see Figures C.1 and C.2), but eventually decided to focus on a range between $\mu = 0.0$ and $\mu = 0.005$. In this context, μ is the probability that each bit in the information-encoding strings changes from 1 to 0, or from 0 to 1. During the preliminary assessment, we used a generic evolutionary system in order to evaluate the effect of variation rates (i.e., rates of mutation or imitation error) in isolation (i.e., without including the complex processes of the dual-inheritance system in our simulation model).

We selected the range $\mu = [0.0, 0.005]$ taking into consideration the following: (1) the explored rates of mutation and imitation error must be low enough for the evolutionary system to achieve long-term maintenance of high a proportion of optimal phenotypes. High values of μ, which maintain the phenotypes of the population at fitness values considerably below the optimal sequence, are not considered in our analyses (see plot A in Figure C.1, which shows a simulation run with a μ value that prevents phenotypes from reaching the optimal sequence). The highest value in our explored range (i.e., $\mu = 0.005$), does achieve long-term maintenance of a high-proportion of optimal phenotypes (see plot B of Figure C.1). (2) For all rates explored, the evolutionary system should be able to reach the optimal sequence before 5×10^4 iterations, with the obvious exception of $\mu = 0.0$ (see plot C in Figure C.1, which shows the smallest non-zero value of μ explored in our results: $\mu = 1 \times 10^-4$). (3) The explored range should focus on combinations of mutation and imitation error rates for which substantial differences, in the proportion of simulations that converge to culture, are observed. In Figure C.2, we show how the range $[0.0, 0.005]$ was selected after a preliminary assessment of a larger range (i.e, $\mu = [0.0, 0.015]$).

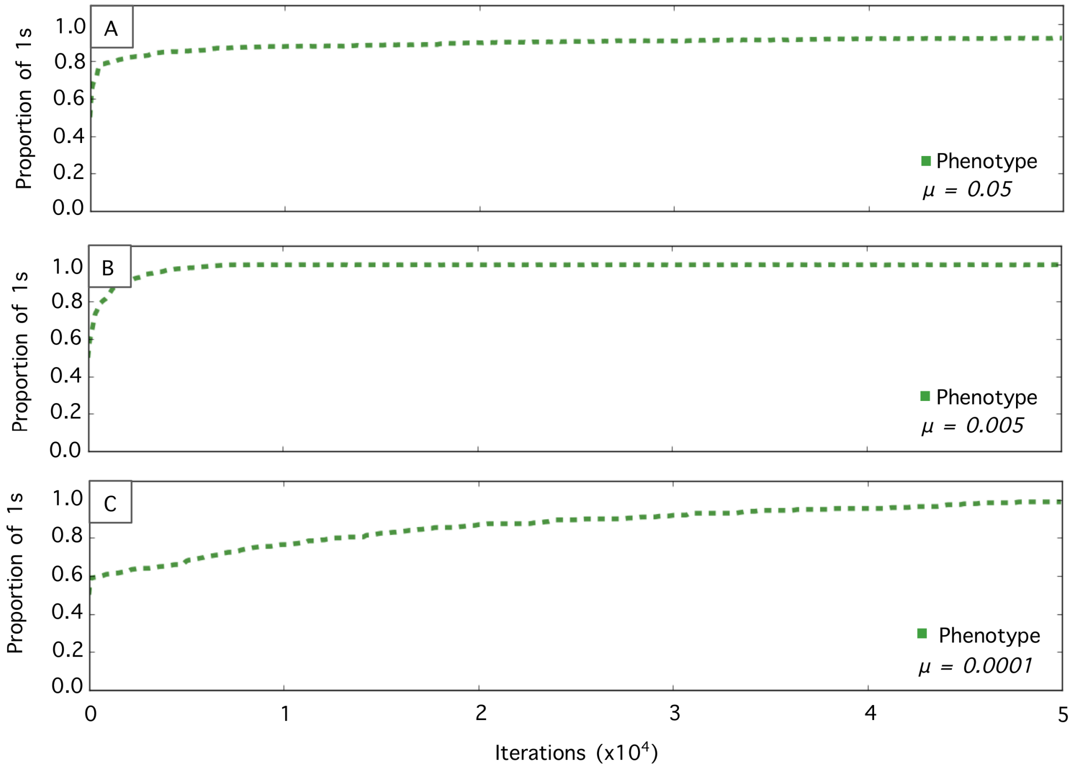

Figure C.1: Simulation runs for a standard genetic evolutionary system (without social learners). Each plot depicts the change over evolutionary time in the average proportion of 1s in the phenotypes (green line). Plot A shows results for a simulation with mutation rate $\mu = 0.05$. In this plot, the phenotype sequences cannot reach optima due to the lack of fidelity in genetic replication. Plot B shows results for a simulation with mutation rate $\mu = 0.005$. In this plot, the phenotype sequences are maintained close to the optima. Plot C shows results for a simulation with mutation rate: $\mu = 1 \times 10^-4$, in this simulation the phenotypes reach the optima after 5×10^4 iterations. In all plots, the strength of reproductive and survival selection is fixed at $x_d = 1$ and $x_r = 1$.

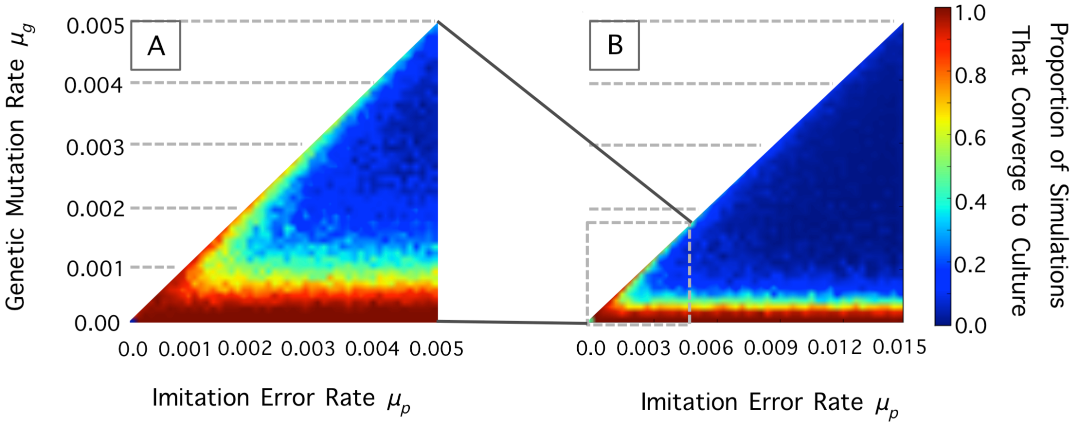

FIGURE C.2: Two heat maps depicting how the tendency for social learners to evolve varies with the genetic mutation rate, μ_g, and imitation error rate, μ_p. For all points in these plots, $\mu_p > \mu_g$. The colour of each pixel represents the proportion of replicates that converge to culture for simulations with parameterisations that corresponds to the "co-ordinates" in which the pixel is located. Map A shows a subset of the range that was initially explored in our preliminary analyses. This range subset was used for all analyses in this thesis. Map B shows the entire range explored during preliminary analyses. We selected the $\mu = [0.0, 0.005]$ range due to its high concentration of simulations that converge to culture. In all simulations, the strength of reproductive and survival selection is fixed at $x_d = 1$ and $x_r = 10$.

C.2 Selecting the Mutation Rate for the Imitation Switch (μ_c)

We chose to fix $\mu_c = 0.01$ to incentivise the formation of *chains of imitation* (i.e., sequential imitation among imitators) in the early stages of our simulation runs. For lower values of μ_c, the density of social learners, produced by mutation alone (i.e., before any form of selection acts on the population), is too low for imitation among imitators to occur with high enough frequency as to kick-start an evolutionary system (i.e., culture). On the other hand, values of μ_c larger than those minimally required to start chains of imitation, will produce a large initial density of social learners without the intervention of any evolutionary process. For obvious reasons, a large initial density of social learners would make our model unsuitable to describe the emergence of culture, as this process is, precisely, the outcome of the evolution of social learning. The value of $\mu_c = 0.01$ was selected as a compromise between these two constraints.

When $\mu_c = 0.01$, on average, once in one hundred iterations, an instinctive actuator turns into an imitator or vice versa. Analyses of our model show that a critical density of imitators is required to start chains of imitation (i.e., imitators copying imitators).

In other words, sequential events of high-error transmission are required to start a non-genetic evolutionary system that is competitive compared to the genetic one. The lower the chance of these chains forming, the less likely it is for social learners to fix, even under otherwise favourable conditions. Fixing the mutation rate for the imitation switch ensures that a constant "background" density of imitators is created by mutation alone, independent of the fidelity of both genetic and social transmission, and that this density is resistant to stochastic fluctuations throughout the simulation. In Figure C.3, we compare two individual simulation runs: in plot A, μ_c is equal to μ_g, which is equal to 0.0005. In plot B, $\mu_c = 0.01$. In plot A, genotype-phenotype disengagement does not occur, and the proportion of social learners in the population remains close to zero throughout the entire simulation run. In plot B, culture emerges during the initial stages of the simulation run. These results are in agreement with our explanation, and show how the emergence of culture is hampered by a low density of social learners in the initial stages of the simulation run.

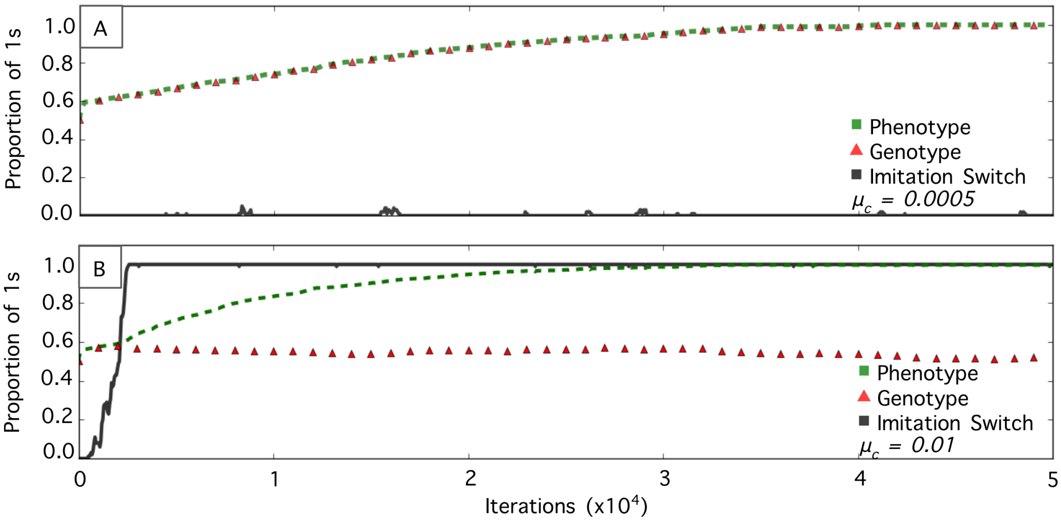

Figure C.3: Comparison of two simulation runs with different values of μ_c. Each plot depicts the change over evolutionary time in the average proportion of 1s in the population's phenotypes (green) and genotypes (red), and the proportion of imitators in the population (black). In plot A, $\mu_c = 0.0005$. In this plot, genotype-phenotype disengagement does not occur, and the proportion of social learners remains close to zero throughout the simulation run. In plot B, $\mu_c = 0.01$. In this plot, genotype-phenotype disengagement occurs early in the simulation run. These runs show how a small background-density of social learners can be crucial for the occurrence of genotype-phenotype disengagement. In both simulation runs, the strength of reproductive and survival selection is fixed at $x_d = 1$ and $x_r = 10$. Also, $\mu_g = 0.0005$ and $\mu_p = 0.005$.

C.3 Selecting the Population Size (N)

When selecting the population size of our simulation setup, we took into consideration the following factors: (1) the population needs to be large enough for critical components of our simulation model to operate, e.g., social learning and horizontal imitation cannot realistically operate in a "population" of a single individual, (2) the population size needs to be sufficiently small, so that in less than a second, a single simulation run could be set up, and the phenotype strings could reach the optimal sequence. Simulations with large populations take longer to reach optima, and (3) the size of the population should be similar, or at least have the same order of magnitude, to the size of natural populations in species that support basic forms of culture. Our selected population size of 100 individuals, falls within the estimates for group sizes of existing social primates and extinct hominids (Markham et al., 2015; Dunbar, 1993). Note that optimal and effective group sizes in natural populations are the result of numerous evolutionary and environmental processes (Dunbar, 1993). None of these processes are considered in this thesis. However, for the scope of our research questions, we consider $N = 100$ as an adequate value to demonstrate the processes described in Chapters 3, 4 and 5.

Our results can be reproduced with simulations that use larger population sizes (see Figure C.4). However, simulations with larger populations, require longer computational times to optimise the phenotype sequences. For example: a simulation with $N = 100$ takes 0.4 seconds to reach optima, in contrast, a simulation with $N = 1000$ takes 16.5 seconds. For results that require a single run, this difference is not significant. However, to produce arrays of maps like those shown in Figure 3.9 and 5.4, more than two million simulations runs are required. Considering that: (1) the computational times required to produce these maps are too large for repeated experimentation and testing, and (2) there are not differences between the results produced with $N = 100$ and the results produced with larger population sizes, we selected $N = 100$ for all our simulation runs in this thesis.

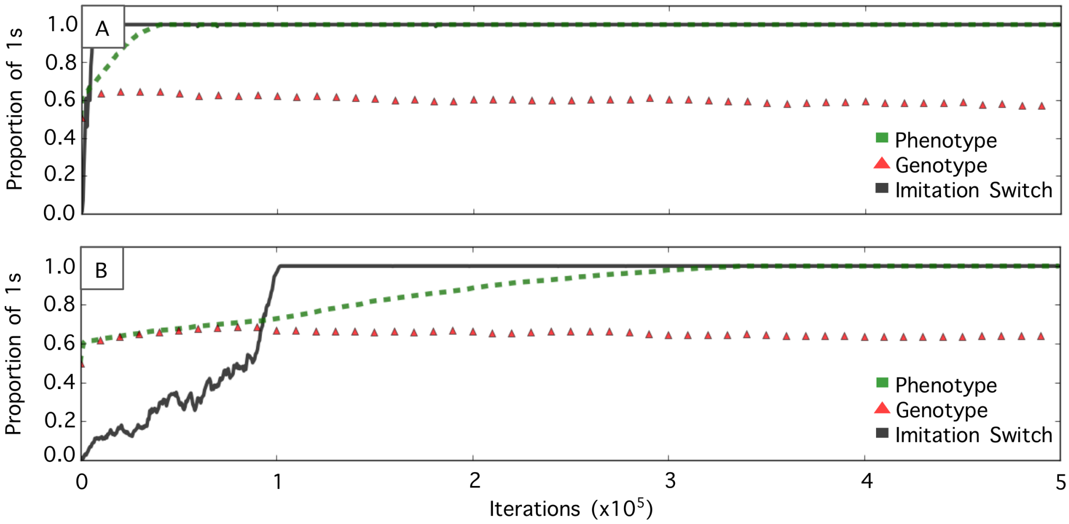

Figure C.4: Simulation runs for our basic model in Chapter 3. Each plot depicts the change over evolutionary time in the average proportion of 1s in the population's phenotypes (green) and genotypes (red), and the proportion of imitators in the population (black). Plot A shows results for a simulation with a population of 100 individuals. Plot B shows similar results for a simulation with a population of 1000 individuals. The simulation with 1000 individuals takes 3.5×10^5 iterations to reach the optimal sequence, whereas the simulation with 100 individuals only takes 4×10^4 iterations. For both runs in this figure: $\mu_g = 0.0001$ and $\mu_p = 0.0005$. The strength of reproductive and survival selection is fixed at $x_d = 1$ and $x_r = 10$.

C.4 Selecting the Length of the Bit Sequences (L)

In our simulation model, the emergence of culture is intrinsically entangled with the process of genotype-phenotype disengagement. For this process to occur, the optimal sequence needs to be sufficiently large; otherwise, the process of genotype-phenotype disengagement could be reverted. Note that in a disengaged population, genotypes are not subject to selection; therefore, they tend to encode sequences with lower fitness values than the population mean. However, when L (i.e., the bit-string length of phenotypes and genotypes) is small, some genotypes under no selective pressure can encode relatively high-fitness sequences, simply as a result of neutral drift. If this situation occurs, and the high-fitness sequence is expressed in the phenotype of an instinctive actuator, the process of genotype-phenotype disengagement could be reverted. The likelihood of this reversion decreases for large values of L.

Furthermore, to the extent to which L is small, the likelihood of the genetic evolutionary system finding the optimal solution, before social learners invade, increases. Note that for small values of L, the number of bit-flip mutations required for the initialised sequences to evolve into the optimal sequence, is also small. As discussed in this thesis, if the optimal sequence is evolved through genetic evolution, culture will not emerge.

After a preliminary analysis (see Figure C.5), we decided to use an L value of 200. This value is well within the range of L values for which genotype-phenotype disengagement is practically irreversible. Most of our simulations setups start with randomly assigned bits; therefore, at the start of the simulation, it is expected that, on average, 50% of the bits will already match the optimal sequence. Therefore, selecting a length of 200 bits, leaves, on average, 100 bits for the cultural and genetic evolutionary systems to evolve/optimise.

In Figure C.5, we present results for our preliminary analyses. Plot A shows a standard simulation run where irreversible genotype-phenotype disengagement occurs in a simulation where $L = 50$. Plot B shows a simulation run where $L = 6$. In this run, instinctive actuators reinvade after disengagement has occurred. In plot C, we show the tendency for genotype-phenotype disengagement to persist, for simulations with different values of L (black line). Also in plot C, the average proportion of 1s in the phenotypes and genotypes of the populations, after 5×10^4 iterations, for 100 replicate simulations, is shown in green and red, respectively. This plot is in agreement with our explanation, and demonstrates that for increasing values of L, the likelihood for the irreversible fixation of social learners also increases.

The results in figure C.5 also serve to support the claim that an extended strategy space is a key element of our simulation model, and that results shown in this thesis could not be obtained with the classical two-state strategy space approach (i.e., where $L = 1$) Cavalli-Sforza and Feldman (1983); Rogers (1988); Richerson and Boyd (2006).

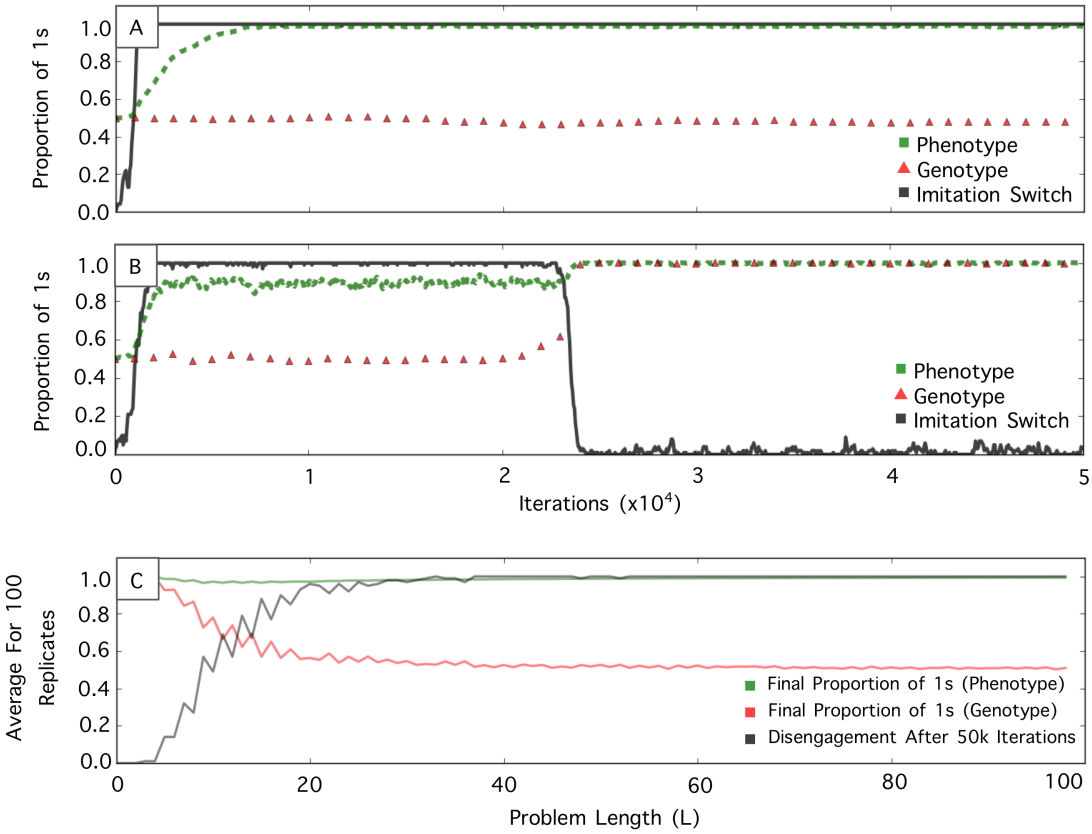

Figure C.5: The effect of problem length, L, on the likelihood of persistent genotype-phenotype disengagement is shown using two representative simulation runs (plots A and B) and a summary of multiple runs (plot C). Plot A shows the result of a simulation run with sequence length $L = 50$, where irreversible genotype-phenotype disengagement occurs during the early stages of optimisation. Plot B shows the results for a representative simulation run in which the length of the problem size is $L = 6$. For this value of L, genotype-phenotype disengagement is reversible. Plot C shows the tendency for genotype-phenotype disengagement to persist (i.e., it shows the number of simulations where the proportion of social learners is above 0.95 after 5×10^4 iterations), in 100 replicate runs, for each value of L between 1 and 100 (black line). In plot C, we also show the average proportion of 1s for the phenotype and genotype after 5×10^4 iterations, for 100 simulation runs, for each value of L between 1 and 100 (green and red lines, respectively). For all simulations shown in this figure $\mu_g = 1/(100L)$ and $\mu_p = 1/L$. The strength of reproductive and survival selection is fixed at $x_d = 1$ and $x_r = 10$.

C.5 Selecting the Range of Reproductive Coefficients (x_r) and Survival Coefficients (x_d)

During the initial development stages of our model, we explored a range of x_d and x_r values, larger than that presented in our results (i.e., $x_d = [1, 5, 10]$ and $x_r = [1, 5, 10]$).

Specifically, the wider range included cases for which $x_d = 100$ and $x_r = 100$. Considering the scope of our research questions in Chapters 3 and 5, we decided to discard the results of simulations with $x_d = 100$ and $x_r = 100$, as these were not substantially different to those produced by simulations in which $x_d = 10$ and $x_r = 10$.

Note that for simulations with $x_d = 100$, the effect of survival selection is relatively negligible compared to the effect of reproductive selection (for cases where $x_r = 1$). As x_d approaches infinity (∞), the death selection function effectively becomes a random selection function (see Equation 3.2), in which each individual has the same probability of being selected for death, regardless of its phenotype. As mentioned in the previous paragraph, results produced with $x_d = 100$ were not substantially different from those produced with $x_d = 10$. Therefore, within the scope of our work, further attenuation of the strength of selection, by making x_d larger than 10, was not considered necessary. This preliminary assessment indicated that a value of $x_d = 10$ was already large enough to make the effect of survival selection negligible. Conversely, the same logic was used when restricting the max value of the reproductive selection coefficient (x_r). An exception to this restriction was made in Chapter 4. The argument for this exemption will be explained in the next section.

C.6 Setting $x_r = 1000$ to Effectively Eliminate Reproductive Selection

In Chapter 4 we showed results for simulations in which $x_r = 1000$ (e.g., Figure 4.6, 4.7, 4.8). This value of x_r reduces the reproductive component of selection, compared to simulations in Chapter 3 and 5, where the largest explored value of x_r is 10. In Chapter 4, the scope of our research question is focused on the potential advantage of horizontal imitation over vertical imitation, if any. As explained in Subsection 4.2.2, to understand this advantage, it is important to remove the effect of reproductive selection, almost entirely. This is achieved by increasing the value of x_r.

In contrast to models explored in Chapter 3 and 5, the model explored in Chapter 4 produces recurrent invasions of vertical and horizontal imitators (see Figure 4.6). In simulations with both types of imitators, the dominance of one type over the other is highly sensitive to the relative strengths of survival and reproductive selection. For this reason, our criteria to select x_r is different in Chapter 4, compared to Chapter 3 and 5.

As explained in Section 3.2 and Subsection 4.2.2, vertical imitators are affected by both kinds of selections: survival selection and reproductive selection; whereas, horizontal imitators are only affected by survival selection. Therefore, we set $x_r = 1000$ in order to evaluate the relative dominance of vertical and horizontal imitators, in a simulation environment in which both types of imitators are under the same selective pressure. To guarantee that this value of x_r was high enough to remove the effect of reproductive selection, we tested increasingly higher values of x_r until the response observed in our results stopped changing.

Appendix D

Extended Results for Chapter Five

The following eight heat map arrays show additional results for Chapter 5. In each array, the rate of environmental change δ_e is different. During this analysis, increasingly high rates of environmental change were explored. The selected range shown in Chapter 5, contains the most contrasting results. Rates of environmental change higher than $\delta_e = 0.4$ do not produce different results in terms of the proportion of simulations that converge to culture, compared to results for $\delta_e = 0.4$.

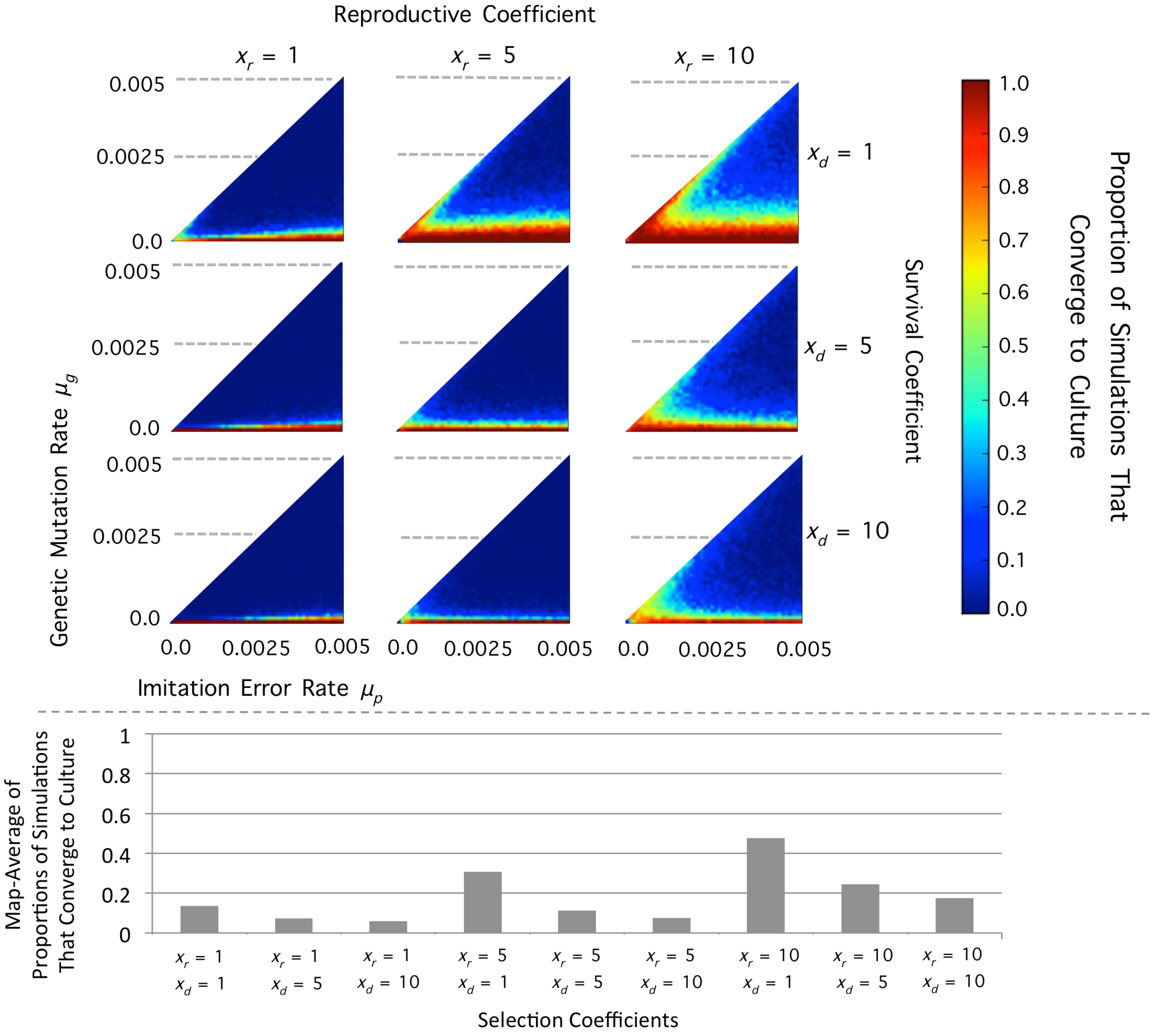

FIGURE D.1: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right columns of maps have x_r values of 1, 5 and 10. From top to bottom rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs, the rate of environmental change is set to $\delta_e = 0.0$ (i.e., $\mu_e = 0.0$).

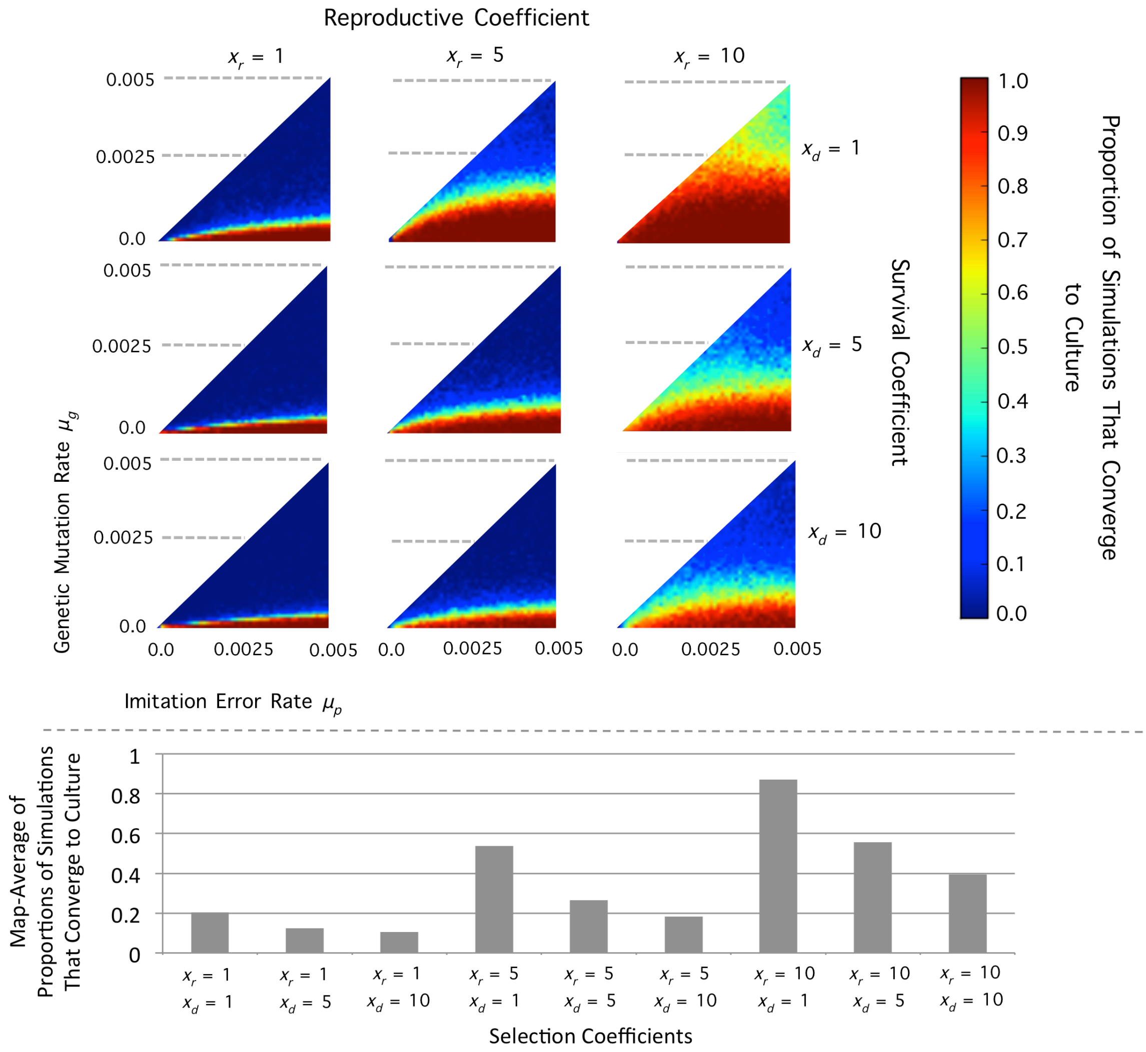

FIGURE D.2: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right columns of maps have x_r values of 1, 5 and 10. From top to bottom rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs in this figure, the rate of environmental change is set to $\delta_e = 0.02$ (i.e., $\mu_e = 1 \times 10^{-4}$); very mild environmental change.

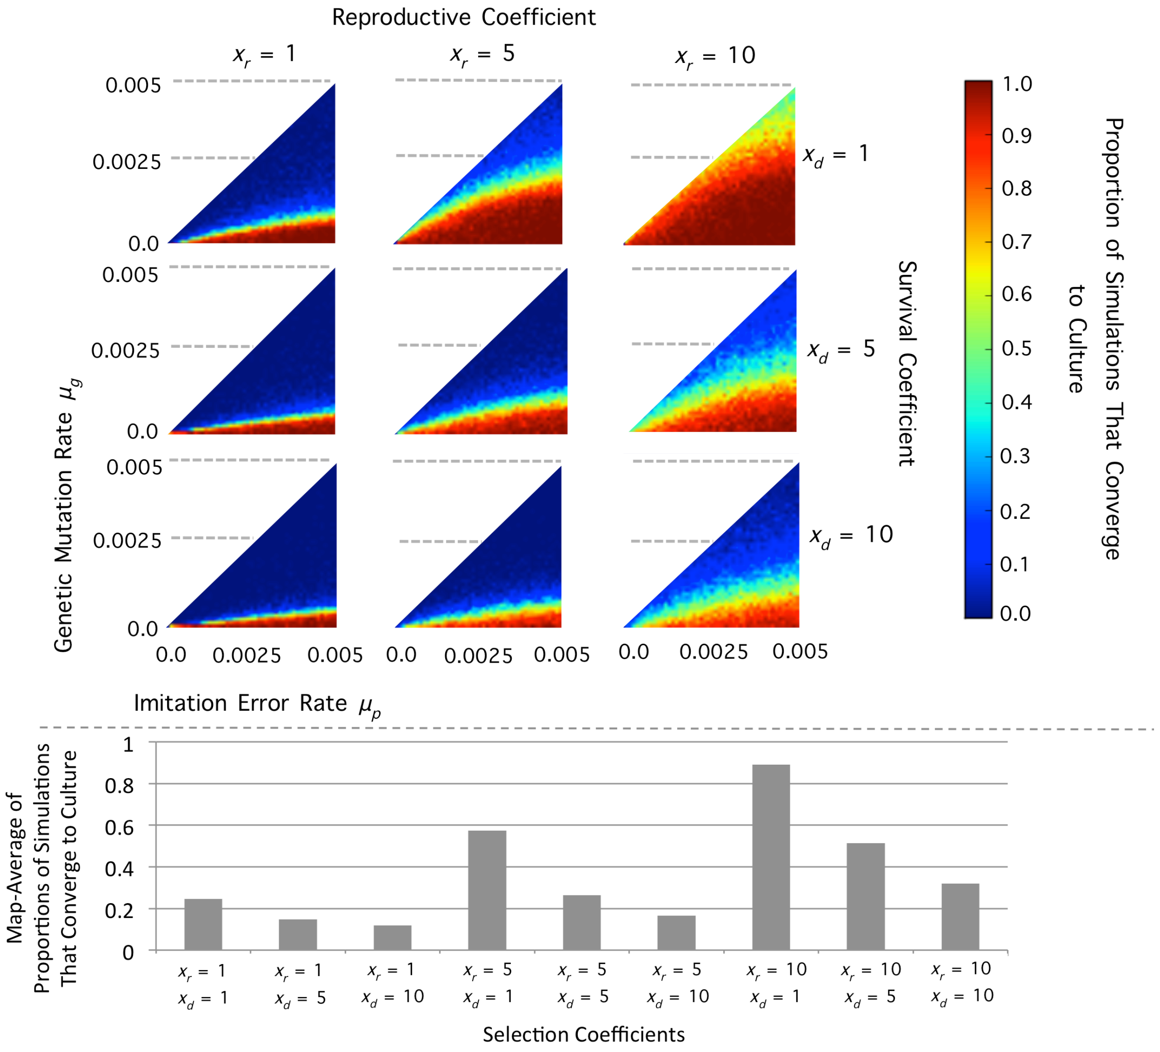

FIGURE D.3: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right columns of maps have x_r values of 1, 5 and 10. From top to bottom rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs in this figure, the rate of environmental change is set to $\delta_e = 0.04$ (i.e., $\mu_e = 2 \times 10^{-4}$); moderate environmental change.

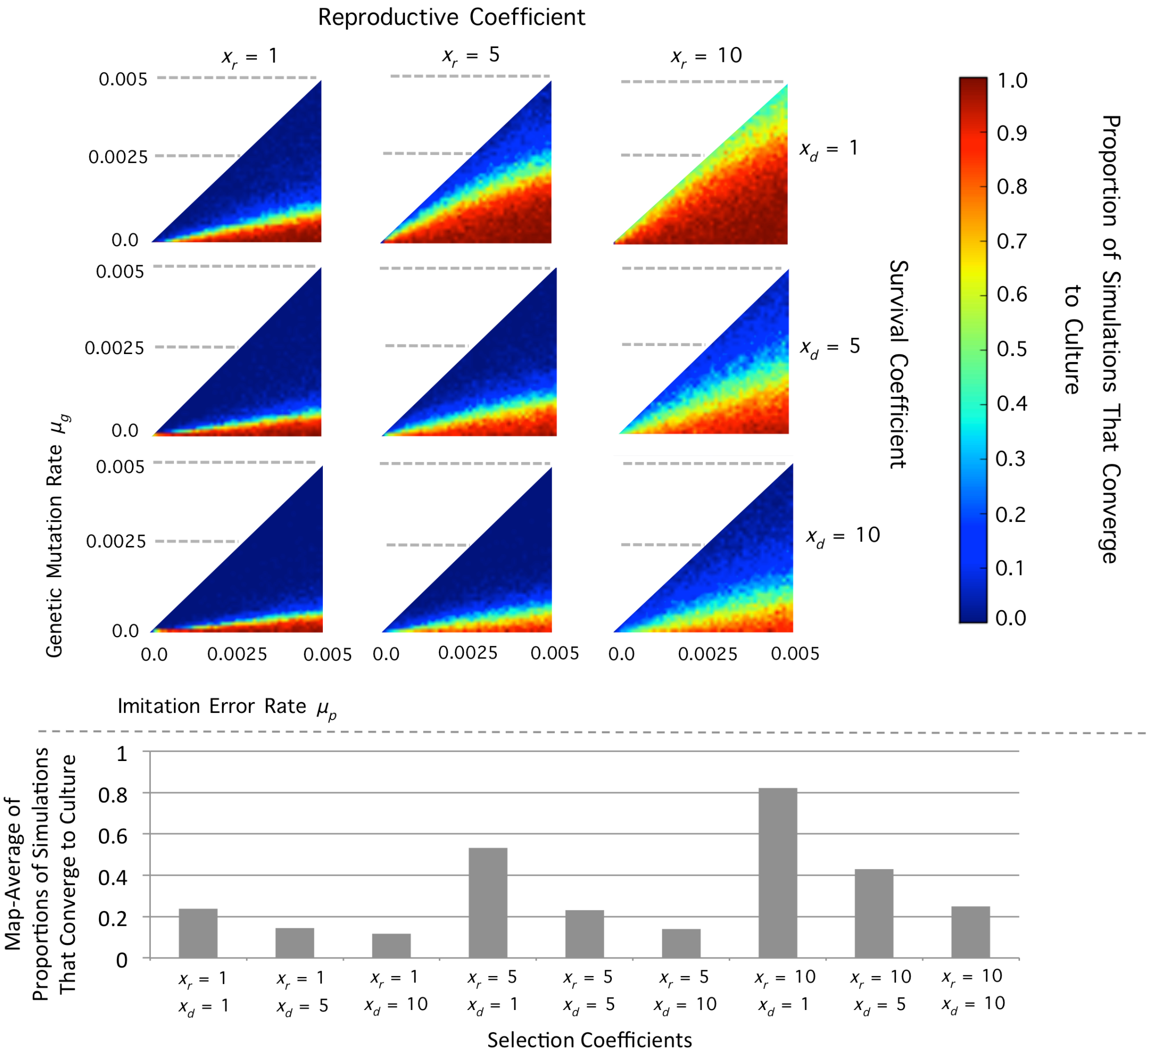

FIGURE D.4: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right columns of maps have x_r values of 1, 5 and 10. From top to bottom rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs in this figure, the rate of environmental change is set to $\delta_e = 0.06$ (i.e., $\mu_e = 3 \times 10^{-4}$); moderate environmental change.

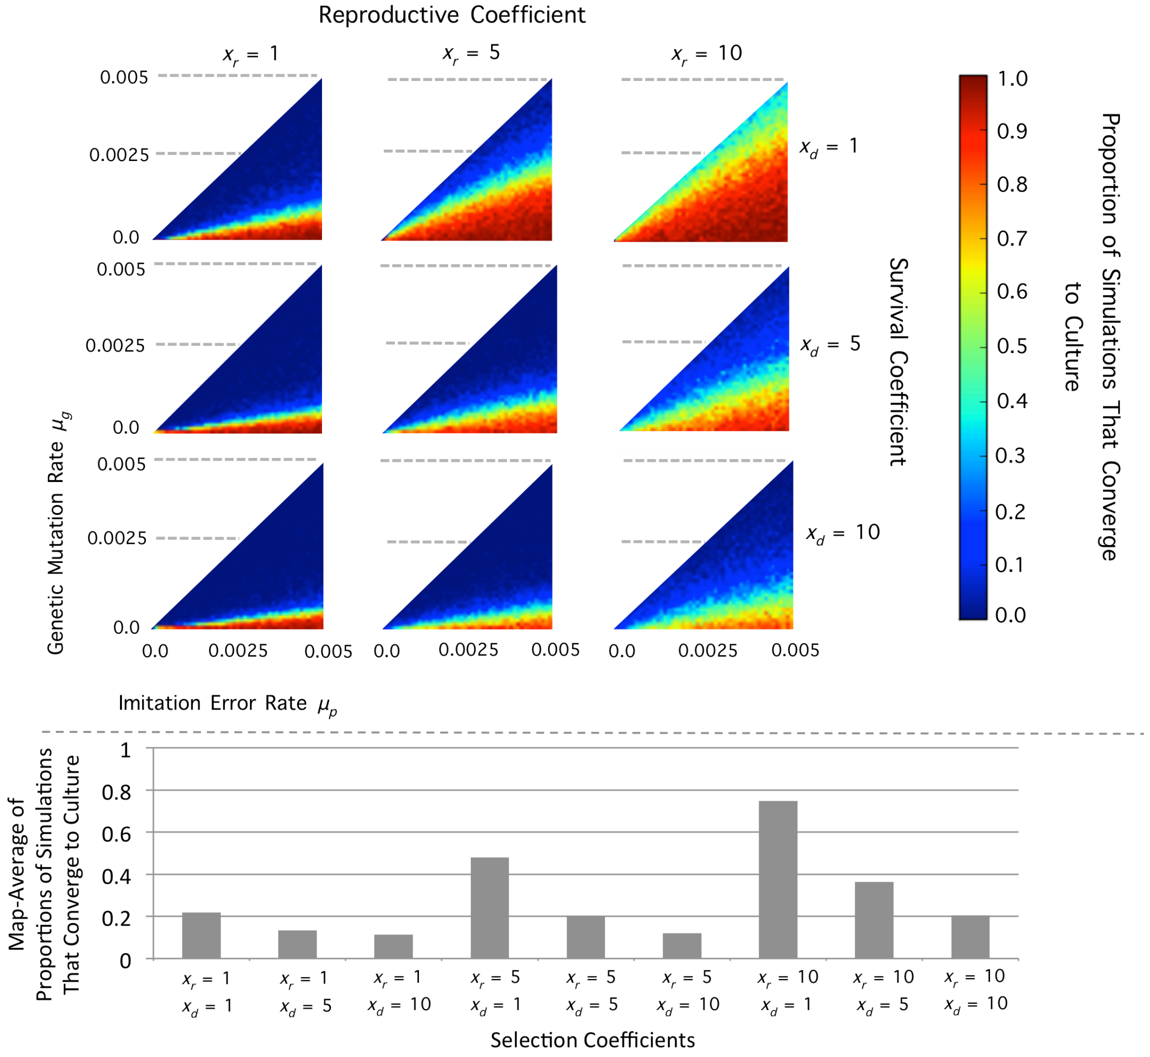

FIGURE D.5: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right columns of maps have x_r values of 1, 5 and 10. From top to bottom rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs in this figure, the rate of environmental change is set to $\delta_e = 0.08$ (i.e., $\mu_e = 4 \times 10^{-4}$); strong environmental change.

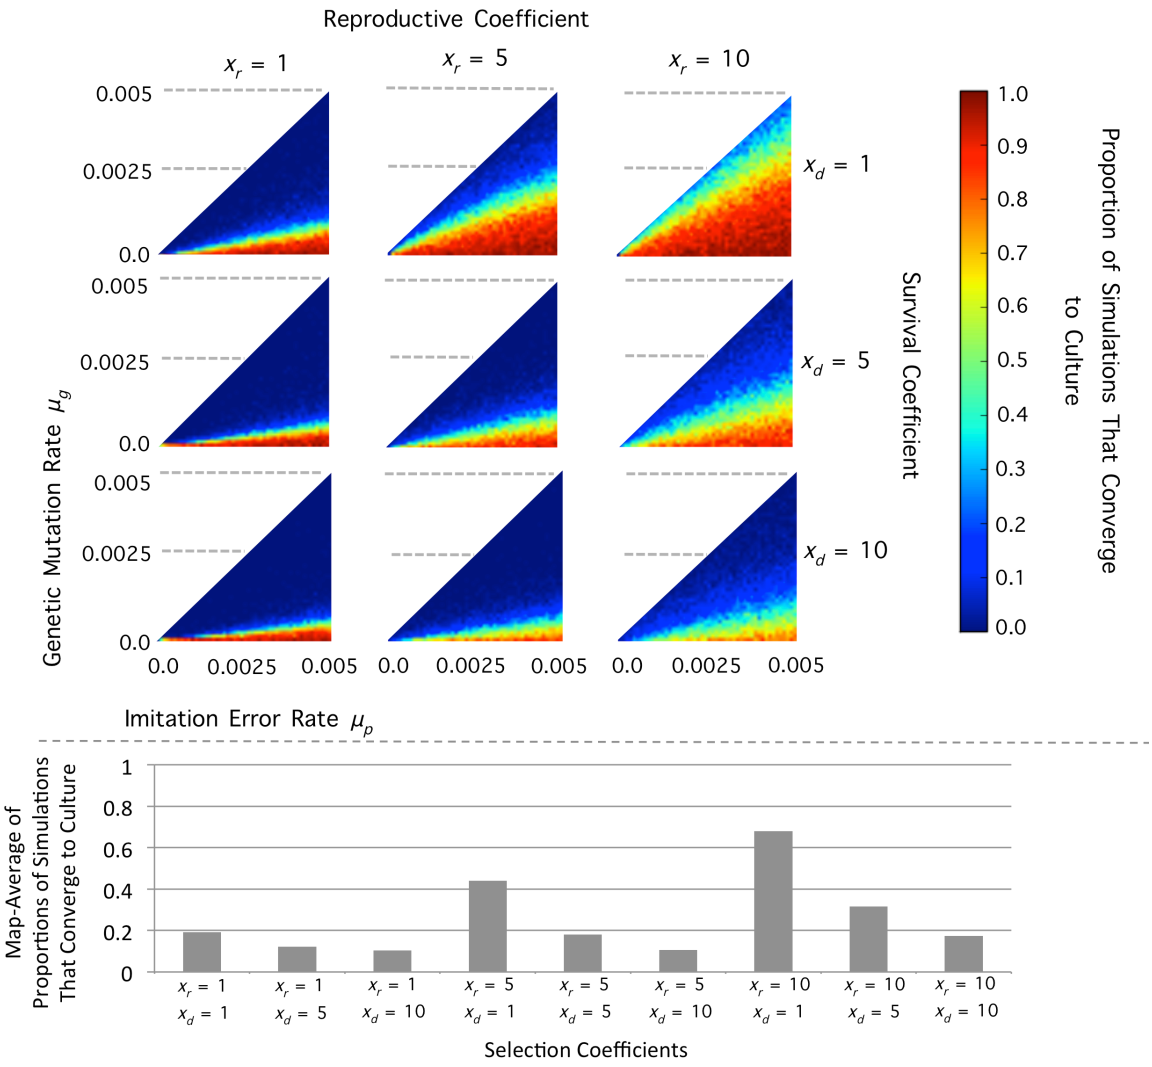

FIGURE D.6: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right columns of maps have x_r values of 1, 5 and 10. From top to bottom rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs in this figure, the rate of environmental change is set to $\delta_e = 0.1$ (i.e., $\mu_e = 5 \times 10^{-4}$); very strong environmental change.

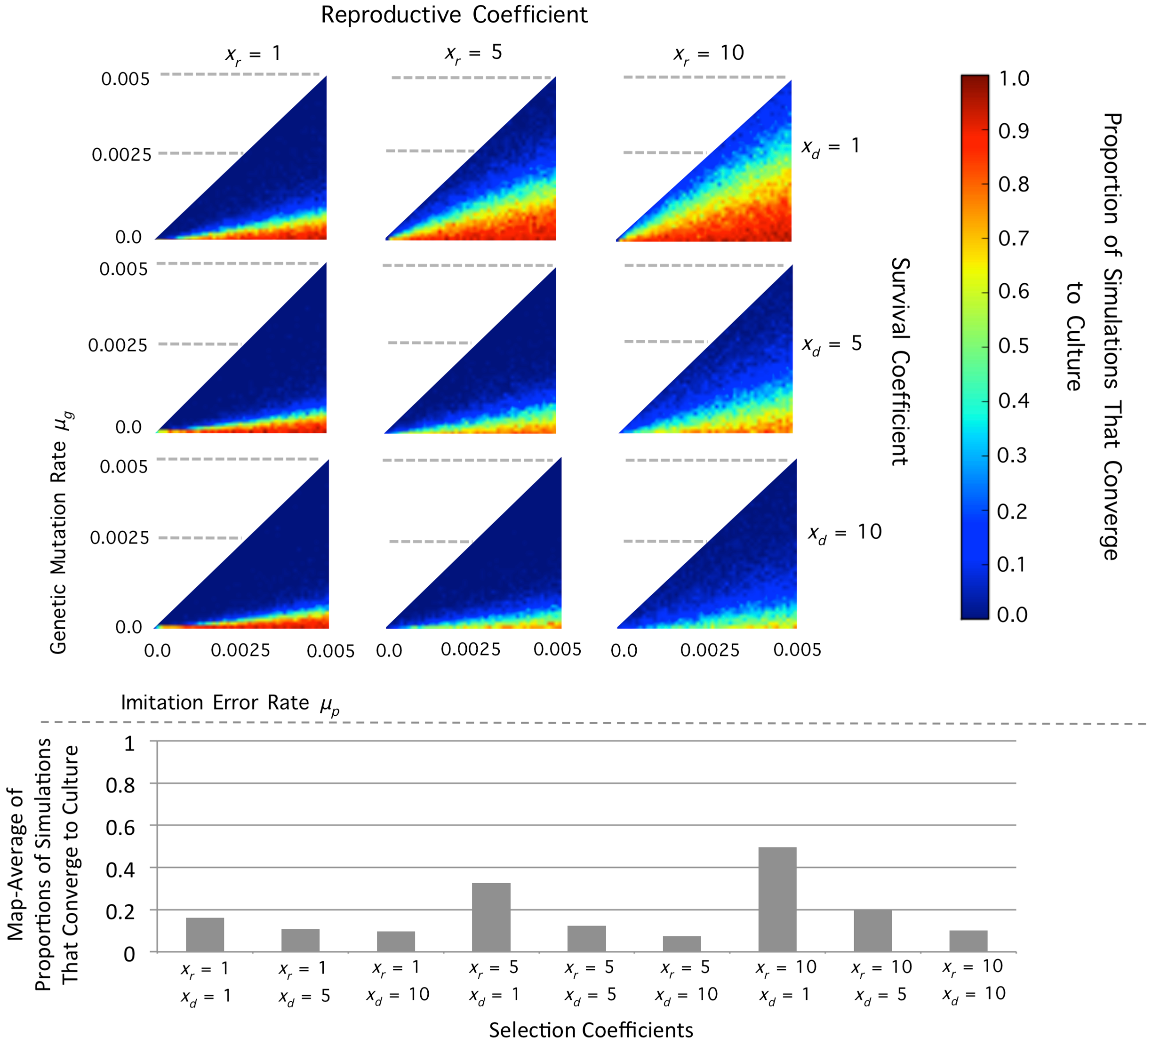

FIGURE D.7: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right columns of maps have x_r values of 1, 5 and 10. From top to bottom rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs in this figure, the rate of environmental change is set to $\delta_e = 0.2$ (i.e., $\mu_e = 1 \times 10^{-3}$); extreme environmental change.

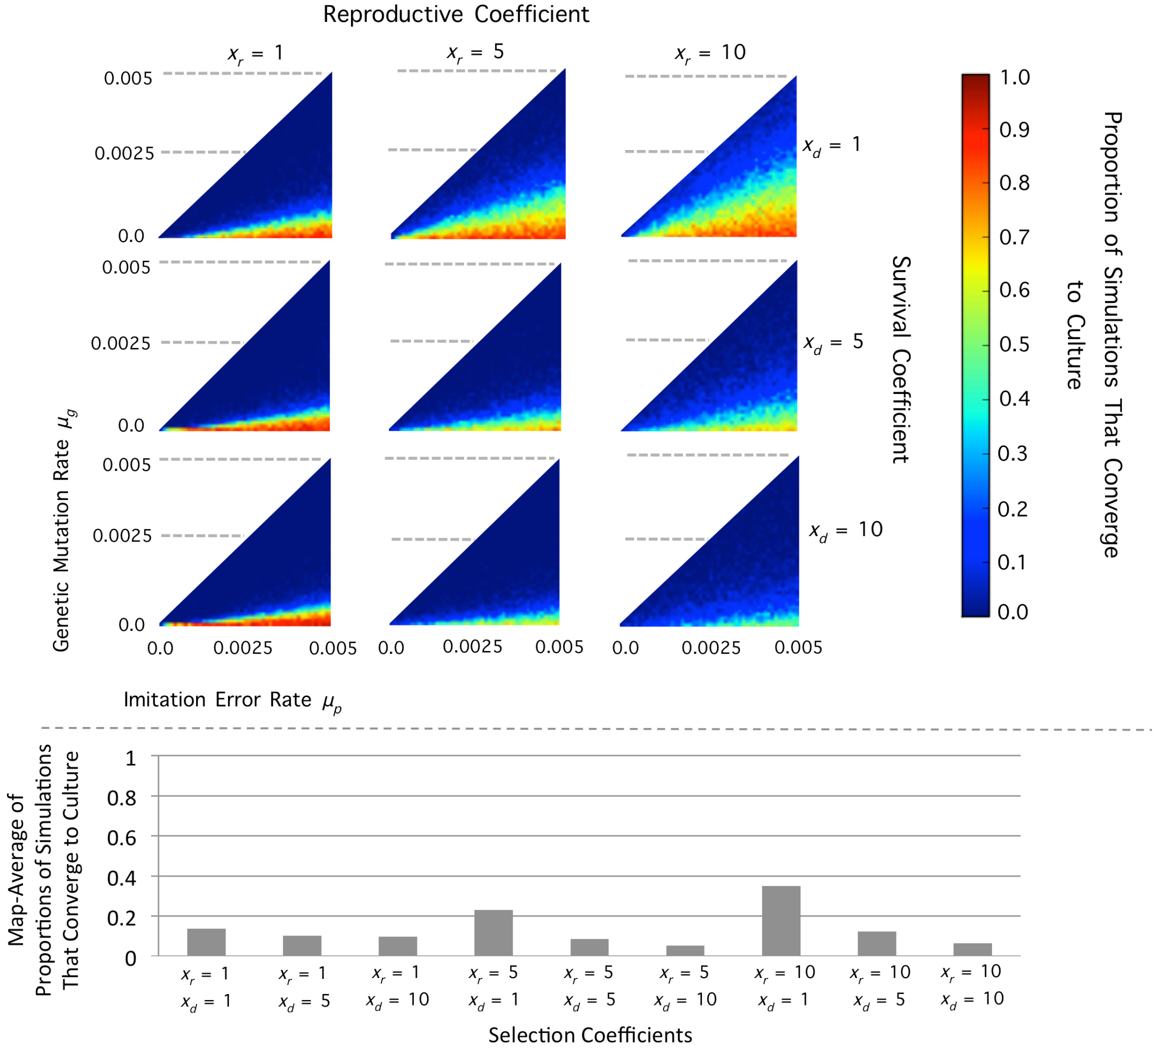

FIGURE D.8: Array of heat maps with different combinations of genetic mutation rates and imitation error rates, each using fixed reproductive and survival coefficients. Within each plot, red areas represent parameter sets where a high proportion of replicates evolved culture through genotype-phenotype disengagement. From left to right columns of maps have x_r values of 1, 5 and 10. From top to bottom rows of maps have x_d values of 1, 5 and 10. Higher values represent lower selection strength. All maps show the same area of mutation rates between 0.0 and 0.005 for combinations where $\mu_p > \mu_g$. A bar plot of the map-average proportion of simulations that converge to culture is provided under the array of heat maps. Each bar corresponds to one map in the array. In all runs in this figure, the rate of environmental change is set to $\delta_e = 0.4$ (i.e., $\mu_e = 2 \times 10^{-3}$); very extreme environmental change.

Bibliography

Acerbi, A., Ghirlanda, S., and Enquist, M. (2014). Regulatory traits: Cultural influences on cultural evolution. In Cagnoni, S., Mirolli, M., and Villani, M., editors, *Evolution, Complexity and Artificial Life*, pages 135–147. Springer-Verlag.

Aoki, K., Wakano, J. Y., and Feldman, M. W. (2005). The emergence of social learning in a temporally changing environment: A theoretical model. *Current Anthropology*, 46(2):334–340.

Aunger, R. (2000). The life history of culture learning in a face-to-face society. *Ethos*, 28(3):445–481.

Avital, E. and Jablonka, E. (2005). *Animal Traditions*. Cambridge Univeristy Press.

Axelrod, R. (1997a). The dissemination of culture a model with local convergence and global polarization. *The Journal of Conflict Resolution*, 41(2):203–2226.

Axelrod, R. M. (1997b). Advancing the art of simulation in the social sciences. In Conte, R., Hegselmann, R., and Terna, P., editors, *Simulating Social Phenomena, Lecture Notes in Economics and Mathematical Systems*, pages 21–40. Springer-Verlag, Berlin.

Baldwin, J. R., Faulkner, S. L., Hecht, M. L., and Lindsey, S. L. (2005). *Redefining Culture*. Laurence Erlbaum Associates.

Bell, A. V., Richerson, P. J., and McElreath, R. (2009). Culture rather than genes provides greater scope for the evolution of large-scale human prosociality. *Proceedings of the National Academy of Sciences*, 106(42):17671–17674.

Bentley, R. A., Hahn, M. W., and Shennan, S. (2004). Random drift and culture change. *Proceedings of the Royal Society B.*, 271(1):1443–1450.

Blackmore, S. (1998). Imitation and the definition of a meme. *Journal of Memetics*, 2(2):159–170.

Blackmore, S. (2000). *The Meme Machine*. Oxford University Press.

Blute, M. (1987). Biologists on sociocultural evolution: A critical analysis. *Sociological Theory*, 5(2):185–193.

Borenstein, E., Feldman, M. W., and Aoki, K. (2007). Evolution of learning in fluctuating environments: When selection favors both social and exploratory individual learning. *Evolution*, 62(3):586–602.

Borg, J. M. and Channon, A. (2012). Testing the variability hypothesis: The adoption of social learning in increasingly variable environments. In Adami, C., Bryson, D. M., Ofria, C., and Pennock, R. T., editors, *Proceedings of the Thirteenth International Conference on the Simulation and Synthesis of Living Systems*, pages 317–324.

Bouchard, T. J., Lykken, D., McGue, M., Segal, N., and Tellegen, A. (1990). Sources of human psychological differences: The minnesota study of twins reared apart. *Science*, 250(4978):223–228.

Boyd, R. and Richerson, P. (1996). Why culture is common but cultural evolution is rare. *Proceedings of the British Academy*, 88(1):73–93.

Boyd, R. and Richerson, P. J. (1988a). *Culture and the Evolutionary Process.* University of Chicago Press.

Boyd, R. and Richerson, P. J. (1988b). An evolutionary model of social learning: The effects of spatial and temporal variation. In Zentall, T. and Galef, B. G., editors, *Social Learning: Psychological and Biological Perspectives*, pages 29–48. Lawrence Erlbaum Associates.

Boyd, R. and Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16(1):125–143.

Boyd, R., Richerson, P. J., and Henrich, J. (2013). The cultural evolution of technology: Facts and theories. In Richerson, P. and Christiansen, M. H., editors, *Cultural Evolution: Society, Technology, Language, and Religion.* The MIT Press.

Brown, C. and Laland, K. N. (2003). Social learning in fishes: a review. *Fish and Fisheries*, 4(3):280–288.

Case, C. M. (1927). Culture as a distinctive human trait. *American Journal of Sociology*, 32(6):906–920.

Castro, L. and Toro, M. A. (2004). The evolution of culture: From primate social learning to human culture. *Procedings of the National Academy of Sciences USA*, 101(27):10235–10240.

Cavalli-Sforza, L. and Feldman, M. W. (1981). *Cultural Transmission and Evolution: A Quantitative Approach.* Princeton University Press.

Cavalli-Sforza, L. and Feldman, M. W. (1983). Cultural versus genetic adaptation. *Procedings of the National Academy of Sciences USA*, 80(16):4993–4996.

Cole, M. (1996). *Cultural psychology: A once and future discipline.* Harvard University Press, Cambridge, MA.

Cole, M. and Parker, M. (2011). Culture and cognition. In Keith, K., editor, *Cross-cultural psychology: Contemporary themes and perspectives*, pages 133–159. Wiley-Blackwell, Chichester, UK.

Dawkins, R. (2006). *The Selfish Gene.* Oxford University Press.

DeJong, K. (2002). *Evolutionary Computation: A Unified Approach.* Bradford.

Dignum, V. and Dignum, F. (2014). *Perspectives on Culture and Agent-based Simulations*, volume 3. Springer International Publishing.

Dugatkin, L. A. (2000). *The imitation factor: evolution beyond the gene.* NY: Free Press, New York.

Dunbar, R. I. M. (1993). Coevolution of neocortical size, group size and language in humans. *Behavioural and Brain Sciences*, 16:681–735.

Durham, W. H. (1976). The adaptive significance of cultural behaviour. *Human Ecology*, 4(2):89–121.

Edmonds, B. (2001). The use of models - making mabs actually work. In Moss, S. and Davidsson, P., editors, *Multi-Agent-Based Simulation, Lecture Notes in Artificial Intelligence*, pages 15–32. Springer-Verlag, Berlin.

Eiben, A., Raue, P.-E., and Ruttkay, Z. (1994). Genetic algorithms with multi-parent recombination. In Davidor, Y., Schwefel, H.-P., and Manner, R., editors, *Proceedings of the 3rd Conference on Parallel Problem Solving from Nature*, volume 866, pages 78–87.

Enquist, M., Eriksson, K., and Ghirlanda, S. (2008). Critical social learning: A solution to rogers' paradox of nonadaptive culture. *American Anthropologist*, 109(4):727–734.

Enquist, M. and Ghirlanda, S. (2007). Evolution of social learning does not explain the origin of human cumulative culture. *Journal of Theoretical Biology*, 246(1):129–135.

Ferrari, P. F., Gallese, V., Rizzolatti, G., and Fogassi, L. (2003). Mirror neurons responding to the observation of ingestive and communicative mouth actions in the monkey ventral premotor cortex. *European Journal of Neuroscience*, 17(8):1703–1714.

Findlay, C. S., Lumsden, C. J., and Hansell, R. I. C. (1989). Behavioral evolution and biocultural games: Vertical cultural transmission. *Procedings of the National Academy of Sciences USA*, 86(1):568–572.

Futuyma, D. A. (2009). *Evolution.* Sinauer.

Galán, J. M., Izquierdo, L. R., Izquierdo, S. S., Santos, J. I., del Olmo, R., López-Paredes, A., and Edmonds, B. (2009). Errors and artefacts in agent-based modelling. *Journal of Artificial Societies and Social Simulation*, 12(1):1.

Gilbert, N. (2007). *Agent-Based Models. Quantitative Applications in the Social Sciences*. SAGE Publications, London.

Harris, J. R. (1999). *The nurture asssumption: Why children turn out the way they do.* New York Free Press.

Hartfield, M. and Keightley, P. (2012). Current hyphotheses for the evolution of sex and recombination. *Integrated Zoology*, 7(2):192–209.

Henrich, J. (2015). *The Secret of Our Success: How Culture is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press.

Henrich, J. and Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and human Behaviour*, 19(1):215–241.

Henrich, J. and Boyd, R. (2002). On modeling cognition and culture: Why cultural evolution does not required replication of representations. *Journal of Cognitive culture*, 2(2):87–112.

Henrich, J., Boyd, R., and Richerson, P. J. (2008). Five misunderstandings about cultural evolution. *Human Nature*, 19(1):119–137.

Henrich, J. and Gil-White, F. J. (2001). The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, 22(3):165–196.

Hewlett, B. and Cavalli-Sforza, L. (1986). Cultural transmission among the aka pygmies. *American Anthropologist*, 88(4):922–934.

Heyes, C. (1994). Social learning in animals: Categories and mechanism. *Biological Reviews*, 69(1):207–231.

Heyes, C. M. and Galef, B. G. (1996). *Social Learning in Animals: The Roots of Culture.* Elsevier.

Holden, C. and Mace, R. (1997). Phylogenetic analisys of the evolution of lactose digestion in adults. *Human Biology*, 69(5):605–628.

Holyst, J. A., Kacperski, K., and Schweitzer, F. (2000). Phase transitions in social impact models of opinion formation. *Physica A: Statistical Mechanics and its Applications*, 285(1):199–210.

Huber, L., Range, F., Voelkl, B., Szucsich, A., Viranyi, Z., and Miklosi, A. (2009). The evolution of imitation: What do the capacities of non-human animals tell us about the mechanisms of imitation? *Philos Trans R Soc Lond B Biol Sci*, 364(1):2299–2309.

Ingold, T. (1986). *Evolution and Social Life.* Cambridge University Press.

Jablonka, E. and Lamb, M. J. (2006). *Evolution in Four Dimensions.* MIT Press.

Jahoda, G. (2012). Critical reflections on some recent definitions of "culture". *Culture and Psychology*, 18(3).

Jenkins, P. (1978). Cultural transmission of song patterns and dialect development in a free-living bird population. *Animal Behaviour*, 26(1):50–78.

Kadushin, C. (2011). *Understanding Social Networks: Theories, Concepts, and Findings.* Oxford University Press.

Kameda, T. and Nakanishi, D. (2002). Cost-benefit analysis of social/cultural learning in a nonstationary uncertain environment. an evolutionary simulation and an experiment with human subjects. *Evolution and Human Behaviour*, 23(5):373–393.

Kauffman, S. A. (1993). *The Origins of Order.* Oxford Univeristy Press, Oxford.

Knuf, L., Aschersleben, G., and Prinz, W. (2001). An analysis of ideomotor action. *Journal of Experimental Psychology General*, 130(4):779–798.

Kohler, T. A. and Gummerman, G. J. (2001). *Dynamics of Human and Primate Societies: Agent-Based Modeling of Social and Spatial Processes.* Oxford University Press, Oxford.

Kolodny, O., Creanza, N., and Feldman, M. W. (2016). Game changing innovations: How culture can change the parameters of its own evolution and induce abrupt cultural shifts. *PLoS Computational Biology.*

Kroeber, A. L. (1948). *Anthropology: Race, Language, culture, psychology, pre-history.* Harcourt, Brace and World.

Kroeber, A. L. and Kluckhohn, C. (1952). Culture: A critical review of concepts and definitions. *Papers of the Peabody Museum of American Archaeology and Ethnology*, 47(1):223.

Laland, K. and Galef, B. G. (2009). *The Question of Animal Culture.* Harvard University Press.

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

Laland, K. N., Atton, N., and Webster, M. M. (2011). From fish to fashion: experimental and theoretical insights into the evolution of culture. *Philos Trans R Soc Lond B Biol Sci*, 366(1567):958–968.

Laland, K. N. and Hoppitt, W. (2003). Do animals have culture? *Evolutionary Anthropology*, 12(1):150–159.

Laland, K. N. and Janik, V. M. (2006). The animal cultures debate. *Trends in Ecology and Evolution*, 21(10):542–547.

Lazer, D. and Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quaterly*, 52(4):667–694.

Macy, M. W. and Svoretz, J. (1998). The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63(5):638–660.

Markham, A. C., Gesquiere, L. R., Alberts, S. C., and Altmann, J. (2015). Optimal group size in a highly social mammal. *Proceedings of the National Academy of the United States of America*, 112(48).

Marriot, C. and Chebib, J. (2014). The effect of social learning on individual learning and evolution. In Sayama, H., Rieffel, J., Risi, S., Doursat, R., and Lipson, H., editors, *Proceedings for the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 736–743. MIT press.

Massen, C. and Prinz, W. (2009). Movements, actions and tool-use actions: An ideomotor approach to imitation. *Philos Trans R Soc Lond B Biol Sci*, 364(1528):2349–2358.

McElreath, R. and Henrich, J. (2007). Dual inheritance theory: The evolution of human cultural capacities and cultural evolution. In Dunbar, R. I. M. and Barrett, L., editors, *Oxford Handbook of Evolutionary Psychology*, pages 555–570. Oxford University Press.

McElreath, R. and Strimling, P. (2008). When natural selection favors imitation of parents. *Current Anthropology*, 49(2):307–316.

McGrew, W. (1998). Culture in nonhuman primates. *Annual Review of Anthropology*, 27(1):301–328.

Mead, M. and Metraux, R. (1953). *The Study of Culture at a Distance.* University of Chicago Press.

Mesoudi, A. (2011). *Cultural Evolution: How Darwinian Theory can Explain Human Culture and Synthesize the Social Sciences.* The University of Chicago Press.

Monfardini, E., Hadj-Bouziane, F., and Meunier, M. (2014). Model-observer similarity, error modeling and social learning in rhesus macaques. *PLOS ONE*, 9(2):e89825.

Morris, A., Ross, W., Hosseini, H., and Ulieru, M. (2014). Modelling culture with complex, multi-dimensional, multi-agent systems. In Dignum, V. and Dignum, F., editors, *Perspectives on Culture and Agent-based Simulations*, pages 13–30. Springer International Publishing, Switzerland.

Moussaid, M., Kammet, J. E., Analytis, P. P., and Neth, H. (2013). Social influence and the collective dynamics of opinion formation. *PLOS ONE*, 8(11):e78433.

Mulder, M. B., Nunn, C. L., and Towner, M. C. (2006). Cultural macroevolution and the transmission of traits. *Evolutionary Anthropology*, 15(2):52–64.

Nelson, D. A. (2000). A preference for own-subspecies' song guides vocal learning in a song bird. *Procedings of the National Academy of Sciences USA*, 97(24):13348–13353.

Norenzayan, A., Smith, E. E., Kim, B. J., and Nisbett, R. E. (2002). Cultural preferences for formal versus intuitive reasoning. *Cognitive Science*, 26(1):653–684.

Perreault, C. (2012). The pace of cultural evolution. *PLOS ONE*, 7(1):e45150.

Pfau, J., Kirley, M., and Kashima, Y. (2013). The coevolution of cultures, social network communities, and agent locations in an extension of axelrod's model of cultural dissemination. *Physica A: Statistical Mechanics and its Applications*, 392(2):381–391.

Pigliucci, M. (2010). Genotype–phenotype mapping and the end of the 'genes as blueprint' metaphor. *Philosophical Transactions of the Royal Society B*, 365(1540):557–566.

Plomin, R., DeFries, J., Craig, I., and McGuffin, P. (2003). *Behavioral Genetics in the postgenomic era.* American Psychological Association.

Potts, R. (1996). *Humanity's descent : the consequences of ecological instability.* New York: Morrow.

Pulliam, H. and Dunford, C. (1980). *Programmed to Learn: An Essay on the Evolution of Culture.* Columbia University Press.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., and Laland, K. (2010a). Why copy others? Insights from the social learning strategies tournament. *Science*, 328(5975):208–213.

Rendell, L., Forgarty, L., and Lanland, K. (2010b). Rogers' paradox recast and resolved: population structure and the evolution of social learning. *Evolution*, 64(2):534–548.

Rendell, L. and Whitehead, H. (2001). Culture in whales and dolphins. *Behavioral and Brain Sciences*, 24(2):309–382.

Resnick, M. (1997). *Turtles, Termites, and Traffic Jams.* MIT Press, Cambridge MA.

Richerson, P. and Boyd., R. (2001). Culture is part of human biology: Why the superorganic concept serves the human sciences badly. In Maasen, S. and Winterhager, M., editors, *Science Studies: Probing the Dynamics of Scientific Knowledge*, pages 147–177. Bielefeld: Verlag.

Richerson, P. J. and Boyd, R. (2006). *Not By Genes Alone: How Culture Transformed Human Evolution.* University of Chicago Press.

Rogers, A. (1988). Does biology constrain culture? *American Anthropologist*, 90(4):819–831.

Rosenberg, N., Pritchard, J. K., Werber, J. L., Cann, H. M., Kidd, K. K., Zhivotovsky, L. A., and Feldman, M. W. (2002). Genetic structure of human populations. *Science*, 298(1):2381–2385.

Shea, N. (2009). Imitation as an inheritance system. *Philos Trans R Soc Lond B Biol Sci*, 364(1528):2429–2443.

Shennan, S. J. and Steele, J. (1999). Cultural learning in hominids: A behavioural ecological approach. In Box, H. O. and Gibson, K. R., editors, *Mammalian social learning: Comparative and ecological perspectives*, pages 367–88. Cambridge Univeristy Press.

Simpson, G. G. (1953). The Baldwin effect. *Evolution*, 7(2):110–117.

Smith, K., Brighton, H., and Kirby, S. (2003). Complex systems in language evolution: The cultural emergence of compositional structure. *Advances in Complex Systems*, 6(4):537–558.

Takahashi, K. (1999). Theoretical aspects of the mode of transmission in cultural inheritance. *Theoretical Population Biology*, 55(2):208–225.

Tomasello, M. (1999). The human adaptation for culture. *Anual Review of Anthropology*, 28(1):509–529.

Tsutsui, S. and Jain, L. (1998). On the effect of multi-parents recombinaton in binary coded genetic algorithms. *Knowledge-Based Intelligent Electronic Systems*, 3(1):155–160.

Tylor, E. (1920). *Primitive culture : researches into the development of mythology, philosophy, religion, language, art, and custom.* London:Murray, 6th edition.

Tylor, E. B. (1871). *Primitive Culture: Researches in the Development of Mythology, Philosophy, Religion, Language, Art and Custom.* John Murray.

van de Waal, E., Borgeaud, C., and Whiten, A. (2013). Potent social learning and conformity shape a wild primate's foraging decisions. *Science*, 340(6131):483–485.

Van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, 1:1–30.

Vernon, P. A., Jang, K. L., Harris, J. A., and McCarthy, J. M. (1997). Environmental predictors of personality differences: A twin and sibling study. *Journal of Personality and Social Psychology*, 72(1):177–183.

Wakano, J. and Aoki, K. (2006). A mixed strategy model for the emergence and intensification of social learning in a periodically changing environment. *Theoretical Population Biology*, 70(4):486–497.

Watson, R. and Szathmary, E. (2016). How can evolution learn? - a reply to responses. *Trends in Ecology and Evolution*, 31(12):896–898.

Watts, D. J. (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research*, 43(1):441–458.

Weissman, D. B., Desai, M. M., Fisher, D. S., and Feldman, M. W. (2009). The rate at which asexual populations cross fitness valleys. *Theoretical Population Biology*, 75(4):286–300.

White, L. (1959). *The Evolution of Culture: The Development of Civilization to the Fall of Rome*. McGraw-Hill.

Whiten, A. and van Schaik, C. P. (2007). The evolution of animal “cultures” and social intelligence. *Philosophical Transactions of the Royal Society B*, 362(1480):603–620.