


Theoretical and Computational Models of Cohesion, Competition and Maladaptation in the Evolution of Human Musical Behaviour

Oliver Bown

Goldsmiths, University of London



Submitted in satisfaction of the requirements for the degree of
Doctor of Philosophy in the University of London

2008

Declaration

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

Signature

Date

Abstract

This thesis is concerned with the evolution of musical behaviour in humans, considered on a broad theoretical level. Drawing on key theoretical developments in the evolution of human cultural behaviour, and the potential interaction between culture and biological evolution, I propose that music has evolved through a feedback process between human cultural systems and human physiology, that places its origin not in any adaptive function but in the non-functional consequences of complex dynamical systems—in particular, systems in which cultural dynamics strongly determine relative individual fitness. In the context of other approaches to the theory of the evolution of musical behaviour in humans, the primary goal of this research is to elucidate the potential for biocultural evolutionary feedback processes to lead to the cultural and biological stability of activities such as music.

In the first half of the thesis, I review existing theories that are concerned with the evolution of musical behaviour and attempt to classify them according to their basic evolutionary theoretical position. I evaluate these theories in terms of broad issues in evolutionary theory and the theory of biocultural coevolution, and compare views from the natural and social sciences on the process of cultural evolution. I also review the role of computer simulation modelling in understanding the general dynamics of biocultural coevolution and emergent behaviour.

In the second half of the thesis, I present a modelling framework specifically dedicated to performing multi-agent simulations concerned with the evolution of musical behaviour, and experimenting with a number of agent models. The simulations presented in this thesis demonstrate how, in the context of cultural behaviour, individual behaviour is capable of tending towards behaviours that are not functional for individuals or for groups. This presents an alternative to sexual selection as a hypothesis for the evolution of human behaviours not based on direct utility. I conclude that these simulations demonstrate how cohesive behaviour could emerge from competitive behaviour in the context of a system of stylistic interaction, and could be extended with the aim of developing a more precise formulation of the ways in which biological evolution can be effected by systems of cultural behaviour.

Acknowledgements

I thank my supervisors, Geraint Wiggins and Tim Blackwell, for their enthusiastic support and encouragement and their many talents. This thesis was also shaped by enjoyable discussions with Alice Eldridge, Mark d'Inverno, Rob Saunders, Joao Martins, Chris Knight, Tak-Shing Chan, Dan Bratton and all of the members of the ISMS group at Goldsmiths College. It was made easier in more indirect ways by Estelle Hoen, Sam Britton, and numerous Bowns.

Contents

1	Introduction	11
1.1	The Problem of Music	11
1.2	Scope	13
1.3	Hypothesis and Objectives	14
1.4	Overview	15
2	Research on the Evolution of Human Musical Behaviour and Related Fields	17
2.1	Music as a Perceptual and Biological Phenomenon	17
2.2	The Core Debate on the Evolution of Human Musical Behaviour	24
2.2.1	Emergence and current popularity of the debate	25
2.2.2	Key views on the evolution of human musical behaviour	29
2.2.2.1	The early debate	29
2.2.2.2	Recent development of themes in the debate	39
2.3	Contextualising the Evolution of Human Musical Behaviour Within Issues in Evolutionary Theory	41
2.3.1	The fundamental debates of Darwinism	42
2.3.2	The evolution of cooperative behaviour	46
2.3.3	Theories of animal communication	50
2.3.4	Major transitions and evolution beyond Darwinism	53
2.4	Considering the Evolution of Human Musical Behaviour from the Point of View of Cultural Evolution and Gene-Culture Coevolution	58
2.4.1	Human evolution and the social brain	61

<i>CONTENTS</i>	5
2.4.2 Evolutionary approaches to culture	65
2.4.2.1 Boyd and Richerson's theory of gene-culture coevolution . . .	65
2.4.2.2 Niche construction	71
2.4.2.3 Memetics	74
2.4.3 Anthropological and sociological views of culture	77
2.4.4 Creativity as a cultural driving force	85
2.5 Summary	91
3 A Framework for Investigation	93
3.1 Theoretical development of existing models	93
3.2 Extension of cohesionist models	95
3.2.1 The handicap principle applied to a changing-style system	96
3.3 Introducing new approaches to competitiveness	98
3.3.1 Social learning leads to the social generation of fitness	99
3.4 The equivalence between stylistic bonding and competitive social learning . . .	103
3.4.1 Competition leads to cohesion	105
3.4.2 Cohesion leads to competition	106
3.5 The gene-culture feedback cycle of social-learning models	106
3.5.1 Runaway biocultural coevolution drives the evolution of innate musical capacities	109
3.6 Additional factors in social-learning models	111
3.6.1 Social lotteries	111
3.6.2 Identity	113
3.6.3 Creativity	115
3.6.4 The relationship between vertical and horizontal transmission	116
3.6.5 Sex	117
3.6.6 Group singing and mother-infant interactions	118
3.7 Theoretical significance of social-learning models	119
3.7.1 The overlap between sexual selection and social selection models . . .	120

3.7.2	Relationship with general trends in evolutionary theory	123
3.7.3	Exploring the competitive style learning hypothesis	125
4	A Simulation Modelling Approach to the Study of the Evolution of Human Musical Behaviour	127
4.1	Theoretical perspectives	127
4.1.1	Models of human behaviour	128
4.2	Using simulation models to explore novel evolutionary processes	130
4.3	Multi-agent models of musical behaviour and evolution	135
4.3.1	Werner and Todd's sexual selection model	136
4.3.2	Managing the complexity of the evolution of human musical behaviour in a model	140
4.4	Simulating the evolution of human musical behaviour in light of current theoretical debates	143
4.4.1	What evolves when music evolves?	146
5	Design of Simulation Model	148
5.1	Aims	148
5.2	Theoretical questions	150
5.3	Approach to building a simulation framework	153
5.3.1	Evolution	153
5.3.2	Interaction	155
5.3.3	Limits of simulating music	156
5.4	Implementation	158
5.4.1	Simulation updates	158
5.4.2	Rules for death and reproduction	159
5.4.3	Environment and energy	159
5.4.4	Modelling the effects of the social determination of fitness	161
5.4.5	Musical representation	161
5.5	Modelling investigations	166

5.5.1	Introducing genetic variables	166
5.6	Summary	170
6	Experiments	173
6.1	Experimental set up	174
6.2	Experiment 1: Initial study of $ P $ and S evolution	175
6.2.1	Further exploration of increasing S	179
6.2.2	Discussion	183
6.3	Experiment 2: The coevolution of S , P and p_{sing}	186
6.3.1	Rationale	186
6.3.2	Predictions	187
6.3.3	Results	188
6.4	Experiment 3: The coevolution of S , P , C and L	189
6.4.1	Rationale	189
6.4.2	Predictions	189
6.4.3	Simulations	191
6.4.4	Results	191
6.4.5	Summary of results	194
7	Discussion of Results and Conclusions	197
7.1	Interpretation of simulation results	197
7.1.1	Summary of conclusions	197
7.2	Comparison of simulations with existing modelling approaches	200
7.2.1	Criticism of these results	201
7.2.2	Conclusions regarding methodology	203
7.3	Proposition of future models	204
7.3.1	Development of the changing style system models	204
7.3.2	Exploration of more complex musical capacities	206
7.3.3	Comparative approaches to models of the evolution of human musical behaviour	207

CONTENTS

8

7.4	Relevance to theories of the evolution of human musical behaviour	208
7.4.1	Towards alternative runaway selection processes	209
7.4.2	Cultural behaviour as a major transition in evolution	210
7.4.3	Relevance to modern understanding of musical and artistic behaviour in society	211

List of Figures

5.1	A simple linear hedonic function relating novelty (the Euclidian distance between a perceiver's location in style space and the location of the perceived stimulus in style space) to perceived value. A variety of novelty preference values is shown. In these simulations the novelty preference value was set to 2 and the susceptibility to enchantment value was used as a genetic variable.	163
5.2	Agent's evaluation and learning methods. On the left, agent A's evaluation function is shown, non-zero within the dark circle with a maximum value at the point of the dotted circle. Thus agent C receives the highest reward from A, and agents B and D receive similar non-zero rewards. On the right, agent A learns from its set of physical neighbours by adjusting its position towards their centroid by a fraction $L/200$ of the total distance.	164
5.3	Flow diagram of update for each agent. This sequence is repeated for each member of the population at each time step. Changes to agents' states, as well as additions to and removals from the population, are enacted after the entire population has completed this sequence, so that agent updates are synchronous.	166
6.1	Final average $ P $ values for several runs of the model with different learning strengths.	176
6.2	Final average S values for several runs of the model with different learning strengths.	178
6.3	Average S values for several runs with a learning strength of 0.1.	178
6.4	S values for one single run with a learning strength of 0.1.	179

6.5	Detail of one dimension of the agent style space during a run with $L = 0.1$ and vertical transmission between $t = 19,000$ and $t = 20,000$. Different individuals are drawn as lines with different colours for clarity. The graph shows the grouping of agents in the style space and occasional sudden leaps to different parts of the style space.	180
6.6	Final average S values for several runs of the model with no vertical transmission over different learning strengths.	181
6.7	Average total pay-outs made by agents mapped against age in the vertical transmission case (dotted lines) and non-vertical transmission case (solid lines). The lower of each pair of lines shows the average of the pay-outs made only to parents.	184
6.8	Averaged pay-outs made by agents to various different categories of kin during the vertical transmission model. The results were gathered by sampling payments made by all members of the population each 1000 timesteps. The payment values are calculated for each agent and each category by dividing the average reward paid to that category, divided by the agent's S value. This gives the payout to that category as a fraction of the agent's maximum payment ability. These payments are then averaged across the population for each category. The graphs shows that payments between parents and children are significantly higher than between other categories of kin and non-kin and that payments between grandparents and grandchildren are significantly higher than the average.	184
6.9	Final average values of C , L , $ P $ and S for sets of ten simulations each with different initial values of L (x-axis) and $\beta = 0.1$ (left), $\beta = 1$ (middle) and $\beta = 10$ (right). These simulations used vertical transmission.	193
6.10	Final average values of C , L , $ P $ and S for sets of ten simulations each with different initial values of L (x-axis) and $\beta = 0.1$ (left), $\beta = 1$ (middle) and $\beta = 10$ (right). These simulations used non-vertical transmission.	194

Chapter 1

Introduction

1.1 The Problem of Music

Over the past century and a half, Darwinism has risen to a level of great importance in the understanding of human behaviour and continues to pervade this subject. At the same time, theoretical biology continues to mature, contending with and drawing on views from other disciplines, and nowhere is this more the case than in the domain of culture. This thesis is concerned with the problem of the existence of musical behaviour in humans and its relation to a scientific evolutionary view of human behaviour. In the spirit of Darwinian evolutionary theory, a philosophically materialistic view of mind is a necessary starting point for this study, forbidding the notion of music residing in a Platonic realm: music must be seen entirely in terms of its relationship with human behaviour. This human behaviour must be understood in relation to a process, natural selection, that is understood to relentlessly constrain the form of biological systems. Yet, whilst constrained by natural selection, human behaviour is often seen as defying its logic, and continues to present considerable challenges to evolutionary theory. Humans are shaped not only by their genetic evolution, but also by the cultural systems they grow up in, and musical behaviour epitomises this contrast. It appears to be influenced by dedicated genetic components, but arguably it does not represent a domain of behaviour that evolved due to its survival value on either an individual or a group level.

Although it was a problem considered by Darwin himself (1883), only in recent years has the theory of the evolution of human musical behaviour become a serious area of scientific en-

quiry, with a number of different points of view competing for acceptance (*e.g.*, Pinker, 1998; Miller, 2000a; Cross, 2001, 2003a; Wallin et al., 2000). The remarkable number of contributing disciplines and overlaps makes for a complicated debate, with an overwhelming number of factors to consider and a great deal of uncertainty amidst a conspicuous shortage of evidence. This thesis focuses strictly on the theoretical side of the study of the evolution of human musical behaviour. It deals with specific interpretations of Darwinian evolutionary theory, and models that explore the efficacy of hypothetical evolutionary processes that may have contributed to the evolution of human musical behaviour. The value of such work lies in the fact that a number of key theories cannot be evaluated fully on the basis of current or impending evidence, and need to be developed further in order to find critical aspects that may produce genuinely testable hypotheses.

This thesis stems largely from a basis in artificial life and theoretical biology (*e.g.*, Di Paolo et al., 2000), and uses computer simulation modelling as a means to explore the implications of theories of the evolution of human musical behaviour. Whilst computer simulation models have enjoyed success in these fields (*e.g.*, Hinton and Nowlan, 1987; Axelrod, 1987; Kauffman, 1993), human musical behaviour, like language, poses a specific problem in which an interaction system evolves within a predominantly cultural niche (Huron, 2001). Both the communication system and the cultural niche are arguably highly complex (Miller, 2000a; Plotkin, 1997). One of the main challenges faced in this thesis has therefore been to attempt to break down aspects of these theories into simplifications that are manageable from the point of view of simulation. In doing so, it has been hard to address or absorb the rich literature discussing the evolution of language and the relationship between music and language (Bickerton, 2001), and the specific focus on notions of music as an essentially functionless cultural activity draw it further from any comparisons with language (Pinker, 1998).

This thesis also embodies another broad interest. Studies of the evolution of human musical behaviour are a particularly exciting area of interdisciplinary research in which theoretical biology collides with the study of cultural behaviour. The 20th Century has witnessed conspicuous battles across these respective academic domains (*e.g.*, the nature nature debate, addressed in (Elman et al., 1996)), which continue to operate to some extent in isolation from each

other (one possible account is Snow, 1993). There is great scope for theoretical biological approaches to cultural behaviour which give culture some autonomy over genetically determined behaviour, and also over environmental determinism. Dawkins' (1976) theory of memes and Boyd and Richerson's (1985) theory of biocultural coevolution both provide such approaches. New approaches to theoretical biology that go beyond Darwinism, such as those of Maturana and Varela (1980), Kauffman (1993) and Odling-Smee et al. (2003) may be even more amenable to the study of cultural behaviour. To consider theoretical approaches to the evolution of human musical behaviour, in the light of these views, addresses the general relationship between science, the humanities and the arts. This thesis was initiated with that exciting area of development in mind.

1.2 Scope

In this dissertation, I treat music as a system of sonic interaction unique to humans (although it may share many features with forms of sonic interaction in other animals) which is distinct from language in that it does not convey truth-functional or referential information (*c.f.*, Cross, 2007a) (although, according to this definition, language and music may still have much in common). This definition avoids the term communication, which, despite being a word with numerous and quite varied definitions itself (Owings and Morton, 1998), often implies a transfer of meaningful information. Music does, of course, communicate (Cross, 2007a), but the approach taken here attempts to side-step the paradigm of communication by talking about interaction in a broad sense. Some definitions of communication already push it so far as to be at risk of being synonymous with any form of interaction (Dawkins and Krebs, 1978). By avoiding a notion of communication and adopting a minimal cognition paradigm (*e.g.*, Quinn, 2001) I attempt to renew a notion of music as a substrate of cultural behaviour defined empirically by certain contexts of use and a specific physical domain. Notions of music are typically limited to a sonic domain (*e.g.*, Blacking, 1995), and by virtue of this, music is necessarily a temporal activity, and in this sense can be grouped with dance. The models in this dissertation do not strictly address the temporal domain of music and can be considered applicable to a wider range of cultural behaviour. Thus, although contributing to an understanding of music is the starting point for this

thesis, the proposed contribution may not be music-specific. However, an argument is made in Chapter 3 for why music might be more fundamental than other cultural domains for the kinds of evolutionary processes proposed.

I take an approach to the theory of the evolution of human musical behaviour which focuses on the possibility that interaction in cultural systems is the driving force for that evolution. There are many different definitions of culture (Plotkin, 2002). For the purpose of this thesis, I will treat culture as a system of behaviours held by a group of individuals which has become shared through social learning (Boyd and Richerson, 1985). Since it may be problematic to separate this set of learnt behaviours from the processes that change these behaviours over time, and the learning process itself (for example see Odling-Smee et al., 2003), I consider a cultural system as potentially including these additional elements. The term *bio-cultural system* is used to refer to such a system whilst explicitly including biological adaptations. I define a cultural domain as a domain of behaviour that is determined by a cultural system. If a cultural domain ultimately drives the emergence of biological adaptations, then these biological adaptations are considered part of that domain, and the domain is still defined as being culturally determined.

1.3 Hypothesis and Objectives

The literature review in Chapter 2 and arguments in Chapter 3 lead to the hypothesis that it is possible for maladaptive cultural interactions based on style learning to drive the emergence of a self-sustaining cultural domain, along with the evolution of aspects of perception exhibiting the essential features of music, which I set out to test using a computer simulation model. The purpose of the model is to demonstrate the potential for such a process as a proof of concept. One may doubt the refutability of the stated hypothesis, but I take the view in Chapter 4 that it is valid to use simulation models to demonstrate possible ways in which a hypothesised process can work, as a way of developing concepts and opening doors to potentially fruitful lines of enquiry.

The primary objective of this thesis is to formulate a theory of how such a process could work, framed in terms of the evolutionary concepts of inter-individual competition, group cohesion and maladaptation resulting from the emergence of cultural behaviour. Secondly, I aim to

build a simulation model based on the hypothesised processes set out in Chapter 3 and in exploring the dynamics of this model, to test the stated hypothesis. I expect the models to reflect the complexity of musical interaction as an evolutionary phenomenon, and therefore a third aim is to investigate how fruitful such complex models can be, when analysis of the model's behaviour is hindered by its complexity.

1.4 Overview

This thesis is divided into a theoretical section and an experimental section. In the theoretical section I make an analysis of existing theories of the evolution of human musical behaviour. I define four general orientations of theories of the evolution of human musical behaviour: consequentialist, cohesionist, competitivist, and cognitivist (Chapter 2). Through a consideration of the strengths and weaknesses of these theories I synthesise a new view based on a hypothetical interaction between cohesive and competitive processes (Chapter 3). According to this view, musical interaction is understood as a system in which competition between individuals has the effect of reinforcing aspects of cohesive behaviour. I argue that, given a certain level of cultural behaviour, including the potential for individuals to form strong allegiances and the intensive use of social learning, runaway evolutionary processes can emerge in which competitive interactions in a cultural domain drive the potential for musical behaviour to become a medium for social cohesion. The determiner of this type of runaway process is the kind of social organisation and cultural behaviour in place at its outset. Of particular interest is the possible maladaptive emergence of a system of style-learning. In Chapter 4 I discuss approaches to computer simulation.

In the experimental section I propose a computer simulation model which demonstrates the potential for such a runaway process to take place (Chapter 5). The model consists of a population of artificial agents, each of whose ability to survive and to reproduce is determined by interactions within the group. I consider simple games in which agents reward each other as a result of their respective evaluations of each other's style, and ask how it could be possible for the strength of these rewards to increase under evolutionary pressure, and also what kinds of perceptual effects could result from natural selection in this context. The conclusions of a number of simulation experiments are that such a process is possible as long as, within a local group,

individual gains are at the expense of the whole group and not of individuals giving the rewards (Chapter 6). To explain the simulation results, a process of kin selection was proposed, in which kin groups exploited the rules of the game and the simulation's design, with the overall effect of increasing rewards. This explanation may be seen as a specific model of the emergence of kin-based altruism in the context of musical behaviour, which may be seen as emerging initially from maladaptive runaway cultural processes. It is also shown that, under such circumstances, more complex perceptual systems associated with the mode of interaction could emerge and stabilise. In Chapter 7, I discuss these results and consider ways in which the model can be developed in response to my own initial criticisms of its design, and with the aim of exploring additional ways in which cultural processes can interact with biological natural selection to produce novel evolutionary processes and modes of behaviour.

Chapter 2

Research on the Evolution of Human Musical Behaviour and Related Fields

This literature review covers three main topics: the debate on the evolution of human musical behaviour; relevant issues in evolutionary theory; and approaches to the study of culture and its relation to evolutionary processes. I present the key views of the evolution of human musical behaviour first, in order to provide the context for the discussion of evolutionary theory. The purpose of the latter two sections is to provide material which will be useful for a critical discussion of current debates on the evolution of human musical behaviour and a formulation of a new theoretical approach which will take place in the following chapter. As a prelude to the discussion outlined here, however, I begin by briefly reviewing recent notions of music as a perceptual and biological phenomenon.

2.1 Music as a Perceptual and Biological Phenomenon

The notion of music as a biological phenomenon goes back at least as far as Darwin (1883) and has passed through periods of popularity as an object of evolutionary enquiry with both biological and cultural orientations (*e.g.*, Rehding, 2000). Scientific approaches to music that manage to incorporate both measurable facts about musical scores or recordings, and individual perceptual experiences of music, mark an important development in the potential for grounding evolutionary theories of music (*e.g.*, Clarke, 1987; Sloboda, 1988; Aiello and Sloboda, 1994; Cook and Everist, 1999). Lerdahl and Jackendoff (1983) stress the importance of a structural approach to

music by setting it apart from mathematical theories, which can describe anything but ultimately say nothing, as well as from artistic intuition-led theories, which “[sever] questions of art from deeper rational inquiry; [they] treat music as though it had nothing to do with any other aspect of the world” (Lerdahl and Jackendoff, 1983, p. 2). Characterising this important shift of perspective, Lerdahl and Jackendoff view a piece of music as a ‘mentally constructed entity’ rather than anything embodied in a score or an acoustic signal:

“One speaks of music as segmented into units of all sizes, of patterns of strong and weak beats, of thematic relationships, of pitches as ornamental or structurally important, of tension and repose, and so forth. Insofar as one wishes to ascribe some sort of ‘reality’ to these kinds of structure, one must ultimately treat them as mental products imposed on or inferred from the physical signal” (Lerdahl and Jackendoff, 1983, p. 2).

Here Lerdahl and Jackendoff concentrate on the *sense* that we make of music and view it as an imposition of order on the physical stimulus that is a musical performance. Treating perception as primary, they see composition as working with this process of experience as its subject. In this sense they actively emphasise the similarity between aspects of music perception and of perception in other domains, such as vision, as a *making-sense* of the world. Considering the variability of music across different cultural idioms, they conclude that some of these aspects of perception are likely to be innate. Their explanation for this innateness is that aspects of music perception are largely borrowed from other aspects of cognitive ability. However, they do not show a strong commitment to the proposition, and even include in parenthesis the suggestion that “there might, nevertheless, be a residue of musical cognition that is a product of nonadaptive evolutionary accident” (Lerdahl and Jackendoff, 1983, p. 283). Perception-based approaches like Lerdahl and Jackendoff’s allow music to be treated as an object of biological enquiry, at which point it can be viewed as one instance of a general class of interaction systems studied by ethologists and anthropologists alike. In this connection lies a genuine scientific basis for the question of why we make music.

Music constitutes numerous areas of enquiry and here I provide a focused account of one aspect of music, rhythm, rather than a broad overview. Rhythmic skills mark a key difference between humans and other animals (Bispham, 2006a), and form the basis for definitions of human music which focus on the temporal organisation of sound (such as Blacking, 1995). My focus on rhythm, however, is not intended to exclude the importance of other aspects of music, in particular pitch and harmony. This focus serves to guide later discussions of the potential use or effect of musical behaviour in human social interaction. However, I do propose that a broad notion of rhythm, as an analytic framework for studying music, is in some sense more fundamental (*i.e.*, less ethnocentric) than pitch or timbre. All music, as performed, exists *in time* and has temporal features that relate to human perception. In terms of fundamentally human aspects of rhythmic perception, it has been observed that humans, including non-musicians, can tap in time with a periodic auditory rhythmic stimulus and are competent at adapting to changes in the rate of this stimulus (London, 2004). This apparently differs from other animals, some of which (such as fireflies) can synchronise rhythmically, but using a very simple mechanism which is not robust to changes in the period of the activity (Bispham, 2006a). Bispham (2006a) provides a compelling argument that robust phase and period entrainment is a distinctive characteristic of human musical ability. But could mechanisms for rhythmic perception and production have direct primary utility in musical activities? Possible explanations for our unique rhythmic abilities include: the perception and production of language or other non-linguistic acts of interaction in humans (Pinker, 1998); the analysis of environmental sounds (for example, the rhythmic regularity of the sound of an animal's footsteps may help identify each sound as coming from the same source) (London, 2004); and the ability to learn sequences of movement, and to plan such sequences, which would have direct practical value such as in tool-making (Donald, 1991). These are open problems in the psychology of music. Separately, there is also the question of the innateness of such capacities. The cross-cultural universality of rhythmic musical behaviour (see Nettl, 2000) provides circumstantial evidence for the theory of the innateness of rhythmic abilities in humans, but can never constitute conclusive evidence. A particularly rich area of investigation lies in studying infant-caregiver interactions, in which babies are seen to rely strongly on rhythmic cues to establish appropriate communicative relations with the caregiver. Little is known

about the innate role of rhythmic behaviour in this process, but experiments show that infants develop temporal expectations about their interactions with caregivers (Trevarthen, 1986, 1993) indicating that rhythmic entrainment is involved.

The question of the relationship between musical rhythm and the use of rhythm in language remains open. Brown (2007) and Cross (2007b) stress the importance of the different interactive contexts in which musical and linguistic timing are embedded. Linguistic interaction centres around turn taking, and involves rhythmic cues which participants use to coordinate themselves, whilst musical interaction is usually simultaneous and often employs entrainment to a common rhythmic pulse (Cross, 2007b). In that sense musical and linguistic rhythmic perception require different mechanisms, but there is also the open question of whether linguistic perception requires rhythmic entrainment, in the musical sense (discussed below) or in a purely linguistic sense, depending only on stress. The latter view is supported by the majority of evidence (Patel, 2006), but Patel (2006) argues that the process of vocal learning in children, as opposed to vocal perception in adults, does require cognitive skills akin to the perception of musical beat and metre.

What, then, is rhythm, and how do we perceive it? Here I define rhythm as the temporal organisation of discretely perceived events in a stimulus. This includes the timing of events, but also their qualitative nature. An important distinction to make in studying rhythmic perception is that between the stimulus itself and an individual's perception of that stimulus (Bolton, 1894). Whilst rhythm generally describes the patterns actually present in the stimulus, accenting, beat and metrical structure are three important resulting perceptual features which cannot be said to be contained by the stimulus, but induced in the listener entirely, or in part, by the stimulus (Cooper and Meyer (1960).

Three experimentally significant aspects of rhythmic perception are beat induction, metre induction and the categorical perception of rhythmic patterns (see, for example, Clarke, 1987; Large et al., 2002; Honing, 2002; Desain and Honing, 2003). Beat induction, or beat tracking, involves the perception of beats in an auditory stimulus. When a sequence of uniform regularly-spaced sounds is heard we *entrain* to their regularity and can predict with reasonable accuracy when the next beat will occur (Clarke, 1987). As stated, the beat is not in the stimulus itself, but

is a perceptual event that is *induced* by the listener. Thus, when individual stimuli are removed from a regularly repeating sequence, one can still perceive beats at the points where they should have occurred.

Entrainment is a more general concept applicable to many dynamical systems. Huygens observed that two pendulum clocks placed on the same mantelpiece may synchronise with each other as a result of the very slight physical forces that are transmitted through the material of the mantelpiece (Clayton et al., 2005). This can be understood in terms of the stability of loosely coupled systems. In such a case the only state variable we need to keep track of is the relative phase between the clocks. The synchronised state, where the relative phase is zero, is a stable attractor in the space of all possible states. When the clocks are not synchronised, the forces passing through the mantelpiece have some effect on their relative phase, pushing it towards zero. When they are synchronised, these forces balance out and have no effect. Theoretically, when one oscillatory system is loosely coupled to another (so that the influence is one way only) it may entrain at a simple integer ratio, or it may fail to adopt a fixed cyclical behaviour. Human beat induction can be seen as a very robust adaptation of entrainment (Clayton et al., 2005), possibly using a complex version of the basic model of entrainment given by Huygens' clocks. When we entrain to musical rhythms we flexibly vary the rate of our internal timing mechanism as well as its relative phase (London, 2004). We can induce missing beats and handle large gaps in the stimulus, and we can adapt to varying tempos (Bispham, 2006a).

Metre induction involves the structuring of beats into repeating cycles. This requires a concept of *accent*: a property of a beat that makes it more salient. Cooper and Meyer (1960) suggest that duration, intensity, melodic contour and regularity, as well as other factors, can all contribute to the perception of accent. They define accent as an *axiomatic experiential concept*. *Regularly occurring* accents establish a sense of metrical structure which is defined by the number of beats over which this repetition takes place.

Accents are also significant in the perception of rhythmic patterning. Cooper and Meyer define rhythm as “the way in which one or more unaccented beats are grouped in relation to an accented one” (Cooper and Meyer, 1960, p. 6). Whilst London (2004) describes rhythm as “patterns of duration that are phenomenally present in the music” (London, 2004, p. 4), Cooper

and Meyer's definition, resting as it does on the perception of accents, establishes rhythm as a perceptual entity too. Beats, accents, metre and rhythm may thus be mutually interacting perceptual elements that are all induced from an auditory stimulus.

Experiencing certain aspects of rhythmic perception only takes place within a certain specific temporal range and is strongly influenced by the tempo of the stimulus (London, 2004). Speeding up or slowing down the stimulus does not simply produce the same perceptual image at different rates (Honing, 2002). Timing perturbations in isochronous stimuli become unidentifiable at intervals faster than around 100ms (which is sometimes referred to as the *trill threshold*) (London, 2004), and the same is true over intervals longer than 1.5-2 seconds. This defines a range within which rhythmic and metrical perception are naturally constrained. Towards the geometric centre of that range lies what is known as the *tactus*, a preferred perceptual tempo, which has a period of around 600ms but varies from person to person (London, 2004). London (2004) adds that tracking beats becomes difficult at a lower period threshold of about 200ms, and proposes that the relationship between the 100ms timing threshold and the 200ms beat tracking threshold is indicative of an important aspect of rhythmic perception: that in order to perceive regular beats we need to be able to time at least the most elementary subdivisions of those beats (*i.e.* that divide the beat in half). This is further supported by evidence that the beat tracking threshold increases for more complex patterns which imply a subdivision into thirds. London adds that "if hearing a beat means hearing the potential for its subdivision, and if metric attending minimally involves an ordering of beats (that is, a *tactus* and at least one superior level), then metric attending in almost all cases will involve three levels of structure" (London, 2004, p. 46).

Naturally extending the notion that a musical stimulus invokes perceptual phenomena in the listener, there is strong evidence that the listener actually *imposes* rhythmic structure on auditory stimuli. Bolton (1894) showed that listeners impose a metrical structure of two, four, or sometimes three beats on a uniform regular auditory stimulus, an effect termed *subjective rhythmization* (although as London 2004 points out, a more fitting term would be *subjective metricization*).

Plausible biological mechanisms underlying entrainment involve nonlinear adaptive oscillators. Following Large (Large and Palmer, 2002; Large et al., 1995; Large and Kolen, 1994;

Large and Jones, 1999; Large et al., 2002), such models begin by considering the behaviour of a single adaptive oscillator responding to a uniform pulse. The oscillator has two state variables: a period and a phase. The phase defines a point at which the oscillator *expects* a beat to occur. Around that point lies a bell-shaped expectation window within which the oscillator is sensitive to stimuli. In keeping with the basic theory of entrainment to a regular stimulus, if a pulse has just recently occurred when the oscillator reaches its zero point, then the oscillator both speeds up and nudges itself forward a bit (increases both rate and phase). If a pulse occurs just after the oscillator has passed its zero point, then the oscillator slows down and holds back a bit (decreases both rate and phase). However, if a beat occurs whilst the oscillator is not near its zero position (the adaptive window is closed), then it will not cause any adaptation. Likewise if the oscillator passes its zero point with no beat occurring, then nothing changes either. In this way, both period and phase are simultaneously modified, but only in keeping with the oscillator's expectation. Whilst Large's models are predominantly presented as abstract but physically plausible, Eck (2002) suggests similar oscillator models which are more directly grounded in models of neural firing behaviour. In simple conditions, Large and Kolen (1994) suggests that single adaptive oscillators are successful at beat induction, but there is a pay-off in their sensitivity which is not in keeping with human beat induction: too insensitive and they adapt far too slowly, too sensitive and they can be thrown right off by small variations in timing.

For more complex rhythms Large and Jones (1999) and Large and Palmer (2002) introduce a multiple-oscillator model. An early version of this model consisted of a number of oscillators with different natural periods. In response to a rhythmic input some will entrain and others will remain unstable. A combined spectrum of entrainment thus points to a best hypothesis for the timing of beats, but also gives additional information about metrical structure, since the stability and response of oscillators at different rates show up as clear traces. Interestingly, this allows for the perception of polyrhythmic metrical structure, which theorists such as Lerdahl and Jackendoff do not treat as perceptually equivalent to regular metrical structures. Lerdahl and Jackendoff's (1983) generative theory assumes that a polyrhythm would be perceived as either one or another of its constituent metrical groupings rather than *as* a polyrhythm. This may be true at the 'higher level' of sense-making that Lerdahl and Jackendoff are concerned with, but

less true at a lower level of rhythmic pre-processing. Large and Kolen (1994) also considered how the shape of the oscillators' adaptive windows could be individually varied according to the strength of their predictions. This has the effect of controlling their sensitivity and hence their rate of adaptation. This has two important consequences for later, higher level, stages of perception: it gives metrical levels a *strength*, placing metrical perception in a continuous realm, and it provides a low-level notion of expectation, which is increased through repetition. If such oscillators are connected in a network which maintains integer-ratio period relationships amongst them, then stronger candidates can drive the less stable oscillators to the correct integer-scaled period.

This work demonstrates a more general view of how multiple hypothesis or networked models can lead to the generation of higher-level structures from singular input stimuli. Although less concerned with cognitive modelling than with the challenge of engineering beat induction by machines, Dixon (2000) similarly uses multiple adaptive beat-inducing agents, and chooses the most likely estimate from a range of hypotheses. Similar results can be obtained from looking for patterns of autocorrelation in audio signals. Eck (2002) has also looked at the beat inducing potential of networks of loosely coupled relaxation oscillators with similar results to Large's work. Following developments in the neuroscience of vision, Neil Todd (Todd et al., 1999; Todd, 1999, 1994; Todd and Brown, 1996) considers how banks of *receptive fields* can filter auditory signals in order to build up a complex rhythmic image. Receptive fields are neural networks that perform transformations on a signal in order to extract specific features. Such features include percepts such as pitch, but also, through differentiation, the movement in these percepts. Todd's model grounds the neurology of rhythm perception in an established computational neuroscience of perception, which may have a significant impact on how more abstract models such as those of nonlinear adaptive oscillators are framed in the future.

This brief review of current research into the perception of rhythm provides the basis for asking how we come to possess such perceptual mechanisms that appear to be so specific to music. As discussed here, we are still not sure whether they *are* specific to music, or, independently, whether they are innate. Nevertheless, the study of such mechanisms have contributed to the growing sense that there could be evolved innate aspects of musical behaviour, although

this comes hand in hand with a greater sense of the role of music in social behaviour, and arguments for the adaptive function of music. This aspect of the debate will form the bulk of further discussion.

2.2 The Core Debate on the Evolution of Human Musical Behaviour

The problem of the evolution of human musical behaviour is a specific problem embedded within the theory of humankind's evolution, but one that has until recently been considered either insignificant or quite simply baffling. In order to understand how music managed so suddenly to take on importance to evolutionary theorists over roughly the past ten years (and similarly why Darwinian evolutionary theory took on importance to musicologists, ethnomusicologists, and so on), one must first follow major developments in the understanding of human evolution.

2.2.1 Emergence and current popularity of the debate

Music and other creative behaviours belong to a collective set of concepts underlying the distinction between humans and other animals: culture, including social learning (Mithen, 1996; Plotkin, 1997, 2002), symbolic culture (Donald, 1991), cumulative culture and a strong dependence on enculturation (Richerson and Boyd, 2005); a very large group size for primates (Dunbar, 1998); complex social organisation including group formation at multiple levels, and a division of roles and of labour (Knight, 1991); altrification (the early birth of highly dependant infants) (Trevvarthen, 1986; Dissanayake, 1995); bipedalism; language (Pinker and Jackendoff, 2005); a highly developed innate social intelligence, or 'theory of mind' (Dunbar, 1998); a reflexivity or conscious awareness (Knight, 2002); and numerous other aspects of intelligent behaviour such as a capacity for design and planning, the use of tools and a general purpose problem-solving ability and capacity for creativity (Boden, 1990).

Other factors have also been noted: for example, Knight (1991) points out that human females are different from other primate females in that they have highly visible menstruation periods but hidden ovulation periods (the opposite of most primates), pointing to a change in female social organisation which may have driven other social changes. There are also many ways to carve up and frame the various aspects of human psychology and cultural behaviour.

For example, in the context of the evolution of human musical behaviour, Tolbert (2007), taking a semiotic approach, considers referentiality important, whilst Whitehead (2007) focuses on pretend play. Furthermore, the key functional application of human intelligence is also commonly debated. For some time a developed hunting ability was considered the main fitness gain resulting from human intelligence, whilst in recent years this view has lost favour to a theory of social intelligence under increased predatory pressure.

A significant factor for contemporary theory is that of social intelligence and the notion of a theory of mind (Whiten and Byrne, 1997). It is now widely believed that human intelligence is at least as much about social intelligence as it is about general purpose intelligence involved in manipulating the physical environment. This supports a view of human evolution in which an arms race of social intelligence was initiated, possibly under other selection pressures that favoured increasingly large and cohesive social groups. This evolutionary process can be seen as a struggle between altruistic and selfish forces in social interaction, and has itself been informed by theoretical developments in the study of the evolution of altruism and of communication (for example Axelrod, 1984; Owings and Morton, 1998), and numerous studies of primate social behaviour that underline the basic dynamics of primate social cohesion (for example Richman, 2001). Questions of the dynamics of social organisation and cultural behaviour have taken centre stage in this most recent set of theoretical developments. Importantly, language, as with other forms of animal communication, can be viewed as a way to selfishly manipulate other individuals, rather than simply to work cooperatively with them (Dawkins and Krebs, 1978; Krebs and Dawkins, 1984; Owings and Morton, 1998).

Social brain, or Machiavellian intelligence, theory (Dunbar, 1994, 1998; Whiten and Byrne, 1997; Dunbar, 2004, 2006) views the evolution of a theory of mind as a consequence of the competition that comes about through life in a close group. As a natural consequence of Darwinian theory, if the ability to manipulate others, especially in the formation of short term alliances, has immediate fitness gains, then this ability will be selected for. Theory of mind is the ultimate such ability and can be viewed as the natural outcome of such a process (Dunbar, 1994).

Social brain theorists have proposed hypotheses for the evolution of language in which language is understood primarily as a means to manipulate others' behaviour, rather than as a

medium for conveying accurate information between individuals with shared interests (Dunbar, 2004; Knight, 2002). This is based on widely recognised theoretical problems associated with a view of language simply emerging under the pressure of its own functional utility for groups of individuals (Krebs and Dawkins, 1984).

This is concordant with a broader theory of animal vocal communication (Owings and Morton, 1998), in which the dynamics of communication are viewed in terms of the costs and benefits to the various interacting participants, whose interests are not necessarily shared. In some cases there may be a direct mutual benefit to the communication, whilst in other cases a more complex set of relations may take place (see simulations by Noble, 2000), possibly including the manipulation of one individual by another, or some complex combination of costs and benefits over the long term.

Thus we can view human language as exhibiting a combination of mutual benefit and manipulation in modern human behaviour and also varying degrees of referentiality and meaning. The question is what route language took to this point. Within the context of these issues, music appears to take on increased significance as a participant in the evolution of human behaviour, rather than just an oddity. However, the same can be said given developments in the understanding of signalling in sexual selection, which shares a similar historical development to that of social interaction (Ridley, 1994).

Alongside developments in human evolutionary theory, psychological research on the perception of music has also made inroads into the exploration of aspects of music in the brain, with increasing value following the development of techniques for analysing neural activity (for example Peretz, 2001; Altenmüller, 2001; Koelsch et al., 2004). A basic dissociation between music and language has been investigated by numerous researchers and this continues to be an active area of research and debate (for example Patel, 2006). Although this research is beginning to throw up significant results, the implications for an evolutionary theory of music are further complicated by the problems of determining dissociation, functional independence and innateness for musical capacities (McDermott and Hauser, 2005). For the moment, this uncertainty at least lends credibility to the potential for an innate, dissociated capacity for music in the brain, although, even if this were not the case, this would not rule out the possibility of a leading

role for music in other aspects of human evolution such as language.

The pioneering work of Colwyn Trevarthen (Trevarthen, 1986, 1993) in mother-infant interactions in the 1980s and 1990s demonstrated a basic innate mode of rhythmic interaction being used and understood by infants, which various commentators have argued bear the essential traits of both musical and linguistic behaviour. This interaction has been viewed both as a mechanism for learning language and as a basic human capacity to engage in social interaction using various aspects of entrainment, alternation, prosody, gesture and imitation. Notions such as attention and referentiality bear relevance to this perspective and have informed discussions of music. Tomasello (1999) presents a complete theory of the evolution of human cognitive abilities stemming directly from capacities for sharing activity.

The contribution of ethnomusicological research to this debate has been relatively minimal, although the theories of John Blacking have been influential in grounding a discussion of music's essential role in social life. Blacking (1995) argued that music is a human universal if not in its production then at least in terms of the essential grasp all normal human individuals have of music. He treated this as a psychological given upon which to build an approach to ethnomusicological research into the use of music in context, and the social construction of musical contexts. In general, it is probably fair to say that ethnomusicologists, following trends in anthropology (Kuper, 1996), have focused on the heterogeneity of music rather than its universal features (Nettl, 2005). Cross (2007b) points out that, with the exception of Blacking (1995), Nettl (2005) and Merriam (1964), ethnomusicology has not concerned itself greatly with definitions of music, an indicator of its extreme cross cultural flexibility. A significant contribution from ethnomusicology, relating naturally to this focus on heterogeneity, lies in demonstrating the diverse range of social contexts in which music is put to use (Nettl, 2005), allaying any suspicion that music has a highly focused and universal domain of application for modern humans, such as courting, or communication with or amongst children (Cross, 2007b).

A dramatic development came in 1995 with the discovery of an object at a Mousterian site in Slovenia believed to be the earliest known musical instrument: a small bone flute aged at around 44,000 years (see Wallin et al., 2000). More exciting still, this flute was believed to have belonged to a Neanderthal, rather than to one of our own direct ancestors. Unfortunately, the

excitement surrounding this artefact was premature, as recent evidence suggests that the appearance of a flute is merely coincidental, the holes in fact being tooth marks. Thus at least some of the driving force behind the emerging interest in the evolution of human musical behaviour was actually based on an unfortunate misclassification (see Cross, 2006).

Around the beginning of the 1990s, various theorists on the evolution of humankind were beginning to allude to the significance of musical behaviour in a number of ways. Dunbar (1994, 1998, 2004, 2006), one of the pioneers of social brain theory, pointed out the possible musical nature of early human language. Donald (1991) proposed that early human communication was multimodal, and points to possible relationships between musical rhythm and other forms of coordinated movement. Knight (1991), meanwhile, proposed a controversial menstruation theory of the evolution of culture based on the problem of establishing the basic social conditions for modern human communicative strategies. Richman's (2001) studies of gelada baboon calls point to the similarities between these vocalisations and human language and music structure and attempt to outline the social context and evolutionary dynamic of these calls. Meanwhile, Miller (2000a) has put forward a sexual selection approach to music embedded in a general sexual selection view of human psychology, whilst Dissanayake (1988, 1995, 2000a,b) has written extensively on the notion that aesthetic practices play a fundamental role in human social interaction. Nils Wallin's (1991) publication "Biomusicology" defined a new discipline covering evolutionary approaches to music as well as psychological, neurophysiological, cultural and ecological aspects of music. Wallin's approach was relatively independent of any one academic lineage and possibly suffered as a result from being perceived as overly convoluted. The term is yet to be widely accepted.

A collection of articles was published under the title "The Origins of Music" (Wallin et al., 2000). In the introduction the authors, Steven Brown, Nils Wallin and Björn Merker, propose that "just as music brings us in touch with the very deepest levels of our emotions, so too the study of music evolution has the potential to bring us in touch with the very deepest aspects of our humanity, our origins, our reasons for being" (Brown, 2000, p. 21). "The Origins of Music" can be seen as summing up the first explicit round of debate on the evolution of human musical behaviour, representing the key theoretical positions and gathering a broad pool of knowledge

from an impressive number of disciplines. At this stage, arguments tended to fall reasonably squarely into accordance with one category of evolutionary theory or another. Justus and Hutsler (2005) and McDermott and Hauser (2005) argue for a more rigorous approach to the problem of whether musical behaviour does indeed constitute a genuine human adaptation, sparking off a more thorough and focussed debate. In this round, views appeared more open to the subtleties of evolutionary processes, let alone the subtleties of trying to make inferences from psychological, archaeological, primateological, or other forms of data. This debate still largely focuses around the task of gathering evidence that directly supports the statement that human musical behaviour is a direct adaptation, rather than engaging in a theoretical debate about the range of processes musical behaviour may have been involved in during human evolution.

This provides an outline of some of the main theoretical trends underlying the recent development of a number of competing theories of the evolution of human musical behaviour. I will now consider those theories in turn.

2.2.2 Key views on the evolution of human musical behaviour

2.2.2.1 The early debate

An objective of this thesis is to analyse existing theories of the evolution of human musical behaviour according to their essential nature as abstract evolutionary processes. I therefore begin this section by proposing four categories of evolutionary theory that can apply to music; *consequentialist*, *competitivist*, *cohesionist* and *cognitivist*. Consequentialist theories are those that deny music an *adaptive function* that could be used to explain its evolution; music is either an offshoot of other evolutionary forces or a technology that is built upon existing non-musical aspects of our physiology and psychology. Competitivist theories view music as a system that has evolved through some form of intraspecies competitive evolutionary process, the typical case being sexual selection. Cohesionist theories view music as a device for forming strong bonds between individuals in a group, and generally tend to argue that such types of interaction dedicated to mediating social bonds are a benefit to individuals in the way that they shape social organisation. Cognitivist theories place an emphasis on the role music plays in facilitating or mediating certain types of thought, such as the ability to manipulate abstract cognitive structures.

This categorisation leads to a set of idealised mutually independent theories. I argue that these are the four most prominent categories of evolutionary approach as far as the evolution of human musical behaviour is concerned. Existing theories are more often combinations of these categories, and might involve periods of evolution that build upon each other under shifting environmental pressures. I consider each of these categories in turn, beginning with consequentialist theories, which act as the null hypothesis for a theory of the evolution of human behaviour by natural selection.

Pinker (1998) takes the consequentialist position that human music is a social development built upon existing cognitive capacities. Music is viewed, first and foremost, as a behaviour that lacks adaptive function, both for selfish individuals and for cooperative groups. Huron (2001) describes Pinker's theory as the NAPS (nonadaptive pleasure-seeking) theory of music. According to a now infamous analogy, music can be likened to cheesecake, sharing the feature that we enjoy it despite it being of no use to us: cheesecake contains a combination of nutritional ingredients, but their arrangement in the form of cheesecake itself serves no purpose, and the excess of each ingredient is ultimately bad for us. Pinker's point is that complex evolutionary systems such as human brains will have inevitable but unnecessary side effects: "Some parts of the mind register the attainment of increments of fitness by giving us a sensation of pleasure. Other parts use a knowledge of cause and effect to bring about goals. Put them together and you get a mind that rises to a biologically pointless challenge: figuring out how to get the pleasure circuits of the brain and deliver little jolts of enjoyment without the inconvenience of wringing bona fide fitness increments from the harsh world" (Pinker, 1998, p. 524). Thus for Pinker a bundle of aesthetic factors and mental capabilities come together to produce our love of music and our musical ability. Each factor is a separately evolved human cognitive trait: the faculty for language; auditory scene analysis; emotional calls; habitat selection; motor control; and other stuff (Pinker's term). Each of these factors can be ascribed an independent adaptive function which explains their evolution. But music itself is something that, over the course of human history, has been adapted and improved so as to maximise the aesthetic experience brought together by the various biological components that make up what is known as music perception. Since adaptive development in a given environment is hugely instrumental in the formation of

brains, Pinker's theory is, to some extent, able to account for the development of the dedicated neural networks for music perception and production that have been observed to emerge in the brains of skilled musicians (*e.g.*, Altenmüller, 2001; Schlaug, 2001), although evidence from musical perception in infants, (*e.g.* Trehub, 2001), is less convincingly explained away by the influence of other adaptive functions.

Miller (2000a) takes the competitivist view that musical behaviour evolved through sexual selection. He agrees with Pinker that music is of no use to individuals or to groups in their relationship with the extraspecies environment, but argues that musical behaviour is sufficiently complex that it requires an evolutionary explanation: "Human music shows all the classic features of a complex biological adaptation" (Miller, 2000a, p. 330). But, as he points out, given that a general form of altruistic information sharing is not evolutionarily viable, animal signalling systems have only been shown to evolve for a small handful of evolutionary reasons, including threats, altruistic warning calls between kin, contact calls, dominance and submission signals, and courtship displays. From this list, he argues, sexual selection is the only viable candidate, and befits music very well.

Sexual selection is a process of natural selection in which adaptation occurs in response to the task of attracting a choosy mate (Darwin, 1883; Andersson, 1994). This task typically falls upon males, since females are the choosier sex due to their greater investment in having children, which is in turn associated with the greater certainty that the children are in fact their own (Bateman, 1948). Males who are better at attracting mates are considerably more successful at producing offspring, while less attractive males are denied opportunities to mate and are disproportionately unsuccessful. Due to their promiscuity and the low cost of reproduction for males, successful males don't need to be so choosy (Bateman, 1948). Under these conditions it is possible that changes in males are adaptive due to their effects on sexual attractiveness even if this change is less adaptive in other respects: the benefits of a certain change outweigh the costs (Darwin, 1883). Accordingly, males will evolve according to females' attractiveness criteria. However, the reason that females are attracted to some males more than others in the first place is that evolution has endowed them with means to attempt to indirectly determine the best genes for their children (Andersson, 1994). Females who chose more attractive males despite the

fact that they are less fit in other respects will not ultimately have the reproductive success of females who are attracted to the fittest males. Accordingly, female attractiveness criteria evolve to be more successful at finding truly fit males (Zahavi, 1975).

Sexual selection can therefore be seen as a coevolutionary “arms race” with males competing to score high on female attractiveness measures (which is always balanced with other fitness criteria), and females refining their attractiveness measures to get the true best males. The result is *runaway selection* (Fisher, 1958) of attractiveness traits limited by an important *handicap principle* (Zahavi, 1975; Zahavi and Zahavi, 1997), which states that an evolved attractive feature must have a higher relative cost to an unfit animal than it does to a highly fit animal. In the language of the handicap principle, attractive features must be honest indicators; they reliably indicate fit individuals (see Section 2.3.3). Sexually selected traits are therefore potentially good indicators of fitness, despite the apparent contradiction that they in fact reduce fitness. In Miller’s theory, the handicap principle is provided by the fact that only a correctly functioning, healthy brain can achieve certain cognitive feats. He proposes that “the more important brains become in human survival and reproduction, the more incentive mate choice would have had to focus on brain-specific indicators” (Miller, 2000a, p. 340). Musical ability in males could therefore plausibly act as an indicator of fitness to females. The more impressive a male’s song, the more a choosy female can be sure that his brain is in order. Musical performance bundles together a number of more specific indicator functions. In particular, Miller argues that creativity and learning ability could be indicators of different types of intelligence, as well as the more traditional showing-off that one can waste time on fruitless activities. He accepts that these are speculations, but argues that they are testable in modern humans.

As well as indicator traits, Miller argues that music can also be subject to runaway sexual selection based on initial biases, in which case nothing is being indicated. Variation in female sensitivity to male musical performances will continue to ensure that, potentially, increasingly positive responses to music will arise. He adds, “The aesthetic and emotional power of music is exactly what we would expect from sexual selection’s arms race to impress minds like ours” (Miller, 2000a, p. 331). Miller’s view that “human music shows all the classic features of a complex biological adaptation” is strongly supported by his presentation of the facts of sexual

selection, although his effort to bring supporting evidence to the argument is weak: the use of music to serenade, the predominance of males in music production, the promiscuity of rock stars, and the similarities between human music and avian song do not appear to be scrutinised as they are incorporated into this theory. Most importantly of all, however, Miller pins his analysis on the assumption that a signalling system can only emerge in any species if it belongs to one of the categories listed above (threats, altruistic warning calls between kin, contact calls, dominance and submission signals, and courtship displays). This is peculiar given that, even if a sexual selection theory of language was vindicated, this would not change the fact that language is used by modern humans in ways that extend this list, unless Miller's very own arguments are to be interpreted as dominance signals. If these are unusual new uses of signalling systems unique to modern humans, then the ancestors who helped drive the evolution of music and language could also arguably have possessed such unique traits.

I describe Miller's theory as *competitivist* because it hinges on competition for goals measured entirely within the species rather than with respect to the environment. All species compete amongst themselves for limited resources within an environmental niche, and this competition drives evolution. Thus natural selection, one could argue, is always competitive. The specific focus of *competitivist* theories in the sense intended here is that they involve competition in which the species' behaviour is both the competitor and the evaluator of that competition: it is competition with respect to intraspecies evaluation. Due to the implications of the handicap principle, Miller's (2000a) theory is also *cognitivist* in that it draws on a relationship between musical behaviour and aspects of our general intelligence: music indicates intelligence. Other examples of competitive evolution include the development of traits used directly for combat between individuals (Fisher, 1958), and psychological traits such as a theory of mind capacity (Dunbar, 1998) that has been touched on already in this section. Theory of mind capacities facilitate direct social competitive behaviour by giving individuals the ability to outwit their conspecifics in games of alliance formation (Dunbar, 1998). However, we will also look at possible less direct forms of competitive interaction in later sections.

Merker (2001) looks at a more specific aspect of music in its role of increasing reproductive success. He argues that synchronous chorusing by males may have evolved in order to call

to potential mates over greater distances. This does not mean that synchronous calling is a sexually selected trait, since it does not depend on judgement by the female, only on whether she hears the males' calls or not. Groups of males who can synchronise their calls can mutually increase their reproductive success by potentially attracting females over a wider area. However, if sexual selection is also operating on calling behaviour, if the attractiveness of the call is as important as its audibility, then Merker's perspective implies a process of sexual selection in which synchronous chorusing is grounded. The suggestion evokes a rich set of considerations for how musical behaviour subsequently evolved through the interplay between synchrony and the patterns generated during chorusing. The synchrony element can be seen as a cooperative component setting up a context within which competition takes place. This is analogous to the phenomenon of similar shops gathering together in distinct areas of a city.

Cohesionist theories add to the debate the notion that musical behaviour has an adaptive function on a group level or works to the immediate mutual benefit of a number of participating individuals. According to Dunbar (1994, 1998, 2004), human brain size evolved under the pressure to live in larger groups, which was kick-started by a dramatic forced change in habitat, before which there was greater pressure to act selfishly. Moving from forest to savannah environment, the theory states, our ancestors were exposed to a greater threat of predatory attack. Under such circumstances the advantages associated with forming large cooperative groups began to outweigh the advantages of selfish behaviour. This explains the emergence of group level devices for binding potentially selfish conspecifics into alliances. The aesthetic and emotional power of music alluded to by Miller, is also used as evidence of an adaptive function for music as a cohesive device, and is taken by others to argue this point of view (*e.g.*, Mithen, 2005; Cross, 2001, 2003a,b; Huron, 2001; McDermott and Hauser, 2005; Dissanayake, 1988).

It should be noted that there is a direct overlap and potential confusion between cohesionist and competitivist theories. As mentioned in Section 2.2.1, Dunbar's model describes the evolution of a theory of mind through a competitive arms race under the pressure of living in large social groups. Theory of mind is not seen as a cohesive mechanism, as such, but as an evolutionary consequence of a cohesive requirement, actually driven by competitive interaction. On the other hand, the establishment of a theory of mind in a population alters the dynamics of in-

teraction, facilitating genuinely cohesive behaviours based on mutual benefits and the increased capacity to identify cheats and punish them. In the following discussion, this interrelation between these two processes should be borne in mind.

Dissanayake (1988, 1995) was one of the first authors to consider a cohesionist perspective, which she applies to artistic behaviour in general, or more precisely, to a general cross-cultural category of behaviour that she describes as *making special*. Dissanayake's earlier work was radical in proposing this point of view, which did not fit well with established evolutionary theory at the time, although she was considerably more concerned with challenging the dominant post-structuralist and post-modernist understanding of art in the West, which she believed was politically problematic.

Mithen (2005) is one of the many theorists particularly concerned with this cohesive role, and is intrigued by the relationship between music and language over evolutionary history and during ontogenic development. Of particular importance to the evolution of music is whether language evolved through the formation of a compositional system, in which mechanisms emerged to combine existing meaningful words into more powerful meaningful utterances, or as the refinement of an earlier holistic language system, in which meaningless complex utterances were refined over time, attached to meanings, and decomposed into meaningful components (Wray, 2002). The latter model suggests that proto-language was more like music than modern language, and suggests a theory of the evolution of language and music from a common origin. Mithen (2005) argues that holistic proto-language was predominantly manipulative: it was used by individuals not to share information about the environment amongst others (as modern language can do), but to manipulate others' actions, in particular through the expression, and perhaps direct control, of emotional states. He proposes that the ability to communicate complex emotions provides a mechanism for increased social cohesion, and that music's relationship with the emotions is such that it encourages a state of *boundary loss*, a term introduced by McNeill (1995), in which individuals increase their sense of unity with the group.

Another important aspect of music's role in this process is found in systems of mother-infant interaction. These too could be said to be musical and multi-modal. Infant-directed speech (IDS) appears to pass through four stages during a child's development, and each stage is associated

with a phase of learning. Stage one is directed towards engaging attention, and uses just three or four basic vocal patterns which have a predictable effect on the infant. Stage two modulates arousal and emotion. Stage three involves the communication of emotion proper. Stage four is geared towards the acquisition of language. Stages one, two and three are as much musical as they are linguistic (Mithen, 2005).

Mithen's theory is not especially well supported by the limited archaeological evidence at hand and mainly focuses on outlining rough scenarios of social behaviour at various stages of human evolution and considering the likely level of development of proto-musical interaction and its domain of application. It does not adequately address the evolutionary mechanism by which a cohesive role for music should come about. Bispham (2006b) provides a thorough criticism of this work.

Richman (2001), a strong supporter of the holistic view of the origins of language, provides an equally speculative, but more theoretically rich, account of the evolution of music which is far more concerned with a specific evolutionary process than with an ultimate adaptive function as such. His theory can be viewed as competitive in this sense, because it does not provide an adaptive function for music, but does suggest that musical behaviour evolved through the natural selection of certain individuals over others in a social realm. Richman argues that gelada monkey 'friendly' sounds resemble speech and music in important ways, both in their structure, and in their use. Friendly vocalisations are "produced in units averaging a total length of about nine or ten syllables, produced at a rate of about five syllables per second, organised by differentiation of strong and weak beats, with about three or four strong beats per unit, and all under an intonation contour (melodic contour) where the end of the unit is signalled by tonal changes" (Richman, 2001, p. 302). Gelada monkeys use these vocalisations to negotiate the formation of bonds with conspecifics (members of the same species), and identify different vocalisations with different social contexts, displaying "a great variety of syllables with different consonant and vowel-like features, as well as a great variety of rhythmic and melodic variation" (Richman, 2001, p. 302).

Richman argues that a key difference between gelada and human vocalisations is that for geladas "a given sequence does not tend to be repeated as a vocal formula. In other words, we have no evidence that geladas are capable of repeating the same, exact succession of vocal

features that would mark two vocal units as the same. At most, they can repeat or vocally match specific phrases, up to three syllables long, that they have just heard” (Richman, 2001, p. 302). He adds that holistic approaches to real speech in humans suggests that in spoken languages we commonly draw on a huge set of learned vocal formulae, rather than constructing original grammatical structures on the fly. The question for Richman, then, is how to explain an evolution from gelada *meaningless* vocalisations to human meaningful linguistic formulae. His answer lies in “regular expectancy based on repetition and a regular beat”, pointing out that these are essentially musical dimensions (Richman, 2001, p. 304).

Richman proposes three *redundancy devices* in music: *repetition*, *formulaicness*, and *expectancy*. He argues that “our ancient ancestors, at the beginnings of music, familiar with only a few fixed formulas, must have depended hugely on the other two redundancy factors. They needed much constant repetition by everyone and a lot of expectancy of what was to follow to hear sequences as recognisable and hence repeatable” (Richman, 2001, p. 304). What must also be needed, he proceeds, in order for the formulaic aspect of music to evolve, is for individuals to be driven to try to imitate each other, to synchronise with each other, and to tie newly found formulae into “intensely felt, highly particular, multi-modal scenes of real life” (Richman, 2001, p. 306). The latter process would depend on an associative learning that connects ritualised behaviour to salient events. This is another big leap, but one that is entirely possible given a culture of such proto-linguistic interaction.

Taking issue with the theories of Pinker and Miller, Cross (1999, 2001, 2003a,b, 2004) argues that greater consideration needs to be given to the assumption that music has no adaptive function. In addition, he insists on taking the perspective of anthropologists and ethnomusicologists as the crucial starting point for understanding how music is used and how it is perceived. Only these areas of research avoid bias towards the isolated individual in a subject that is so clearly a social phenomenon, and emphasise the incommensurability of musical experience from one context to another, leading to the assertion that musics, in the plural, “are musics only in their cultural contexts” (Cross, 2003b, p. 19). With this in mind, Cross attempts to redefine music in terms that acknowledge the diversity of use and significance in musical contexts: “music is not only sonic, embodied and interactive, it is bound to its contexts of occurrence in ways that enable

it to derive meaning from, and interactively to confer meaning on, the experiential contexts in which it occurs, these meanings being variable and transposable” (Cross, 2003a, p. 108).

The important fact about music is that it “can appear to be *about* something, but its *aboutness* can vary from context to context, within a context, and from individual to individual” (Cross, 2003a, p. 3). This characteristic is exploited in our earliest social learning experiences, as described by studies of mother infant interaction by Travarthen (1986; 1993), Papousek (1996) and Dissanayake (2000a), where it is adaptively functional in terms of group cohesion. This unique interaction is closely associated with human altrification: the increase in the length of time a child is dependent on his or her parents. This in turn is closely related to the existence of culture, since it is the point at which the great variety in adult behaviour, described as cultural, is established. A central concern for Cross, then, is that music is functional in making cultural individuals. Just how it is functional, and how tightly intertwined it is with the rest of human behaviour remains little understood.

On the surface, Cross, Dissanayake and Mithen’s theories all appear squarely cohesionist, allowing for the caveat that this may indeed be a competitiveness process in disguise. But in all three there is an inkling that musical behaviour goes beyond social intelligence. For example, for Dissanayake (2000a), early interactions provide benefits for infants including, amongst many other things, the development of “cognitive abilities for recognising agency, object, goal and instrumentality, a narrative-like mode of thought and perception ... and [the predisposition of] the infant generally to intellectual and social competence, including intentionality, reciprocity, and expansion beyond the present situation” (Dissanayake, 2000a, p. 393). Most of the elements mentioned here, as well as all of the other benefits mentioned by Dissanayake, are to do with social intelligence, and this is clearly where the emphasis lies. But this passage also suggests that the interactions bear equally on non-social aspects of our intelligence.

Cross also discusses the relationship between musical thought, or the capacity developed through protomusical interaction, and general flexible thought that is characteristic of human intelligence: “protomusical activities provide mechanisms for acquiring cognitive flexibility as well as consequence-free means of exploring and achieving competence in social interaction.” (Cross, 2003b, p. 27). Like Dissanayake, Cross frames this discussion in terms of social in-

telligence, but this cognitive flexibility is not explicitly distinguished from intelligence that has conferred advantages to humans in their relationship with the environment.

Thus both Cross and Dissanayake indicate that this learning is also beneficial to a general intelligence, which is itself adaptively functional with respect to our manipulation of the physical environment. For Cross music is a powerful model of versatile human cognition. It is for this reason that I categorise Cross' theory as being a combination of cohesionist and cognitivist.

With reference to earlier work by Mithen, Cross elaborates on this relationship between music and cognitive flexibility. Mithen's (1996) theory of the evolution of the human mind proposes that, in accordance with contemporary evolutionary psychology (see Barkow et al., 1992) the mind is modular, consisting of evolved domain-specific competences, with some degree of general connectivity that accounts for our flexible metaphorical and analogical thought. However, in Mithen's more recent work on the evolution of music, the same claim is not explicitly made, nor is it explicitly denied. Whilst I have distinguished between cohesionist and cognitivist orientations, it is quite possible that in practice they are rarely independent of one another. I maintain that the distinction is useful, because it allows us to identify where each theory is placed on the spectrum between the two pure categories.

2.2.2.2 Recent development of themes in the debate

There is at least good suggestive evidence for each of the four theory-types that I have defined. For each theorist's convictions, there is also a plausible theory as to why they might have been *misled* by the poverty (or enormity) of data at hand. Plotkin (1997) reasons that if the human brain can be considered the most complex system in the known universe, then human social systems can be considered more complex still, consisting as they do of human brains. These discussions indicate a consensus that musical behaviour must lie towards the more complex end of human social behaviour incorporating a vast number of interacting behavioural and evolutionary phenomena. It is no surprise that competing theories generally fail to incorporate every contributing factor.

One fact that all theorists treat as critical is the importance of music for modern humans. In the case of Western cultures, this fact is hard to avoid. For non-western cultures, anthropological

and ethnomusical data generally supports the claim: for example according to Huron (2001), “In the Atlas mountains of Morocco, full-time Jujuka mountain musicians are supported by the local villagers. That is, there is an entire caste of people whose principal productive activity is music-making. A ready index of the importance of music in such a society, ” he suggests, “may be the ratio of the number of musicians to the number of farmers and herders.” (Huron, 2001, p. 46).

For Pinker (1998), this is a phenomenon that needs to be explained in spite of evolutionary theory predicting that we should not be spending time making music. For the large camp who oppose the Pinkerian view, it is a phenomenon that needs to be understood in evolutionary terms, although not wholly in genetic evolutionary terms, but with respect to culture.

The maturation of the debate following this basic staking out of interests has two main strands. Firstly, several theorists and experimenters have argued over the psychological and comparative data supporting the cognitive innateness and domain specificity of musical behaviour in the brain. How can we determine once and for all whether or not human musical behaviour is a biological adaptation? Secondly, the potential ambiguity of *any* outcome from this problem of evidence has driven a more theoretically open discussion concerning the nature of adaptation and evolutionary process in contexts relating to human musical behaviour. This can be generalised as the problem of trying to explain musical behaviour as part of a complex system of social interaction that exists against a backdrop of evolution by natural selection.

In response to Justus and Hutsler (2005) and McDermott and Hauser (2005), who present the various problems in experimentally testing the adaptiveness of musical behaviour, Dean and Bailes (2006) argue for the importance of focusing on the “maintenance and transformation of musical activity, rather than on origins” (Dean and Bailes, 2006). Dean and Bailes advocate a combined approach to cohesion phenomena and sexual selection, proposing a role for studies such as experimental economics in gaining an understanding of potentially relevant social dynamics.

Similarly, Fitch (2006) stresses the overemphasis on the question of adaptation and states a preference for investigation into ‘intuition pumps’ for the generation of testable hypotheses. For Fitch, the problem of music’s existence in an evolutionary context is straightforward, since it is heritable (either culturally or genetically); if it were “truly useless” then it should be “a

target of selection to disappear, like a cave fish's eyes" (Fitch, 2006, p. 86). He questions the distinction between culture and biology, and between innate and learnt behaviour, arguing that the human musical capacity should be seen as an "instinct to learn" (Fitch, 2006). More important, he suggests, is the task of trying to ascertain how music acts in social interactions and whether musical abilities or behaviour have any relation to reproductive success. Trainor (2006) also takes the view that identifying innateness is problematic, and that an understanding of how music could have exerted selection pressures is desirable.

Livingstone and Thompson (2006) add that "such strategies allow researchers to constrain theories of music as an adaptation to environmental pressures; they provide less guidance for theories that implicate exaptations or secondary adaptations" (Livingstone and Thompson, 2006, p. 89). Merker (2006) argues that vocal learning is a fundamental aspect of music that needs further study and points to iterated learner models that can be used to explain some of the emergence of structure in language and music entirely through ordering processes in cultural transmission, suggesting additional dimensions to the context in which we understand the evolution of human musical behaviour to have taken place. Likewise, Patel (2006) argues that vocal learning is an uncommon trait in primates, and that auditory rhythmic entrainment in humans is striking given that it is so much stronger than visual rhythmic entrainment.

This thesis stems directly from these lines of enquiry, particularly from Fitch's view that we need to explain music's non-erasure regardless of its relation to any innate genetic determination in the brain, and it is therefore not necessary to consider the details of the debate on neurological innateness in this review. Instead, I review the theoretical content of these approaches to the evolution of human musical behaviour. In particular, cohesionist theories face the challenge of explaining how music comes about as a system of interaction that holds greater benefits for those that adopt it than to those that do not.

2.3 Contextualising the Evolution of Human Musical Behaviour Within Issues in Evolutionary Theory

In this section I introduce Darwin's theory and some major themes in the modern theory of evolution, in particular the emergence of cooperative, coordinated behaviour. Given the various

models of the evolution of music discussed in Section 2.2.2, this discussion is geared towards evaluating the construction of these models in general theoretical terms. My review of the literature focuses on the relationship between natural selection and the existence of structures at different levels of organisation, including problems of group selection, Maynard Smith and Szathmáry's major transitions (Maynard Smith and Szathmáry, 1995), and Maturana and Varela's concepts of autopoiesis and structural coupling (Maturana and Varela, 1980, 1987).

2.3.1 The fundamental debates of Darwinism

The theory of evolution predominates biological explanations of the natural world, and has been predominated itself, throughout most of the 20th Century, by the *modern synthesis* (Huxley, 1975) – a combination of Darwin's theory of natural selection, Mendel's theory of genetics, and the mathematics of population genetics. Darwin's (1860) original theory postulates three ingredients for evolution to occur in a population of reproducing individuals: heredity; variation; and selection. A reproducing population produces new individuals which are identical to their ancestors (heredity) except for slight differences due to mutations during the copying process (variation). The existence of minute differences in physiology caused by this variation results in minute differences in reproductive success (which incorporates the ability to survive and the ability to produce lots of children), and those individuals with greater reproductive success will have a greater chance of ensuring the survival of their line of descent than will those with diminished reproductive success (selection). Natural selection emerges from the constraints on survival that result from competition for resources in a limited environment, which is most pronounced between similar individuals since they have similar needs. In this way, Darwin's (1860) theory of evolution takes a giant step towards explaining the relationship between all biological species (as connected points on the *tree of life*) and, at the same time, their apparent adaptivity, their often staggering suitedness to a given environmental niche. Darwinian evolution is described as *process* rather than *mechanism* (Dennett, 1996) on account of the fact that Darwin had no idea how heredity and variation were actually manifest in biological systems. Natural selection can be viewed as a powerful algorithmic process regardless of its application to biology, and this fact has been reinforced by the successful application of genetic algorithms to engineering problems

(*e.g.*, Holland, 1975; Goldberg, 1989).

The modern synthesis sees Darwin's theory as applying predominantly to strings of discrete units of information – genes – which determine how individual organisms will develop (Huxley, 1975). In almost all cases, these strings, called genotypes, are understood only to vary through mutation and sexual combination, and cannot be changed by the action of the individuals that carry them (Jablonka and Lamb, 1995). Although Darwin himself accepted the possibility that individuals could change their genetic makeup through their own actions, known as Lamarckian evolution (Darwin, 1868), with the hindsight that Lamarckian evolution cannot occur in genes because it is not possible to re-encode adaptive changes back into genetic information (the flow of information between genotypes and phenotypes is one-way), the process of Darwinian evolution has taken on a hugely more creative role in our understanding of the design of biological systems (Huxley, 1975). Evolution's most successful public spokesperson of recent years, Richard Dawkins (*e.g.*, Dawkins, 1976), has sought to emphasise the extent to which science can reduce much of biological nature to the process of the natural selection of genes. For example, Dawkins (1976, 1986) tackles the charge that Darwinian evolution alone could not have led to such complex, modular systems as the vertebrate eye. The complaint about natural selection's applicability to such systems is that they are made of hugely interdependent components, none of which would have any use without the others, and are therefore irreducibly complex. He argues plainly that this is a misconception, that one can interpolate between simple light-sensitive patches of skin and full vertebrate eyes, and that each stage in this interpolation involves a step towards improved vision (Dawkins, 1986). Thus in Dawkins' view, extremely gradual changes can indeed lead to complex and sophisticated systems that have the appearance of intelligent design. Dawkins' approach implies that the ceaseless action of natural selection is such that scientists should have strong reservations about calling other forces into play when explaining complex physiologies and behaviours. In short Dawkins' view upholds two great but controversial tenets of modern Darwinian theory: gradualism and adaptationism.

Others argue for the increased significance of additional factors in the evolutionary process, especially when viewed on a macroscopic scale. Owing to its elegant and pervasive simplicity, Darwinian evolution runs the risk of inspiring an image of mutable organisms adapting to fixed

environmental niches. But it was important to Darwin and to any of his better-known intellectual descendants, Dawkins included, that environmental niches are themselves composed of evolving entities. This implies that all evolution is in fact open-ended *coevolution* (see Ridley, 1994, for a popular account) and that additional processes need to be evoked in order to reach a more global understanding of the biological phenomena that surround us. Whilst this is well understood, it is argued from one camp that theorists who lean too heavily on natural selection, practising what is known as naïve adaptationism, can only explain processes at a certain level of detail. Gould and Lewontin (1979) argue that chance has a much greater role to play in a truer explanation of biological phenomena than naïve adaptationism permits, and that the process of natural selection is not reducible to a process acting solely on genes, since the role individual genes play in the formation of phenotypes is so strongly affected by other factors. Gould is generally sceptical of both gradualism and adaptationism (Gould, 1977; Gould and Lewontin, 1979; Gould, 2002) (see also Lewontin, 1979). In the first case he cites evidence that many species simply did not undergo change for huge periods of evolutionary time, and when they did they did so very suddenly. Thus they were not involved in the sort of complexity-increasing arms race that Dawkins describes. Gould refers to this mode of macroscopic evolution as *punctuated equilibrium*. In the second case, he argues that adaptationist explanations all too easily obscure simpler processes such as the structural requirements of building a certain body shape. The complexity of structure and behaviour exhibited by all evolved organisms is therefore seen to throw up artefacts which contribute to the context in which further change takes place. A key concept for Gould then is *exaptation* (Gould and Vrba, 1982), the process by which an existing phenotypic trait (adaptive or not) becomes adapted for a new purpose. Gould's view of evolution is neatly represented in the analogy of Galton's polyhedron. Imagine a convex polyhedron resting on a table top, the orientation of which represents the state of an evolving organism. Selection pressures are represented by gentle lateral forces acting on the polyhedron. As the forces change, the polyhedron may shift from one face to another, but this always involves a rapid change followed by a stable state. A modern version of this perspective uses the notion of basins of attraction. As well as exaptation, Gould and Lewontin (1979) also introduces the term *spandrel* to describe traits whose original evolutionary origin is obscured by their subsequent evolutionary development.

Sterelny (2001) provides an informative summary of the classic debate between Dawkins and Gould.

The two sides agree, on the other hand, that the process of natural selection itself cannot be simplified. As Plotkin (1997) points out, evolutionary theory can be seen to fall upon *final cause* explanations rather than science's typical *efficient cause* explanations (we owe both terms to Aristotle (Aristotle, 2007)); the first provides an explanation in terms of what an object's purpose is, whilst the second explains in terms of what process brought the object about). Darwin's theory cannot actually be used in this way; final cause explanation is at best a neat (or sloppy, depending on who you ask (Sterelny, 2001)) shorthand for efficient cause explanation: in science, final causes don't explain unless they are related to efficient causes. Natural selection can act as the process that mediates this relationship, but the details pertaining to how natural selection operates in real complex environments can never be glossed over and edged out to leave a final cause form of science. Popper, inspired by influential biologists of the time, described the theory of natural selection as tautological and untestable, and explained that scientific interest in natural selection lay in its value as a metaphysical research programme. He later retracted the statement that it is tautological and untestable (see Popper, 1978).

Heeding Gould's warning, it is unwise to evoke adaptation by natural selection without weighing up the effects of other processes (Orzack and Sober, 2001). A common heuristic then is that if a phenotypic trait is either improbably complex or demonstrably near-optimal in the light of some function then natural selection is likely to have had some part in its design (Dawkins, 1976). However, both of these conditions can be problematic. It is widely believed that complexity can emerge from simple processes (of course natural selection is one such process) (Highfield, 1996), therefore one must consider other routes to complexity when considering the origins of a complex system. Optimality was famously brought into doubt by Lovelock's Daisyworld model (Lovelock, 1979), which showed that organic systems can regulate their environments in such a way that an observer may be led to conclude that the organism has adapted whilst the environment remains fixed. Furthermore, in increasingly complex environments, and with increasingly complex structure, the factors that need to be taken into account when working out what effect any variation will have on reproductive success become increasingly numerous

(Kauffman, 2000). Phenotypic changes will inevitably lead to both advantages and disadvantages in different domains of life, according to different aspects of longevity and reproductive success. For example, greater visibility may attract both mates and predators; it is not a foregone conclusion that any such change should lead to greater reproductive success. Thus adaptation needs to be considered in terms of sets of costs and benefits (Huxley, 1975), and a mutation can be evaluated according to the ratio between the benefits and costs associated with that change. This makes the task of singling out specific evolutionary processes fraught with difficulty.

2.3.2 The evolution of cooperative behaviour

One of the greatest battlegrounds for evolutionary theory, which continues to the present, is the matter of seemingly altruistic behaviour between individuals either of the same or different species. Selfish-gene theory (Dawkins, 1976) defines genes as the key protagonists of evolution, and phenotypes as the vehicles that they have built over successive generations of natural selection. Existing individual genes will always tend to behave in their own best interests because they are the descendants of those genes which were best predisposed to survive. Selfishness is nothing more or less than the maximisation of one's own chances of survival, and so the notion of 'survival of the fittest' can only predict ubiquitous selfishness amongst genes. In most cases, this logic of selfishness extends to the organisms that these genes inhabit, and in countless cases this selfishness is exactly what is observed in nature. But there are also counter-examples in which individual organisms do not act to maximise their own fitness. The theory of inclusive fitness (Hamilton, 1963) was one of the earliest attempts to explain this phenomenon, and extends the basic principle of equations of costs and benefits to cases in which it is advantageous for an individual to act to the benefit not only of itself but of its close relatives. Since the genes that inhabit one vehicle are likely also to inhabit a close relative, it is of benefit to these genes that the relative has a good chance of survival, based on the assumption that the survival of either individual will serve the survival of the gene. On the other hand, there is some chance that these genes do not exist in the relative. The result is a strong correlation between the degree of relatedness between two individuals and the net benefit gained from acting altruistically towards each other. If r is the relatedness of two individuals, expressed as a fraction, then we can ex-

press inclusive fitness in the form of Hamilton's inequality: a change in behaviour with altruistic consequences is beneficial if the benefits to the relative, b , the costs to the individual, c , and the degree of relatedness are such that $b/c > 1/r$. However, as Maynard Smith and Szathmáry (1995) point out, this can only operate between direct descendants where the relevant mutation actually occurred: if the relative survives to reproduce but the altruistic individual does not, then the relevant adaptation towards altruism only survives if it is also possessed by the relative, thus inclusive fitness theory only works in cases where an adaptation towards altruism already exists in more than one individual, which may mean that other explanations become feasible (Maynard Smith and Szathmáry, 1995). Inclusive fitness can also lead to more complex sets of interactions such as the process of enforcement, in which one individual coerces another into acting altruistically based on their relative investment in certain relatives. For example, since sisters have an average relatedness of 0.5 and mothers and daughters also have an average relatedness of 0.5, then a female, A, would do equally well to care for her sisters as for her daughters. However, A's mother is related to A's sisters with an average relatedness of 0.5, but to A's daughters with an average degree of only 0.25. According to the theory of inclusive fitness, since it makes no difference to A but it does make a difference to A's mother, A's mother is likely to be successful in coercing A to devote time and energy to caring for her sisters rather than her daughters (Maynard Smith and Szathmáry, 1995).

Inclusive fitness clearly does not explain certain other forms of apparently altruistic social behaviour in less closely related individuals, and the only other generally accepted explanation is that of *mutual benefit* or *reciprocal altruism* based on game theoretic models of interaction (Maynard Smith, 1982; Maynard Smith and Szathmáry, 1995). In many cases, what we call cooperative behaviour simply acts to the benefit of both participants. However, a particularly interesting evolutionary possibility arises in a situation known as the prisoner's dilemma (Flood, 1952). The prisoner's dilemma is a generalisation of a common situation in which two individuals would do well to cooperate with each other, but one would do even better if they were able to act selfishly while taking advantage of the other's cooperation (Dawkins, 1976; Axelrod, 1984). The example that gives this situation its name is as follows: two criminals are caught red-handed at the scene of their crime. They are taken away and interrogated in separate rooms.

Each is given the choice to plead innocent or guilty, but their punishments will depend on the combination of both pleas. If they both plead guilty they each get a one year sentence. If they both plead innocent they each get a two year sentence. If one pleads guilty but the other pleads innocent, then the former gets three years and the latter is freed. Here pleading innocent is the selfish option: the participant is banking on being able to take advantage of the other, and in taking this option eliminates the outcome which is best overall for the pair (one year each = two years total). Pleading guilty is the altruistic option: the participant realises that the other could take advantage of them, but is willing to take this risk in order to allow the possibility of the overall best outcome. In fact, strictly speaking, altruism defines behaviour that benefits others without benefits to oneself. This is only altruism if the participant has no expectation of any gain, and in that sense the term reciprocal altruism is slightly misleading.

As discussed above, selfish behaviour is often treated as a given, a baseline upon which subsequent cooperative behaviour may be seen to emerge under specific socio-evolutionary dynamics (Dawkins, 1976). In a single blind prisoner's dilemma game, with no communication between the players before or during the game, there is no better choice than to plead innocent (to act selfishly). Note that the selfish option is best whatever your opponent's decision turns out to be. If your opponent turns out to have pled innocent, you would have done better to plead innocent than guilty, and if they turn out to have pled guilty you would also have done better to plead innocent than guilty. In nature, however, where the prisoner's dilemma presents a common situation in which individuals must decide whether to cooperate with or defect from another individual or group, there are all sorts of conditions that affect the way that participants enter into such a situation (Maynard Smith and Szathmary, 1995). With additional knowledge about an opponent's strategy, the dynamics of the game change considerably. The simplest case is with the hindsight of previous prisoner's dilemma situations. Axelrod (1984, 1987) investigated how simple agents could develop successful strategies to what is known as the *iterated prisoner's dilemma*, a form in which the game is repeated a number of times by the same two individuals, who have a memory of the outcome of one or more of the previous games. If individual strategies are linked to fitness, and fitter individuals have greater reproductive success within a population, then we can begin to understand the evolutionary dynamics of cooperative

behaviour. Axelrod's results show that always defecting remains a good strategy, since hindsight of previous situations doesn't necessarily lead to foresight in the current situation, whereas always cooperating is very risky if there are any defectors about. A surprising result from this study is that another extremely simple strategy, known as "tit-for-tat", also proves very effective. The tit-for-tat strategy always starts an iterated prisoner's dilemma round with cooperation, and then continues by copying the opponent's previous choice *ad infinitum*. In a world populated by variously selfish and altruistic agents, this represents a minimal form of adaptation. Against a purely selfish opponent, a tit-for-tat agent does as well as possible (ignoring the first cooperative move the tit-for-tat agent becomes a defector); against a purely cooperative opponent it doesn't do as well as a selfish strategy, but against another tit-for-tat agent it does better than the selfish strategy, because it 'convinces' the other to cooperate rather than defect. Tit-for-tat and pure selfishness are both evolutionary stable states (ESS), meaning that a process of simulated natural selection on a population of tit-for-tat individuals, or a population of selfish individuals, won't trigger a major change in the population's makeup. If you introduce a mutant into either of these populations, then the mutant will be unsuccessful and will soon die out. However, pure altruism is not an ESS. If you introduce a selfish mutant into an altruistic population, then the selfish mutant will be successful. In Maynard-Smith's terminology, the population can be successfully *invaded* (Maynard Smith, 1982).

In the above ESSs the tit-for-tat population scores higher overall than the selfish population, because whilst the tit-for-tat individuals are scoring two points each time a game is played, the selfish individuals are scoring only one. It follows that under the right conditions altruistic behaviour might evolve (Dawkins, 1976; Axelrod, 1984). The main condition is the efficacy of the tit-for-tat behaviour itself: individuals must not be naively altruistic, but use information about their opponents to determine whether risking altruism is worthwhile. Given a sufficient number of tit-for-tat individuals in a population otherwise dominated by defecting individuals, the tit-for-tat strategy is the more successful. If tit-for-tat individuals get to interact with each other more often than with defectors, a tit-for-tat strategy can evolve out of a defector strategy (Dawkins, 1976). Kin selection again steps in to provide a suitable context for this to happen, on the basis that close kin are likely to interact with each other more than distant kin. The

complex social psychologies of many animal species appear to be testament to this dynamic coevolutionary arms race, in which the ability to cheat without getting caught is pitted against the ability to catch cheats.

The value of the relative costs and benefits in the prisoner's dilemma games also affects their outcomes. Recall that the prisoner's dilemma is itself only a special case of game theory systems. If cc is the net benefit awarded to mutual cooperators, dd to mutual defectors, cd to cooperators whose opponents defect, and dc to defectors whose opponents cooperate, then we have $dc > cc > dd > cd$ as the condition for a prisoner's dilemma situation. However, Maynard Smith and Szathmáry (1995) propose that in many realistic situations $dc = dd = cd = 0$ and $cc > 0$. That is, there are situations (like moving a rock which is too heavy for a single person to carry, but light enough for two people to carry) in which only reciprocal altruism is beneficial. This is evidently the case with mutual benefits between species, such as between bees and flowers, but Maynard Smith and Szathmáry (1995) also propose that it can explain the evolution from selfishness to cooperation within a single species. This requires the establishment of situations in which mutual benefit occur.

In a similar way, Sober and Wilson (1998) argue that various approaches to the problem of groups selection can be brought together in a legitimate pluralist view of cooperative or altruistic behaviour, stating that proposed alternatives to the problem of altruism actually embrace group selection principles. Sober and Wilson propose a hierarchical view of selection, with natural selection potentially occurring at different levels of organisation, from genes to organisms to social groups, along with a methodology for identifying the various strengths of selection at these various levels.

2.3.3 Theories of animal communication

In an influential paper, Krebs and Dawkins (1984) argue that communication systems can emerge as a result of both situations of mutual benefit or as a result of situations in which one individual benefits from the manipulation of another individual. Since Dawkins and Krebs' (1978) earlier proposal that selfish manipulation is a more important factor in communication than the mutually beneficial sharing of information, in keeping with selfish gene theory, the study of

communication has been primarily concerned with the goal of understanding how communication systems evolve if only the producer of signals benefits. The situation is straightforward in the case of communication between different species. For example, an animal may evolve the ability to deceive a predator into believing that it is more dangerous than it actually is, and this clearly represents a beneficial adaptation for that species (Owings and Morton, 1998). However, communication between individuals of the same species is slightly more complex because an individual that is successful at manipulating other individuals pushes the population towards that kind of behaviour. Owings and Morton (1998) formulate the developments of this body of theory in an assessor/manipulator model, which essentially describes a coevolutionary arms race between abilities on both sides of this relationship, the senders and receivers of communicative signals.

A second key principle underlying the development of the assessor/manager model is the theory of costly signalling developed by Zahavi (1975), otherwise known as the handicap principle. Zahavi argues that signals *can* be of use to the receiver if they are inherently reliable. An inherently reliable signal is one that, by virtue of its means of production, guarantees some information about the signaller. A peacock's tail reliably indicates the genetic quality and good health of a peacock because a poor quality or unhealthy peacock would not be able to survive with such a costly burden (Zahavi, 1975). Low frequency vocalisations often honestly communicate the size of an animal, because the ability to generate low frequency sounds is partly dependent on size (Owings and Morton, 1998). Gazelles leap energetically in the air in the presence of a lion, signalling their superior fitness over other nearby gazelles (Owings and Morton, 1998). The lion can use this as a reliable indicator of which animal to chase, to the mutual benefit of both the lion and the fitter of the gazelles. Following Dawkins and Krebs (1978), Owings and Moreton argue that a coevolutionary arms-race can occur within the same species between *managers* who manage to squeeze fitness gains out of the exploitation of conspecific *assessors*, and assessors who manage to accurately distinguish reliable signals from exploitative ones in conspecific managers. The efficacy of this evolutionary process relies on the finer details regarding how individuals do indeed perceive and evaluate signals, and also on the fact that one successful modification, as it proliferates through the species, alters the environment in which further potential modifications

will be selected for. Given the sharp asymmetry between the processes of sending and assessing signals, the relative time lags in evolutionary adaptations in different domains, and the wide variety of very different contexts in which managing and assessing may take place, it is easy to see how this situation leads towards the positive generation of complex systems of interaction, rather than simply neutralising any change.

Krebs and Dawkins (1984) predicted that costly signalling would therefore arise in situations in which only the signaller was really the beneficiary, and that low-cost ‘conspiratorial whispers’ could emerge in cases where there was a mutual benefit to both parties. Noble (2000) explored these possibilities in computer simulation models but could not find situations in which communication emerged in the absence of mutual benefit. However, his simulation results suggested that costly signals could emerge in cases where the benefit to the recipient was negligible, suggesting a kind of ‘sales resistance’ which signallers need to overcome. Searcy and Nowicki (2005) review recent work in this area and conclude, based on the most recent evidence, that animal signalling is more likely to be honest than deceptive. Cross argues (personal communication) that signalling theory avoids the cognitive context of signalling and is therefore less applicable to music.

Theories of the evolution of language, music, or a combined musilanguage precursor have drawn on this theoretical process to varying degrees, but surprisingly its implications for the evolution of music have not been explored in any great depth. I discuss Dunbar’s (2004) theory of the evolution of language in Section 2.4.1. This approach considers the way in which what he called ‘vocal grooming’ (which could be either music or language) could be used as a sustainable tool for managing social relations, but does not account in any great detail for the kinds of handicap costs that we would expect to be associated with vocal grooming. Very recently, however, Cross (2007a) has developed his theory of music as possessing floating intentionality to incorporate the assessor/manager principle. He therefore proposes how music can be seen as a domain of social interaction in which coordination can be achieved *despite* conflicting interpretations held by different individuals. In doing so I believe that Cross makes a huge leap towards a theory that can be accepted and understood by ethnomusicologists and evolutionary theorists alike, but does not go to the full extent of exploring the way in which the evolution of

musical behaviour itself is subject to the battles between individuals in everyday interaction.

Meanwhile, sexual selection theories of music, mostly in the hands of Miller (2000a), have made far greater capital out of the kinds of arms-race principles presented here. Sexual selection theory was also strongly influenced by Zahavi's (1975) handicap principle, as already suggested by the example of the peacocks tail. Darwin's (1883) original formulation of sexual selection simply stated that as well as being useful in general survival, traits could also be useful in attracting mates. If mates were attracted to a certain feature, then that feature would become firmly established, perhaps exaggerated over time. The theory has been developed more recently as follows (Fisher, 1958). If female preferences and male traits vary in a population, then it can be assumed that partnerships will form which represent well-matched pairs of traits and preferences. We assume here that females are the choosy sex and males compete for the opportunity to reproduce with as many females as possible, which is likely due to the higher cost of reproduction for females. According to this reasoning, females are more likely to have similar numbers of mates, whilst male access to mates may be hugely imbalanced (Miller, 2000b). Thus many such partnerships might involve the same male. The offspring of one such partnership will share both the traits and the preferences of mother and father. Assuming that the genetic contributions to preferences and traits are subject to further random variation in subsequent generations, then this pairing process will continue. In some cases preferences and traits may become rather extreme, especially if female preferences run ahead of male traits, becoming effectively insatiable. This is where the runaway process begins (Fisher, 1958). Furthermore, in some of *these* cases, the females will be attracted to traits that actually happen to indicate the fittest males (Zahavi, 1975; Miller, 2000b). Some highly salient male traits are directly selected for on their ability to satisfy certain female preferences, regardless of whether they are good for the male, but female preferences are indirectly selected for on their ability to discriminate truly fit males from less fit males. If the female is a bad judge of fitness, her offspring are likely to be either of poor genetic quality, or poorly looked after. The kind of exaggerated traits that should tend to survive in the long term, therefore, are traits that reliably indicate fitness (see Miller, 2000b). The best known honest signals are those that demonstrate that an individual is energetic, healthy, coordinated and capable of survival despite wasteful behaviour (Zahavi, 1975). The notion that a fit peacock

would be even fitter *without* the burden of its big tail is neither here nor there with regard to the evolutionary history of the peacock: what good is a fit peacock without a peahen with whom to mate?

For these reasons, sexual selection should be understood as a very powerful evolutionary force, as well as a stunningly creative one. Miller (2000b) has even suggested that sexual selection can be relied upon to actively accelerate the evolution of fitness-enhancing traits. In other words, once sexual selection has found reliable indicators of fitness it maximises those indicators, thus maximising the underlying fitness. Sexual selection clearly presides over highly salient traits: traits that are visibly bright and large, or audibly loud, or involve movement. Sexually selected traits are also often complex, intricate, and highly idiosyncratic between different species. This superficially makes music a strong candidate for being a sexually selected trait. However, the question of whether sexual selection is respectively a strong candidate for explaining musical behaviour depends primarily on the social context in which our ancestors struggled for existence during the period of music's evolution. This will be discussed in Section 2.4.

2.3.4 Major transitions and evolution beyond Darwinism

According to the reasoning established in Sections 2.3.2, Darwinian evolutionary theory requires that the emergence of music as a mechanism for fostering social bonds involves a process of either kin selection or mutual benefit, assuming, that is, that a social bond refers to some established basis for cooperation. The important catch is that a mutual benefit argument cannot simply posit as an explanation the fact that a more cohesive group will be more successful than a less cohesive group. It must go further and provide a plausible mechanism by which the more cohesion-prone individuals within the group are likely to be more successful than the less cohesion-prone individuals within the group, and how music itself functions in making this happen. We must also bear in mind the difference between cooperation, as discussed in Section 2.3.2, and cohesion, as discussed in Section 2.2.2.1. Cooperation refers explicitly to benefits with respect to separate individuals, whereas cohesion refers to the coming together of elements to form a whole. Cohesion implies a more profound process than the mere choice of individuals to cooperate under certain circumstances that facilitate the survival of their genes, and involves the

emergence of new orders of structure which may be coordinated in such a way as to be defined as individual units themselves. An important question then is the degree to which the formation of closely bound groups necessarily implies the cooperation of individuals within those groups. To the extent that it does, kin selection or mutual benefit are required to explain how cohesion becomes established. To the extent that it does not, one must explain in what other sense a group is cohesive. It is therefore necessary to consider how theoretical biology has approached the problem of organisation at different levels.

Maynard Smith and Szathmáry (1995) discuss the evolution of social behaviour as one of the major transitions in evolution. They observe that evolution exhibits an overall increase in complexity over time, which can be categorised according to a small number of major transitions: from replicating molecules to populations of molecules in compartments; from independent replicators to chromosomes; from RNA as gene and enzyme to DNA and protein; from prokaryote to eukaryote; from asexual clones to sexual populations; from protists to animals, plants and fungi; from solitary individuals to colonies; and from primate societies to human, language-using societies. Most of these transitions appear to have been one-off events in the Earth's evolutionary history (except for the transition to multicellular organisms, and to colonial species) and each constitutes a step up in complexity and a rewrite of the rules of evolution to take into account a major new context, with new actors. They suggest that natural selection on its own does not account for these leaps in complexity, and that the mechanics of individual situations, such as the one discussed above, need to be investigated as specific transitional events. As to whether general biological laws apply across all of these levels of scale they are not sure.

Concentrating on the earliest stages of this process, a number of theorists such as Szathmáry (1997) and Kauffman (1993) have been successful in demonstrating computationally how life may have originated by showing how the essential feature of replication emerges naturally under the right circumstances in complexes of interacting chemicals. Since replication is a prerequisite of Darwinian evolution, this cannot be a Darwinian process as such. Kauffman accordingly describes his work as aiming toward a *general biology* (Kauffman, 2000). Central to his theory is the notion of 'order for free' which implies that the self-organising properties of complex collections of molecules provide much of the explanation for biological systems, which if taken

further essentially relieves the emphasis on natural selection in evolutionary theory and amounts to the kind of process that Maynard Smith and Szathmáry seek in their analysis of the major transitions. The question remains as to the significance of this process of self-organisation in other transitions. According to Gould (2002), Kauffman is the most notable modern intellectual descendant of D'Arcy Thompson, whose book *On Growth and Form* (Thompson, 1992) revitalised the pre-Darwinian idea that a great deal of the *form* we observe in nature is a result of classical physical processes as much as it is of natural selection. In a similar vein, Kauffman (1993) introduces complex systems approaches to the emergence of form that are not Darwinian. However, the notion of self-organisation through feedback introduces a chaotic form to the dynamics of evolution which is not a part of Thompson's theory. He argues that we cannot "prestate the configuration space of a biosphere" (Kauffman, 2000, p. 130). That is, no analysis of a complex evolving system can genuinely tell us what's going to happen next. Although we can understand how aspects of organisms were shaped by natural selection in the short term, we cannot be sure which of the organism's interactions with its environment will be of any great importance in the future, because the possible interactions are limitless. Although selfish gene theory (Dawkins, 1976) can precisely define processes that can emerge under certain conditions, it is less well equipped to clarify how those conditions emerge.

In each of Maynard Smith and Szathmáry's (1995) major transitions, competition at one level becomes overwhelmed by the unity required by structures forming at a higher level. For example, they consider how it is that competition between cells did not interfere with the formation of multicellular organisms, and conclude that due to the way multicellular organisms organise themselves (through ontogenic development from a single cell) there simply is no such competition: all of the cells contain identical DNA and are thus bound by a mutual investment in the survival of the same essential entity. Individual cells are bound by a common mechanism for reproduction. Thus the emergence of multicellular organisms from cellular organisms is strangely unproblematic, involving no struggle for existence or conflict between the two levels, only a new way for entities to produce reproductions of themselves. However, human societies are certainly not bound together by a unified process of reproduction. Many (*e.g.*, Boyd and Richerson, 1985; Laland et al., 1999; Dawkins, 1976; Blackmore, 1999; Knight, 2002) have ar-

gued that social learning provides a new means of information transmission which does, in a similar way, bind humans into cooperative groups. These arguments will be considered in Section 2.4.2. But this is also not the only way that the formation of more complex structures can be understood.

Maturana and Varela's (1980; 1987) theoretical biology explicitly approaches this process by treating natural selection not as a given pervasive force but as a process that is integrated into a range of other less imminently conflicting relationships between entities. Their theory builds from their notion of *unities* "the description, invention and manipulation of [which]" they suggest, "is at the base of all scientific enquiry" (Maturana and Varela, 1987, p. 73). In this way they define a mode of description which holds generality across a range of biological forms (both physical and abstract): "There is an organization that is common to all living systems, whichever the nature of their components" (Maturana and Varela, 1987, p. 74). Defining this organization is the process of autopoiesis, by which a unity produces (and transforms and destroys) the components which: "(i) through their interactions and transformations regenerate and realize the network of processes . . . that produced them; and (ii) constitute it as a concrete unity in the space in which they exist by specifying the topological domain of its realization as such a network" (Maturana and Varela, 1987, p. 135). The two parts of Maturana and Varela's definition correspond to two crucial properties of autopoietic systems: a metabolism, which is the set of interactions that the system uses to maintain its unity; and a membrane, the maintained boundary that delimits the space in which the unity exists from the outside. Autopoiesis is considered to apply to systems existing at different hierarchical levels, from the cell, to the multicellular organism, and possibly to the social group.

Another concept discussed by Maturana and Varela (1987) is *structural coupling*, which is defined in the establishment of autopoiesis. Structural coupling describes the interaction between the constituent elements of a unity, or elements that have the potential to form a unity. "We speak of structural coupling whenever there is a history of recurrent interactions leading to the structural congruence between two (or more) systems" (Maturana and Varela, 1987, p. 75). These recurrent interactions trigger, rather than specify, structural changes in the autopoietic unities that engage in them. Structural congruence is the result of a set of structurally coupled relations

between unities. Since structural coupling between a unity and aspects of its environment has an effect on the behaviour of that unity, it will have an effect on the evolution of the unity, as well as of those other elements of the environment. It is easy to picture collections of unities, in the extremes of such circumstances, coming to behave as new larger unities upon which natural selection operates.

Maturana and Varela's view provides a valuable alternative to traditional biological theory, and poses a set of processes beyond the reach of natural selection. Importantly, whilst promising to concur with most selfish gene theory, it redefines the problems of the emergence of cooperative groups discussed in Section 2.3.2 by explicitly addressing the full set of relations between organisms and their environments. I interpret their concept of structural congruence as suggesting that sustained interactions between individuals ultimately provide reliable environmental contexts which are likely to become utilised by those individuals in beneficial ways *even if* they are not cooperative interactions. This has been explored experimentally in an artificial life context by Quick et al. (1999), demonstrating the analytical value of looking at the long-term relationship between a system and its environment, in this case, a simulation of bacteria interacting with the World Wide Web.

Following Maturana and Varela, coordination and coupling are offered as more general terms subsuming cooperation as the only way in which individuals can come together to form higher levels of organisation. Similarly, aspects of the work of D'Arcy Thompson (1992), Gould (1985) and Kauffman (2000) all echo this critical point. Therefore, in response to the question posed at the beginning of this section, these points of view propose that cooperation is not the only means by which it is sensible to talk about social cohesion.

There is in fact little conflict between this point and the principles of kin selection and mutual benefit. They all address the problem of understanding how the notion of selfish genes, shaped by Darwinian evolution, are reconcilable with the formation of more complex coordinated structures. In the case of musical behaviour, however, it would be a mistake to reduce musical interactions to games of cooperation and defection only to gloss over the potential subtlety of musical interaction between individuals and its effects on their relative fitness. In particular, this would ignore the *structure* of social groups (which is something more than their mere ex-

istence and influence). Theories of the evolution of human musical behaviour need to address the problems posed by a game theoretical view of music as a cohesive medium, but can do so by drawing on notions such as structural coupling as a means for forming coordinated unities. It is important to reiterate Kauffman's (2000) point that any number of interactions may direct the future course of evolution. Thus we can consider cases in which individuals' adaptations to detrimental situations take the form of unusual compromises – quick fixes – and from this compromise arises the basis for greater coupling of behaviour.

One final view that also echoes these concerns is that of niche construction (Odling-Smee et al., 2003). Niche construction is the modification of environmental niches by their inhabitants. The authors of niche construction theory argue that this process is as important to understanding natural evolutionary systems as is the process of natural selection, and set up niche construction and natural selection as an interacting pair of processes with certain significant dynamics not present in the study of natural selection alone. I leave a further discussion of niche construction to the next section, since the niche constructionists' debate is particularly salient in the case of human evolution.

2.4 Considering the Evolution of Human Musical Behaviour from the Point of View of Cultural Evolution and Gene-Culture Co-evolution

Human cognition unambiguously constitutes a major evolutionary departure from the nature of other animal behaviour. Humankind's distinct cognitive power is associated with the use of language (Burling, 2005), the ability to use abstract and metaphorical thought (Donald, 1991), the ability to comprehend complex social relations (Dunbar, 2004), the experience of consciousness and free will (Dennett, 1993), and the ability to learn structured sequences such as movements of the body.

It has been proposed that our cognitive evolution occurred in two relatively rapid phases, separated by extended periods of only gradual change (Mithen, 1996). The first of these occurred 2 million years ago and is associated with the development of very basic tools, partial bipedalism, larger social groups, but very little else that could be described as cultural. The

second occurred about 100,000 years ago and is associated with the emergence of language and numerous cultural activities, and the dispersion of modern humans out of Africa. In the intervening period archaeological evidence suggests that very little cultural change may have taken place in hominin cultural practices (Mithen, 1996). However, this later figure is a topic of lively debate and Mellars (2006) suggests a more recent figure of around 50,000 years for dispersal from Africa of genetically and culturally modern humans. There is one primary form of artefact representing hominin creativity throughout this vast period; chiselled, often symmetrical handaxes which various authors have proposed various uses for, the most obvious use being as a tool, although a serious competing hypothesis is that the handaxes were used for display (see Kohn and Mithen, 1999).

A tenet held by many evolutionary psychologists is that brains have a modular structure (Barkow et al., 1992), with distinct modules for distinct cognitive capacities (stemming from the notion of modular perceptual systems (Fodor, 1981)). These distinct modules imply the possibility of being able to break down specific evolutionary histories for specific cognitive capacities, which is one of the main bases for research in evolutionary psychology. At the same time, the human mind appears to have a powerful general intelligence capacity associated with abstract and metaphorical thought, as well as a high degree of plasticity. Evolutionary psychologists Barkow et al. (1992) argue that human minds need to be understood as complexes of specialist adaptive capacities that have been shaped by natural selection. They assert the importance of this modular view of cognition – that the brain has evolved numerous context-specific modules in response to distinct evolutionary pressures – as compared to what they call the Standard Social Scientific Model Barkow et al. (1992), which treats the brain as a uniform, content-free device that allows behaviour to be determined entirely as a result of cultural influences. The Standard Social Scientific Model, they argue, “rests on a faulty analysis of nature-nurture issues, stemming from a failure to appreciate the role that the evolutionary process plays in organising the relationship between our species-universal genetic endowment, our evolved developmental processes, and the recurring features of developmental environments” (Barkow et al., 1992, p. 33). The modular view also has a very practical relevance, in that it facilitates the division of behavioural capacities in order to study specific evolutionary situations.

Barkow et al's search for the origins of our evolved psychology emphasises the key fact that "the development of increasing cultural variation throughout the Pleistocene was made possible by the evolution of psychological specialisation that exploited the regularities of human meta-culture in order to learn the variable features of culture" (Barkow et al., 1992, p. 93). That is, cultural variation is above all a consequence of psychological capacities.

For example, Barkow et al. (1992) present a study of the evolution of human social contract behaviour based on game-theoretic evolutionary approaches. They emphasise that one of the main cognitive adaptations for human social behaviour is in the identification of cheats in situations of social exchange (reciprocal altruism such as food sharing and the division of labour). Identification of cheats is a must-have capacity if species-wide social behaviour is to evolve, this being one of the cognitive developments that would allow competitive interaction to become established at a new level. Evidence for this cognitive ability comes from data on a special kind of psychological test, the Wason test (Wason, 1966), which shows that in modern humans logical problems are more easily solved, more transparent, when placed in a social context. As well as showing that social context matters, Cosmides and Tooby (1992) present a further experiment in which the difference between detecting altruistic behaviour and selfish behaviour in an exchange situation in an imaginary culture is shown to be significant: we are less proficient at identifying altruists than we are at identifying cheats. However, the implications of results from the Wason test on innate biological traits are not clear, as the authors do not after all distinguish between a learnt skill and an innate one.

Such arguments have had an important impact on the way we understand cultural behaviour, and so-called followers of the Standard Social Scientific Model have been urged to tighten up the implicit assumption that cultural influences could ultimately arrive at any human behaviour imaginable, regardless of our genetic history (Barkow et al., 1992). This also suggests, but does not guarantee, that universal human activities such as music are likely to be associated with specific cognitive modules, each of which requires a Darwinian explanation. Here I have presented the more extreme arguments from evolutionary psychology but naturally a number of more moderate views are also worth considering. Shore (1996) defines the debate in terms of the notions of *psychic unity* and *psychic diversity* (humans having universal cognitive mechanisms

versus culturally determined cognitive mechanisms) and argues that the distinction is based on a false dichotomy and an overly essentialistic biology (Shore, 1996, p. 312). This synthesis is similar to the *Rethinking Innateness* argument of Elman et al. (1996).

2.4.1 Human evolution and the social brain

At present, the dominant theory for the emergence of human intelligence is known as Machiavelian intelligence theory (Whiten and Byrne, 1997), or the social brain hypothesis (Dunbar, 1998). Dunbar (1994, 1998, 2004) explains humankind's evolution in terms of the evolutionary advantage of living in larger groups and the evolutionary pressures of handling social relationships in those groups. According to his theory (2004), this was a pressure that became predominant following change in the forest environments of our earliest ancestors, which drove them into more open environments where predatory attacks were a more serious threat. Assuming that larger cooperative groups are more secure against attacks from predators, groups that can address the challenges of establishing cooperative behaviour, namely that selfish individuals may be able to take advantage of that cooperative spirit, will be more successful than those that do not. Dunbar (2004) argues that the social grooming behaviour of apes provides exactly that defensive mechanism against selfish behaviour. Social grooming is a means by which strong, emotional social bonds are formed in ape societies, and is typified by the fact that the development of each interpersonal relationship requires a great deal of time and effort. Why, Dunbar asks, would such a costly behavioural ritual evolve that distracts these apes from other useful behaviours, rather than a very short economical one that frees up their time? His answer, following the handicap principle (see Section 2.3.3), is that the high cost of the behaviour is the key to forming a social system which is relatively robust against selfish behaviour: in a society where acceptance into an alliance involves a high cost, the cost of selfish behaviour itself is higher still, since the selfish individual, expelled from one group, would have to start afresh to get accepted into a new group. In such a system, the best strategy for each individual is to stick with their group and behave altruistically. It is important to reiterate that, according to Dunbar, the high cost of this bonding behaviour is a necessary feature of it, rather than an unfortunate side effect.

The social brain hypothesis is presented in opposition to other theories which aim to explain

the rapid expansion in size of the human brain, such as foraging memory or intelligent tool use (Dunbar, 1998). Dunbar supports his claim with data that suggests a correlation existing across primate species whereby group size is directly proportional to brain size, or specifically, the size of the hippocampus, which is the area most commonly associated with our theory of mind capability (the ability to comprehend other individuals as thinking, perceiving beings like ourselves, and to infer things about their behaviour as a result) (Dunbar, 1998). In non-human primates, time spent grooming is also proportional to these two factors. By extrapolating from these relationships, Dunbar predicts that humans should live in groups with a natural size of about 150 and groom 35% of the day, which he argues is an unmanageable commitment, given other environmental demands (Dunbar, 2007). Dunbar proposes that our ancestors shifted from physical grooming to ‘vocal grooming’, allowing grooming to take place in larger groups (between 2 to 7 individuals communicating vocally at the same time) and whilst engaged in other activities. Vocal grooming actually applies equally to song, and since Dunbar’s theory does not require that the vocal signals contain any referential content, but are emotionally stimulating, this can essentially be understood as another cohesionist theory of the evolution of human musical behaviour.

The important problem, as discussed with reference to the prisoner’s dilemma and the evolution of altruism in Section 2.3.2, is to explain how the vocal system manages to maintain the honesty of the signal (“I am committed to this relationship”) in light of its potential abuse by selfish behaviour. According to Dunbar:

In many ways, primates’ success from an evolutionary perspective is a direct consequence of [their intense] sociality. Primate societies are implicit social contracts which allow some of the problems of survival and reproduction to be solved cooperatively. Social contracts of this kind work because they allow relevant problems to be solved more efficiently. However, social contracts require individuals to be willing to forego some of their more immediate personal interests in order to benefit from greater returns later through group-level cooperation. If too many individuals act in their own selfish interests, the cohesion of the group will be threatened simply because too many others will end up paying the costs of sociality. Group stability is

quickly threatened, leading to the rapid collapse of the contract.

The real issue here seems to be the cognitive demands of maintaining the stability of relationships through time. . . . In the limit, and certainly in the human case, [this] requires individuals to be able to understand another's perspective sufficiently well to appreciate what kinds of adjustments are necessary to create the levels of 'bondingness' required to keep a group together. One element of that is knowing when to trust another individual.

However . . . freeriders . . . are an inevitable by-product of a biological system that is built on cognitive flexibility rather than genetic hardwiring or endocrinological determinism (as is the case with many social insects, for example).

(Dunbar, 2006)

The evolutionary problems associated with cooperation are a central part of the puzzle that Dunbar addresses, and yet his description of theory of mind dynamics at times implies a system entirely predicated on the selfish use of social relations. In this he suggests a possible scenario in which complex social systems, and complex social brains, emerge only from the runaway creation of conditions within which increasingly sophisticated forms of selfish behaviour can be performed. In other words, cooperation is actually an indirect consequence of that evolutionary process. Conversely, the evolution of cooperation, a functionally adaptive behaviour, as Dunbar defines it, can be seen as being driven by the coevolution of suitably watertight systems of interaction, which attempt to guarantee honesty of signalling, along with new tricks for cheating this system.

Other work focuses more closely on the trajectory of incremental developments along which human cognition may have evolved. Donald (1991) describes human cognitive evolution in three major phases. The first phase is characterised by the *episodic mind*, which is seen in modern apes. The episodic mind is adept at storing, recalling and analysing literal episodic perceptual information, but has no capacity to abstract and to imagine. The second phase is characterised by the *mimetic mind*, which apes do not possess, and which humans have surpassed. The mimetic mind has greater control over mental representations and can visualise movements of the body.

It can plan movement and it can structure sequences of movement. Donald proposes that this cognitive development was multimodal, or easily extended through different modes once a general capacity emerged. Members of a mimetic society, as the name suggests, learn movements and patterns from each other. The final phase is characterised by the mythic mind, which represents the beginnings of metaphorical thought, and thus the fundamental cognitive development upon which language evolution rests.

Much weight has been given to such an outline by the discovery of *mirror neurons* in humans and other primates (for example Arbib, 2005). Mirror neurons are sets of neurons that are seen to fire both when an action is performed and when the same action is observed, implying a neural device for learning action from observation, as in the development of Donald's mimetic mind. Arbib (2005) proposes the mirror system hypothesis for the evolution of language as constituted by the following sequence:

1. Grasping.
2. A mirror system for grasping shared with the common ancestor of human and monkey.
3. A simple imitation system for object-directed grasping through much repeated exposure. This is shared with common ancestor of human and chimpanzee.
4. A complex imitation system for grasping.
5. *Protosign*, a manual-based communication system, breaking through the fixed repertoire of primate vocalizations to yield an open repertoire.
6. *Proto-speech*, resulting from the ability of control mechanisms evolved for protosign coming to control the vocal apparatus with increasing flexibility.
7. *Language*: the change from action-object frames to verb-argument structures to syntax and semantics; the co-evolution of cognitive and linguistic complexity.

This ties language closely to *theory of mind*, grounding its evolution in a dedicated learning mechanism that ultimately applies to a range of social learning behaviour.

I have only briefly sketched out a number of key social intelligence theories of the evolution of language in this section. Their common thread is the runaway evolution of increasingly sophisticated tools for social interaction. On the one hand, theories such as Dunbar's (2004) provide a basis for understanding how vocal communication could have had a functional utility which was based on the manipulation of social cohesion. This requires some kind of positive hedonic response in individuals in response to hearing vocalisations, which also inherently enforces the requirement for reliability (Dunbar, 1998). On the other hand, theories such as Donald's (1991) emphasise how social learning provided building blocks for the evolution of language as we know it: a system largely dependent on vocal imitation (Pinker, 1998). Neither require specific features for language at the outset and can equally apply to music: more so in the case of Dunbar's (2004) theory.

2.4.2 Evolutionary approaches to culture

The essence of culture, from a scientific point of view, traditionally lies both in its role as a non-genetic mode for the inheritance of useful behaviour (an adaptive behaviour leading to further adaptation), and in its essential capacity to generate maladaptive behaviour. One view amongst biologists is that modern culture does indeed produce many human behaviours that are detrimental, maladaptive, but that these are side-effects of a generally adaptive development: *cumulative social learning* (Boyd and Richerson, 1985). This view is discussed below. The study of culture *in situ* is hard to reconcile with a study of individual psychology grounded in biology (Geertz, 1973), and provides only a tenuous link to a theory of what culture may have been like in the hominin lineage (Barkow et al., 1992). However, it is imperative for a combined evolutionary theory of biology and culture that this gap is closed. In the remainder of this section these difficulties will be discussed. I will consider notions of dual-inheritance in bio-cultural evolution and compare these views with those of social anthropologists. I will also discuss creativity in the context of cultural processes.

2.4.2.1 Boyd and Richerson's theory of gene-culture coevolution

Boyd and Richerson's (1985) theory of gene-culture coevolution represents a key view of culture as an aspect of human adaptive behaviour. As for many other theorists of cultural evolution, for Boyd and Richerson, learning behaviour from other individuals (social learning) is viewed as a primary adaptive behaviour. Boyd and Richerson's view is that social learning evolved because it allowed a population to adapt to a changing environment without each individual having to learn how to survive directly from the environment, ultimately introducing a new and more efficient mode of adaptation. They reason that effective social learning strategies require individuals to adopt behaviours based on limited information about whether the behaviours they are learning are in fact useful, and that these strategies have detrimental consequences which are probably unique to humans; in short, behaviours that are not best suited to a particular environment, or of any use to the individuals that adopt it, can become distributed throughout the population. Boyd and Richerson argue that a form of cultural runaway selection can emerge which is essentially similar in character to runaway sexual selection: individuals imitate other successful individuals, adopting useful behaviours, but also various behavioural red herrings which are of no actual use, and can therefore be considered detrimental to fitness. As these detrimental behaviours get copied, they increase their chance of being copied, spreading throughout the population.

Two behavioural notions that Boyd and Richerson (1985) develop for this model are frequency-dependant bias and indirect bias. Underlying these behaviours is the central notion of blending: an individual brings together (blends) information about different role models in order to decide what behaviour to adopt. In a changing environment, where individuals learn directly from the environment to some extent, but also learn from each other, it would be a reasonable assumption that the most common behaviours are the best behaviours. That is, if a large number of individuals are learning directly from the environment, rather than imitating, and assuming that they have a reasonable success rate, then most individuals will have adopted a suitable behaviour with respect to the current environment. Therefore, an efficient, albeit slightly risky, rule for choosing a behaviour to imitate would be to adopt the most common one, or to adopt behaviours with a probability based on their frequency. This is known as frequency-dependent bias. Assuming that imitation is in some way less costly than direct learning, then a small number of

frequency-dependent imitators in a population of direct learners will have far greater fitness. Imitation will therefore increase in frequency through natural selection. As the number of imitators increases, the fitness advantages of imitation decreases: imitators require that those they imitate are doing the right thing, and frequency-dependent imitators require that the majority of the population are doing the right thing. In a slow-moving environment this is more likely, but as environmental change increases, imitation becomes more error-prone. An evolutionarily stable equilibrium is therefore reached between direct learning and frequency-dependent imitation. In Boyd and Richerson's (1985) basic model, the average fitness of the population at this state is the same as with no imitation. But in a model where individuals choose between learning and imitation, based on the strength of environmental cues that could indicate the correct behaviour, the equilibrium learning/imitating strategy produces higher fitness. Boyd suggests that introducing a notion of cumulative culture into this model further secures the evolutionary advantage of imitation.

An important consequence of frequency-dependent bias discussed by Boyd and Richerson is that it causes behaviour to vary on a group level. That is, it forces homogeneity of behaviour within groups, reducing the effect of selection between individual behaviours. This provides a basis for a much-disputed process of group selection. This is a very powerful argument with respect to Maynard Smith and Szathmáry's major transitions (Maynard Smith and Szathmáry, 1995). Maynard Smith and Szathmáry argue that in multicellular organisms, competition between cells is rendered irrelevant by the fact that all cells share identical DNA. Similarly, social insects are bound together by their common descent. Boyd and Richerson (1985) evoke a similar, albeit weaker, version of the same principle: effective behavioural variation between individuals is reduced because all individuals share the same culture. In this way, groups of individuals, including their accumulated cultural systems, may become units of selection. This is widely accepted by many as being the only feasible argument for the possibility of group selection, although many who accept its feasibility also do not consider it likely.

The second principle investigated by Boyd and Richerson (1985) is indirect bias, which emphasises the likely inaccuracies that imitation will suffer. In order to be a successful imitator, any individual must know both who to copy and what to copy from them. The question of who,

according to Boyd and Richerson, comes down to identifying what are known as indicator traits, identifiable behaviours that indicate that the candidate makes a suitable role model. But as well as determining who would make a suitable role model, an imitator must decide what it is about their role model that they want to copy. In keeping with the principle that simple heuristics are more likely to have evolved than complex analysis, Boyd and Richerson's answer to this question is that imitators will do well to copy as much as possible from their role models. That is, crude forms of imitation are particularly likely to evolve because they reduce the cost of direct learning from the environment. This leads to a situation where individuals fix on certain indicator traits and then adopt a plethora of extra traits from those models who display the required indicator traits. The result of this process is that insignificant traits become indirectly biased and can spread through a population. If the model also allows that individuals also make choices about *which* indicator traits they consider important, then it is possible that quite arbitrary behaviours become indicator traits and a runaway selection of cultural traits becomes possible. That is, individuals chose models based on useless indicator traits, and become models themselves based on those traits.

Henrich and Gil-White (2001) propose additional factors that could consolidate the runaway nature of this process. They highlight a distinction between dominance and prestige, stating that dominance is the result of powerful individuals asserting their authority over the group, whilst prestige comes from a process of model selection, following the notions established by Boyd and Richerson. In the acquisition of prestige, individuals who are identified as good models are given rewards by those who want to spend time around them, a necessary step towards learning their valuable behavioural knowledge. Prestige is the accumulation of this social attention by popular individuals, and can also be an advertisement of model quality.

The combination of frequency-dependent bias and indirect bias provides solid ground for the view that culture promotes detrimental behaviours, despite its having arisen because of the overall benefits of social learning in a population. However, these behaviours are not arbitrary but are the result of a new kind of evolutionary process based on the learning strategies of individuals. In Boyd and Richerson's words: "Under some assumptions, the model is consistent with the sociobiological expectation that human culture is adaptive, but under others it allows

cultural processes a causal autonomy... Often it seems that culture can be autonomous and adaptive” (Boyd and Richerson, 1985, p. 242). Boyd and Richerson’s view of culture as autonomous clearly indicates how cultural systems can generate new selection pressures, which introduces the notion that *useless* behaviours might become foci not only for continued cultural evolution, but also for continued biological evolution.

In summary of their theoretical orientation, (Richerson and Boyd, 2005) stress that, a powerful adaptation that has various maladaptive consequences does not necessarily die out (since these costs do not outweigh the benefits of the adaptation) and, most importantly, does not necessarily fix the problems associated with this maladaptation. This apparent imperfection can occur either because the organism’s phylogenetic history and current situation does not permit it to find improvements through genetic variation and natural selection, or because the ideal solution simply isn’t possible. Boyd and Richerson refer to the latter state of affairs with respect to human social learning. “No herbivore,” goes their analogy, “can be as fleet as a gazelle, as tall as a giraffe, and as powerful as an elephant. Inescapable biophysical tradeoffs ground magical organisms such as gigantic, flying, fire-breathing dragons. Pigs can’t fly; even if they had optimally designed wings, they’d be too heavy” (Richerson and Boyd, 2005, p. 155). Likewise, you simply can’t have the powerful benefits of social learning without its inherent uncertainty and the potential for behaviour which does not optimise individual chances of survival. Again, this is based on an emphasis on specific ‘fast and frugal’ cognitive mechanisms that actually implement social learning, for which we have no reason to believe nature could provide improved replacements.

This argument is yet another cautionary tale in adaptationist reasoning. Boyd and Richerson uphold the principle of adaptationist reasoning themselves, but not its extension towards indiscriminate assumptions about how ‘much more functional’ human behaviour could be. Such assumptions are based on the implicit inference from the principles of natural selection that nonoptimal behaviours result in extinction. There are two equally satisfactory corrections that can be made to this mistaken view: firstly that human behaviour is as optimal as it can be, and secondly that human behaviour is good enough to ensure human survival, and that is all that matters. In fact, this point has an obvious repercussion for the theory of evolution in general.

The same argument could have been applied to genetic inheritance: natural selection should have favoured organisms whose systems of heredity were perfect, the most accurate replicators. This leads to a *reductio ad absurdum*; we are faced with the evolution of systems that don't evolve, and therefore the stagnation of evolution. Replication with mutation results in all sorts of disasters, dead-ends, and death, but it is unquestionably adaptive.

Boyd and Richerson's model can be applied to the evolution of human musical behaviour in two ways. Firstly, the possibility that human cultural behaviour facilitates group selection opens the door to group cohesion arguments that were stuck with the problem of explaining mechanisms of group cohesion in terms of individual fitness. Brown (2000) follows this route. But, as stated, Boyd and Richerson's group selection hypothesis is still widely accepted only as an *in principle* argument (Blackmore, 2006). Secondly, the possibility that social learning can drive the runaway evolution of a variety of maladaptive behaviours provides a new explanation for the cultural emergence of musical behaviour and *could* explain how it became established as a stable aspect of human social interactions. The question, then, is how the slow march of biological evolution was altered by the sudden new context invoked by cultural behaviour. Did it drive the evolution of individuals who were defended against maladaptive musical pursuits by being increasingly numb to musical interaction? Or did it drive the evolution of individuals who thrived on the benefits gained from living in a population of other musical individuals? In Chapter 3, I will develop an approach to answering these questions.

Other significant approaches to the evolution of culture that will not be considered in significant detail here include the works of Sober and Wilson (1998), Tomasello (1999) and Shennan (2002). Sober and Wilson (1998) present arguments essentially in support of the value of group selection, by proposing a pluralist approach that draws in supposed alternatives to, or counter-arguments to, group selection. They propose a hierarchical approach to selection, arguing that selection can operate on units of organisation at various levels, and to various degrees simultaneously. These include units of selection such as the gene, the organism and the cultural group. In this way Sober and Wilson (1998) incorporates the kind of mechanisms proposed by Boyd and Richerson (1985) along with other proposed fundamental changes to the organisation of evolutionary units, such as those discussed by Maynard Smith and Szathmáry (1995).

Tomasello (1999) presents a theory of the evolution of culture and human cognition which focuses primarily on the emergence of cognitive capacities for directing attention, the attribution of intentions, and the sharing of intentions. This approach is also extended more specifically to language by Tomasello (2003). As with Boyd and Richerson (1985) and Sober and Wilson (1998) focus on sharing intentions focuses on collective aspects of human behaviour and, like Boyd and Richerson (1985), on the potential for culture to accumulate content through the collective actions of individuals. Tomasello et al. (2005) propose that the evolution of collaborative skills may have driven the main departure for human cognition, and that this may have resulted from selection at a group level, for example in situations where cohesive groups can collectively defend a food source whilst maintaining the benefits of this collective action to all of its members. Collaborative action, they argue, requires shared intentionality, whilst competitive interaction does not require an understanding of intentions at all. Current work by Tomasello and his team involves comparative research into the capacities for humans and other primates to engage in shared intentions.

Shennan (2002) develops gene-culture coevolutionary theory (referred to as dual inheritance theory) focusing largely on the tools of ecological modelling and the analysis of adaptation and optimisation, and also supporting the utility of the concept of a *meme*. Sterelny (2004) praises Shennan's commitment to confronting the problem of the transition from primate social life to human enculturation, but criticises his confidence in the applicability of ecological optimisation models to human behaviour. One problem that Sterelny flags up is the problem of cognitive constraints for evolving humans, altering the conditions for optimality in as yet unclear ways. Sterelny argues that "if the archaeological record is too poor to support strong tests, that would seriously limit the value of behavioural ecology for archaeology" (Sterelny, 2004, p. 253). If Sterelny's complaint is justified, then whilst ecological models adapted for the context of dual inheritance are likely to be a valuable starting point for models of the evolution of human musical behaviour, the notion of optimality itself may be more problematic, at least in the context of comparisons with empirical data. We know that evolution optimises, but the problem of what constraints guide the process of optimisation in biocultural systems may render this focus untenable.

2.4.2.2 Niche construction

More recently, a perspective has emerged that attempts to achieve a greater sense of continuity between animal and human behaviour, and their relationships to environments. In a definitive introduction to this perspective, known as *niche construction*, Laland et al. (1999) present the relationship between genetic evolution and culture as one of mutual interaction, pointing to “empirical data and theoretical arguments that human cultural activities have influenced human genetic evolution by modifying sources of natural selection and probably altering genotype frequencies in some human populations” (Laland et al., 1999). “Cultural traits, such as the use of tools, weapons, fire, cooking, symbols, language and trade,” they continue, “may also have played powerful roles in driving hominin evolution in general, and the evolution of the human brain in particular” (Laland et al., 1999). Rather than accept the sociobiological tenet that culture contributes to evolution only in that it is an extension of the phenotype, they posit that cultural systems can be seen as a particularly extreme case of what they claim to be a universal phenomenon of complex organisms modifying their environments. Culture is therefore seen as an aspect of both phenotype and environment. Laland et al. (1999) consider the notion that many other organisms can act to modify their environments and change the conditions for their own selection in cases far less complex than that of human culture.

Breaking down a perceived qualitative barrier between human cultural behaviour and the behaviour of other animals with respect to their environments implies a gradualist development of culture, in particular, the specific human phenomenon of cumulative culture, from the basic niche construction process present in all biological activity. Accordingly, Laland et al. exemplify the niche construction hypothesis with the case of hominin evolution. In a significant borrowing of terms from the *order-for-free* perspective of dynamical systems theorists such as Kauffman (Kauffman, 1993), they refer to the possibility that the distinctly human process of vertical cultural transmission (the transmission of cultural information from parents to children) “may become an autocatalytic process: greater culturally generated environmental regulation leading to increasing homogeneity of environment as experienced by parent and offspring, favouring further vertical transmission” (Laland et al., 1999). This is not dissimilar to Boyd and Richerson’s model of emerging imitation, except in the way that an environment is viewed. Niche

constructionism breaks down the distinction between social and ‘natural’ environment, that has been used for the convenience of this discussion. In this way, vertical transmission becomes somewhat more relevant than horizontal transmission because the benefits of social transmission are mutually intertwined with the process of parents modifying their environments for their own sake and for that of their offspring.

In their discussion of hominin evolution, Laland et al. (1999) discuss the possibility that niche construction can both speed up and restrain evolution. For example, Wilson (1985) argues that evolution occurs more rapidly in larger-brained organisms whose behaviour has a more dramatic effect on their environment. On the other hand, niche constructing traits can shield other traits from selection, meaning that they can be more variable. “For instance, improved levels of health care and sanitation are examples of culturally mediated counteractive niche construction that damp out selection against individuals with some gene-related disorders, who may survive and reproduce in the modified environment” (Laland et al., 1999). They argue that the greater variation allowed at these points might then be significant in future genetic change, such as under new environmental pressures. “The particular significance of this for human evolution is that,” they claim, “as unusually potent niche constructors, hominids should be particularly resistant to genetic evolution in response to changing environments, while at the same time capable of dramatic evolutionary change following major innovations” (Laland et al., 1999). This is perhaps the most explicit statement quoted so far of the notion that culture can introduce completely novel evolutionary pressures that would explain the evolutionary importance of certain aspects of human behaviour. There is a subtle difference in how the niche constructionists present this statement, however, compared with the other approaches that we have seen. In one sense it is obvious that the current behaviour of a species is the starting point and context for its future evolution, but for many this is subsumed by the more significant principle that it is imperative for any species to adapt to a physical environment. Therefore, cultural behaviour is first and foremost adaptive, and although it may have many consequences that are not directly adaptive, these are only consequences; they cannot be the source for new evolutionary processes. The niche construction view, however, invites the possibility that cultural behaviour truly generates environments for future adaptation, built upon layers of other cumulative cultural behaviour.

Niche constructionism also takes an insightful step away from adaptationist reasoning, one that Boyd and Richerson continue to be reluctant to do. As will be discussed below, the simplifying nature of standard evolutionary explanation is less favoured by social scientists, especially in its tendency to reduce cultural behaviour to one or another understood evolutionary principle. Niche construction, its proponents argue, presents an immediately more complex view of evolutionary process that should appeal to social scientists sceptical of evolutionary reductionism. Although Boyd and Richerson do, to some extent, present a complex view of culture which attributes to it some level of autonomy, they continue to insist on framing that in terms of adaptationist explanations (albeit a cautious and open-minded form of adaptationism). The step away from adaptationism is a giant leap. We have seen it in Maturana and Varela's (1980) concepts (see Section 2.3.4), here in niche constructionism, and will look at similar principles in anthropological thinking in Section 2.4.3. In the following section on memetics, these same problems are present again.

2.4.2.3 Memetics

Pursuing the basic fact that elements of culture persevere through a process of inheritance between individuals, it has been proposed that an analogy between cultural evolution and genetic evolution could be far reaching. Evolutionary and *diffusionist* theories of cultural change, which shared such a goal, were enthusiastically debunked by social anthropologists during the first half of the 20th Century (these remain convincing theories of technological change, nevertheless), but the central importance of the gene in the modern synthesis of evolutionary theory introduces a new perspective on this recurring debate. It is from this perspective that Dawkins (1976) suggests that genetic evolution had a cultural analogue in the *meme*. Looking at human behaviour, Dawkins complains that gene-centric explanations provided by biologists, largely concerning the advantages of altruistic behaviour, are dubious, in particular because of the great differences that describe fundamental aspects of human behaviour throughout the world. He rejects the centrality of genes in the explanation of human behaviour by holding onto the pivotal concept of his own analysis of genetic evolution: replication. He proposes that there may be other replicating structures that share the gene's passive, survival-driven character, and in the case of culture these

might collectively be called memes. Memes are assumed to exist in the environment of sufficiently complex organisms, such as humans, who can learn and exchange information. Whilst meme-exchanging organisms are crucial to the survival of the memes, the memes are nevertheless autonomous, just as genes are in selfish gene theory Dawkins (1976). In fact, these organisms are seen as survival vehicles for memes, just as they are for genes (although memes may have alternative survival vehicles such as books and buildings). “Examples of memes,” Dawkins suggests, “are tunes, ideas, catch-phrases, clothes fashions, ways of making pots or building arches” (Dawkins, 1976, p. 192).

Dawkin’s suggestion that a tune is a meme is intuitively attractive. We hear them, they get ‘stuck’ in our heads, and we reproduce them. In this way they propagate, possibly with mutations due to copying errors. Unlike genes, we can also create, or accidentally discover new ones. Tunes in general propagate because they are so well suited to our musical tendencies – to hear, remember, reproduce – and some specific tunes propagate because something about them makes them even more well suited to the copying process, perhaps how memorable they are, or how exciting they are at the moment (only to become boring in time). Presumably certain musical tendencies, such as the tendency to whistle whilst walking down the street, are also memes, which facilitate the transmission of the tune memes. This view also echoes Boyd and Richerson’s (1985) view that social learning can lead to behaviours that are not necessarily as useful to us as the behaviours that have been carefully crafted and handed down through our genes over centuries of evolution. Meme theory makes more of a point of this arbitrariness. Social learning, and other aspects of human psychology, set up a context in which any number of unusual behaviours could emerge, but not only emerge; they also combine, compete for space, adapt, pollute, cause episodes of rapid widespread extinction, and so on. It is as conceivable that musical behaviour emerged in this complex domain as it did in the domain of genes. But despite the fact that such processes may well occur, memetics has been plagued by a general lack of interest because it doesn’t actually provide any powerful tools for analysis. For many, the problem begins with the fact that there is no such thing as a meme: at least not a useful thing.

In the milestone volume *Darwinizing Culture*, Aunger (2000) notes that the meme does not appear to have moved beyond ‘analogy’ status and, given that “the lack of subsequent devel-

opment of the meme has been conspicuous” (Aunger, 2000, p. 2), asks whether the concept of meme has any future. Whilst scientists of culture are generally not averse to the meme concept, anthropologists, tentatively welcoming scientific research into culture, write it off. Laland and Odling-Smee take the pro-meme position: “In our view, biologists and human scientists alike will not be able to understand the evolution of culture unless they are prepared to break down the ‘complex whole’ into conceptually and analytically manageable units. To this end, we regard memes as a valuable scientific tool. For us, the pertinent question is whether they are a useful theoretical expedient.” (Laland and Odling-Smee, 2000, p. 122). Conte (2000) delves deeper into the relationship between memes and cognition. Her central interest is the formalization of rules of social interaction for the purpose of social simulation. To this end, she draws attention to “the various mental states (including social goals, motivations, obligations) and operations, such as (social) reasoning and decision-making, necessary for an intelligent social system to act in some domain and to influence other agents (through social learning, influence and control)” (Conte, 2000, pp. 83–84). In this sense, Conte does not succumb to a meme-centric view of culture at the expense of the individual. On the contrary, her view is that cultural evolution occurs through the interactions of memes and (social) minds.

Opposition to memetics is unanimous amongst the traditional social scientists contributing to this collection (Aunger, 2000). Three specific cases are summarised here in the authors’ own words. Firstly the analogy is weak: “cultural items are ‘re-produced’ in the sense that they are being produced again and again [but not] in the sense of being copied from one another. Hence they are not memes, even when they are close ‘copies’ of one another” (Sperber, 2000, pp. 164–165). Secondly, the explanatory power of the meme is weak: “evolutionary – or Darwinian – approaches to culture, or society, or humanity are not to be reduced to a single question, let alone a single type of answer. The Darwinian programme in the human sciences should be open, eclectic, and multi-faceted.” (Kuper, 2000, p. 186). Thirdly, the distinction between meme and behaviour is not possible: “as the American critics of the diffusionists showed, memes, like traits, will continually be integrated and transformed by the receiver of information... This is where life is, not in the bits... [and] culture would not be made up of a single isolable type of coded information, which, even for the sake of analysis, could usefully be understood as separate

from social life” (Bloch, 2000, p. 201). The overwhelming opinion of those whose expertise lies in the study of culture *in situ*, then, is that memes do not exist.

In sharp contrast to Kuper’s (2000) call for an eclectic approach to evolutionary theories of culture, Blackmore, the key proponent of memetics, expresses astonishment that other theorists, such as Richerson and Boyd (2005), do not accept the view that social learning involves a process of replication, although without offering a defence of the utility of her view beyond a subjective conceptual framing of the facts (Blackmore, 2006). However, Blackmore’s approach to memetics is strongly geared towards the autonomy of cultural evolutionary processes, relieved of the constraint of ‘Wilson’s leash’: the principle that evolutionary outcomes are always ultimately constrained by the genetic determination of behaviour. Following Dawkins, Blackmore proposes that cultural processes are not necessarily adaptive, but may be quite the opposite; humans have survived to this day despite their wayward cultural behaviour, not because of it, and this cultural behaviour may have also driven genetic evolution in new directions (Blackmore, 1999). This minority opinion is an incredibly important proposal to acknowledge, and has rarely been stated with such clarity, but it does not itself depend in any way on the central argument of memetics, that cultural units of information exist that are replicated by learning individuals. Indeed, many views of anthropology, in particular those pitted in opposition to functional and ecological explanations of cultural systems, view cultural processes as potentially autonomous and not driven by adaptive pressures.

According to Bloch (2000), the debate on memetics would have been short lived if there had been a properly established communication between anthropologists and scientists of culture. He condemns the ignorance of scientists over the major issues in anthropological theory, but his condemnation of anthropologists is greater still for their unwillingness to even acknowledge scientific approaches to culture. This remains the general state of affairs. Whilst it is the principle of Occam’s Razor that guides scientific approaches to cultural dynamics, critical theory, it could be said, has become the guiding principle of social anthropological discourse in managing the chaos of ethnographic and global political reality. Taking Bloch’s (2000) advice, it is necessary for a scientific approach to the evolution of culture to understand the knowledge and motivations guiding an anthropological view of culture. I consider only a handful of very broad issues

regarding the general state of theory in social anthropology, and how these can influence our general understanding of the kinds of contexts for social interaction that we would expect to have existed during the evolution of human musical behaviour.

2.4.3 Anthropological and sociological views of culture

In 1895, the prominent sociologist Emile Durkheim wrote:

“whenever any elements combine and, by the fact of their combination produce new phenomena, it is evident that these phenomena are not given in the elements, but in the totality formed by their union. The living cell contains nothing but mineral particles, just as society contains nothing but individuals; it is obviously impossible, however, for the phenomena characteristic of life to exist in the atoms of hydrogen, oxygen, carbon, and nitrogen. For how could the properties of life exist within inanimate elements? ... Let us apply this principle to sociology. If, as we may accept, the synthesis *sui generis* which every society constitutes yields new phenomena, differing from those which take place in the individual consciousness, we must also admit that these facts reside exclusively in the very society itself which produces them, and not in its parts.”

(Durkheim, 1972, p. 69)

With this emergentist stance, Durkheim was at the centre of a revolution in the social sciences leading to the European traditions of sociology and social anthropology. “Thus we have a new justification,” he argued, “for the separation which we have established between psychology proper, which is the science of the individual mind, and sociology” (Durkheim, 1972, p. 70). He continues, “social facts do not differ from psychological facts in quality only: *they have a different substratum*. They do not develop in the same milieu and they do not depend on the same conditions” (original emphasis). In Durkheimian tradition, sociologists and anthropologists throughout the 20th Century have studied social systems as phenomena that are in some sense irreducible to human psychology. Clearly this premise can assume many different guises, and indeed it has been a central point of debate for anthropologists over the past century. It is

therefore a peculiar fact that, from the perspective of onlookers in the natural sciences (Barkow et al., 1992), anthropologists have stood up in favour of quite extreme versions of this view in spite of so much evidence that suggests a more compromised perspective.

Ingold (1996) sheds light on the character of the modern discipline of anthropology by emphasising its two quite separate driving forces: “anthropology is at once the most resolutely academic and the most fiercely anti-academic of disciplines. Its commitment is to human understanding of a very fundamental kind, and it continues to exist and thrive only thanks to a university system which – at least in principle, if no longer in practice – is dedicated to the production of knowledge for its own sake. Yet at the same time, anthropologists have been foremost in challenging the claims of academia to deliver authoritative accounts of the manifold ways of the world, along with the implicit ranking of such accounts above those that might be offered by ‘ordinary folk’ whose powers of observation and reason have supposedly not been cultivated to the same degree” (Ingold, 1996, p. 1). Although this appears to offer *carte blanche* for scientists to stop reading and conclude that anthropology has nothing to offer the *scientific* study of culture, a lesson can be learnt from studying how such a state of affairs came about.

Anthropology’s drift from the scientific mainstream is epitomised in its relationship to evolutionary theory. Although popularised by Herbert Spencer in the 19th Century (Spencer, 1867), in the early 20th Century an evolutionary theory of culture was proving problematic in two ways. The first, and most commonly cited reason was that it was politically dubious, particularly in light of Spencer’s progressivist formulation. The suggestion that European and North American cultures were further along an evolutionary journey than other cultures encouraged a sense of Western superiority and justified colonial rule and domination. If the dangers of this supremacist attitude have dwindled through a greater understanding of evolutionary theory, this is in part thanks to the progress of anthropology, in which detailed ethnography has uncovered a great deal of sophistication and complexity in cultures that had always been tagged primitive or simple (non-literate or non-industrial are common alternative descriptions, although they are both negative definitions). Whilst it should be possible to maintain an evolutionary view of culture that does not permit supremacist overtones, in practice the modern anthropologist is justifiably sensitive to the fact that academic theories can profoundly influence social practice.

This factor is central to the second view of anthropology mentioned by Ingold (1996).

A second reason for adopting a non-evolutionary viewpoint is that the evolutionary study of human cultures has proved to be methodologically intractable, for good scientific reasons. In non-literate cultures, whose practices and structure were of greatest importance to 20th Century anthropologists, evidence for the diffusion or inheritance of cultural traits was scant and of questionable meaning, and historical information was nonexistent by definition (inseparable from myth, since only transmissible through oral communication). In other words, there were serious methodological reasons to abandon the study of cultural change over time and to devise an alternative problem-specific approach to the study of human cultures in which the structures currently in place were to be described in terms of their immediate function. This functionalism was necessarily couched almost entirely in understanding the stability of the present, rather than speculations about the dynamics that led to this point, and in this sense could be viewed as loosely Darwinist in its organic view of cultural systems.

The Polish-born British anthropologist Bronislaw Malinowski is largely credited for devising the modern approach to ethnographic fieldwork during a visit to the Trobriand Islands at the time of the First World War (Malinowski, 1932). His extended period of close involvement with the Trobriand Islanders illustrated the value of immersive first-hand study of a social group, known subsequently as *participant observation*. It demonstrated the imperative need to develop an analytical approach that recognized the difference between what people did, said and actually believed, as well as subjective differences between individuals within the group. These three levels of ethnographic reality define a sophisticated framework for studying cultural systems, in which cultural rules and norms are fluid entities that can be broken by free-thinking agents, and recorded statements needn't correlate strictly to the intricacies of an individual's beliefs (Malinowski, 1932).

The history of social anthropology has, with notable exceptions (see discussion of Lévi-Strauss in the following section), continued to expand upon the intractability of such theorising about culture. The most significant major shift in anthropological theory has been to enthusiastically embrace post-modernism. In its reflective nature this exposed a lack of transparency in the use of language itself as a medium for academic discourse. Clifford Geertz famously proclaimed

that “man is an animal suspended in webs of significance he himself has spun” (Geertz, 1973, pp. 4–5), arguing for the study of culture as an interpretative search for meaning within specific cultural systems. Geertz pioneered a style of ethnographic writing which attempted to capture the ethos of cultures, rather than to lay bare the structures and systems that constituted a culture, as a scientist might. The dissection of a cultural system into theoretical fragments has become treated with increasing suspicion. Instead, culture came to be understood as resembling a text, with complex and overlapping layers of meaning, and myriad interpretative possibilities. The aim of the modern anthropologist is to understand culture as the genuine participant in a cultural context would see it, rather than achieving a bird’s-eye systems view. This can be understood in terms of the distinction between emic and etic ways of describing culture (Pike, 1967). In these terms, general systems-level questions remain important, in particular, how individual acts ultimately generate established cultural systems, an issue that Tomasello (1999) addresses from an evolutionary point of view whilst managing to acknowledge the importance of the specificity of human enculturation.

Following this post-modern turn, the act of writing ethnographic material has itself gone under the microscope (Clifford and Marcus, 1986), and anthropologists move increasingly towards studying themselves and their methods as part of the system in question (Strathern, 1995). Contemporary anthropology recognizes the importance of analyzing earlier ethnography as contextually bound literature, uncovering theoretical biases that have fed back into analysis through fieldwork practices. Gellner (1987) captures a fundamental issue at the heart of this problem:

“The situation facing a social anthropologist who wishes to interpret a concept, assertion or doctrine in an alien culture, is basically simple. He is, say, faced with an assertion S in the local language. He has at his disposal the large or infinite set of possible sentences in his own language. His task is to locate the nearest equivalent or equivalents of S in his own language”. [There is no mediating third language.] “Naively, people sometimes think that reality itself could be this kind of mediator... that equivalences between expressions in different languages could be established by locating just which objects in the world they referred to”. [But] “language functions in a variety of ways other than ‘referring to objects’ ”. “Many

‘objects’ are... created... by the manner in which [language] carves up the world of experience”.

(Gellner, 1987, pp. 22–23)

This analysis captures a problem that is clear to social anthropologists, but to which many scientists studying human evolution seem to be oblivious. Gellner manages to express the problem in a way that is clear to scientists, but that has presumably been obscured by the incomprehensibility of much anthropological analysis.

Returning to the notion of the function of cultural systems in anthropology, it is fitting to briefly introduce one more perspective that sheds light on notions of evolving cultural structures. This is a perspective that emphasises power struggles within social systems rather than assuming homogeneity or merely stochastic variance, and is most commonly associated with the social theory of Karl Marx (*e.g.*, Marx, 1995). Marx’s legacy is a good example of why such debates are so heavily politicised and why it is often inappropriate to reduce such theory to pure science: the social and cultural systems that naturally occur, whose structures naturally maintain themselves through reinforcement, possibly due to an autonomously evolved functionality, are often systems that maintain great inequality and suffering throughout the world. Science grinds up against culture at this point, and is complicated by it.

Marx’s view of social structure is one of relationships of dominance based on mechanisms of exploitation (Marx, 1995). In Western societies, for example, women were seen as being subordinated by men’s economic and political dominance, and in capitalist societies, workers were seen as being subordinated by notions of ownership reinforced by the dominant classes. In response to capitalist doctrine which assumed that capital could generate value in and of itself, Marx argued that the social meaning of value could only really be a property of the labour, rather than the capital, invested in a product. Workers were alienated from the value of their production by a capitalist system that reduced their labour to a commodity, whilst maintaining the value-generating power of capital, the financial investment in an enterprise. One might take from this analysis the simple rule that a certain social organisation can be reinforced by a minority to whom it presents a benefit. The reinforcement incorporates a feedback process in that the power

gained by the dominant individuals is channelled back into the maintenance of that organisation. In this way we can describe a system of domination as an evolutionarily stable social organisation, without having to agree with some supporters of capitalism that its existence is evidence of its overall evolutionary ‘fitness’. These kinds of power relations are characteristically human and can be seen as yet another elaboration on the basic notion that cultural entities, such as a conception of money, dominate our perception of the world quite extensively. Marx’s labour theory of value (Marx, 1995) promoted a new world view that became embodied, for better or for worse, in an alternative social organisation, communism, from which new forms of exploitation naturally emerged. Simple models of gene-culture coevolution, however, necessarily gloss over this complexity of internal social structure in a population. Traditionally, theories of human evolution assume that such complex relations only emerged, or reached significance, during the cultural explosion that followed the bulk of human biological evolution (Knight, 1991, see, for example), but the view that this kind of relational structural complexity is commonplace in animal behaviour is gaining pace (*e.g.*, Dunbar, 2004).

There are essentially two important and closely related views that are defined in the anthropological debates that have been only briefly covered here. The first, characterised by Geertz’ “webs of significance” (Geertz, 1973), is that according to anthropologists culture bears heavily on the way that individuals understand and behave in the world. This may be truer in some domains than in others. Thus in some sense cheats and altruists are no different across cultures, across the animal kingdom and in imaginary logical games such as the iterated prisoner’s dilemma. But, according to this view, culturally constructed entities such as gods and magical powers dramatically alter individual behaviour by effectively altering the environment in which those individuals behave. The contexts in which acts of selfishness and altruism are performed, then, are not seen as being limited to situations in which the outcome actually determines an ecologically tangible fitness gain. In some cases these culturally constructed phenomena may only affect an individual’s perception of the world, quite possibly to the detriment of their survival, but in the sense that these phenomena affect the behaviour of other individuals – they define the social environment – they constitute a new context for survival, potentially benefiting some at the expense of others. This ties in with the second important view, epitomised by Marxism, that

sees cultural systems as embodying and reinforcing inequalities. In reality, it is not only individual privileges that differ in a social group, but world views themselves; different religions, ethnicities and classes coexist in modern societies playing out various power struggles. World views even differ dramatically from one family to the next. This is not necessarily due to cultural variation operating on an otherwise homogenous cultural system, it may be a stably reinforced cultural structure that perseveres through time.

One evolutionary theory that has taken into account this more complex view of human social relations, in which power relations between identifiable groups is identified, is Knight's (1991) development of a menstruation theory of the evolution of culture, built on earlier work by Turke (1988). Like Dunbar (2004), Knight is interested in the conditions allowing for the evolution of cooperative behaviours such as the establishment of an information bearing linguistic system. Following the same prisoner's dilemma reasoning as Dunbar, Knight (2002) argues that information bearing language simply could not evolve in a non cohesive system of social interaction, and that one needs to understand what changes took place that allowed the right evolutionary conditions to become established. But Knight's focus is on the relationship between males and females in primate groups, and the power struggles between these groups. He argues that typically, in primate societies, female investment in child rearing is prohibitive and subject to massive exploitation by promiscuous but unhelpful males who have little certainty of fatherhood of their children, and therefore little motivation to invest in their upbringing. A dramatic change came about in hominin evolution, however, through a change in female behaviour, from a state of individual competition for the best mates to cooperation through a mutually beneficial action that Knight terms *sex strike*: "some unsuspected pattern [of mating system] resting on inter-female solidarity on a level unknown among non-human primates" (Knight, 1991, p. 204). Under selective pressures resulting from new predation risks and increasing infant dependency, sex striking emerged as females began to coordinate their menstrual periods, which, according to Turke (1988), drove, and subsequently coevolved with, selection for more cooperative provisioning males, who provide food for individual females with whom they have established sexual relations. It can be seen how this pair of coevolving phenomena could also have coevolved with increased altrification and bipedal development.

The details of the argument supporting this theory are complex and beyond the scope of this thesis, resting on extensive discussions on the sets of costs and benefits of particular behaviours within particular social contexts. In short, the theory convincingly addresses a number of observations about human behaviour that fall outside of the issues discussed so far, including the unique sexual characteristics of human females amongst primates (continuous but discriminating sexual receptivity along with concealed ovulation, the opposite of most primates), the widely observed capacity for passive menstrual synchrony amongst women, and the near universal existence of ritualised offerings of meat from hunting males to females in hunter gatherer societies. This approach is of interest from the point of view of this thesis not because of the specific details of the argument, but because it represents a level of complexity of evolutionary explanation that is not typical of evolutionary psychologists, but approaches the complexity of anthropological discourse, a key aspect of which is the formation of individual and group roles and identities in heterogeneous societies, and their economic power relations. This complexity is embodied in the set of relations and concurrent processes required for the menstrual theory to be established. Complex social relations themselves could be capable of driving new organisations of individuals that radically alter evolutionary environments, as Knight's theory can be broadly understood to propose. As Knight states, "although answers at this stage must remain tentative, merely to ask this question is to begin to glimpse the possibilities in a new way" (Knight, 1991, p. 204).

Furthermore, the theory addresses in detail the specific theoretical consequences for dynamic sets of social relations of a biologically grounded system with a reasonably limited set of parameters: a signalling element, a handful of basic costs and benefits, and constraints derived from time management, cognitive complexity and the strength of social networks. This holds promise as a general methodology for approaches to other aspects of hominin evolution, which may identify similar processes operating at different levels. Knight suggests that human cultural evolution involved the establishment of a collectively imagined 'virtual reality' (Knight, 2002) which radically altered power relations between groups of individuals. Such a view forges a strong and sought-after link between trends in contemporary social anthropology and Darwinist thinking about human evolution, allowing anthropological thought to become embedded into Darwinian

explanations of human evolution, rather than just established as a consequence of Darwinian processes. But for the Darwinian demands on such a view to be satisfied, the most important challenge lies in understanding how this virtual reality became established and maintained to the point of influencing human behaviour. Social intelligence theories of human evolution touch on such notions, but do not propose clear lines of investigation in this area.

2.4.4 Creativity as a cultural driving force

This final section does not follow neatly from the discussion so far, because the discussion of human creativity presented here does not take into account either the evolutionary origins of creativity or the role of creativity in evolutionary processes. Likewise, little of the discussion so far has touched explicitly on human creativity, and yet it is hard to think of human musical behaviour without considering creativity. The discussion of culture has focused firmly on the notion of the transmission or inheritance of cultural information through a population, the significance of cultural contexts to the individuals who inhabit them, and the structure and function of cultural systems. But it is not only external environments and chance events that drive change in cultural systems. In this section, I consider approaches to human creativity. Although this discussion is independent of evolutionary debates, some of the ideas discussed here will inform decisions made in Chapter 5 regarding the use of a simple notion of creativity as a driving force for cultural change.

Sternberg and Lubart (1999) define creativity as the ability to produce work that is both novel and appropriate. This view is widely shared by contemporary creativity theorists. They point to a conspicuous lack of study in creativity by psychologists, two reasons for which are “[the appearance that it lacks] a basis in psychological theory or verification through psychological research”, and “problems with the definition of and criteria for creativity that seemed to render the phenomenon either elusive or trivial” (Sternberg and Lubart, 1999, p. 4).

An early attempt at a psychological theory of creativity is that of Koestler (1967). Koestler proposed that the creativity of the human mind arose as the interconnectivity of the evolving human brain increased across a threshold beyond which the potential for analogical and metaphorical thought emerged. He introduces the word *bisociation* to describe the process by which

human brains make such connections: “Bisociation means combining two hitherto unrelated cognitive matrices in such a way that a new level is added to the hierarchy, which contains the previously separate structures as its members” (Koestler, 1967, p. 183). Introspective evidence for the creative thought processes in a group of American mathematicians, gathered by Hadamard (1954), is taken by Koestler to support the notion that creativity itself occurs at a level lower than rational thought, and not represented in linguistic terms.

Koestler (1967) also proposes three different emotional responses to the creative process of bisociation, referred to as the HAHA, AHA, and AH reactions. The HAHA reaction refers to the humour that we commonly find accompanying bisociation; an expression of joy at the stimulation offered by the new combination of concepts. The AHA reaction refers to the moment of intellectual discovery (eureka!), in which newly combined concepts prove to open up a new perspective on reality. Koestler refers to the common historical fact that radical, albeit correct, scientific theories are often met with laughter – this was exactly the response of Galileo when he heard Kepler’s theory that the tides were caused by the movement of the moon. Finally the AH reaction refers primarily to the emotional response to art (or indeed science), which involves a certain experience of self-transcendence caused by the mystery and ineffability that can arise in bisociation. Koestler summaries thus: “The HAHA reaction signals the collision of bisociated contexts, the AHA reaction signals their fusion, the AH reaction their juxtaposition” (Koestler, 1967, p. 192).

Boden (1990) presents an analytical framework for a cognitive theory of creativity. She dismisses *combinatorial* views such as Koestler’s bisociation, arguing that this cannot account for the open-ended nature of creativity, and describes creativity instead as a form of search through a *conceptual space*. Search, she suggests, may be either confined to this space – in which case it is dubbed exploratory creativity – or, more rarely, have the effect of enlarging it or mutating it – in which case it is dubbed transformational creativity. Whilst individuals inhabiting similar social contexts will have developed roughly the same conceptual space, it is possible that one arrives at a part of the space, or opens up a new space that had not been arrived at before. Boden distinguishes this H-creativity (for historical) from P-creativity (for psychological).

Boden’s exploratory-transformational view of creativity defines the role of *heuristics* in cre-

ative search. Heuristics can inform exploratory search, and the discovery of new heuristics coincides with a transformational event. Heuristics can also be hierarchically structured: some heuristics order the processes by which other heuristics are generated and manipulated. And so it is possible that creative individuals possess (learned or innate) high-level heuristics that can be called into practice during any creative search, for example “consider the opposite”. The formulation of empirically observed heuristics, such as those provided in the mathematician G. Polya’s book *How To Solve It* (1957), allows for computer simulation of the creative process within restricted domains.

The development of this kit of concept-handling tools in humans is explored by psychologist Karmiloff-Smith (1990). Karmiloff-Smith presents experiments that involve asking children of different ages to draw a house that does not exist (that had something unusual about it). She groups the various concepts that are explicit in the drawings according to their generality, and shows a gradual progression of conceptual versatility with increasing age. Younger children’s drawings tend to show a rearrangement of the features present in ordinary drawing: doors and windows exchanged, individual elements resized, and so on. Older children are able to manipulate their reality with far greater dexterity, showing numerous ways in which they had reorganized their set of concepts pertaining to the form of houses at a higher level. These experiments support the hypothesis that creativity takes shape during ontogenetic development.

Boden also focuses on the state of reverie which is known to have preceded many well-known H-creative events. For example, when Kekulé discovered the structure of benzene rings he was drifting off by the fireplace and began dreaming of snakes biting their own tails. This led him to realise that the structures he assumed had been strings were indeed rings. This associates creativity again with a subconscious process, and one in which metaphor is instrumental, or as Koestler describes it, “some far-fetched, reckless combination of ideas, which would be beyond the reach of, or seem to be unacceptable to, the sober, rational mind” (Koestler, 1967, p. 181).

Creativeness, as a general capacity, is viewed in higher-level psychological and social terms. Barron and Harrington suggest a number of personality traits that come hand in hand with creative ability: independence of judgment; self-confidence; attraction to complexity; aesthetic orientation; and risk taking (Barron and Harrington, 1981, cited in Sternberg and Lubart 1999).

Sternberg also cites other ways in which creativity interacts with other aspects of social life. One is that recognition of creativity can increase motivation for creative individuals. Creativity can then be seen as operating in a feedback relationship with other personality traits. Another is that creativity varies with respect to environmental variables such as cultural diversity, war, availability of role models, and so on (Simonton, 1984, 1988, 1994a, 1994b, cited in Sternberg and Lubart 1999). Sternberg goes on to suggest a confluence model in which creativity is understood to occur due to the convergence of a number of components. For example, Amabile (1983, cited in Sternberg and Lubart 1999) suggests that creativity is the confluence of intrinsic motivation, domain-relevant knowledge and creativity-relevant skills, which include cognitive skills, heuristic knowledge, and a work style. Similarly, Gruber (1981, 1988; Gruber and Davis, 1988, cited in Sternberg and Lubart 1999) considers the development in the individual over time of knowledge, purpose and affect, leading to creative undertakings. Sternberg's own approach, meanwhile, centres on the economic nature of creativity, suggesting that creative individuals are those who are able to, as he puts it, "buy low and sell high". It is not clear to what extent any of the above definitions suffer from an ethnocentric bias. The centrality of the notion of creativity itself, largely unquestioned in these accounts, also potentially suffers from a lack of cultural comparison.

A biocultural coevolutionary approach to creativity is provided by Lumsden (1999), who asks whether creativity is an adapted trait specific to humans or smoothly varying across the animal world. Reminding us that neo-Darwinian theory has by no means dealt with the problems of macroevolution, he suggests that the study of creativity strikes at the heart of the concerns of contemporary neo-Darwinism with questions of play and innovation behaviours in non-human species. "Human creativity" he states, "is the fire that drives gene-culture coevolution" (Lumsden, 1999, p. 160). This drive, he argues, is also closely integrated with the nature of cultural diversity.

"It is important not to underestimate the role in our evolution of innovated social ephemera like the earliest chants and narratives. Our childhood encounters with them build our minds. There can be no serious evolutionary thesis about creativity until this paleoarchaeology of ephemera has been developed to its limits and inte-

grated fully with the data on ‘endurables’, and the fieldwork on play and innovation in animals integrated with human data.”

(Lumsden, 1999, p. 162, with reference to Kriendler and Lumsden, 1994.)

Lumsden also asks to what extent great minds matter in the empirical study of creativity. He suggests that the occurrence of creative events has a chaotic nature and asks whether the specific creative mind matters as much as the social conditions that give rise to the event. This point of view questions the effectiveness of psychological explanations of creativity.

Likewise, Csikszentmihalyi’s (1999) theory derives directly from the frustration experienced in his early research in trying to pin down the psychological attributes of a creative person. His conclusion from this experience is that creative personalities and real creative success simply do not correlate. Csikszentmihalyi rejects the solution to this problem that states that the researcher should focus on the creative process rather than the creative success: “If creativity is to retain a useful meaning, it must refer to a process that results in an idea or product that is recognized and adopted by others” (Csikszentmihalyi, 1999, p. 314). And so he argues: “what we call creative is a phenomenon that is constructed through an interaction between producer and audience. Creativity is not the product of single individuals, but of social systems making judgments about individuals’ products” (Csikszentmihalyi, 1999, p. 314, original emphasis). Similarly, it is impossible to separate creativity from the act of persuading others about that creativity.

Csikszentmihalyi’s systems theory approach therefore considers creativity to occur at the intersection between three entities: the individual; the domain (or cultural aspect); and the field (or social aspect). The domain refers to the particular activity that is the object of the creativity: music composition, for example, or something more or less specific. The field refers to the set of individuals who will provide feedback through their judgment of the creative act, or their acceptance of its product.

These social perspectives on creativity place less emphasis on the power of the creative mind and allow that the automatic reinterpretation and rethinking of domains by new minds can act as a creative force in its own right. Structuring, categorising and making sense of the world appear to play a vital role in the process and can be seen as the backbone of creativity.

The study of creativity is a critical factor in understanding how culturally evolving systems behave, although creativity may not itself be a suitably universal notion and may itself be re-defined by cultural contexts. Again, in general a more universal and systemic perspective on creativity treats it as a simple force that drives cultural change from within cultural systems. Although social learning is the main component of such systems, we would expect little cultural change if learning acted in the absence of creative behaviour, except if errors are the cause of change. Our understanding of the extent of cumulative culture throughout our evolutionary history therefore depends on an understanding of how creative individuals were, and also how geared the system was towards driving specific individuals towards creative behaviour. Meanwhile, Boden's (1990) analysis leaves little doubt about the practical problem-solving utility of creativity, and its close relationship to learning. We saw in Miller's (2000a) theory of the sexual selection of music the proposal that musical creativity was an indicator of a good mind, and that this ultimately accelerated human intelligence. It is equally possible that social competition in activities such as music could have had the same effect. In Boyd and Richerson's model the conditions are established for individuals to compete for the prestige associated with being a good model of behaviour. A feedback process kicks in because the rewards associated with being a good model actually constitute fitness; that is, the model's behaviour does not need to be fit with respect to the environment because the behaviour itself generates its own value. In the sexual selection case, the driving force behind the emergence of this creative behaviour is easily explained through the combination of handicap principle and runaway process, whereas in the case of social competition, the role of creativity would have been less indirect; creativity and natural talent become the most critical qualities in becoming the most successful individuals.

In the previous section, I discussed the study of how humans inhabit rich cultural worlds involving a great number of false beliefs. How can these beliefs be acceptable, let alone useful? Sexual selection theory can explain such beliefs, and the creativity that may have led to them, as part of the elaborate system for displaying intelligence. But the intelligence in question is almost certainly an intelligence judged through social interaction – the ability to understand and take advantage of a system of beliefs – otherwise we should expect truer beliefs to have won out over more extraordinary ones; outside of a social domain, a belief in rational analysis of

evidence would have been more useful than a belief in magic and witchcraft, whereas in a social domain, belief can be interpreted as the process of inhabiting the same cultural world as one's neighbours in order to successfully interact with them. In that case the indicator-trait nature of such beliefs is a redundant addition to their immediate fitness value. Creativity is directly applicable to survival in these dynamic cultural worlds.

2.5 Summary

In this chapter, I have set out the key debates in the evolution of human musical behaviour, and moved on to consider contributing factors to these debates in general evolutionary theory, biocultural coevolutionary theory, and related disciplines involving the study of culture. With regard to the theory of the evolution of human musical behaviour, the key issue that I have been concerned with is the fact that, whilst a cohesionist account of music provides a compelling final cause account of how music functions in modern humans, it lacks a detailed development of evolutionary processes, which could have led to this outcome in terms of the selection of gradually fitter individuals. Dunbar's (2004) theory comes closest to explaining such a process, but is slightly unclear as to the mechanism. If vocal grooming took over from physical grooming using the same principle of honest signalling, then the time dedicated to talking, and its directedness at a specific individual, or small group, is presumably the primary reliable indicator of this individual's commitment to the relationship. Mere exposure to that individual would not do because primate societies are too complex; your nearest neighbours or closest kin may not be your preferred allies (Dunbar, 2004). Dunbar (2004) proposes that the brain is pleasurablely excited by this stimulation, just as the body is in the case of physical grooming, and this is the mechanism by which individuals actually *feel* the bonds that are being formed. If communication is suitably *directed*, but also manages to allow individuals to carry on with other activities, and to communicate in small groups rather than just in pairs, then it is conceivable that a pleasurable response to communication that clearly satisfies these criteria could emerge. To reiterate, reliable communication is reliable for the same reason that grooming is, because if an individual is forming bonds here, he can't be forming bonds there. This means that the mechanism for evaluating incoming communication needs to be particularly good at discriminating between

communication that satisfies this condition, and communication that fails to. The pleasurable response to communication leads to the risk of exploitation (Dawkins and Krebs, 1978; Owings and Morton, 1998). Even very subtle exploitation could lead to evolutionary benefits over time. The conclusion of this chapter is that this approach provides one of the most convincing starting points for an evolutionary theory of music. However, Boyd and Richerson's (1985) theories of cultural behaviour also provide a body of hypothetical principles describing the nature of social learning, also with convincing evolutionary origins, which introduce additional potentially important cultural dynamics. In the next chapter, I will propose a relationship between Dunbar's cohesionist approach to musical behaviour and a proposal for emergence of musical behaviour through Boyd and Richerson's cultural dynamics.

Chapter 3

A Framework for Investigation

3.1 Theoretical development of existing models

In the previous chapter, I proposed four evolutionary orientations for theories of the evolution of human musical behaviour, all of which are supported by empirical evidence and compelling theoretical arguments. To begin this chapter, I contend that further purely theoretical developments can be made in fleshing out the scope of these orientations and their interactions, pointing to the development of relevant testable hypotheses. Such a purely theoretical approach, which does little to address existing empirical evidence, is not readily recognised amongst evolutionary musicologists. Purely theoretical development, for example, was not mentioned in the list of methodological approaches discussed by Wallin, Merker and Brown in the introduction to *The Origins of Music* (Brown, 2000). And yet a great deal of work goes on in theoretical biology and gene-culture coevolutionary theory in developing abstract models that define potential processes in evolution, much of which stems from theoretical debate as well as empirical observation. Given the puzzling nature of human music, its recent rise to prominence as a topic of evolutionary theory, and the massive limitations in gathering empirical data about the actual process of the evolution of human musical behaviour, it seems sensible that purely theoretical discussions about evolutionary processes in such a domain would be informative, especially in light of problematic notions such as adaptation and innateness, and their applicability in this domain. It is particularly valuable because although there are a number of open questions that are currently being tested experimentally regarding the evolution of human musical behaviour, for example

whether musical processes of cognition uniquely inhabit certain parts of the brain, the experimental results may turn out to have limited direct implications for the theory of the evolution of human musical behaviour in general, as illustrated by various recent discussions (Justus and Hutsler, 2005; McDermott and Hauser, 2005; Cross, 2006; Fitch, 2006; Dean and Bailes, 2006; Livingstone and Thompson, 2006; Merker, 2006; Trainor, 2006). Whilst experimental processes chisel away at the mass of available theories, the extension and development of these theories also provide new questions which can be tested experimentally. Furthermore, speculative evolutionary explanations, especially those related to human evolution, inevitably run the risk of being classed as elaborate just-so stories. A common criticism of such hypothetical scenarios is that they seek justification through their plausibility, rather than through experimentation. Unguided speculation aside, it should be noted that plausibility (as well as acceptability, which may take on other dimensions such as political views) is a strong contributor to the evaluation of any theory, and abstract theoretical study can reveal new perspectives on a scenario that radically alter this level of plausibility, and thus our preference of one theory over another.

In the light of the danger of teleology in evolutionary theory (Lewontin, 1979; Plotkin, 1997), I aim to treat the evolution of human musical behaviour primarily as a question of process, sidestepping problems of explanations that rest on final causes, and grounding this theory in solid Darwinian reasoning. But, although theories of the evolution of human musical behaviour are all clearly working within a Darwinian paradigm, it is important to accept that the evolution of human behaviour is just as likely to throw up novel evolutionary processes as it is to exemplify familiar ones, as in the theory of memes (Dawkins, 1976; Blackmore, 1999; Aunger, 2000). Humans differ from other animals in more than one way, and it is clear that some differences would necessarily have contributed to the context for other differences, as in a number of popular accounts such as those of Mithen (1996), Donald (1991) and (Dunbar, 2004). That being the case, theories of the evolution of human musical behaviour should be cautious about relying on out-of-the-box Darwinian processes.

In this chapter, I will propose that various theoretical perspectives described in Sections 2.2.2.2 and 2.4 point to an extended theory which views the evolution of human musical behaviour as intertwined with processes of social learning. I will argue that cohesionist theories,

in addressing the problem of honest signalling, point to the potential significance of music's variability, its nature as a changing-style system. This is unorthodox in the sense that a changing signal might generally appear to be a less reliable indicator of an individual's state, but nevertheless, it is, in principle, capable of indicating an investment. I will also propose a novel approach to the evolution of human musical behaviour, which takes Boyd and Richerson's (1985) view of social learning at its starting point. This theory states that musical behaviour became amplified through the effects of social learning, possibly involving some of the consequentialist cognitive elements pointed to by Pinker (1998). I argue that a strong version of the runaway cultural evolution proposed by Boyd and Richerson (1985) could drive the establishment and reinforcement of musical behaviour in human social dynamics to the extent that musical interaction becomes a domain for selection. This is ultimately a competitivist view, given that individual interactions are based on competition for socially determined fitness. I then propose that both the social cohesionist and social competitivist models lead to a similar process of evolutionary feedback, ultimately based on the tension between a process of genetic take-over, a cohesion-based handicap principle, and the effects of social-learning. Possibly most importantly for this discussion, the critical aspect of music from an evolutionary point of view is that it is a salient, immediately mutable signal allowing a one-to-many or many-to-many interaction.

3.2 Extension of cohesionist models

I begin by considering cohesionist theories and the problem of what musical behaviour contributes to a process of cohesion. Following arguments such as those of Dunbar (2004), social cohesion is driven by external environmental pressures favouring larger groups and the challenge of overcoming local competition between conspecifics. Dunbar's own grooming theory points to an honest sign of allegiance; since you cannot groom everyone all of the time, you are necessarily bound to allegiances with those you do groom, who remember your efforts. This is assuming that grooming implies allegiance at all, an assumption supported by the pleasure-inducing endorphin release that grooming causes, an evolutionary affirmation of the fitness rewards available to those who play by the rules of social grooming. Primates are therefore assumed to have evolved a sort of innate set of ground rules for interaction in which the complex dynamics of social

cohesion can be played out. Dunbar (2004) proposes that laughter, and perhaps music, have a similar effect, with the advantage that they are not limited to one-to-one contact and can be engaged in whilst multitasking. They are therefore more efficient at establishing bonds in larger groups, for which grooming would be too costly. But what of that cost? If the time and energy cost involved in grooming is actually a meaningful handicap, not just a mere hindrance, then is there a similar such time cost associated with vocal interaction? Also, since grooming is an unambiguously one-to-one interaction there is no question about which individuals the bonding involves, but with vocal interaction the specificity of the participants needs to be established. By what mechanism could music establish the honesty of this interaction?

3.2.1 The handicap principle applied to a changing-style system

To answer this question I propose another cost: the cost of accurately learning style. It is a quintessential fact of human musical behaviour that musical style variation correlates well with social group organisation. Hagen and Bryant (2003) point out that this would make an excellent indicator of coalition strength, since it would indicate how long a group has been together. This is an inversion of the above argument: time itself is enough to solidify bonds within a group, and music functions to communicate the strength of that bond to enemy groups. If they hear complex well-rehearsed musical performances, they can be sure that the coalition strength of that group is high. Necessary to Hagen and Bryant's (2003) theory is the notion that the musical performance takes time and effort to learn. It is therefore necessary that performances are not the same across time and across cultural groups. Varying style establishes a context in which time-intensive learning is the only way to achieve a goal; it sets up the context for a handicap. If this is true between groups, then it could also be true within groups, defining a way of establishing bonds, as well as indicating them to others. Thus, individuals develop their allegiances by spending time correlating their musical style with others, what we might call 'stylistic bonding'. Small groups of style are driven to change, to intensify the uniqueness of the style that needs to be learnt. This still satisfies the one-to-many advantages of musical behaviour over grooming, but with an important bounding effect; you can't bond stylistically with everyone if there are too many different styles to learn. Nettle (1999) has also developed a theory of the group-cohesive

value of dialect in spoken language which is very similar to the theory being proposed here.

But there are additional consequences of treating stylistic variation as the key to music's role in social behaviour: it actually sets up the possibility that bonds become established across larger groups in a hierarchical manner. For example, individual *A* bonds with his immediate neighbour *B* by developing a close stylistic proximity to him, *B* bonds with his immediate neighbour *C*, and so on. *A* and *C* turn out to have a strong bond, if measured in stylistic proximity, although not as strong as *A* and *B*, or *B* and *C*. According to this view, style becomes a marker for relatedness, and social groups can use these behavioural markers as indicators of group cohesion at different levels, if appropriate opportunities arise. Indeed, although Dunbar's number, 150, (Dunbar, 2004) arguably defines the predominant group size in many human organisations, modern humans perceive themselves as embedded in variably sized groups ranging from a necessary minimum of 2 people to a maximum that is determined by the planet's human population.

In human society, therefore, we can use behavioural trends as markers for group cohesion, as well as the mechanism by which group cohesion is achieved. This could explain why we are so stimulated by specific learnt styles. The ape's grooming mechanism is less powerful. If *A* and *B* form an alliance, and *B* and *C* also form an alliance, there is no implicit marker for the alliance between *A* and *C*.

This extended cohesion theory of music suggests that music functions initially through stylistic learning as an honest signal, and consequently takes on additional novel utility through the abstraction of this process; style proximity becomes a flexible marker of cohesion. This places two critical requirements on musical style: that it does indeed vary perceptibly from group to group and individual to individual, and that it is hard to learn (note that there is no reason why this should not apply to language as well). With regard to the first requirement, whilst learning drives the convergence of musical style, something must be driving its divergence. Thus, there should be little doubt that this theory depends on some kind of change. With regard to the second requirement, it is clear that individuals who are quicker and more flexible learners in that domain should in principle be more successful. We need to ask whether this should ultimately cause the collapse of the proposed system, and if not, why not? I will address these two issues in later sections, but before moving on it is necessary to consider an alternative

route to this same point.

3.3 Introducing new approaches to competitiveness

The system described above depends crucially on a capacity for social learning, and many theorists view the origins of social learning in terms of its utility in acquiring useful behaviours at a low cost. Boyd and Richerson (1985) view social learning as capable of producing undesirable but affordable runaway cultural evolutionary effects. Therefore, I propose that we may find the origins of stylistic bonding through music in such a runaway process. Such runaway processes can be seen as competitiveness processes that share many of the characteristics of sexual selection, but ultimately constitute a novel and more powerful process of evolution. I will argue that it is possible for these processes to embed themselves in particular domains of social behaviour, and subsequently contribute to the direction of biological evolution. Therefore, it may be possible to discuss the changing-style system proposed above in the absence of the cohesionist utility associated with it in that section.

According to Boyd and Richerson (1985), powerful cognitive social learning heuristics evolved in human behaviour, leading individuals to attempt to identify successful models in the population, and try to imitate their behaviour. Maladaptation could easily occur if individuals were to accidentally copy the ‘wrong’ behaviours, especially since frequency-dependent biases might then amplify and consolidate these wrong behaviours. Perhaps, as Boyd and Richerson suggest, further cultural selection might be able to weed out these maladaptive behaviours, but this is more of a tentative hypothesis on their part, based on the self-evident fact that there are, by definition, limits to maladaptive behaviour. Their main point is that the benefits of social learning, when it works, far outweigh most of these slight distractions from useful behaviour. Thus, I begin by proposing a scenario in which Boyd and Richerson’s social learning heuristics lead to the erroneous emergence of music as a significant, highly regarded, social behaviour. I propose that music is a likely contender for runaway cultural evolution for two reasons. Firstly, music is highly salient, or, to be precise, music is the evolutionary *result* of highly salient aspects of behaviour getting accidentally copied during social learning; things such as walking styles and vocal sounds. These are aspects of behaviour that are the most visible, audible, and

identifiable as being unique to an individual. When a learner copies a model, then, unless the copying process is very rigorous and extremely successful at discarding irrelevant behaviours (and the thrust of Boyd and Richerson's (1985) argument is that this is pretty much impossible), these aspects of behaviour are strong candidates for inadvertently getting copied. Secondly, it is conceivable that the collision of cognitive attributes that Pinker proposes provides an alternative starting point for music's accidental rise to prominence: the cheesecake theory transposed back in time. Essentially, any of Pinker's (1998) attributes might be seen as providing the basis for the salience of music as a focus for copying.

In fact, the salience argument is problematic. Assuming we evolved to copy useful behaviours, such as using tools, or washing food, what would cause us to copy walking styles and vocal sounds, even if they are salient? I will propose below that this is not a problem because, very rapidly, the social determination of the fitness of certain behaviours overtook their *actual* fitness. At that point, all salient areas of behaviour are open to runaway cultural evolution.

To reiterate, following Boyd and Richerson (1985), I favour an understanding of process over the temptation to think in adaptationist terms. Naïve adaptationist thinking would have us believe that a tendency for maladaptive behaviour would either be weeded out or cause the extinction of our species, but this is not so. Boyd and Richerson's point is that the social learning mechanism really can't do any better; it is optimal, and more than satisfactory to ensure the survival of the species, even once we factor in the time wasted on useless behaviour along with its disastrous effects. In fact, Boyd and Richerson (1985) describe culture as having the occasional appearance of autonomy over environmental selection pressures (although they are at pains not to concede too much autonomy). This is a critical point in understanding the evolution of human behaviour and has been expressed in different ways by various other theorists, in particular by numerous anthropologists over the past century who have opposed arguments of both the environmental and genetic determinism of cultural behaviour. Taking this view, we can now begin to consider the processes by which music became established and maintained as a result of social learning heuristics and existing behavioural adaptations.

3.3.1 Social learning leads to the social generation of fitness

The above approach relieves us from finding an adaptive origin for music in genetic evolutionary terms, but still leaves us with questions of evolutionary stability, in particular: how does music as maladaptive behaviour defend itself from ultimately becoming weeded out? That is, we may accept that arbitrary maladaptive behaviours come and go, but how do specific behaviours become entrenched? Boyd and Richerson's (1985) model provides a mechanism by which a behaviour such as music might propagate through a social network and become established in a population, but so far this provides no more than a drift effect. If stronger effects were implicated, they would point to a more plausible theory. But Boyd and Richerson's (1985) models assume that the external, non-social environment is the ultimate determinant of individual fitness. In our daily lives we know this not to be the case. An individual who is severely disabled can survive through the altruistic behaviour of others. A person with musical talents has a good chance of surviving and passing on their genes even if they adopt many other fitness-reducing behaviours. This may be the payoff of thousands of years of successful cultural evolution, but the possibility remains that the cultural determination of survival took over to some extent from the environmental determination of survival at an early stage in human evolution. This can be referred to as a state of ecological dominance, and many have argued that at this point intraspecies competitive factors, including sexual selection and Machiavellian intelligence, are likely to increase their role in determining survival. Thus, we consider ways in which Boyd and Richerson's (1985) cultural dynamics can be taken further, to the point at which culture is not only autonomous, but contributes to biological evolution.

Henrich and Gil-White (2001) propose one process which may contribute to a stronger effect. As Boyd and Richerson (1985) state, once individuals have identified other successful individuals as good models, they need to learn from these models. In order to learn from them they need exposure to these individuals, and that may require placating the model in order to gain access time. Successful models should therefore expect to become increasingly successful through this kind of respect payment. Furthermore, it is unsurprising that, with individuals looking for other individuals to learn from, frequency-dependent biases feed back on themselves; we look towards those that others are looking towards. Thus, successful models should find that

they attract increasingly large numbers of followers, and also become increasingly successful through implicit rewards that arise from this attention. These factors conspire towards a winner-takes-all cultural process, possibly built on an arbitrary and environmentally useless behaviour. In short, cultural niches emerge in which success can be achieved. In this way, it may no longer matter *which types* of behaviour are being copied; at an earlier stage, learning was limited, at least roughly, to the domain of genuinely useful actions. Later, once behaviours had the capacity to generate their own fitness rewards, learning could branch out in any number of directions. Furthermore, domains of learning coevolved with domains of social evaluation.

But let us consider in greater detail what needs to happen for this system to become autonomous and independent of environmental factors. Imagine that individual *A* is successful and has arbitrary behaviour *a*. *A*'s success is identified, he becomes a popular model, and increasingly successful as a result. Numerous individuals copy his style, producing lots of variants of behaviour *a*. But if *A*'s success also points to the value of *a*, then anyone else who possesses behaviour *a*, or a variant of *a*, could, in turn, become successful. This requires that *A*'s success results in the identification of *a* as being the reason for his success, or alternatively, that *A*'s success drives other individuals to infer that *a* is a successful behaviour. A circular argument ensues. *A*'s success depends on the value of *a*, and the value of *a* depends on *A*'s success. However, it appears possible that, in principle, all of the elements in this puzzle could indeed reinforce each other. Individual *B* is a typical social learner. He surveys his social environment looking for good behaviours to absorb. He notices that *A* is successful, this is reinforced by the fact that other individuals also think that *A* is successful. He suspects that *a* is the cause of *A*'s success. He strives to learn from *A*. This requires time, effort and the payment of rewards, but eventually *B* gains competence at behaviour *a'*, a variant on *a*. Individual *C* also notices that *A* is successful, and also suspects this is because of behaviour *a*. *C* notices that *a'* and *a* are similar, and also that gaining access to *B* is easier than gaining access to *A*. *C* strives to learn from *B*. *B* gains rewards from *C*, and more attention is drawn to *B* as a result of *B*'s increased success. In such a way, it is clear that although both *B* and *C*'s suspicions about the value of behaviour *a* were potentially wrong, the reiteration of the mistake by numerous individuals redefined it as *de facto* being correct. Thus, some kind of inherent bias towards certain behaviours during learning

could lead to the reification of those behaviours as important, thus reinforcing those behaviours as the ones worth paying attention to. As a final step, we can replace the term ‘suspect’, in this description, with the notion that *A*, *B* and *C* have evolved innate tendencies to treat behaviour *a* as significant. This requires that over genetic evolutionary time the preference for this mode of behaviour becomes innately embedded.

There is no externally defined need for a handicap in this case, and therefore no functional reason why this would need to be a changing-style system, as compared with the case of stylistic bonding, where the value of cohesion drives a need for style learning to be genuinely costly. Nevertheless, if this system was not a changing-style system it would converge onto a single behaviour, and little variation would ensue. Therefore, this system only makes *sense* if it is possible to identify something that also drives changing style. This may sound like a weaker case for the existence of a phenomenon, but is, in fact, satisfied by an anthropic principle: if that particular condition were not in place, there would be no such thing as musical behaviour for researchers to be puzzling over. I consider three possible driving forces for a changing style system. The first is that those towards the top of this system, the models, rather than the learners, may gain additional benefit from changing their styles so as not to converge with those learning from them. If it could be shown that divergent behaviour was of value to successful individuals then this would provide a basis for stylistic change. The second possibility is that the stylistic centres are simply relatively ephemeral and quite noisy, so individuals with very different styles also occasionally become successful models. If this noise is enough to stop successful behaviours collapsing onto the same style then the system maintains its tendency to change, and the dynamic aspect of the system becomes one of its most consistent and reliable features. Since Boyd and Richerson’s view of learning focuses on adaptation to a changing environment, it should be equally well-suited for adaptation to a culturally constructed environment. The third possibility is that having started out as a changing-style system, the momentum of stylistic change is kept up because individuals actually begin to anticipate future trends, not aiming to copy successful models perfectly, but instead to slightly modify their learnt behaviour. Establishing the plausibility of a competitive style-learning system requires modelling how one of these change-inducing processes, or something similar, fits into the system in a stable manner.

There is a sense in which the champions of these culturally constructed niches are, in fact, freeloaders, on the basis that they extract benefits from established systems of social interaction without reciprocating those benefits. Those systems may not have been established for the purposes of mutual benefit, but, nevertheless, other individuals could have been more successful without being bounded by them to the benefit of the winners. However, this is actually the opposite of freeloading on cohesive behaviour discussed by Dunbar, because it does not promote the downfall of the behavioural system. The success of these individuals is predicated on a set of cultural behaviours that they themselves uphold as much as any other individual, perhaps even more so. These behaviours involve selecting individuals in the group who are successful or who others think are successful, attempting to spend time with them, inadvertently rewarding them, and copying their behaviour. With cohesion-based freeloading, the freeloaders who are successful do not uphold the altruistic behaviour which they themselves exploited to such great effect. The champions of cultural selection are also the supporters and maintainers of cultural selection. Similarly, this applies to the emotional response associated with music. Assuming that musical behaviour is established in the system, those with stronger positive emotional responses to musical stimuli may be the ones more driven to succeed in cultural competition, thus propagating that emotional response. Thus, the system is a self-sustaining one. Even if it is a risky waste of time for the group as a whole, it is an unavoidable situation. Again, social learning is as optimal as it can be with respect to the fitness benefits it brings to the population.

3.4 The equivalence between stylistic bonding and competitive social learning

The process described above depends on the notion that learning from an individual either requires or implicitly involves enhancing their fitness in some way. On this point, social-learning competition and cohesive stylistic bonding are closely related. In both cases, the process of learning is associated with benefits to one or other of the individuals. In the social-learning competition approach, learning goes one way, as do the benefits. In the stylistic bonding approach, learning and benefits are mutual. In all cases, learning incurs a cost, and additional costs will be associated with the bonding level, *i.e.*, in the social-learning competition case,

the learner pays to learn (the learner is more bonded to the model than vice versa), and in the stylistic bonding case, both pay the cost of bonding behaviour (assuming they play by the rules). Thus, in the simplest case, these two views can be seen as two variations of the same model of interaction. According to that model, in any interaction between a pair of individuals, each participant decides the extent to which they wish to learn from the other. They then learn a certain amount, which is proportional to a learning cost, and they also pay a certain cost to their model. We might determine coefficients for the cost of learning, and for both the cost to the learner and the benefit to the receiver associated with the payment (it is conceivable that the receiver's benefit is greater than the giver's cost, for example if the giver has a surplus of food, or is dealing in an imaginary quantity such as status, which may have the effect of offsetting or distributing the cost). In the competitive social learning case, learning would be entirely one way; in the cohesion case, learning would be mutual. The key to distinguishing these variations of this model, then, are the factors that influence the decisions made by each individual about whom to learn from and how much (or what) to learn from them.

Another difference between the stylistic bonding view and the competitive social learning view is that stylistic bonding implies that an innate positive response to musical sound is necessary to consolidate its function as a bonding agent, whereas in the competitive social learning case there is no functional utility associated with music and therefore no need to infer the emergence of such innate traits, even if innate traits do indeed emerge from the system. In fact, this difference may be mostly superficial, since in either case the evolution of innate musical behaviours (including musical pleasure) needs to be explained in terms of a process of emergence, rather than in terms of a functional utility.

It is also possible that cohesive and competitive aspects of style-learning are tightly intertwined. On the one hand the cohesive explanation for human musical behaviour can be seen as a gloss for a more convoluted process in which the dynamics of social learning establish the conditions under which the evolutionary step towards cohesion is taken. On the other hand, we can view a competitive social learning model as emerging out of a cohesionist system: once music becomes important in establishing social bonds, with an innate component actually making individuals susceptible to emotional manipulation by music, it becomes a domain which is open

to exploitation, and therefore to competition.

3.4.1 Competition leads to cohesion

The possibility that the products of cultural behaviour could feed back into genetic evolution is supported by various gene-culture coevolutionary theorists, including Boyd and Richerson (1985), niche constructionists such as Odling-Smee et al. (2003), and memeticists such as Blackmore (1999). Blackmore has argued that humankind survived despite behaviours such as music, rather than because of them. She describes musical and linguistic behaviour as parasites for which human behaviour is the host. Parasites are a burden to their hosts, but necessarily survive because their host species also survives in the long term. The result is a coevolution between the two, where future evolution may indeed make adaptive use of the parasite, possibly ultimately towards an overall increase in fitness. Having established the equivalence between competitive style learning and stylistic bonding above it is possible to conceive of processes which would drive the emergence of the latter from the former. For example, the rewards associated with being a good stylistic model might become formalised in a concept such as status. Once a subgroup of individuals with similar behaviour and high status emerged they may begin to engage in mutual stylistic learning towards a common target style. But mutual competitive style learning *is equivalent to* stylistic bonding, due to the bidirectional nature of the bonds. Therefore group cohesion becomes feasible.

Alternatively, we can view this principle in terms of Maturana and Varela's (1987) notion of structural coupling. Assuming the system of interaction is underway and stably established, this system becomes a *reliable* part of each individual's environment. Like any aspect of the environment, each individual uses it to the best advantage it can, and in some cases this results in more interaction, not less. More interaction means more coupling, and this means greater structural congruence; *even if* the interactions involve competition in certain domains of activity, greater levels of interaction will create stronger pressures for cooperation within these domains.

Finally, we can also view this principle in terms of niche construction (Odling-Smee et al., 2003). Once a niche of musical interaction is culturally constructed, it is up to individuals to survive in this niche, which involves success at competition within the specified domain. If

their success depends in some way on aspect of behaviour that reinforce the domain, then their survival will enhance and sustain it in future generations. As the niche becomes increasingly consolidated, it introduces new situations in which mutual benefit may operate.

3.4.2 Cohesion leads to competition

As a basic principle of the cohesion argument, certain qualities of a signal must promote an innate positive emotional response. This requires a non-conscious cognitive apparatus for identifying the signal as appropriate, and triggering the positive emotional response. Trivial though it may seem, if we did not discriminate at some level, then all sound would trigger the same response. Minimally, we would expect at least to distinguish sound that is the vocalisation of a conspecific, which is not of any other category: an alarm call, aggressive growl, and so on. More likely, we would expect socially cohesive vocalisations to express some basic intent and convey additional emotional complexity, possibly taking advantage of existing responses to sounds in the environment. Either way, musical sound is already located along a continuum according to its evaluation through emotional response. Even from the point of view that stylistic variation is not a critical feature of music, then, nevertheless, some communicative sound will be more effective than others at triggering responses, meaning that some individuals will be more successful than others at pushing their conspecifics' musical buttons, either innately or through learning, or even by sheer chance. This invites the potential for an evolution of the capacity to manipulate others' emotions through musical behaviour.

Much of this reasoning has already been outlined by Miller (2000a) in his sexual selection theory of musical behaviour. If we have any such pleasure-inducing biases in our behaviour, sexual selection could take advantage of them: "aesthetic displays play on psychological foibles" (Miller, 2000a, p. 341). But in a social system in which there is equal value in attracting conspecifics into alliances or gaining greater mileage out of existing alliances, as there is in attracting mates, then social competition is also capable of playing on these foibles. Miller argues that attracting mates is all-important. Many support an alternative view that this is simply not true of humans, who are so careful to regulate sexual access.

3.5 The gene-culture feedback cycle of social-learning models

The interaction between cohesive and competitive interactions based on style-learning can therefore be seen to branch in many directions, from a purely cohesionist model to a purely competitive model, with a number of possible cause and effect processes connecting the two. Either way, once musical behaviour becomes established as a context for human interaction, I argue that it is necessarily subject to a process of genetic take-over, which can be seen to kick in at any point along this continuum of possible theories. In the cohesionist case, consider a social group identified by a musical style. Each new member of the group has to learn that style in order to become safely established in the group, and it is a necessary consequence of the handicap principle that it should be hard, but not impossible, for individuals to learn this local style. Individuals who, due to new genetic endowments, find it easier to learn that particular style, do well out of the situation, gaining stronger acceptance in the group with less effort. This is the basic principle of genetic take-over (Baldwin, 1896; Hinton and Nowlan, 1987): a *specific* behaviour that needs to be consistently learnt will tend increasingly towards an innate behaviour, due to the benefits rewarded to those individuals for whom the burden of learning is least.

But according to the principle that the need to learn provides the handicap, then as the behaviour becomes increasingly innate, the handicap principle associated with it diminishes. This is a problem for the cohesionist view. I have already proposed that a changing-style system, rather than a static style system, provides a convincing handicap for music to work as a cohesive agent, because it takes time, effort and exposure to learn a style. I proposed that styles needed to vary across groups so that individuals could identify their allegiance to a specific group, in distinction to any other group, by having to learn that *particular* style. So if genetic take-over erodes the effect of stylistic bonding in one domain, then either the system collapses, or it is able to shift into a slightly new domain in which the costs of learning are still high. It is not clear that this would automatically be the case, so, as with competitive style learning, something driving the existence of a changing style system across changing domains would help ensure that the system wouldn't collapse.

Therefore it could be said that only a changing style system would be a sustainable cohesive

system (see Nettle, 1999), otherwise genetic take-over would ensure that learning would be replaced by innate behaviour and undermine the handicap of stylistic-learning; a changing style system evades genetic take-over and the erosion of the handicap principle. However, if genetic take-over was a true threat to the handicap principle, then its effects would have been felt to some degree: some learnt aspects of musical behaviour, probably those most easily learnt, would have become innate. Whatever the rate of adaptation for the changing style system to dodge genetic take-over, some drift must have occurred.

Furthermore, genetic take-over doesn't stop causing trouble for the cohesion hypothesis just on the basis of specific behaviours. Individuals with greater *general* learning abilities are always likely to be able to take greater advantage of the system, thus driving the evolution of increasingly sophisticated general learning abilities. Again, the handicap principle is in danger of exploitation.

The possibility that genetic take-over conspires to erode the honesty of stylistic bonding leads us back to the potential of the driving force of competitive style learning, and I argue on this basis that the cohesionist theory is only feasible if a process of competitive style learning counteracts the eroding nature of genetic take-over. The result is a runaway evolutionary process which one would expect to have a powerful effect on social learning abilities and on the autonomy of cultural processes.

It should be noted, however, that, beginning from a competitivist point of view, it is not necessary to construct the reciprocal argument. Competitive style learning can also be seen to induce a runaway process of genetic take-over, without requiring social cohesion to facilitate the process. It is still possible that genetic take-over could cause competitive style learning to collapse, but this depends on the social and psychological structure of that competition. From Boden's (1990) analysis of creativity discussed in Section 2.4.4, it seems that the human mind strives to construct new concepts from those around it. This notion of creativity comes across as the missing element in Boyd and Richerson's model, and would be necessary to establish the kind of changing style system considered here. I argue that the feedback process that embeds competitive style learning into social behaviour must do so in such a way that a need for change always outpaces genetic evolution, although I have not attempted to incorporate the evolution of

creativity into this picture of runaway biocultural coevolution.

Thus, returning to the equivalence between competitivist and cohesionist processes of style learning, it seems that the symmetry is broken. Cohesionism requires competitivism to make sense. Competitivism does not require cohesionism, and indeed may set up the conditions under which cohesive processes can sensibly arise. Nevertheless, the competitivist approach has a number of question marks over how such a process can sustain itself. It may be that cohesionist arguments do ultimately come the rescue of a competitivist theory.

3.5.1 Runaway biocultural coevolution drives the evolution of innate musical capacities

According to this view, the process by which specific musical capacities evolve would be a genetic take-over process based on competitive style-learning. Such innate musical capacities could include actual innate emotional, and perhaps physical, responses to musical behaviour, specific perceptual capacities associated with musical structure, and the strong neural plasticity associated with music learning.

The emotional response would be an adaptation to the need to identify and strive to learn successful styles in individuals. Innate preferences for styles would be the result of genetic take-over of the common properties of musical structure during the period of this take-over (imagining a similar genetic take-over process to the abstract model given by Hinton and Nowlan, 1987). They would not be entirely arbitrary, but may have involved and restructured arbitrary elements. Otherwise, emotional responses would be style-relative and depend in part on what style is learnt. Plasticity is in part a prerequisite of the whole process, since the process has social learning as its starting point, but neural plasticity associated with music, as well as other aspects of cultural behaviour, could be seen as being optimised for rapidly soaking up stylistic information. This could be a result of the need to be able to adapt to a changing style system. Together these elements could provide both an innate capacity and a driving force to focus in on successful styles.

If music evolved through socially determined runaway evolution, then specific perceptual capacities associated with musical structure might be explained through the layering of genetic

take-over effects on top of one another in the context of competitive style-learning. According to such a process, new, fitter perceptual layers could emerge and be propagated through the population, and the system as a whole could subsequently interact according to a slightly different set of rules, a slightly altered context within which further successful perceptual systems may emerge, similar to predator-prey arms races such as the simulation models of Cliff and Miller (1996). As individual perceptual systems vary due to slight genetic differences, then, which variations might we imagine to produce successful results?

One suggestion is that perception evolves towards increasingly structuring musical stimuli. Structuring allows us to more accurately distinguish stimuli by increasing the number of dimensions required to express (and therefore perceive) them, as in multi-agent beat-tracking models such as those of Large and Kolen (1994) and Dixon (2000). But is an individual with greater capacity to categorise actually any better off than one who is not, in an entirely socially constructed domain? Clearly this is a context-dependent question, but I will propose a scenario in which this advantage is manifest in cases where a small group of individuals share this greater capacity in a social domain. One way that this could work is if, once a small group has such an ability, the internal dynamics of that group establishes the importance of that ability amongst the wider population. For example, the subgroup could develop a distinct style of behaviour, which they value, and which other members of the population value, but which drives them to view other styles as less valuable. Note that, in such a case, the increased perceptual capacity is not better in any absolute sense, but only in the sense that it drives a new cultural context which is advantageous; it is niche constructing in a cultural context (c.f., Laland et al., 1999). This notion will be explored in the simulation models in Chapter 5. Such a niche-constructing process, due to the feedback inherent in the niche construction model (Odling-Smee, 1994), has the potential to drive exotic and divergent evolutionary results similar to sexual selection and predator-prey arms races (Ridley, 1994).

A potentially important structuring process of human musical perception is the perception of an isochronous pulse in a musical stimulus, which is also closely related to the perception of metrical structure (e.g., London, 2004). Numerous functional explanations have been proposed for the role of rhythmic entrainment between individuals: as a means to maximise the amplitude

of a group signal (Merker, 2001); as a way of establishing expectation and turn-taking in mother-infant interactions (Dissanayake, 2000a); as a way to establish joint attention in individuals (Cross, 2004), or a feeling of boundary loss or one-ness in a group (Mithen, 2005). Few have considered the role of rhythm as structuring the content of musical sound due to its categorical qualities (*c.f.*, van den Broek and Todd, 2003). Rhythmic ability is the subject of great interest in theories of the evolution of human musical behaviour, and a dominant view is that rhythmic entrainment allows individuals to share joint attention (*e.g.*, Cross, 2003a). A genetic take-over evolutionary process could equally well provide an explanation for the perception of rhythm, including beat induction and metrical perception, and their interaction. In fact, the two views are reconcilable if one considers the possibility that rhythmic entrainment became a means to channel the learning of styles, such that learnability became enhanced both by the rhythmic structure of the stimulus, and by the possibility to learn by mirroring actions with a degree of reliable timing expectations. According to this view, intersubjectivity could then be a product of the emergence of mirror neurons for performing such tasks (*c.f.*, Arbib, 2006).

3.6 Additional factors in social-learning models

I have now described in some length a framework for investigation into the evolution of human musical behaviour involving the interaction between competitive and cohesive processes, and certain proposed evolutionary consequences of these scenarios. This view, taking the competitive starting point defined above, is summarised in the hypothesis that it is possible for maladaptive cultural interactions based on style learning to drive the emergence of a self-sustaining cultural domain, along with the evolution of aspects of perception exhibiting the essential features of music. In doing so, I have focused as much as possible on the elaboration of the focal points of this theory at the expense of other secondary but also very important phenomena that are peripheral to it. I will discuss these points here. With the exception of creativity, I will not consider these elements elsewhere in this thesis. This is by no means to suggest that they are irrelevant, and I hope that they will play an increasing part in the development of this theory in later work. In fact, the first three factors – social lotteries, identity and creativity – all contribute immediate extensions to the theory.

3.6.1 Social lotteries

Given that a competitive style learning process is not assumed to impose function on musical behaviour, it may nevertheless be worth considering potential functions associated with the subsequent social organisation. This organisation resembles the kind of winner-takes-all processes seen in the sexual competition of many polygamous species, and yet it is a process that is socially determined. The difference is that, in the case of culture, social value is generated, allowing for a form of competition in which the winner is subsequently supported by the group, as opposed to singular competition in which the winner is simply no longer troubled by competitors (see, for example, Gellner, 1987; Knight, 1991; Henrich and Gil-White, 2001). Successful individuals become super-successful not only by championing the local group but also by gaining their support. This is therefore a social system that supports the super-fitness of individuals, who may have a massively enhanced survival capacity over others, extending beyond sexual access to include numerous other benefits (see, for example, Burling, 2005). Super-individuals, as defined here, may exist in systems that have any form of material wealth, collective notions of status, or mythical notions, whereby the structuring of the social system could be geared towards defending the importance of that individual. The principle of the social lottery, I suggest, is that it is only the super-individuals who have enhanced survival fitness, and they may be a very small proportion of the population, whilst, everyone else actually has diminished fitness. The cultural dynamics of the population create what is essentially a lottery: arbitrary members of the group get super-fit through the rewards of others, but importantly, not because they exhibit innate selfish behaviour (if so, the effect of this process would rapidly diminish). This may have some net survival advantage for the species. In a sense it is a very biased form of cohesion.

This notion of a social lottery involves the implicit exploitation of individuals who are not super-fit, by those who are, and if the super-fit are the ancestors of future generations, how is this exploitation maintained? Should the principles of the selfish-gene undermine the existence of such a system, with individuals developing resistance against social exploitation? In sophisticated modern societies, super-fit individuals establish ways of passing more than just their genes and their capacity for social success on to their offspring; they also pass down material wealth and power. Such power-building does of course happen in human culture (e.g., Dumont, 1980),

but it is also clearly dynamic and unstable (e.g., Hobsbawm and Ranger, 1992). We could assume then that specific periods of familial dominance are short lived enough that these structures have a minimal impact on genetic selection itself, and that the population remains well-mixed over all, although there is unlikely to be evidence for this over suitable evolutionary time scales. Life in a dynamically structured social system would be a very specific niche to adapt to, but also a very variable one, and it may be that survival depends on an individual's social context and, importantly, their ability to adapt to that context; given enough potential for social mobility (over sub-genetic time scales), genes would have to become well adapted to survival at different levels of existence within such a system, not just one, and could therefore adaptivity feed back into the maintenance of the system. It is likely that in the earliest phase of the evolution of human musical behaviour there was no material wealth, no mythical thought (Donald, 1991), but just maybe the emergence of notions of status, based on the processes invoked by Boyd and Richerson (1985), and Henrich and Gil-White (2001). As various commentators have suggested (e.g., Lévi-Strauss, 1966; Donald, 1991), musical interaction resembles an embryonic form of mythical behaviour. Likewise, here I suggest that through the establishment of small, ephemeral social power centres musical interaction could have set the context for the emergence of mythical behaviour, defined as a form of collectively generated reality around which social organisation takes place.

Social lotteries, or social power structures, may have been social organisations that had functional benefits for our ancestors, but that our ancestors also successfully maintained through their own behaviour (c.f., Dunbar, 2004; Burling, 2005). I mention them here because they appear to be an implicit outcome of the competitive style learning theory and provide adaptive functions that count as bonus points for the theory. However, there is no clear way in which they directly contribute to the evolutionary process itself, thus I would argue that the social lottery effect is a fortunate consequence of competitive style learning, which could explain why we are not extinct, but not how we came to be the way we are. Put this way, it also provides another view on the notion of cohesion, which could be described in the same way.

3.6.2 Identity

In modern human behaviour we typically think of the proliferation and establishment of stylistic traits throughout a group in terms of identity (Anderson, 1983; Hobsbawm and Ranger, 1992). That is, certain behaviours are not even mistaken to be actually useful with respect to an external environment (Cannadine, 1992), they merely communicate the identity of an individual, which may be more or less similar to other individuals and thus more or less associated with a group. I discussed the possibility that stylistic group identity would be an emergent consequence of a stylistic-bonding system, and argued that this had the additional property that it allowed for a continuous hierarchical evaluation of group identity. Kinship, assuming we can recognise it, is a more literal marker for group identity, but it cannot be adapted according to new contexts, although kinship can be perceived as being more or less flexible in different cultures (Barnard, 1994). Stylistic identity marking would allow family groups and other groups to fragment, for individuals to become assimilated into existing groups, and for individuals to use a measure of cultural proximity in evaluating situations beyond their regular local environment (although this is a principle that can be as easily exploited as it can be trusted) (Gellner, 1983). If the time cost of stylistic learning can be trusted as an honest signal (Hagen and Bryant, 2003), then arguably it can be relied upon to indicate a great deal about an individual's life history. An individual's life history is likely to give some indication of allegiance, although of this we can be less certain. They may, after all, be fleeing a dispute from life-long allies, intent on revenge. Nevertheless, a possible rule to emerge, simplistic though it is, would be to trust those more similar to you, and to distrust those less similar. This is a rule that could potentially reinforce itself once reasonably prevalent, following the reasoning that, if there's a good chance that other individuals around you follow that rule, then it would be sensible to follow the rule yourself. After all, there's no point in trying to make friends with people who think you are their enemy. Thus, the emergence of identity actually enforces upon its subjects a new notion of the group, which is hierarchical and continuous (Gellner, 1983), which drives those who are similar closer together, whilst pushing those who are different further apart (c.f., Shelling, 1978), and which is at the same time flexible, allowing creative individuals to devise new identities as part of everyday social competitive dynamics (Geertz, 1973). This argument has already been

developed by Nettle (1999) as applied to dialects in spoken language, which could provide the basis for the emergence of separate languages across human cultures.

As I have already argued, the most salient aspects of behaviour, the most individualisable aspects (capable of variation and style changing), the most copyable behaviours, and arguably the most arbitrary (the most explicitly symbolic) are strong contenders to become the currency of identity, if identity is based on the process of copying style. The learnable behaviour as identity marker is an aspect of modern human behaviour (see, for example Durkheim, 1972; Anderson, 1983). It need not be dependent on language, but it is highly cultural, and implies the use of symbols (see, for example Feld, 1982).

Although I have suggested that identity marking emerges from stylistic bonding, it can also be argued to have emerged from competitive stylistic learning, if, again, the reliability of the time cost of stylistic learning is assumed. If identity marking could be shown to reinforce cohesion, it could arguably be treated as an element that ties stylistic bonding and competitive style learning closer together. Identity marking is therefore another aspect of human behaviour that I argue provides additional support to the plausibility of a competitive style learning theory. Unlike the notion of social lotteries, identity marking also seems to suggest local evolutionary mechanisms that may contribute to the process. However, I defer engaging with the complexity of this very rich area of theory without looking first at simpler processes of competitive style learning.

3.6.3 Creativity

I have argued that a changing-style system is important in sustaining a competitive style learning process and proposed a number of ways in which the competitive style learning process may maintain a changing-style system. Unless driven by noise, a changing style system must come about because individuals search out new styles, in other words, are creative. Creativity tends to refer to a powerful human cognitive capacity. In this context, however, it needn't be a powerful form of intelligence as long as it is a cognitive capacity for driving cultural change. This kind of creativity is commonly understood in terms of a desire for novelty or surprise (for example in Boden, 1990; Wiggins, 2003; McCormack, 2007). Typically, however, the desire for novelty is counterbalanced by an equal desire for recognisability, and the result is a preference for a

modest degree of novelty (as discussed in Saunders, 2001; Saunders and Gero, 2001a,b). In humans, as well as other animals, it has been observed that visual familiarity plays an important part in the process of incest inhibitions. Thus, childhood friends are less likely to become sexual partners even if they are not closely related (see Lieberman, 2006). The reduction of arousal from familiar things is associated with a notion of boredom, and the opposite sometimes with a notion of curiosity. Thus, for the present purposes, curious individuals who create cultural artefacts can be described as innately driven to modify their own behaviour without external influence.

A minimal notion of creativity such as this provides one mechanism by which a changing-style system continues to change, and as discussed above, we can consider how this might contribute to successful behaviour in individuals, and therefore reinforce itself.

One approach to this that I considered in section 3.3.1 was that, assuming that a changing-style system exists to begin with, it is reasonable to infer that intelligent individuals will strive to innovate new styles, slightly different from those around them. This can be seen as the ultimate development of a Boyd-and-Richerson (1985) style social learning process: Boyd and Richerson (1985) view social learning as an adaptive trait built upon the capacity for direct learning from the environment. This original learning from the environment can be seen as being creative in a stronger sense of the word than I use above: novel *and* valuable behaviours are innovated. But, following from the autonomy of a culturally-generated environment, I argue, the individual innovation of behaviours can arise not only because they have a direct payoff from interaction with the environment, but because they have a strong social payoff.

3.6.4 The relationship between vertical and horizontal transmission

An important distinction that is widely discussed in gene-cultural coevolutionary theory (Durham, 1991; Laland et al., 1999; Odling-Smee et al., 2003) but that has been glossed over in this chapter is that between vertical and horizontal (and also oblique) cultural transmission. Vertical transmission passes from parents to offspring, thus following exactly the same path as genetic transmission. Any kind of social learning other than vertical transmission may traverse lines of genetic heredity (Durham, 1991). Horizontal transmission describes transmission between

individuals who may not be directly related, and as the name implies, of similar ages. Oblique transmission describes transmission from one generation to the next between unrelated individuals. Thus, horizontal and oblique transmission are the processes that Boyd and Richerson (1985) are primarily concerned with, whereas niche constructionism emphasises the role of vertical transmission (Laland et al., 1999), in which parents are able to construct the niches, including cultural niches, that their offspring inherit. This distinction underlines some important notions in gene-cultural coevolutionary theory.

From the point of view of the theories discussed in this chapter, the vertical-horizontal distinction is significant in two ways. Firstly, we must consider how vertical transmission may interact with a Boyd and Richerson (1985) style process based primarily on horizontal transmission. It may be reasonable to assume that the notion of pay-to-learn suggested by Henrich and Gil-White (2001) is not strong between parents and their offspring, who automatically have access to their parents as models. Offspring would therefore be capable of inheriting their parent's style through social learning with greater ease, due to this increased access. If individuals are successful because of their style, then not only are they capable of having more and better funded-for offspring, but, under such conditions, these offspring would also have ready access to successful styles, a solid basis for the emergence of a strong vertical transmission of style. In addition, I propose, having many children with similar style implicitly increases the attention towards oneself. This can be seen as flooding the cultural market, so to speak, with a specific family style, and could equally be a hindrance for the offspring of individuals with less success. Secondly, since vertical transmission effectively tracks genetic transmission (*i.e.*, according to its definition, it passes from parents to their offspring) it is capable of becoming involved in processes of kin selection, and kin competition. Kin groups are likely to be more cohesive for reasons of kin selection (Dawkins, 1976; Maynard Smith and Szathmáry, 1995). The vertical transmission of style therefore points to a certain potential for musical style to convey cohesion, and yet, as discussed, this is a cohesive device that has flexibility, since culturally learnt behaviour remains adaptive during an individual's lifetime, whilst behavioural adaptations cannot pass through genetic inheritance (Dawkins, 1976). Musical and cultural behaviour is not bound to vertical transmission, and individuals can manipulate it creatively during their lifetimes, to

their selective advantage.

These issues suggest again how closely competitiveness and cohesionism could potentially interact. For example, the suggestion of flooding the cultural market through vertical transmission to one's offspring can be reinterpreted as the establishment of a large and cohesive kin group within a wider cultural system. This returns to the notion that, even without cohesion as a starting point, competitive style learning leads to the situation in which a stylistic proximity represents a necessary social bond.

3.6.5 Sex

The theories discussed here also do not take into account various aspects of sexual interaction. Sexual selection hypotheses (for example Miller, 2000a) are a competitivists approach to the evolution of human musical behaviour which share a strong overlap with the competitive style learning theory described here. I will discuss this relationship specifically in Section 3.7.1 below. However, sexual phenomena also include sexual aspects of social organisation (*e.g.* Merker, 2001), which provide a different focus from sexual selection theories. Menstrual synchrony theories such as those of Turke (Turke, 1988) and Knight (Knight, 1991, 2002) constitute one example. Knight (1991) develops an extensive theory of female cohesive behaviour that ties in with earlier concepts developed by the anthropologist Lévi-Strauss (1971) about the cultural centrality of the process of exchange of women as wives between groups of men. Clearly, a social organisation that is this rich confounds the simplified notions of stylistic interaction discussed here, which are abstracted away from sexual dynamics. The richness of primate sexual relations in general are, similarly, at the very least, an obstruction to the clarity of stylistic interaction processes. But, since there is no immediate reconciliation between these areas of focus I argue that it is sensible to carry on in one area until a point is reached where they either conflict or can be reconciled.

3.6.6 Group singing and mother-infant interactions

Two factors that have reinforced the strong emphasis on a cohesionist approach to the evolution of human musical behaviour are group singing and mother-infant interactions (see Dissanayake,

2000a; Cross, 2003a). Singing and playing music in groups is one of the main contexts for musical behaviour, and is widely regarded as an important aspect of musical behaviour (Cross, 2003a). Therefore it is reasonable to hypothesise that this phenomenon constitutes part of the functional utility of music. However, this is not a logical implication. For example, as I have discussed, from a stylistic learning point of view it is salient, individualisable and learnable behaviours that are likely to become culturally ingrained. We might add to this that, as a direct consequence of these properties, such learnt behaviours should also be able to be performed simultaneously, confirming the extent to which they were similar, and possibly reinforcing the learning process themselves. This view of group performance has been little considered, although it is suggested by Hagen and Bryant (2003) and Miller (2000a). In theory, if we consider distinguishing the possible functional or consequential origins of music from the possible function or consequential origins of performing music in groups, we see that it is possible to attribute cohesive functional explanations to group performance which do not explain musical performance in general. I emphasise that this is only a devil's advocate position raised to question the assumption that the groupishness of musical performance is essential to a theory of the evolution of human musical behaviour (Cross, 2003a; Mithen, 2005; Brown, 2007). In fact, the possibility that group performance has a slightly independent set of origins and consequences from stylistic bonding or competitive style learning is intriguing. It is possible that the act of performing in groups represents the intersection of a wide range of learning, competing, collective signalling and cohesive processes, a social melting pot of myriad overlapping individual goals (as suggested by Richman, 2001). On the other hand, I argue that, as with sexual behaviour, it is also simply true that it is hard to reconcile stylistic learning theories with group singing.

Mother-infant interactions provide another perspective on the cohesive nature of musical behaviour (Trevarthen, 1986, 1993; Dissanayake, 2000a,b). The highly rhythmic nature of these interactions has been interpreted as mediating a sense of joint-attention between mother and infant, providing a basis for referentiality, and therefore setting the scene for proper cultural learning, ultimately through language (Trevarthen, 1993). This view provides a more subtle and refined development of the social bonding hypothesis of the evolution of human musical behaviour since, although it refers to interactions that arguably share a great deal with music,

these are mother-infant interactions and not general, kin-independent interactions. The role of rhythmic behaviour in mother-infant interactions may also explain certain musical capacities (Dissanayake, 2000a), including the cohesive power of musical behaviour. To the extent that mother-infant interactions are geared towards the development of sociality (and therefore social learning skills) in infants, they provide a suitable context for biases towards salient, individualisable and learnable behaviours to arise in a vertical transmission context (*c.f.*, similar modelling work in language acquisition by Smith, 2002). Both bonding and stylistic learning, we could say, begin at home, but can flexibly extend outwards over time. Since stylistic learning theories focus on competitive and cohesive learning in a predominantly horizontal transmission context (following Boyd and Richerson, 1985), they do not directly address mother-infant interaction, but do not stand in opposition to the evidence that mother-infant interaction is musical (Trevvarthen, 1993), and establishes basic sociality in the infant, as well as a strong bond with the mother.

3.7 Theoretical significance of social-learning models

In this chapter I have presented an argument that if a purely competitive style learning theory is not fruitful, then at least competitive style-learning is a necessary component of cohesionist theories and needs to be considered in depth. I believe that a combined cohesionist/competitivist theory is the strongest contender for an evolutionary theory of human musical behaviour, but for this reason the development of a theory of cultural evolution driven entirely by the consequences of social learning interaction is by far the most compelling direction in which to develop this view. As I suggested at the beginning of this chapter, human evolution is as likely to shine light on new evolutionary processes as it is to fit neatly into existing ones, and runaway cultural evolution driving further biological evolution is indeed novel. It is particularly important to emphasise the difference between this process and the process of sexual selection, especially because they are so similar and lead to very similar observations. On the other hand, it is also fitting to see how these evolutionary processes fit in with broader issues in evolutionary theory. In particular, can these theories contribute to Maynard Smith and Szathmáry's (1995) notion that the emergence of human cultural behaviour is a major transition in evolution?

3.7.1 The overlap between sexual selection and social selection models

Miller (2000a) argues for a sexual selection approach to the evolution of human musical behaviour on the basis that musical behaviour has the appearance of a complex adaptation, but is in fact useless. He draws on evidence to suggest that music is predominantly performed by men in order to attract women. This particular evidence has been met with much suspicion (Bispham, 2006b) and throws up a basic question about how observations of modern human behaviour in the field can inform theories of the evolution of human musical behaviour (Miller, 2000a; Fitch, 2006; Dean and Bailes, 2006). Is it possible simply to anecdotally support arguments about the evolution of human musical behaviour with selected aspects of modern human behaviour? Miller's evidence is easily balanced with anecdotal exceptions that do not seem to prove the rule, such as largely male followings of male bands. Nevertheless, evidence is problematic in either direction on these issues (Miller (2000a) calls for more quantitative social research himself, also see Todd (2000)), and in truth the energy of Miller's argument derives largely from the fact of music's appearance of complex adaptation combined with its apparent uselessness. In his general approach to sexual selection theories of the evolution of human behaviour, Miller emphasises, following Darwin, that evolutionary processes are strictly divided between two types: natural selection and sexual selection. Miller's analysis begins therefore with the question of whether, assuming that music is an adaptation, it is of the natural or the sexual variety. For Miller, music, along with numerous other aesthetic capacities, is subsequently categorised as a sexually selected trait, a product of runaway selection and an indicator of cognitive fitness. Humans, he argues, inhabit a cognitive niche, and sexual selection has largely fixated on the selection of cognitive traits over our recent evolutionary history. The sexual selection argument is powerful, seductively plausible, and may very well be entirely correct.

But there is a strong overlap between Miller's (2000a) motivations leading to a sexual selection argument and the nature of stylistic learning theories. Competitive stylistic learning also provides a scenario in which a complex behaviour can emerge without function (following Boyd and Richerson, 1985). It can also be argued to be a powerful evolutionary driving force due to its inherent feedback dynamics (Laland et al., 1999). It shares with sexual selection the characteristic of being a runaway process emerging from a powerful set of pre-existing adaptations

(Ridley, 1994). In this way it has the potential to make similar predictions about the adaptive appearance of music. On this basis I would argue that Miller's (2000b) initial division of evolution theory into natural and sexual selection is mistaken, at least to the extent that sexual selection is unique. What is more significant is how unique human social learning is. A reasonable rebuttal would be that competitive style learning *is* sexual selection; styles are sexually attractive, and absolutely nothing is more critical than reproductive success (Todd, personal communication). However, competitive style learning proposes an indirect relationship between style and sexual reproduction in which a population of individuals determine social success, which then determines reproductive success (similar to the argument of Knight, 1991). If sexual selection is powerful because the winner takes all, then competitive style learning is equally powerful, perhaps more so, because the winner takes all socially. If our ancestral societies were polygamous then sexual selection makes sense, if they were social learners then competitive style learning makes sense, and if they were both then both effects are possible, and may even have had a joint effect (Knight, 1991).

Sexual selection theory constrained by the handicap principle also proposes a function associated with complex adaptations: music exists to indicate cognitive fitness (Miller, 2000a). Competitive style learning as proposed in this chapter, on the other hand, promotes the emergence of apparent function around the sustained existence of a certain behaviour; functionless music becomes useful exactly because it is a behaviour rewarded with success. And yet this competitive social learning view proposes the emergence of strong music learning abilities without reference to the development of other more obviously useful learning abilities. One could argue that natural selection surely drives the emergence of the useful behaviours first: the useless behaviour is either seen as a consequence of the useful behaviour, or as serving it as an indicator trait (Miller, 2000a). Competitive social learning promotes the view that, although originally a consequence of existing learning abilities, the learning of useless behaviours follows its own evolutionary path, and it is only fortunate that they resulted in a species that survived (*c.f.* Blackmore, 1999). Prior to the acceptance of the handicap principle (Zahavi, 1975), sexual selection followed the logic that completely arbitrary traits could evolve as a result of their attractiveness to the opposite sex (this continues to be validated by models such as Todd, 2000). The handicap

principle placed a sensible restraint on that notion: attractive traits will tend to be indicators of useful traits. But runaway processes are still runaway processes. Indicator traits are big and salient, rather than subtle, for that reason (Zahavi, 1975). So could it actually be fluke that the runaway process of competitive style learning led to a successful species? In the broadest possible sense yes, but then learning and creativity skills *are* useful for a species, and these two aspects of behaviour are presumably selected for during a competitive style learning process, amongst many other genuinely useless behaviours. However, this does not mean that competitive style learning evolved *because* humans needed learning and creativity skills in order to survive with respect to their non-social environment. Useless runaway selection can happen to drive behaviour that is ultimately useful.

Along similar lines, one must ask when a trait is an indicator, or simply a more literally attractive quality. The runaway nature of sexual selection is a result of the fact that genuine genetic goodness is largely hidden, and honest indicators need to emerge (Zahavi, 1975). However, Zahavi (1975) also gives examples of human wealth indicators such as large houses. Clearly a large house does communicate a measure of wealth to others, including potential sexual partners. However, a literal view of large houses is that since we like large houses, anyone who could afford a large house has one, and potential mates, and other social individuals are attracted literally to the large house itself (“I want to live in a large house”), and not because of the thing it supposedly indicates. Ultimately these literal and communicative perspectives flow seamlessly into one another. In human behaviour, cultural contexts can turn things that may be useless or only indicative of something else into things that are directly useful. According to competitive social learning, sexual attraction towards musical behaviour can be interpreted as existing because the perceiver understands that the musical behaviour is directly responsible for social success, it is neither an indicator of the potential for social success, nor for cognitive abilities, and needn’t actually *be* attractive.

3.7.2 Relationship with general trends in evolutionary theory

The competitive style learning theory suggests that from a certain point in humankind’s evolution, where social learning had developed to a certain degree of sophistication and other factors

were in place, a runaway process occurred, in other words a process that took over from gradual adaptation to a changing environment. This is a process commonly cited by Gould in his theory of punctuated equilibrium (Eldridge and Gould, 1972; Gould, 2002). One change invokes a cascade of changes that ultimately settle on a new stable form. Competitive style learning can be seen as a punctuated equilibrium theory in which an adaptation towards social learning inevitably triggered the competitive-style-learning cascade. From the point of view of a cascading process, it does not matter whether individuals are on average fitter, less fit, much the same, or extinct. This is no more than an outcome of the process, and all of these outcomes are possible. The interesting thing about our evolutionary cascade is that it potentially led to a new organisation of individuals, one that was not just new in the lineage of this species, but completely new to nature. Since positive feedback is the driving force for such runaway processes, they are not necessarily adaptive and as such are susceptible to negative outcomes for the species concerned. Viewed in this light, these dramatic episodes can be seen as similar to the minute process of genetic mutation associated with every act of reproduction, the probabilistic nature of which is discussed by Dawkins (*e.g.*, Dawkins, 1986), but more risky, and more rarely, resulting in successful behaviour. Feedback processes are high-stakes evolutionary changes. These changes are so dramatic that they barely resemble adaptation, because, as Gould argues, they simply do not have the property of gradual change (Gould, 1977). Instead, by chance the evolving population survives its violent cascading change, perhaps massively diminished in size, even conceivably less fit as far as each individual is concerned, but its new mode of organisation is, nevertheless fundamentally capable of survival and future adaptation.

Similarly, from Maynard Smith and Szathmáry's (1995) point of view, human cultural behaviour is a major transition in evolution associated with the development of cultural behaviour, involving a strong degree of cooperation. Boyd and Richerson (1985) frame this transition in terms of the homogenising effect of social learning, which reduces selection between individuals and enhances it between groups, thereby creating a mechanism by which group selection becomes feasible. Stylistic learning theory provides a mechanism by which individual competition drives the increasing emergence of contexts in which cohesion is increasingly feasible, offering a link between selfish behaviour and cohesion in the evolution of human musical be-

behaviour which does not rest on problems of group selection (Knight, 2002).

Maturana and Varela (1987) introduce a number of new concepts with which to view evolutionary theory. Through the notion of structural coupling, the evolution of stylistic learning can be viewed as a process in which necessary repeated interactions (Boyd and Richerson's (1985) genuine social learning for adaptive reasons) provide a reliable context for the set of participants to begin to be perceived as a unity. Socially learning individuals are loosely coupled, and their collective behaviour takes on a reliable structure that is not a function of the system but a consequence of other aspects of it. The existence and permanence of this structure provides scope for new adaptive change, but importantly, the existence of any element in that structure is bound up with the existence of the structure itself, and its other components. Therefore, the survival of those components is each intimately coupled to the survival of the structure and associated elements. But in this case the feedback of stylistic learning by its nature appears to reinforce the structure: that is, the loose coupling of individuals through social learning. Again, this points to the role of a feedback process that wasn't driven to work under adaptive pressures but was predestined to work by its particular nature.

If the mechanisms for such processes can be understood, then we can consider them in terms of significant evolutionary jumps, like the ladders (rather than the snakes) in the game of snakes and ladders. Then we can view human evolution as a process whereby our ancestors encountered culture, which kick-started a runaway evolutionary process. The result could have been collapse, but it wasn't, by chance. In the language of evolutionary computation, a fitness landscape is the surface produced by calculating the fitness of individuals in a population as a function of their genetic structure (Mitchell, 1998), and increasing evolutionary fitness means climbing the hills of this landscape, potentially getting stuck at its local peaks and not necessarily finding the highest peak. Neutral networks (Schuster et al., 1994) are defined as evolutionary pathways that populations can travel along without gaining or losing fitness: essentially the contours of the fitness landscapes. Large neutral networks imply a greater potential for an evolving species to jump towards new, higher fitness states in a greater diversity of ways, effectively increasing the search power of evolution by natural selection. One reason why this might be possible is that high-dimensional evolutionary spaces may be more likely to have more neutral networks (Har-

vey and Thompson, 1996). However, since in real evolutionary systems fitness landscapes do not stay still but are constantly changing (Odling-Smee et al., 2003), the analysis of neutral networks based on static evolutionary environments does not necessarily extend towards these real systems. When a population itself generates the conditions for its own members' evolutionary fitness, it no longer makes sense to consider the population as travelling together on a fitness landscape. Thus, whilst the notion of following evolutionary change towards a point of runaway evolutionary departure provides some similarity between the theory of neutral networks and the notion of major transitions, it is inappropriate to think of the process suggested here as a process of random drift around neutral networks by populations of individuals, interspersed by sudden leaps in fitness. A runaway process implies constant evolutionary change under selection pressures, even if this ultimately results in lower chances of survival for individuals within the population. However, the change caused, as with the change caused in neutral networks, can lead to new avenues for evolution to follow.

3.7.3 Exploring the competitive style learning hypothesis

Style learning redirects the focus of the study of the evolution of human musical behaviour in modern humans towards a balance between psychology and the dynamics of style. This new focus is daunting, to say the least: every technological innovation in the last several thousand years has had some effect on human social organisation, and therefore potentially on style dynamics (*c.f.*, Latour, 1993). Style dynamics today don't guarantee anything about style dynamics in the past, and there would certainly be no potential for direct inferences from modern behaviour to theories of the evolution of human musical behaviour. This is true for small-scale hunter-gatherer societies as much as for modern post-industrial societies, whose daily life still involves language, which would also have disturbed the original nature of style dynamics (Ingold, 1994). However, psychological research into the evolution of human musical behaviour is also faced with incredible challenges associated with understanding the current functioning of musical behaviour in the brain, and, from that, attempting to untangle its complex history (Justus and Hutsler, 2005; McDermott and Hauser, 2005).

The approach followed in this thesis is to consider style learning processes in their most

abstract form and thus to study computer simulation models of gene-culture coevolutionary processes involving style learning. This method and associated hypotheses will be outlined in Chapters 4 and 5. But, following the reasoning behind this abstracting approach, it is also worth rethinking modern humans not as archaeological objects that we can use to determine facts about our ancestors' behaviour, but as a lab for studying style-learning behaviour as an abstract phenomenon. In this thesis I only attempt to study the artificial component of this proposed body of research, bearing in mind that computer simulation models can easily be designed to give the appearance of a certain effect, and no simulation model is meaningful outside of a careful description of what assumptions and design considerations have led to the results.

This discussion helps to ground the remaining aim of this thesis. Following the proposal such a process of runaway selection, based on what is effectively an arbitrary consequential behaviour, is possible, what mechanisms could make it work? What kinds of factors are likely to encourage or disrupt such a process? Developing a more detailed theory of how such a process could have worked provides a genuine contribution to the study of the evolution of human musical behaviour through altering the balance of theoretical plausibility between competing theories, in particular lending the power of the sexual selection hypothesis to the convincing role of music in social behaviour. In Chapter 4, I consider how computer simulation modelling can be applied to the problem and, in Chapter 5, I describe a specific set of hypothetical processes, stemming from the discussion in this chapter, that I aim to explore through computer simulation. I propose a number of hypotheses that can be tested using computer simulation models, and describe the design of a model for investigating these processes. Chapter 6 describes specific studies of the modelling framework described in Chapter 5, and Chapter 7 discusses the results and conclusions of those tests with respect to this chapter.

Chapter 4

A Simulation Modelling Approach to the Study of the Evolution of Human Musical Behaviour

4.1 Theoretical perspectives

In the previous chapters, I have placed particular emphasis on a view of the evolution of music that is not well represented in the literature, but that is potentially compatible with theories of human and cultural evolution in other fields. This theory is distinct from consequentialist theories that propose that music is a cultural development based on existing perceptual capacities, cohesionist theories that propose that musical interaction is beneficial to individuals because of the benefits of group behaviour, and cognitivist theories that propose that music-like thought represents a powerful general cognitive mechanism. The central hypothesis of this theory is that it is possible for maladaptive cultural interactions based on style learning to drive the emergence of a self-sustaining cultural domain, along with the evolution of aspects of perception exhibiting the essential features of music. Such a hypothesis is not easily verified or refuted through observation of archaeological, psychological, anthropological or ethological data, although I have discussed approaches to suggestive evidence for and against the theory in the preceding sections. Following the paradigm of artificial life, Todd (2000) argues that computer simulations of evolutionary processes can inform theories of the evolution of human musical behaviour in novel ways, alongside its successful application in other areas such as the evolution of language

(for example Kirby and Hurford, 1997; Kirby, 1998; Smith, 2002). The approach taken in this thesis is to extend an understanding of the ideas presented in the previous section, following this methodology.

4.1.1 Models of human behaviour

In all simulation models, simplicity is necessary for practical reasons, including our ability to program models and the computational time and memory required to run them. Simplification is also desirable from a theoretical point of view, assuming that the goal of such models is to develop our understanding of significant abstract processes that capture, epitomise or characterise the behaviour of a system. We understand simpler things better than complex things. In that sense, to understand is to reduce complexity whilst maintaining accuracy. In the human sciences there is a considerable challenge in appropriately simplifying the behaviours of individual human beings, and thus making valid use of simulation models.

Certain domains of human behaviour have been successfully incorporated into simulation models. Shelling (1978) modelled very simple behavioural biases such as simple frequency-dependent judgements to demonstrate self-organisation in populations. Economic models using agents with very simple rules can likewise demonstrate the properties of complex economic systems. Such research supports the view that although human behaviour as a whole is very complex, certain aspects of human behaviour are simple and predictable, and, as long as they operate in a relatively well-bounded domain, can be successfully modelled.

Gilbert (1995) argues that models of human behaviour should differ from models of other complex systems in one crucial domain: that “people are routinely capable of detecting, reasoning about and acting on the macro-level properties (the emergent features) of the societies of which they form part” (Gilbert, 1995, p. 151). This particular human ability is of course tied in with our approach to living in groups. The relationship between dynamical systems in general and specific dynamical systems in which the components have a reflexive capacity could certainly tell us something about the unique qualities of human social behaviour, but at the same time it is hard to see how this point of view can be extended towards a rigorous categorisation. It is probably untrue that no other animal is capable of intelligently responding to the macroscopic

behaviour of its conspecifics, in which case it is hard to see exactly where one would draw the line. Possibly, in the words of the anthropologist Benedict Anderson (1983), it is more precisely our ability to *imagine communities* that is significant. When we *imagine*, we mentally simulate phenomena that may not exist. A *community* is one such phenomenon, being more than just a group of people, and perhaps something more like a domain of cultural activity. Also, insofar as Gilbert's (1995) shares common ground with theories of shared intentionality (Tomasello, 1999), it may be possible to draw a clear line between human and animal behaviour in terms of the sharing of goals. Other work in modelling human behaviour does not attempt to incorporate this notion of human collective imagination, but instead to construct emergent patterns of human behaviour from simple local interactions (*e.g.*, Axelrod, 1984; Shelling, 1978). The models considered in this thesis will follow from this tradition, on the assumption that musical interaction does not require consideration of social behaviour on a macroscopic level (although some aspects of the proposed models could be argued to). The significance of Gilbert's view to other areas of human behaviour is not questioned, however.

I argue that human cultural behaviour in contexts such as musical interaction at present represents an extreme of complexity, suggesting that a quantitative style of simulation modelling is currently incapable of addressing such problems. The data required to develop accurate simulation models cannot sensibly be gathered, and it is nevertheless unclear where the boundaries lie around musical behaviour. In particular, musical behaviour involves a temporal granularity on the order of milliseconds, whilst cultural processes work over months, years and centuries. These factors are further accentuated when considering the evolutionary process leading to human musical behaviour, which occur over thousands of years, rather than simply understanding human musical behaviour in short-term modern contexts. Although, owing to this complexity, it could be reasonably argued that there is nothing to be gained from attempting to build and study simulation models pertaining to the evolution of human musical behaviour, a qualitative approach which is not primarily concerned with fitting empirical data, still offers the opportunity to begin charting out general categories of evolutionary system in a music-evolutionary domain. A qualitative approach to the evolution of human musical behaviour suggests different aims to a quantitative approach and works almost entirely in the domain of theory development. The

agents of an artificial-life model of human musical behaviour don't even need to be understood as simulations of humans, the system as a whole needs merely to demonstrate a hypothetical evolutionary process, and the abstraction of this process supports its generality. Artificial life has sometimes been referred to as the study of *life as it could be* (Langton, 1988). This is a valuable point of view in justifying a qualitative approach to the evolution of human musical behaviour that is not too strict about modelling accuracy. However, as I will also argue in this Chapter, we can still impose rigorous demands on what kinds of simulations produce relevant results for the evolution of human musical behaviour.

4.2 Using simulation models to explore novel evolutionary processes

What information can computer simulation models contribute to theories about real world processes such as evolution? A model is a formalisation of aspects of a real process that aims to precisely and usefully describe that process. Models are mental and social constructs that are shaped by the direct observation of real systems, sets of concepts that already exist within a sociocultural group and individual contemplation (Kuhn, 1996). Precision of a model can be determined by how accurately it predicts certain behaviours, and at what level of detail. Arguably, a model's usefulness is primarily determined by technological achievements in science: the extent to which the real world can be manipulated towards certain goals.

Kuhn (1996) describes how science develops by working within *paradigms*; gathering experimental data according to a given paradigm during which evidence against that paradigm accumulates, and then breaking out into a new paradigm which poses new methods of observing and gathering data about the world. According to this view, scientists are necessarily bound to work in the idiom of previous closely related work and adopt a great deal of cultural information in order to operate in increasingly specialised domains. In all academic disciplines, dogmatic sets of concepts are a necessary consequence of the collective formulation and communication of models. Wherever sets of concepts are nontrivial, individual arguments about the relationships between them can evoke change in dogma without necessarily being evoked by a change in evidence from experiments. In short, academic thought can be as much internally generated through social-intellectual means as driven by evidence.

Computational models can provide information about real processes that was not apparent before a certain process of simulation was performed. In this sense they constitute observations of the real world, albeit a world that is not material but logical. On the other hand, they result from the formulation of a set of concepts; their consequences stem from the arrangement of those concepts in the mind of the researcher: they are intellectual products. In each case, theoretical sciences can proceed in an experimental manner by investigating the behaviour of interacting sets of concepts. These arrangements of concepts can be refined iteratively with respect to the same set of observations of the real world; most of the information that informs changes to the model comes from studying the model itself. Computer models take this process to a new extreme, facilitating approaches such as artificial life, in which lifelike behaviour (including social behaviour) is simulated, and these simulations inform the evolution of concepts relating to life. Following Kuhn's analysis of thought experiments, Di Paolo et al. (2000) define artificial life models as *opaque thought experiments* in that they explore the natural logical consequences of processes that cannot be thought through. Di Paolo et al. propose that this opacity requires that artificial life researches engage in an interpretative relationship with the qualitative results of a simulation model and iteratively approach new models with their findings.

Popper (1972) proposes as a logical necessity that scientific theories can only be refuted. No amount of supporting experimental evidence constitutes proof of a theory, and yet a single repeatable piece of counter-evidence can require that the theory is developed or replaced. Scientific hypotheses, therefore, should consist only of refutable statements, and only by resisting all attempts to refute it should a theory maintain authority. As an experimental form, then, it is important to ask what kinds of statements computer simulation models can refute. The output of a computer simulation model is necessarily dynamic, or temporal, and thus represents a *mode of behaviour*, which is real; it is repeatable and independent of the observer. Computational processes can therefore clearly be an object of study in their own right and each computer simulation model thus defines a real *possible process*. Of all possible processes, many will be intuitively obvious, some discernible from experimentation, or from mathematical or logical models, and others will be neither intuitive nor predictable in any other way than to run the simulation. It is in this latter case that computer simulation models fit best with Popper's requirement for refuta-

tion: they are able to pose new possible processes and refute the implicit assumption in existing theories that a smaller number of possible processes were available to explain the behaviour of a system. Following Kuhn, the structure of reasoning within a theoretical frame can constrain thinking to the point of asking *how could it possibly be any other way?* In indicating novel possible processes computer simulation models are capable of promoting these other ways of thinking about a problem. A more extensive discussion of the problems of both Popper and Kuhn's approaches is given by Lakatos (1970).

In a classic example, Hinton and Nowlan (1987) model the Baldwin effect, which had been known for years but generally ignored in mainstream evolutionary theory. According to Baldwin, organisms that learn from their environments may repeatedly learn a certain specific behaviour from one generation to the next, thus provoking a selection pressure driving the emergence of the innateness of that behaviour in future generations. Their model, which shows the Baldwin effect to be theoretically viable, was so simple that with hindsight it is hard to believe that a computer was actually required to make the point. However, the historical fact that the Baldwin effect was of less importance before the model than after it speaks for the power of that hindsight: it is as much a struggle to think in new ways as it is to understand how old ways of thinking were so impermeable.

Despite the fact that, in cases such as Hinton and Nowlan's, a computer simulation model can shed light on the possibility of a process, the act of relating that process to real world systems is invariably an act of interpretation: the processes gleaned from computer simulation models can only inform the real world through interpretation and argument, and this mode of interaction will invoke new concepts with which to construct models. In a sense, models confuse matters by constructively adding processes to nature's repertoire, and introduce the task of understanding the relationship between simulation and real world process.

When theory discounts a possible process due to questions over its efficacy, evidence of that efficacy is all the more significant, as in this example. This, the proof of concept case, is the simplest case of how a simulation model can contribute to an understanding of real world processes. In more complex cases, the model throws up new suggestions for how to view real world processes. For example, Di Paolo (1997) investigates potential implementations of Maturana and

Varela's concepts using a model in which the relationship between organisations of individuals and their behaviour is mutual and coupled. Di Paolo considers how spatial social organisation influences the evolution of behaviour in a game that resembles the iterated prisoner's dilemma, containing the potential for the emergence of cooperation. The model consists of a uniform toroidal grid world, with food at each point on the grid, which is replenished at a fixed rate. Individuals have to engage in a prisoner's dilemma-like game with a random neighbour before having access to food. The food is of a certain type, with four possible types, and the agent has to perform one of four actions to obtain the food (the effective component) and one of four actions which is to be made visible to a potential collaborator (the external manifestation). A second agent, the potential collaborator, chosen randomly from the first agent's neighbourhood, then observes the external manifestation of the first agent's action and also has to perform one of the four food-obtaining actions based only on this observation. The amount of food gained by each individual is determined by these two actions. If they both perform the *correct* action for that food type, the total food is shared equally. If only one performs the correct action, a proportion c of the total food ($0.5 \leq c \leq 1$) is given to that individual, and the rest is left behind. If an individual's energy rises above a certain level they get to reproduce (sexual reproduction with a random neighbour). If it falls below a certain level, they die. These rules indicate that the best action for each individual is to perform the correct action whilst one's opponent performs an incorrect action, but unlike the prisoner's dilemma, the second best outcome is that both perform the correct action. In other words, there is no immediate dilemma. The long term evolutionary conflict, however, is that whilst it would be ideal to trick one's opponent by giving a false signal, the homogenising effect of sexual reproduction in a locale implies that all individuals would be equally well deceived.

Accordingly, this model results in a specific structure that demonstrates the evolutionary payoff of these conditions. The structure consists of clusters of individuals, which are increasingly densely populated towards their centres. Individuals on the inside therefore have many neighbours and not a lot of food within their neighbourhoods (food is taken at a greater rate). Thus although they are short of food, they get chosen to be second players more often, and so get more chances to acquire food (from further afield). Individuals on the periphery have more

food but not so many neighbours. Due to homogenisation caused by sexual reproduction, all individuals within a cluster necessarily have roughly the same behaviour. However, behaviour can differ from cluster to cluster. It was noted that cluster size increased with coordination success, but only historical contingency could explain why some clusters stabilised with particular sizes. Clusters can stabilise at a number of sizes. At each size a relationship will maintain between that size, and the coordination success, which is a result of a roughly homogenous genetically determined behaviour.

Following Maturana and Varela, Di Paolo's (1997) model scrutinises final-cause (adaptationist) views of evolution, in particular with respect to communication, complaining that communication is often defined using "the same terms which are used to *explain* it," which, he adds, "is not only confusing, but methodologically very questionable as descriptions of phenomena and descriptions of the generative mechanisms that give rise to such phenomena (explanations) belong necessarily to different domains" (Di Paolo, 1997, p. 2). This view is relevant to the theory of the evolution of music established in previous sections. Music can be seen in an evolutionary light without the need to provide it with a functional explanation. In Di Paolo's (1997) game, communication is first and foremost functional for the observer, and is directly beneficial to the first agent only when it is unsuccessful (*i.e.*, when deceit is successful). The model also implies the emergence of autopoietic *quasi stable* structures. These structures are defined by a spatial organisation and a process that maintains that organisation.

Di Paolo's (1997) model is set up to explicitly explore theories of the evolution of communication, but works best as an abstract exploration of concepts defined internally rather than as a simulation of a specific real world process. Whilst one can talk of the relationship between a structure and a pattern of behaviour in the case of the model and in the case of a particular real evolutionary context, the model does not stand for the real context. Rather it is a means of exploring and generating concepts.

One central theoretical problem that straddles biocultural coevolutionary theory, anthropological theory and this kind of complex dynamical systems research is a opposition between environmental determinism and historical contingency. Di Paolo's model demonstrates a certain form of historical contingency, and yet arguably most models of human social behaviour are

concerned with how contexts such as physical environments determine behaviour.

A good example of this debate and its potential complexity is given by the thesis at the heart of the popular science book *Guns, Germs and Steel* (Diamond, 1997). Diamond argues that the broad East-West spread of the Eurasian continent provided a large area of climatically similar land across which domesticated plants and animals could be diffused through trade and cultural exchange. This is in contrast to Africa and the Americas, where the North-South orientation of the continents meant that exchanged crops and animals entered different, less suitable climates, and other regions such as the Pacific islands, where geographical barriers limited trade. This is a sophisticated environmental determinism argument because it does not depend on microcosmic factors such as the ability for one farmer to successfully grow one crop, but on the potential for a large-scale propagation of objects and ideas.

This hypothesis is extremely compelling, but Diamond's analysis of the history of the world rests, necessarily, on a small number of test cases (on the continental scale, this number is less than ten). Even if the development of a solid base for computer simulation research into such areas is slow, cumbersome and far removed from the real complexity of world history, we can ultimately begin to look at how historical contingency stacks up against environmental determinism as general facets of human behaviour.

4.3 Multi-agent models of musical behaviour and evolution

Todd (2000) and Miranda and Todd (2003) have put forward key arguments for the value of simulation models in the study of the evolution of human musical behaviour. In light of the lack of evidence available regarding the evolution of human musical behaviour, Todd (2000) argues that simulation models can help reconstruct hypothetical scenarios of evolutionary processes. Although I concur with this view and hope to develop directly on this initial work, the approach followed so far by its main advocates, Todd and Miranda, has failed to be convincing to a wider audience, I suggest, because of the lack of any tight integration between theory development and model building in their work. Existing models have tended towards implementing potential processes without at first clearly specifying problems in current theories that need to be developed, and without the additional interpretation of the model's implications for theory, or any

iterative development of models based on these findings. However, I have proposed above that the problem of reducing human musical behaviour to a level of simplification appropriate for simulation modelling is a major initial hurdle, and it is necessary to appreciate this early work as highly exploratory development of a sensible, and ultimately convincing modelling strategy. In developing a powerful framework for computer simulation models of this new domain, it is necessary to freely explore simulation frameworks. These considerations point to a final more powerful argument for the value of computer simulation models in the study of the evolution of human musical behaviour: they force verbal theories into the formal requirements of programming code. By thinking about how a theory can be coded, one is already taking steps towards scrutinising it according to a modelling paradigm, even before running a simulation.

4.3.1 Werner and Todd's sexual selection model

Werner and Todd (1997) develop a sexual selection model of the evolution of human musical behaviour based on the sexual selection hypothesis of Miller (2000a), discussed in Section 2.2.2.1. The model is defined by the sexual selection hypothesis in its runaway form, in other words without the constraints of a handicap principle, with aspects of male and female behaviours co-evolving according to the attractiveness of male songs as evaluated by female song preferences. Genotypes store male songs, which are sequences of 32 notes, each of which can have one of 24 values, and female probability transition tables which define sets of expectations of transitions from one note to the next. Each female is given a random selection of males to choose from (the size of this group being a parameter of the model). They judge male songs according to their transition tables, and mate with their preferred male. Each female has exactly one mate per generation, whereas some males might be chosen by many females to mate with, and others fail to be chosen; a typical sexual selection scenario found in nature that leads to female choosiness. The offspring of this reproduction has its sex determined randomly (such that a 1:1 male-female ratio is maintained) and contains the genes of mother and father, recombined and mutated. Werner and Todd (1997) looked at three ways in which females could judge male songs: a local transition preference, in which each transition of the male song is scored according to the female's expectation of that transition; a global transition preference, in which a transition table is con-

structed from the male's song, and its similarity to the female transition table is measured; and a more complex surprise preference. According to the surprise preference rule, females consider each transition in the male's song in terms of how unexpected it is. They look up this transition in their transition table and subtract that value from the highest value in the transition table for the given start note. So if the female hears C-E, she looks up the transition C-E in her transition table. Assuming that G is the most likely note to follow C in her transition table, she measures the difference between the C-E expectation and the C-G expectation. The greater this difference, the greater the surprise, and the higher score she rewards.

Werner and Todd (1997) compared these three preference types, and also varied the choir size – the number of males each female was given to sample (*i.e.*, the strength of sexual selection versus random mating) – and whether or not female transition tables experienced mutation or not (*i.e.*, comparing coevolution of males and females to evolution of males under fixed conditions of selection). They looked at the effect of these variables on the amount of diversity over time and within the population. They discovered that in support of their expectations, surprise preferences and coevolution led to greater change over time; coevolution also led to greater diversity across the population at any given point in time. But increased choir size (equivalent to an increased level of sexual selection) led to a lower change over time and diversity across the population. However, surprise preference combined with a large choir size, whilst producing a largely converged population most of the time, could lead to the formation of distinct 'species' of song, albeit briefly. Thus increased sexual selection could be said to increase convergence, but allow the runaway divergence of groups. These results point to interesting properties of the coevolution of production and evaluation of musical patterns. Todd (2000) points out that learning during an individual's lifetime could also increase the rate of change within a population.

There are two important criticisms of Werner and Todd's (1997) approach to what Todd describes as "using evolutionary computer simulations to study the origins of music" (Todd, 2000, p. 385). Firstly, the authors are driven towards the assumption of working within a sexual selection framework, but without questioning this framework. Secondly, they treat diversity as such an important *outcome* of a model of the evolution of human musical behaviour. Since Miller's (2000a) sexual selection hypothesis is a valid competing theory of the evolution of

human musical behaviour, it is highly relevant to consider simulation models of this process. Yet the model offered by Werner and Todd is seemingly detached from a plausible scenario for the evolution of musical behaviour in humans (although it may apply well to birds). Synchronous diversity in humans is predominantly a product of human learning capacities: a phenomenon closely associated with a human capacity for culture. The musical preferences that we observe ethnographically are not the products of genetic variation but of cultural change. This is self-evident in the fact that musical styles change during individuals' lifetimes, as Werner and Todd (1997) point out. Even if they were the result of genetic variation, this could also be explained as the result of geographic separation (separation and sexual selection together make a particularly powerful model), and no case is made for identifying diversity as a significant emergent property of the evolution of human musical behaviour. Which type of diversity is it that this sexual selection model is aiming to explain: modern day cultural variation, or a hypothesised genetic variation? To explain modern day cultural variation, the sexual selection model would have to be an analogy. This seems to be the most rational interpretation of the value of Werner and Todd's (1997) model to the evolution of human musical behaviour, in which case, to the extent that this *is* a valuable model, then the implicit suggestion is that there is a valuable analogy between sexual selection processes and processes that drive cultural diversity.

Having argued that it is reasonable to consider hypothetical processes without requiring that they accurately model a certain real world process, it is also important to balance that view with a requirement to produce models that can at least be interpreted as applying to the evolution of human musical behaviour. The fact that the genetic inheritance of songs and song preferences, *in the absence of learning*, seems inappropriately far removed from human behaviour, places Werner and Todd's model in an awkward position given that, besides this, it points to potentially interesting aspects of the evolution of human musical behaviour: the role of group sizes, evaluation methods, and the effect of coevolution. So why ground these aspects in an inappropriate model of sexual selection? On closer inspection, the model is redeemed by the fact that vertical cultural transmission can be operationally identical to genetic inheritance. If the model assumes that the inheritance of musical patterns and preferences is managed through vertical cultural transmission, then the problem of which diversity it is that the model is trying to explain goes

away. It may be controversial to say that vertical cultural transmission follows the same rules of recombination as genetic inheritance, and also problematic to say that offspring inherit culturally from absent fathers as much as they do from dedicated mothers, but these are questions that can be explored, perhaps even used to generate hypotheses about cultural transmission.

It may be that an open approach to potential analogies between cultural and genetic processes was the intention of the authors all along. If so, the emphasis of this point is missing from the discussion of the model. Todd (2000) points out that learning would be an excellent development for the model, but without questioning its sexual selection basis or asking whether equivalent cultural processes could be found to produce similar effects. Of course, the addition of a learning layer would be interesting; so interesting, I would argue, that it should be treated as an invitation to break open and question the sexual selection paradigm that Werner and Todd (1997) embark upon in this research. The authors fail to transpose the arguments about possible processes in the real evolution of human musical behaviour to their own model.

The reason for a focus on sexual selection is implicit from Todd's development of the model's rationale in terms of its properties: sexual selection provides a reasonably well-understood mechanism for the emergence of diversity. The emphasis on diversity, then, becomes a point of contention. Diversity is implicitly bound to what is already known about the power of sexual selection *models* (as opposed to sexual selection hypotheses of the evolution of human musical behaviour). Qualitative computer simulation models become compelling through the demonstration of emergent properties, so a sexual selection model that demonstrates the evolution of an arbitrary trait would be uninteresting, and effectively useless. The dilemma of applying computer simulation models to the study of the evolution of human musical behaviour is that some very convincing theories lead to very dull simulations, and the simulation modeller is driven towards theories which produce interesting models. In this sense, modelling appears to run the risk of being sensationalist. This is not intended as a cynical observation, but one that highlights the difference between supposedly modelling a process, and searching for new process with interesting emergent properties.

However, we can begin to address this issue by scrutinising what emergent properties would be relevant to the study of the evolution of human musical behaviour. Ultimately, such an anal-

ysis depends on one's definition of music. Werner and Todd's (1997) original interest lay in the possibility that audible structure may emerge in the musical patterns and judgements of their evolving agents. Unfortunately this did not occur. If it had then we could say that human judges recognised the musicality emerging from a process of sexual selection on originally random forms. The emergent diversity produced by the model has been taken as an acceptable alternative, as implied in the nature of the model itself. Thus the means to evaluate the model came after the existence of known mechanisms for generating diversity from a process of sexual selection. Furthermore, diversity has no immediate value in developing an understanding of the evolution of human musical behaviour unless this diversity is shown to function in maintaining or driving the emergence of other aspects of musical behaviour. Above all, the very existence of musical interaction is an invariant part of this model. An outcome with no diversity over time or across the population is just as musical as one with high diversity, according to the notion of music that is presented in the design of the model.

4.3.2 Managing the complexity of the evolution of human musical behaviour in a model

How does this analysis clarify the ways in which a model of musical evolution can be useful? Werner and Todd's (1997) model can be viewed in terms of two forms of evolution of music model. According to a dual model of inheritance, musical behaviour is manifest in one or both of two locations: culture and genes. A theory of the evolution of music needs to explain how music establishes itself and stabilises itself in a population in either or both of these locations, according to definitions of the elements that we consider important to defining musical behaviour. A model that proposes to explain how this happens fits the first form of model, which can be described as *complete*. The other form is a model that incorporates from the beginning some kind of musical behaviour. In this case, the existence of musical behaviour in the model is at some level invariant, but some additional features associated with music may emerge. Such models can be described as *incomplete*. On the surface, Werner and Todd's (1997) model is of the latter form. The sexual selection hypothesis is written into the model, along with a description of how music is made and perceived, and the nature of the model is explored. There is almost

no potential for the stability of specific musical characteristics in this system to fail; if there were then the efficacy of the sexual selection hypothesis could become a point of investigation. On the other hand, different settings in the model are compared to determine which conditions produce potentially music-like outcomes. One demonstrates a clustering effect, and this is an important characteristic relating music to social structure (more detailed than diversity, this is a *distribution* of diversity), which is not itself written into the model. If the model defines a process in which musical style groups can emerge and stabilise (or not), then its efficacy with respect to a certain aspect of musical culture can be pinpointed. An important theoretical question is whether or not complete models of the evolution of music are actually possible, not only in terms of the practical challenges of reducing the problem into a programmable simulation, but in terms of the continuum between modes of interaction in humans and our ancestors. Clearly we do not expect to create models in which musical organisms emerge from primordial soup, but on the other hand we *can* expect models in which music emerges as a system of interaction from other forms of behaviour. For example, Quinn (2001) presents a model of the evolution of communication in which two simulated robotic agents begin to communicate using jittering movements; that is, the substrate for their communication, visual signalling, emerges from the existing properties of moving, seeing organisms. Realistically, we are confined to incomplete models, but our goal should be to extend these models through interpretation and more extensive interaction with current debates in the evolution of human musical behaviour. A theoretical approach to building simulation models of the evolution of human musical behaviour needs to keep in mind this notion of what music is, and what existing aspects of behaviour it emerged from.

In conclusion, Werner and Todd's (1997) model does not shed a great deal of light on the application of the sexual selection hypothesis to the evolution of music, but it does propose a starting point for studying in simulation the relationship between a musical interaction regime and macroscopic social phenomena. It also implies, but does not directly state, a tantalising analogy between sexual selection involving genetic heredity and sexual selection involving vertical cultural transmission. This implication emerges from questioning the model's basic assumptions (or ambiguity of its assumptions, as I have discussed above), and attempting to reallocate elements in the design of the model according to that critique. Furthermore, there appears to

be a level of circularity in defining the interesting outcomes of the model. It may be that rather than raising this point as a criticism of Werner and Todd's work, it should be treated as a general problem of reconciling modelling considerations with real-world theory, associated with the necessary incompleteness of models discussed above. The notions of "*amount of synchronic and diachronic diversity*", and "*structure of diversity and organisation of individuals*" must be considered in terms of what they tell us about the totality of musical behaviour. We observe these aspects of music in the world, but could they actually be conditions rather than outcomes of musical behaviour? How do these phenomena interact with other aspects of musical behaviour? Could certain social structures be seen to lead to certain musical structures in extensions of Werner and Todd's model? I believe that the argument that computer simulation models are relevant to the theory of the evolution of human musical behaviour will gain greater strength as we begin to investigate the interaction between multiple evolutionary processes and aspects of musical behaviour.

There may also be good reason for computer modellers of the evolution of human musical behaviour to devise their own paradigm within which comparative analyses can be performed. This goes beyond the notion of simply proposing hypothetical processes, but in some ways it is a natural implication of the argument for extended analysis and design. Although different models may essentially be irreconcilable due to their differences of focus, the factors relevant to theories of the evolution of human musical behaviour could be broken down into a number of categories that could define a modelling framework: reproduction, genetic heritability and variation, sexual interaction, learning mechanisms, vertical and horizontal transmission, properties of sound and rhythm, properties of neural systems, perceptual mechanisms, social evaluation, social intelligence, social organisation, regimes of interaction, the distribution and extraction of environmental resources, and individual fitness with respect to an environment.

With these considerations in mind, I return to the arguments developed in Chapter 3, and propose the design of a computer simulation model, with the aim of exploring the potential processes that could contribute to a competitive style learning theory of the evolution of human musical behaviour.

4.4 Simulating the evolution of human musical behaviour in light of current theoretical debates

As identified in Chapter 2, one of the most important clarifications we can make about the evolution of human musical behaviour concerns the distinction between cooperation and cohesion. Whilst cooperation is a form of interaction clearly defined by rules of costs and benefits, cohesion cannot be defined purely in terms of the relationship between two individuals in a given context. Instead, cohesion must incorporate not only the cost-benefit relations between two individuals but also the scenario within which these individuals come to interact. Measuring the costs and benefits of musical behaviour clearly depends upon gaining a good understanding of the role of music in various contexts in social life. It depends upon the cognitive machinery of the individuals involved, as well as upon the cultural systems in which they find themselves, as well as more specific social arrangements. Understanding music in terms of cooperation, in a traditional game of costs and benefits, will simply not be possible because the far more fragile, variable and flexible entity is the context itself. Rather, a complex social system is one in which modes of interaction are entities in their own right, which specific individuals can manipulate, exploit, promote or erode.

What this means is that, on the one hand, we must maintain that musical behaviour evolved because more musical individuals were in some way fitter than less musical individuals, whilst on the other hand it is difficult to escape the necessity of appreciating that musicality is only a peculiar form of behaviour which is retrospectively defined as having evolved through a certain biocultural evolutionary process.

The problem of cohesion is one of understanding how the interactions between sets of individuals drives new contexts according to which behaviour is selected for. At each point in time we can consider specific interactions in terms of costs and benefits. We can be reasonably sure that individuals (or genes, at least) act in their own self interests, although self-interested behaviour may involve certain statistical approaches to trust and cooperation that allow individuals the potential to be deceived (Dawkins, 1976; Maynard Smith and Szathmary, 1995). However, cohesion refers to a high level of interaction which may be more or less cooperative. There are many social contexts that are highly antagonistic but involve close interaction. The benefits of

being in a large group outweigh the costs of potentially harmful interactions within that group. Often the antagonistic relationship is a direct consequence of the cohesion of the group, such as in a leadership competition (Dunbar, 2006). More subtle situations involve economic relations such as class relations in modern society (*c.f.*, Knight, 1991). Individuals are bound into economic relations – they would be in trouble without them – but compete within that context, often resulting in sustained fitness differences (*e.g.*, Dumont, 1980). I argue that it is important to bear in mind the potential for musical behaviour, and other forms of cultural behaviour, to embody this double-edged sword of cooperation and competition that underpins cohesion.

In modern human behaviour, music provides a focal point for the unity of groups of people, which is very flexible in that it does not appear to be strongly bound by kin relatedness (Cross, 2003a). But it also provides a genuine means for a person to benefit from their cultural context (exploit may be too strong a word) (Huron, 2001). Most importantly, it is possible that musical behaviour binds people to it (in a similar sense to how Dawkins (1976) and Blackmore (1999) describe memes manipulating human behaviour). Interacting in a musical culture entails being musical. It is possible that the only missing link in the runaway process of the evolution of human musical behaviour, then, is the problem of how and why a successful music maker binds other individuals into musical behaviour. This would be to bind them into a cohesive system, and to alter their evolutionary context, but it says nothing of the level of cooperation involved. Recall that there are strong arguments, such as those of Pinker (1998) for an origin of music that is based on the collision of cognitive biases. Therefore, the argument that other individual's perception of music can be exploited is feasible even in the absence of an evolutionary history to music or a cooperative explanation for its existence.

There is some common ground between the kind of social dynamics involved in sexual selection and the ideas presented here. Sexual selection also tackles the problem of how someone who is successful can reinforce the demands upon future individuals for behaviour in that specific domain. However, whilst Miller's (2000b) view of cultural behaviour and social organisation is that it works primarily to enhance the effects of sexual selection, we can view cultural behaviour as creating a whole new selection pressure that emerge alongside sexual selection. Since culture potentially changes too fast for biological change to track it, there are two ways

in which cultural phenomena can influence biological evolution. The first is that what emerges biologically is a general capacity to handle varied cultural contexts; increased flexibility, in the form of creativity and learning abilities, and general social intelligence. The second is that cultural phenomena arrive at some degree of stability through processes of cultural reinforcement, and therefore provide a stable environment in which biological adaptation takes place.

The central phenomenon that underlies this latter process was discussed in Chapter 3. It is the potential for an individual who is successful in some socially determined behavioural domain to go on to reinforce the existence of that same domain. It may not only be the music maker that is successful, but also those that support him or her, and assert the value of their behaviour. From the point of view of evolutionary dynamics, it is clear that this is a kind of evolutionary stable state (ESS) (Maynard Smith, 1982) in a cultural domain. In Chapter 3 I discussed how runaway evolution could occur through the disequilibrium between individuals performing successfully in a culturally established domain and trying to reinforce that domain, and individuals undermining that domain altogether.

A computer simulation modelling approach to the evolution of human musical behaviour can be useful for considering ways in which cohesion and cooperation interact in this kind of domain. Firstly, can cultural contexts, resting on existing biological facts, establish a situation in which individuals are able to enhance their fitness using musical behaviour. Secondly, can this situation create the basis for cooperation. Failing that, can it at least sustain itself as a form of social behaviour and drive the evolution of other aspects of behaviour.

In the following chapters, I attempt to address the problem of modelling such phenomena. However, the difficulties outlined above regarding the complexity of working musical behaviour into a model result in a very roundabout approach to this central problem. In these models a general concern is whether or not it is possible to breach the principle of keeping models simple enough that they are easily understood. The complexity of human musical behaviour embedded in social contexts puts forward a tempting argument for building models that do go beyond our ability to simply see what is going on in the models. It should be clear that, given that the phenomenon in question is a computer simulation model, any hypothesis we have about the behaviour of the model can ultimately be tested. The fact that this may take a long time is not in

itself a scientific justification for sticking with simpler models.

4.4.1 What evolves when music evolves?

In the above criticism of Werner and Todd's (1997) modelling approach, the question was raised as to what relevant phenomena could be counted as aspects of the evolution of human musical behaviour. I argued that the emergence of a general capacity for diversity in a group could not significantly contribute to theories of the evolution of human musical behaviour. Todd's later work (*e.g.*, van den Broek and Todd, 2003) addresses this by looking more closely at an empirically significant aspect of musical structure, the role of isochronous beats. In general, I propose that the evolution of human musical behaviour consists of three areas of *individual* human behaviour in which music can be seen to evolve: a music-specific perceptual system; specific capacities for the production of music; and a tendency for music to evoke a strong emotional reaction, or at least a stimulated or curious reaction. I emphasise the fact that these are aspects of individual behaviour. Human behaviour can also be seen on a group level and it can (and should) be argued that the existence of a music-making culture, regardless of individual behaviours, is the central characteristic of the evolution of human musical behaviour. One could refer to this as the music itself, music independent of human behaviour. Whilst I firmly believe that this is a central concept in understanding the development of human musical behaviours, I find it hard to accept the notion that music is itself an agent on a par with human actors; it is the actions of individuals, human behaviour, that produces and reproduces music, even if this process of reinforcement has qualities that are greater than the sum of its parts. Thus we must understand the evolutionary process as one operating on a macrocosmic, biocultural scale, but we must identify change at the microcosmic scale, in the behaviour of individuals.

More importantly, thanks to cultural inheritance, these three capacities can be stable without having particularly biologically prescribed qualities. We do not fully understand the developmental origins of musical abilities, but we accept that musical behaviour has considerable staying power in human populations, with trends that may not be universal, but are at least extensive. I propose that these three capacities broadly cover the full extent of the meaning of the evolution of human musical behaviour, and that in examining models of the evolution of human musical be-

haviour, it is therefore necessary to focus on these areas. In fact, capacities for production stand out as being considerably less significant than those for perception and appreciation. Production requires that we rise to the challenge of satisfying our own and others' evaluative demands in a specific perceptual context, and could arguably be discounted from a formal model of the evolution of human musical behaviour, based in the assumption that production is secondary to evaluation in the process of actually driving evolution. If, in a context of strong social learning, the most salient and learnable behaviours get picked up and incorporated into stable cultural systems, as discussed in Chapter 3, then certain regular modes of behaviour, such as walking, must necessarily have been the starting points for this evolutionary process. Nevertheless, we can, at least in theory, explain musical production in terms of musical perception, musical evaluation, plus general capacities to learn, reflect and create. It is harder to imagine explaining music perception and evaluation as being natural consequences of an inherent ability to make music. For this reason I argue that perception and evaluation are the things that evolve when we say that music evolves, and that computer simulation models should look at the evolution of perceptual and evaluative systems in artificial musical agents.

Chapter 5

Design of Simulation Model

5.1 Aims

In Chapter 3, I developed a theoretical argument for focusing attention towards the role of competitive style dynamics in the study of the evolution of human musical behaviour. Runaway evolution through competitive style dynamics could be a process that is unique to our species' history because our direct ancestors were unique in their social learning skills. In section 3.7.2, I discussed the notion that this runaway process can be understood as having swept up the evolving species from one environmentally determined situation and deposited it in a new niche. This is not an adaptive process, as such, but is also not necessarily neutral in evolutionary terms. The consequences of such a process could have been fatal, but in this particular case, the qualities to emerge were advanced learning skills, which may even involve the emergence of symbolic thought and imagination: in short, skills easily transferrable to engaging with challenging environments. I argued that it is important to understand in greater detail the interaction between cohesive and competitive processes in the context of stylistic learning, and to see to what extent it would be possible to create a plausible and more detailed model of how competitive stylistic learning might actually be capable of driving the emergence of musical behaviour, including developed perceptual capacities, learning capacities, and absolute as well as culturally relative musical enjoyment.

In Chapter 3 and at the end of Chapter 4, I discussed purely cultural processes that would establish and reinforce musical behaviour in a population. I suggested that it is necessary that

the successful musical individual has the means and inclination to reinforce the value of musical behaviour in the rest of the population, thus sustaining musical behaviour as a stable cultural system. We can therefore ask what biologically innate qualities a musical individual should have from two points of view: firstly, from our understanding of what biologically innate capacities contribute to human musicality; and secondly, following the above argument, by asking what kinds of cultural behaviour would have the property that successful proponents of such behaviour would inherently reinforce the behaviour in others. From the first point of view, I suggested that three broad categories of behaviour are relevant: capacities for music perception; capacities for music production; and capacities for musical motivation. With these capacities in place, human musical behaviour becomes a biological likelihood, if not a cultural one. In fact, capacities for production become a relatively weak component when one considers how readily they are precluded by capacities for perception and motivation.

The aim of the experimental section of this thesis is to explore these phenomena through computer simulation models. Following on from the discussion in Chapter 4, a secondary aim involves contributing to the development of the methodology of computer simulation modelling in this particular field. The evolution of human musical behaviour, as it fits into place with other aspects of the theory of the evolution of humankind, brings along a unique set of concerns that differs slightly from theories of the evolution of language, altruism (and social behaviour in general), theory of mind, intelligence, sexually attractive traits, and all of the computer simulation models associated with these processes. In the broadest sense, all computer simulation models of the evolution of human musical behaviour will fit into the category of multi-agent or population-based simulations, mostly with biocultural elements. But they should also form a more closely related subset. For example, all computer simulation models of the evolution of human musical behaviour will implement means, perhaps evolvable ones, for producing sound and for perceiving sound, as well as genetic evolutionary processes and other rules for interaction. At least in some cases it should be possible to consider a set of models that spans a wide range of evolutionary processes associated with the evolution of human musical behaviour. In part, this aim is motivated by a belief that, although models need to be simple in order to fully explore their behaviour, these same sets of simple models can be seen as subsets of more complex and

more general families of models, with greater numbers of parameters to consider. As computing power increases and grid processing technologies and information retrieval technologies become more powerful, models which are not irreconcilable, but can be related to each other at a greater level of complexity, will become increasingly valuable in collaborative work on theories such as the evolution of human musical behaviour. It will become increasingly possible to engage in parametric searches of these models. Therefore working a level of generality into the design of these models is valuable.

5.2 Theoretical questions

What hypotheses can be developed in order to test the theory that competitive style learning is primarily responsible for the evolution of human musical behaviour? At present this theory is not well established with cohesionists, sexual selection competitivists, and consequentialists, because it is not recognised as a plausible process compared to the existing processes known to stem from Darwinian principles, even if the grounds for such an approach may have growing theoretical support from gene-culture coevolutionary theorists and niche constructionists. Like theories of sexual selection, prisoner's dilemma interactions and genetic take-over, the approach needs to be demonstrated in order to support its feasibility. To reiterate, the logic of Darwinism is forced onto phenomena, which may, through their apparent defiance of this logic, invoke us to propose new principles. From a Popperian perspective, a demonstration of the possibility of competitive style learning refutes the implicit hypothesis that possible evolutionary processes are limited to a finite list excluding this one. This is implicit in many evolutionary explanations because we set out, as far as possible, to build theories from processes that we do know to work. Therefore, we can consider hypothesis of the form: *it is possible for functionless music-like systems of cultural interaction to be self-sustaining and to drive the evolutionary emergence of biologically innate musical capacities and tendencies*. This seems like an awkward stretch of the refutability approach to science, however. From a Kuhnian point of view, contributions to science come from diverse processes, including theory development. Therefore, competitive style learning would, following the demonstration of its efficacy, gain plausibility, and this offers a genuine contribution to theories of the evolution of human musical behaviour.

In particular, a competitive style learning hypothesis is most closely associated with the problem of the emergence of an interaction system that exhibits the capacity to foster group cohesion. Throughout Chapter 3, I focused on a hypothesised positive emotional response to a musical signal, and proposed that the main problem for a cohesionist theory of the evolution of human musical behaviour is to explain how a situation could arise in which this positive emotional response can be selected for. I also discussed the notion that processes of genetic take-over could drive the evolution of innate cognitive capacities in response to the stable establishment of a system of learning-based competitive interaction.

These two factors provide two elements that we can look at in genetic evolutionary terms: the evolution of a positive emotional response to musical stimuli, and the evolution of specific perceptual capacities that result from interaction within a musical domain. Since the positive emotional response is the result of a manipulative stimulus by another individual, I will refer to this process of interaction as *enchantment*. The potential strength of the positive emotional response can therefore be referred to as the *susceptibility to enchantment* (S), and will be assumed to be an innate aspect of an individual's behaviour that varies genetically.

Except at the outset of this thesis, I have avoided being particularly specific about musical structure, cognitive capacities for musical interaction and the functions that music is put to in modern human life. This is open to criticism from various commentators who have insisted on the importance of understanding the phenomenon one is trying to explain (*e.g.*, Bickerton, 2001; Cross, 2006). The problem of how an interaction system emerges in a population is arguably less dependent on the specific mechanisms that it uses than on the social organisation of the individuals in the population and its potential role in this organisation, as I have discussed in Chapter 2 and this also happens to be appropriate from the point of view of designing a model of musical interaction that explores the evolutionary consequences of these social contexts. At the outset of such an attempt, extremely elementary representations of this system of interaction are necessary, so whilst interaction will be the focus, it is still necessary to look at simpler notions than Cross (2006). Furthermore, although in Chapter 3 I discussed a set of reinforcing processes that would establish the social determination of fitness, modelling these processes directly, in a first attempt, would require too much behavioural detail. This is because of the large number

of parameters that would need to be considered in order to design agent behaviours that did a realistic job of choosing models, learning styles, paying for learning access, and so on. In order to simplify this set of processes, I have broken it down into two parts: (a) a dynamical system of stylistic interaction that results in individuals possessing to varying degrees a quality which I will refer to as *status*, and (b) a set of rules determining how individuals' status affects their survival. These components can be referred to respectively as the *interaction game* and the *status game*.

- The interaction game represents the existence of a cultural process of style learning which does not have a direct adaptive function. This is the component that will affect everything to do with agents' styles. It is assumed that all of the aspects of the model to do with the interaction game have no *direct* benefit to individuals' fitness; the interaction game has nothing to do with either mating or extracting resources from the environment to stay alive.
- The status game represents the existence of a social reward system that results in some individuals enjoying privileges which ultimately enhance their fitness.

In the competitive style learning theory, the set of reinforcing processes described above are assumed to explain the emergence of the status game from the interaction game. As stated, the goal of demonstrating this process in a model is not directly attempted. Instead, we can approach questions about the efficacy of a competitive style learning process indirectly, by considering the relationship between the interaction game and the status game. We therefore ignore, for the moment, the notion that an individual pays rewards in order to gain access to a style for the sake of learning it, and consider instead a more elementary process in which an individual pays rewards to another individual simply as a maladaptive result of enchantment, and the extent to which an individual is enchanted by another is a product of the pair's relative styles, as well as the enchanted agent's S .

Therefore, another element from the outline of competitive style learning in Section 3.3.1 that is abstracted in the model is the notion that individuals make judgements about who to interact with based on status. Although these elements are important to a complete theory of

competitive style learning, their contribution can be set in stone for the purpose of exploring other aspects of the model.

In short, we begin by considering multi-agent biocultural models that are purely based on two components: a system that provides socially determined fitness rewards, and a maladaptive, culturally dynamic system of interaction. The variable S , a property of each individual agent, provides the link between these two elements. It determines the extent to which agents' enchantment by other agents results in them paying status rewards to their enchanters. A status reward could be seen as representing a degree of bondedness from one agent to another, and we assume that agents that have attracted many strong bonds from other agents will be fitter than those that haven't. To begin with, we do not assume that 'paying' a status reward is actually directly costly, only that higher status individuals have greater access to resources, imposing an indirect cost on lower status individuals around them.

Thus, although this model ignores some aspects of the process of competitive style learning, and therefore does not directly test the competitive style learning hypothesis, it sets up a context in which it is possible to test the hypothesis that one-way social bonds can become established in a population of agents who are competing in an interaction game.

5.3 Approach to building a simulation framework

5.3.1 Evolution

Genetic algorithms provide a good natural starting point for the design of a multi-agent biocultural model, because they already simulate the process of genetic evolution in a population (Holland, 1975; Mitchell, 1998). Genetic algorithm theory has identified many different ways to recreate the basic evolutionary properties of genetic evolutionary processes, highlighting the fantastic flexibility inherent in Darwin's original conditions for natural selection to occur (Koza, 1993; Holland, 1975; Mitchell, 1998). Genotypes can be implemented as binary strings, strings of real numbers (Mitchell, 1998), strings of program code in the case of genetic programming (Koza, 1993), and anything else that can be gradually mutated and recombined can be treated as a genotype (see, for example, Dawkins, 1986). Following a non-Lamarckian approach (Jablonka and Lamb, 1995; Gould, 2002), what matters is that genes are never influenced by the interac-

tion of their phenotypes except by the process of natural selection: some genes produce fitter phenotypes and those genes are the ones more likely to survive into future generations. These genes are copied from one generation to the next, but may also mutate during this copying process (Darwin, 1860). Sexual recombination is an additional process whereby two genotypes are merged together in some way during reproduction, the important point for natural selection being that, through this process, pairs of distinct good traits which have evolved independently might become combined in one individual, who is subsequently fitter than both parents (of course the opposite can happen too).

Most genetic algorithms work in steps of a generation (Mitchell, 1998). After each generation, many of the less fit individuals die and are replaced by the offspring of the fittest individuals. Between each generation all of the individuals are tested for their fitness. In models of coevolutionary processes, the individuals may be required to interact with each other in order to establish their fitness, but this will still involve one iteration of testing (Mitchell, 1998). Coevolutionary processes result in tests of fitness that are not fixed over time, because other agents in the population contribute to the environment in which they are tested. Under certain conditions, this can result in continual change, and the emergence of complex behaviours and adaptations (see, for example, Cliff and Miller, 1996; Sims, 1994). In order to study biocultural processes effectively, we need to stretch out this between-generation process and fill it with cultural interaction, whilst being able to control the ratio between the rate of cultural interactions and the rate of genetic interactions (*c.f.*, Di Paolo, 1997, 2000). One way of doing so would be to have lots of iterations of agent interactions between each generational update. However, this introduces an artificial discontinuity whereby, all of a sudden, a large group of individuals is replaced with new ones. If cultural learning is involved then we require a smooth process of genetic evolution which does not interrupt cultural evolutionary processes, and instead facilitates a staggered population update (Di Paolo, 1997).

The simplest way to do this is for the internal state of the agents themselves to determine whether or not they should die or be able to reproduce, as in ecological modelling (Odum, 1983), which in turn should be designed to depend sensibly on the results of their interactions with the environment, as well as on internal variables (for another example see Taylor, 2004). This allows

for a dynamic population in which genetic and cultural evolution can carry equal weight. It also means that we relinquish direct control over the population size, since we do not know exactly when any agent will be born or will die. In a sense the rules of birth and death go from being top-down in the case of traditional genetic algorithms, to being bottom-up. Population size is determined dynamically by the birth and death of agents rather than being prespecified.

In the models presented throughout this thesis I will focus only on the mutation aspect of evolution, based on the assumption that the essential results of models of cohesion and competition will not be affected by sexual reproduction. Thus, natural selection is implemented in its most elementary form. Along with many of the other assumptions in this section, this assumption is made under the demands of simplicity (Axelrod, 1997), and may be misplaced. Ultimately it would be desirable to test all such speculations with specific experiments, but this is necessarily deferred due to the time-consuming process of building and exploring such models. Obviously, models of sexual selection would have to have a notion of sexual recombination, but I do not consider these possibilities here.

Following Di Paolo (1997), I have chosen to model agents as having probabilities of death and reproduction at each time step, which can be independently defined according to any other factors in the model, such as an agent's age, or even the overall size of the population. This approach is common in ecosystem modelling (Odum, 1983; Begon et al., 2006) and in artificial life research in which the coexistence and interaction between different organisms is relevant, such as in research into the evolution of organismic complexity (*e.g.*, Werner and Dyer, 1992; Cooper and Ofria, 2002; Richard E. Lenski and Adami, 2003; Taylor, 2004), but is subtly different from co-evolutionary models where agents interactions are evaluated in generational blocks, and then reset (*e.g.*, *Cliff and Miller, 1996; Quinn, 2001*). The probabilities of death and reproduction can be attached to some physically grounded measure of survival capacity such as food, but we can also attach them to culturally defined factors, which is relevant in the case of the evolution of human musical behaviour. Having chosen an asexual model of evolution, the probability of reproduction simply determines the probability of producing a genetic copy of the individual with mutation. We can also attach costs to reproduction, and can allow offspring to inherit non-genetic factors from their single parent, the most obvious being spatial location in the case of a

spatialised model.

5.3.2 Interaction

A model of musical interaction in its most minimal form involves one agent producing a stimulus and at least one other agent perceiving that stimulus. A stimulus is a pattern *placed* in the environment. This is a very broad view and arguably it should not be refined any further in the context of these initial exploratory models. The key restriction that this definition places on musical interaction is that it does not allow a communicative link directly from one agent to the other. An agent is not given direct access to another agent's behaviour, but must perceive it, and in doing so internally represent it, whereas such a limitation does not apply at the production stage, where a different process takes place. Production of behaviour could also be said to require a process of translation from concepts to actions, but we can also consider a simpler subset of behaviour which involves perceiving one's own actions and comparing different arbitrary variations of these actions in order to determine a chosen behaviour. That is not to say that an individual is capable of producing *any* behaviour, only that they are not necessarily involved in a process of conversion, whereas they are at the perception stage. This is a potentially contentious point, and I stress that it is not intended as a view on the difference between musical perception and production, but, again, as a simplifying assumption.

Nowhere in this sequence can we specify a thing we can definitively call music, but we can say that the entity referred to above as a stimulus is an aspect of the environment (it is *external* to its producer, cognitively speaking) that no perceptual system can internalise in its entirety; it is in some sense more absolute than the agent's perception of it. An agent may perceive two slightly different stimuli as the same. Two agents may perceive the same stimulus in different ways, even creatively adding information to their perception. In fact a single agent may even perceive an identical stimulus in two ways at different times. The completeness of the stimulus compared to the incompleteness of the perception of that stimulus is critical. There is no need, at this stage, to consider specific production factors such as bodily constraints on making music, aspects of acoustics, and the temporal nature of musical experience.

5.3.3 Limits of simulating music

It is also far too challenging to consider the difference between musical interactions amongst members of a group and musical behaviour *in groups*. This would require a model of multi-agent interaction *within the time* of musical performance, whereas individual interaction models can abstract that moment of musical performance into an instance, a single signal; hence the notion of producing a stimulus and placing it into the environment. Drawing a boundary around a notion of musical interaction reduces the complexity of having to model how individuals process signals on a moment-by-moment basis, and also does not over-stretch the use of different timeframes, when we already require extreme simplification to achieve a simulation that runs over two timeframes. Biocultural evolutionary models are already straddling the genetic evolutionary time frame and the time frame of daily, monthly or yearly interactions. Similarly, modelling the perception of musical stimuli created by more than one individual in concert is also not attempted.

As a result of these considerations, a very simple geometric model of style is proposed as a means to tap into the essence of a changing style system. According to this model, styles are viewed as elementary particles, which move in a geometric space according to simple dynamical laws determined by the design of agent behaviours. These behaviours are not directly inspired by flock or swarm dynamics, but bear a clear resemblance to them (Reynolds, 1987; Bratton and Blackwell, 2007). This is not to suggest a general theory of style, based on particle dynamics, but is mostly for convenience. It is also intended to emphasise the fact that in studying the evolution of human musical behaviour in simulation models it is necessary not to pre-suppose the existence of the thing you are trying to explain. Therefore, considering elements of human musical behaviour and trying to understand their emergence in models cannot work if we design them into our models. The substrate underlining musical behaviour cannot *be* musical behaviour.

The fact that many of these aspects play a major role in theories of the evolution of human musical behaviour poses a methodological challenge that is applicable to modelling challenges in other areas; if we always attempt to avoid complex design issues, will the modelling approach favour the development of theories that continue to ignore these aspects? The simple answer to this question is yes: models will always be biased towards more *modellable* approaches. But

this is a necessary component of any modelling methodology; our models will inevitably be inaccurate, but they will tell us new things about abstract processes that may be relevant to more accurate models. It is necessary to be aware of the favouritism of modelling towards certain aspects of a real world process and not others. Pointing out missing elements in a model may detract from the model's *realism*, but it does not, however, detract from the *validity* of the model. Some, if not most, arguments about the evolution of human musical behaviour will therefore take place outside of the actual domains that these models define.

5.4 Implementation

All of the models in this thesis have been implemented as programs written using the Java programming language. A library was written to facilitate generic agent-based models which could run over a computing grid and access a remote MySQL database to store the results of the simulation. The aim of this library was to allow models to be designed by specifying agent behaviours in a single Java class, and to explore model parameters using properties files to specify general model parameters (such as the length of the simulation and the initial size of the population), and particular parameters given in the agent file (such as mutation rates and expected longevity). The use of a computing grid with a central repository of results in a MySQL database meant that many simulations could be explored in parallel across a range of parameters. I will not discuss the implementation of this library because its functionality has no impact on the design of any of the models.

Individual agent behaviours were therefore defined for different experiments and a wide variety of approaches were considered during the development of the final models discussed in this thesis. I will discuss these final models here and point out only where necessary any variations that had relevance to the development of these main models.

5.4.1 Simulation updates

For each agent design, a population update method was implemented, which defined how the whole population changed over time, along with an individual update method, which defined how agents performed actions at each time step. Populations were updated synchronously,

meaning that at each time step every agent was updated exactly once. A synchronous update assumes that all agents act simultaneously and is unproblematic as long as the state of any individual is not changed until the end of the iteration through the whole population, *i.e.*, every agent has performed an action at that time step. Thus any values that change for each agent should not actually be set to their new values until all agents have performed their update. Otherwise, the order of the iteration through the population may have an effect on the simulation. The basic sequence for the population update was therefore as follows: (1) individuals perform actions, but temporarily buffer their new states; (2) individuals update their states; (3) individuals are tested to determine if they reproduce or die.

5.4.2 Rules for death and reproduction

Following simulations by Di Paolo (1997), discussed in Chapter 4, I model agents as having a probability of death and reproduction at each time step of the model. This provides the basis for a model that runs over biological and cultural evolutionary time scales with a staggered turnover of new individuals (as compared with an abrupt one in the case of standard genetic algorithms). These probabilities are determined by the agent's fitness, which is in turn determined by other factors described below. If fitness falls below a low threshold, $thresh_{die}$, the probability of death rises from zero to a non-zero probability, P_{die} (Equation 5.1). If fitness rises above a high threshold, $thresh_{rep}$ the probability of reproduction rises from zero to a non-zero probability, P_{rep} (Equation 5.2).

$$p(\text{death}) = \begin{cases} P_{die}, & \text{if } fitness < thresh_{die} \\ 0, & \text{otherwise} \end{cases} \quad (5.1)$$

$$p(\text{reproduction}) = \begin{cases} P_{rep}, & \text{if } fitness > thresh_{rep} \\ 0, & \text{otherwise} \end{cases} \quad (5.2)$$

5.4.3 Environment and energy

Agents are placed in a spatial model in order to allow local divergence of groups both in terms of cultural style and genetic structure, with frequent interaction amongst local, more homogenous

groups, and less frequent interaction between individuals over a greater distance. The local establishment of styles within a broader local context is integral to the discussion of Chapter 3, and the evolutionary processes discussed there depend on the fitness advantages of interaction in local groups with respect to wider social groups, thus social structures is a critical part of the model. Spatial distribution is a simple way to implement this structure.

Agents live in an environment consisting of a toroidal grid of $W \times W$ square cells. Each agent inhabits a specific cell at each time step, with a probability p_{move} of moving to any neighbouring cell (chosen at random) at any given time step. Multiple agents can inhabit the same cell, and, when an agent reproduces, its offspring start life in the same cell as itself. Cells contain energy, and agents feed from a square *foraging region* extending a given number of cells, $R_{foraging}$, in each direction (R stands for *range* here, as opposed to W , which stands for *width*). When an agent feeds, it can extract a fixed amount of energy, $e_{extract}$, from each cell in its foraging region. However, if a cell does not contain enough energy for all of the agents feeding from it at any given time step then the available energy is distributed evenly amongst those agents. This approach treats food sources as continuous and divisible entities equivalent to zones of food-bearing terrain, such that it is realistic to assume that multiple agents taking from the same food will simply share it equally. A constant amount of energy, e_{add} , is added to each cell at each time step, representing the regrowth of food, but each cell has a maximum energy capacity, e_{max} . At each time step, an agent's energy level is also depleted by subtracting a fixed amount, $e_{deplete}$. Whether or not agents feed is determined by the rules of social interaction. An agent's fitness is determined by the equation $f = \text{Max}(0, e \times (1 - a/a_{max}))$, where f is the agent's fitness, e is the agent's current energy, a is the agent's current age, and a_{max} is a provisional maximum age for any agent. The relationship between fitness and age is used to embed a natural age limit on agents. Note that agents can, in fact, be older than a_{max} , because death is determined probabilistically. When an agent reproduces, it shares its current energy equally with its child.

The above rules collectively establish a multi-agent world within which it is possible to study the effects of social interactions. These rules are directly inspired by Di Paolo (1997) with adjustments to fit the goals of the model. Other examples of rules for energy, food, death and

reproduction are discussed in systems ecology literature (*e.g.*, Odum, 1983). We can now specify relationships between how agents interact and their access to energy (and therefore fitness) in order to establish an interaction between processes happening at a cultural time scale and their genetic effects. Meanwhile, the limitation of resources will ensure that the population will naturally be regulated in size across specific locales in the environment.

5.4.4 Modelling the effects of the social determination of fitness

In order to define a context for the social determination of fitness I introduce a socially determined parameter, *status*, which ultimately determines an agent's access to energy. I also introduce a set of rules to determine a dynamical cultural process. As with other modellers of musical behaviour, I consider the variable nature of musical style to be an important phenomenon to establish in a model of musical behaviour. From this point of view, following Boyd and Richerson's (1985) basic position that the significant aspect of culture is the fact that individual behaviours vary dramatically within the lifetime of the individual, the social determination of fitness is a meaningful concept only if there is broad variation of behaviour which is *not* genetically determined. I treat this as the main requirement of the agents' interactions and attempt to devise a simple artificial system of dynamical cultural change based on an arbitrary behavioural *style*. The system is defined as being arbitrary because this style has no direct effect on fitness, but only an indirect effect based on individual interactions.

Whether or not an agent feeds at each time step is determined by a status tournament that the agent plays with another agent. Agents only interact with other agents within their local neighbourhood, which is also a square region extending a fixed number of cells, $R_{interact}$, in each direction from the agent. For status tournaments, a random agent is chosen from the neighbourhood. If the choosing agent has a higher status than the chosen agent it gets a chance to feed from the environment using the process described in the previous section (but not the other way round, so each agent has *at most* one opportunity to feed at each time step).

Status also diminishes over time. This is done using exponential decay with a decay constant, s_{decay} (where $0 < s_{decay} < 1$), which is multiplied to each agents status at each time step.

5.4.5 Musical representation

Status is determined by musical interactions with other agents. Each agent has a style, which should be understood as a point particle in a style, or cultural space. Each agent's style is therefore a real vector. This style is simply defined as a set of D real-valued numbers, where D denotes the *dimensionality* of the style space. Each agent performs in front of each of the other agents in its neighbourhood and receives a status reward from each of these agents based on their evaluation of the performer's style.

I approach evaluation by assuming that the particle space is featureless. That is, no style is fundamentally different from any other style, except in their relative positions. I begin by considering relative Euclidian distances between the styles of performer and evaluator as the sole criterion for evaluation. An evaluation therefore begins by measuring the Euclidian distance between an agent's own style and the perceived style of the performing agent. This measurement can be understood as a novelty rating. We can then consider different ways in which agents respond to novelty, by applying a hedonic function to this novelty rating to get a preference score (see Saunders, 2001; Saunders and Gero, 2001a,b). The agents use a simplified Wundt curve (Berlyne, 1971; Saunders, 2001) for the hedonic function, which increases linearly from $(0,0)$ to a point (x,y) , where x varies between 0 and x_{max} and y varies between 0 and y_{max} , and then decreases linearly to the point $(n_{max},0)$ (see figure 5.1), with $n_{max} > x_{max}$. n_{max} represents the highest novelty value for which a non-zero response is given. The point (x,y) is an attribute of a specific agent; x represents the novelty preference of the agent, and y represents its susceptibility to enchantment. As x and y vary, agents can go from a conservative (low x) to a novelty-seeking (high x) preference behaviour, and can be more or less enchantable (varying y). The left side of figure 5.2 depicts agents' styles as points in a space of possible styles, and shows the peak novelty preference (dashed inner circle) and most distant non-zero novelty preference (dark outer circle) for an agent, A. In this case A would give E and F scores of zero, C its maximum score, and B and D scores between zero and the maximum.

Each such reward is added to the singing agent's status. Thus status rewards are accumulated over time by agents, and status represents the accumulation of social gains that result from enchantment. Status values also decay exponentially with time, with coefficient s_{decay} , repre-

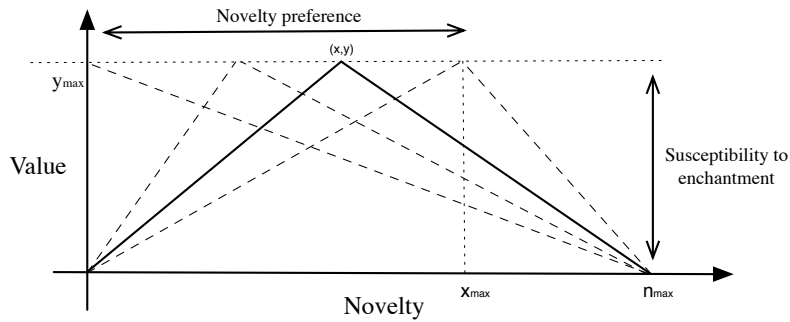


Figure 5.1: A simple linear hedonic function relating novelty (the Euclidian distance between a perceiver's location in style space and the location of the perceived stimulus in style space) to perceived value. A variety of novelty preference values is shown. In these simulations the novelty preference value was set to 2 and the susceptibility to enchantment value was used as a genetic variable.

senting the weakening of bonds in the absence of reinforcement.

The point of defining styles as particles, and setting up a system of cultural interaction operating at a shorter time scale than genetic evolution, is that individual styles change over time. This will mean that inter-individual evaluations will also change over time. Musical styles could be innate, but one of our modelling assumptions is that they stem from an existing dynamical cultural interaction game. As I argued in section 5.2, it is not imminently important that this cultural dynamical system follows the rules of Boyd and Richerson's cultural evolutionary theory precisely. Given the unexplored nature of such a modelling framework, I consider it necessary and sufficient, to begin with, merely to see what happens if we give this system a simple set of rules that drive cultural change.

I have built the dynamics of this game on the combination of two cultural processes: one of creativity (the generation of new behaviours) and one of learning (the convergence of behaviours). In order to model creativity, agents are designed to constantly search for new styles in an undirected fashion by trying out new styles and testing them against their own evaluations. At each time step, an agent generates N new random variations of its current style (the Euclidian distance between original and new style is drawn from an exponential distribution with mean m), and choses the style that scores the highest according to its *own* evaluation. The agent resets its own style to this winning variant. Note that agents use their own styles as reference points

from which to evaluate other agents' styles, and so this change means that the agent is also re-positioning its reference point for future evaluations. For non-zero values of x in the hedonic function described above, this search favours slightly novel styles, and this keeps agent styles constantly on the move in an undirected manner.




Figure 5.2: Agent's evaluation and learning methods. On the left, agent A's evaluation function is shown, non-zero within the dark circle with a maximum value at the point of the dotted circle. Thus agent C receives the highest reward from A, and agents B and D receive similar non-zero rewards. On the right, agent A learns from its set of physical neighbours by adjusting its position towards their centroid by a fraction $L/200$ of the total distance.

In order to model learning, agents are assumed to passively absorb style from their neighbours. Each agent determines the geometric centroid of its complete set of neighbours' styles, and moves its own style a given fraction (determined by a learning strength parameter, L) of the distance towards this centroid (see figure 5.2, right hand side). Thus, as well as agents diverging in this style space due to their own self-evaluating curiosity, this learning process drives agents' styles to converge within their neighbourhoods. The strength of the learning process, the value of L , determines the extent to which groups of neighbours' styles converge, and we are interested in how the results of the model vary with different values of L .

The question also remains as to how the styles of new agents should be initialised in the model. I consider two options: a vertical cultural transmission of styles from parents to their offspring, in which a child's style starts out as an identical copy of its parent's style; and a neutral initialisation, in which children's styles are initialised with all parameters set to zero. The

vertical cultural transition is indistinguishable from genetic inheritance, except that the inherited elements continue to vary culturally.

These behavioural rules define the elementary cultural dynamical system within which agents compete for status, which is the ultimate determinant of their survival. The system established here is highly abstract from the point of view of theories of the evolution of human musical behaviour, but can be analysed and manipulated relatively easily, and this is the main reason I have focused on this representation of behaviour. Any more complex behaviour would have made the system harder to understand and, as it stands, the behaviour of the system is relatively opaque, requiring considerable investigation to clarify. The main point of the design of this aspect of the model is that the behaviour varies as a cultural process with a certain cultural ‘proximity’ between agents who are physically close, but also a certain amount of diversity. Agents evaluate each other’s behaviour relative to their own, and we can therefore say that local cultural behaviour, in particular an agent’s position within an evolving cultural system, has some effect on its, and other agents’ survival.

Figure 5.3 summarises the sequence of events performed for each agent at each time step.

5.5 Modelling investigations

5.5.1 Introducing genetic variables

In this section I define two parameters that are treated as genetically determined aspects of behaviour: the susceptibility to enchantment, S ; and the the number of perceptual parameters, $|P|$.

We can indirectly test the hypothesis that a status-based mediation between cultural behaviour and fitness is sustainable by considering the genetic variation of the parameter which determines the size of the status reward agents are willing to pay to each other. This is the variable y in the hedonic function given above, which is the peak value of the reward an agent is willing to pay out. I termed this value the agent’s *susceptibility to enchantment* (S). If an agent’s S is zero then it never rewards other agents, whereas for any non-zero value it gives a reward

$thresh_{die}$	The fitness threshold below which agents have a given probability (P_{die}) of dying.
P_{die}	The probability of dying at each time step, if an agent's fitness is below $thresh_{die}$.
$thresh_{rep}$	The fitness threshold above which agents have a given probability (P_{rep}) of reproducing.
P_{rep}	The probability of reproducing at each time step, if an agent's fitness is above $thresh_{rep}$.
W	The width (and height) of the toroidal environment (number of cells).
p_{move}	The probability that an agent moves to an adjacent cell at each time step.
p_{mut}	The probability of a mutation during reproduction for any given genetic variable.
$R_{foraging}$	The range within which agents forage for resources.
$e_{extract}$	The maximum energy an agent can extract from a cell.
e_{add}	The amount of energy added to a cell at each time step.
e_{max}	The maximum energy a cell can store.
$e_{deplete}$	The amount of energy agents lose at each time step.
a_{max}	The age at which an agent's fitness is scaled to zero.
$R_{interact}$	The range within which an agent interacts with other agents: its neighbourhood.
D	The dimensionality of 'signals' passed between agents.
s_{decay}	The decay constant of agents' status, multiplied to each agent's status at each time step.

Table 5.1: Model parameters and their descriptions.

based on the perceived novelty of the other agent's style. If, through genetic variation, S tends to decrease throughout the population, then this implies that individuals with lower S are consistently fitter than individuals with higher S . We might expect this to be the case because given two agents living in the same neighbourhood, both rewarding each other, the agent that pays out the higher status reward is likely to be the agent that ends up with the lower relative status, and subsequently the lower fitness. A reasonable strategy in this situation would be not to pay out any status rewards, and therefore not to contribute to any other agent's high fitness. However, actual status rewards are also based on relative evaluations: the Euclidian distances between agents' style parameters. This means that it is possible, in principle, for agents to generally pay out few, or low, rewards, whilst maintaining a high S . It is also possible that subgroups of agents within a population are able to exploit the fact that they can mutually reward status to each other and gain fitness at the expense of other subgroups of lower S agents. In that case, it may be possible for S to increase. Given the uncertainty introduced by the cultural dynamical system, however, it is also possible that no strong selective pressure applies either way to S and it drifts

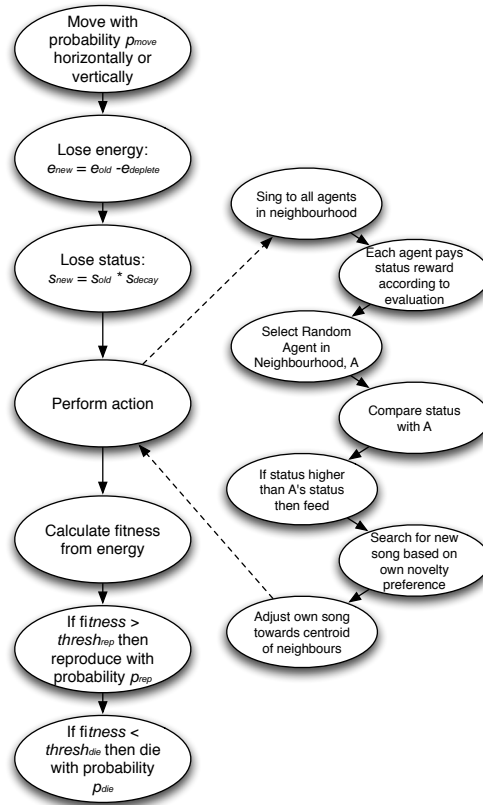


Figure 5.3: Flow diagram of update for each agent. This sequence is repeated for each member of the population at each time step. Changes to agents' states, as well as additions to and removals from the population, are enacted after the entire population has completed this sequence, so that agent updates are synchronous.

randomly.

S is a single value and its genetic variation consists of mutating this value with a probability of p_{mut} when offspring are produced. When mutation occurs, the value is varied by a random number between -1 and 1, with a minimum value of zero and a maximum value of 100.

My main interest lies in studying the evolutionary trajectory of S because it indicates whether or not there is any chance that such a system is sustainable. However, I am also interested in understanding where else we could see the effects of the social determination of fitness operating on an arbitrary cultural dynamical system. I therefore considered how the perception of our agents could evolve in this artificial scenario. As explained in Section 5.3, perception differs from production in that it involves a process of conversion from the observation of a behaviour to a cognitive form. Rather than allowing the agents a perfect representation of their neighbours' style, they were given a limited perceptual capacity: an agent does not perceive all of the D style parameters of another agent, or of itself, but a limited subset. When novelty ratings are made, only the parameters perceivable to the agent are taken into account. The final novelty rating is scaled according to the dimensionality of the space—*i.e.*, if the agent can perceive d parameters, its novelty rating is divided by \sqrt{d} , otherwise stimuli would tend to appear more novel to higher- d agents than to lower- d agents. \sqrt{d} is the maximum distance between two points in a unit D -dimensional hypercube, so scaling by \sqrt{d} scales this maximum back to 1. This limited perceptual capacity affects an agent's evaluation of other agents, its evaluation of itself when deciding new behaviours, and its determination of other agent's behaviours when converging with its neighbourhood. However, in the case of vertical transmission of styles, vertical inheritance of behaviours from parents to offspring is complete; all of the parameters are copied, assuming that the parent-offspring relationship of a stronger nature than other individual relations.

This factor, which I refer to as the agent's set of perceptual parameters P (where $|P|$ is the number of perceptual parameters), was also defined as a genetic attribute of the agents. The set of perceptual parameters is represented by a list of ten binary digits, each of which determines whether or not a certain style parameter is perceivable to that agent. This set is mutated during reproduction with the same probability as with S (p_{mut}). If a mutation occurs, one of the binary digits is flipped.

I predicted that increasingly large sets of perceptual parameters would be selected for in the model, according to the following reasoning. Consider two agents, A and B, one of which, A, has an extra perceptual parameter, p_a , and both of which start off with similar styles. If the agents' styles drift apart along the direction of the axis of p_a , then B will not perceive this

difference, whereas A will. If this distance increases beyond the domain of A's novelty function, then B will no longer be receiving status rewards from A, even though B still might be giving status rewards to A. Also, in general B will fail to converge with a group of agents similar to A, because B will only learn along the axes of all of the perceptual parameters except p_a . This suggests that higher $|P|$ agents are likely to become fitter overall than lower $|P|$ agents and that over time the population will tend towards higher $|P|$. Again, however, since the potential for different social configurations is reasonably broad and unpredictable, we cannot be sure that this process will occur.

An increase in $|P|$ implies a way in which social behaviour could affect specific aspects of perceptual behaviour in a runaway manner. It is a relatively trivial observation that more detailed perception is most likely better for fitness in general interactions with one's environment. The process that I consider here is only slightly different in that what matters is not that the mode of perception is determined by the accurate acquisition of information from the environment, but that modes of perception are correlated between individuals interacting in the same cultural system. In other words, it allows the creative emergence of perceptual capacities through games of social interaction. This provides a way of looking at aspects of musical perception such as the perception of rhythm. Through beat induction we actively cast rhythmic patterns of sound onto a perceptual frame which allows the extraction of various parameters of information such as expressive timing. As suggested in Chapter 3, our capacity to perceive rhythmic sound in this way could be a consequence of sustained social interactions, in which more complex information extractors are inherently more successful than less complex information extractors. In this model I am only interested in seeing how this works in principle using the elementary notion of perception sketched out in Section 5.3. If the principles of such behaviour can be established, I expect that it will be possible to develop the current model towards more complex and music-specific aspects of perception.

To summarise, given the design of this model, it is predicted that S is likely to decrease in most cases (*i.e.*, as the critical parameters defining the behaviour of the model are varied), because it is more likely to be in an individual's benefit *not* to pay status rewards to other agents. Higher status agents will, on average, gain more energy from the environment, within any given

region, than lower status agents, and the best way to ensure access to energy is not to contribute to the increased status of other agents. In this model, evolution has no direct effect on agents' stylistic behaviour, or on the way that individuals distribute environmental resources; these are completely fixed aspect of the agents' behaviour, assumed to be held in place by other evolutionary pressures not defined in the model. Thus the interaction game and the status game are both part of the environmental context in which natural selection works on S . Following this reasoning, S could only increase if the act of paying status rewards was closely correlated with the act of receiving them. Put simply, under some conditions interaction can be cooperative, whist under other conditions it can be competitive. If competitive, then S should decrease. If cooperative, it should increase. However, the environment provides its resources at a fixed rate, and nothing that agents do affects the amount of energy that they can extract. Cooperation is only beneficial due to the existence of agents not participating in the cooperation. Therefore, overall, cooperative behaviour does not implicitly involve mutual benefit; it is necessary for there to be individuals who don't do so well out of the interaction due to the context. Increasing S would therefore show that the model had provided the conditions under which agents could mutually reward each other's status in a structured way.

It is also predicted that $|P|$ is likely to increase. At any point in time, the interaction game between agents is being played within the context of particular modes of perception. There may be many perceptual differences between agents within a population, although we expect the population to remain reasonably closely converged. These modes of perception affect the ways in which agents obey the cultural dynamical rules of the system; they affect agents' creativity (agents are only able to search for preferable new styles according to their own perception of their own style), and learning (agents are only able to learn from those aspects of their neighbours' styles that they are able to perceive). They also affect the evaluations that agents make of each other's styles; additional perceptual parameters provide new dimensions along which agents can perceive difference. This additional difference is likely to mean that agents with higher $|P|$ are more likely to perceive larger differences than agents with lower $|P|$, and sufficiently large perceived differences result, according to the hedonic function, in zero rewards.

5.6 Summary

A summary of the model's design features is as follows:

- The model spans genetic and subgenetic timescales. It is therefore based on a system of food and energy and agents with fixed survival capacities, so that a given population size can be maintained even though the process of birth and reproduction is probabilistic.
- The model does not span the lower timescale of musical action. A single musical interaction is treated as an instantaneous production of a stimulus which is also simultaneously 'consumed' by the recipient agent.
- An agent's access to food is determined entirely by its social interactions. Status is determined by musical evaluation, and relative status determines access to food in a probabilistic manner. Status is accumulated, and fades out over time. Individuals pay status rewards according to both their innate susceptibility to enchantment, and their evaluation of musical stimuli.
- Agent's behaviours, or styles, are understood as being adaptive under two forces: a convergence force, attributed to learning; and a divergence force, attributed to a bias towards novelty, which also determines how agent's evaluate each other.
- Perception is incomplete, meaning that variations in how agents perceive stimuli are possible and subject to evolutionary pressures. The strength of perception affects how agents evaluate and learn.
- Agents can either acquire their parents styles at birth or start out with a default style, depending on the vertical transmission-type parameter of model. The difference corresponds

to a switch in the strength of vertical transmission of style.

- There is a great deal of uncertainty about what constitutes successful behaviour. Although it would not seem sensible to pay out high status rewards at all, it may be possible to accumulate high status in small groups that mutually reward high status.

Chapter 6

Experiments

In the previous chapter, I described a computer simulation model which attempts to set up an artificial society in which the basic notions of cultural imitation and evaluation are established in an evolutionary context. This simulation model does not aim to model exactly the processes discussed in Chapter 3 and a number of simplifying decisions had to be made in order to arrive at a complete, runnable model, rather than allowing the demands of accuracy to drive an increasingly complex and intractable model. Under the influence of the handicap principle and the central problems of selfish gene theory, the main question these models can aim to answer is how it can be evolutionarily possible for individuals to reward musical behaviour (in turn providing an explanation for why it is good to make music).

In this chapter, I begin with a description of the process of calibrating aspects of the simulation model in order to establish suitable working parameters for experimentation. I then present the main experiment involving the evolution of susceptibility to enchantment, S , and perceptual parameters, P , as described at the end of Chapter 5. The experiment shows that it is possible for S to increase under certain conditions and therefore that, in principle, these rules of social interaction allow the evolution of a preference for musical interaction. I discuss possible explanations for this result and observe that the effect is essentially described by the rules of kin selection. This was not an expected outcome of the model and in many ways highlights deficiencies in the model's design. However, it also shows how we can define a complex model in a new domain and go on to observe well-known evolutionary phenomena, and therefore to analyse the model in these terms. I then consider two extensions of the main experiment which introduce

other genetic evolutionary parameters into the model. These fail to produce conclusive outcomes concerning the coevolution of the resulting set of parameters, but in doing so illustrate the uncertainty of selective pressures in the given context with potentially interesting implications in further work.

6.1 Experimental set up

The model design established in the previous chapter has a large number of specifiable parameters. Since the population size is variable, and itself dependent on various parameters, I calibrated the model first by searching for parameters that would allow a 30×30 grid to host 100 agents, with the agents' maximum age, a_{max} , set to 100. I set the cells' maximum energy, e_{max} , to 1, and agents' energy extraction rates, $e_{extract}$, agents' depletion rates, $e_{deplete}$, and cells' gain rates, e_{add} , all to 0.1. Since agents feed from every cell in their foraging region, each agent is capable of gaining more than 0.1 units at each feeding. I also set $P_{die} = P_{rep} = 0.1$, and $thresh_r$ to be twice the size of $thresh_d$, since an agent's energy is shared equally with its offspring. With $thresh_{die} = 5$ (therefore $thresh_{rep} = 10$) and $R_f = 3$, the population tended to maintain a size of approximately 100. This number was found by trial and error. It was not greatly affected by the other parameters in the simulation, although in all cases populations did also have a certain chance of randomly dying out as a consequence of the nature of this model design. These parameters are summed up in Table 6.1.

I ran simulations for 1,000,000 time steps with the initial value of S for all agents set to zero, and with all agents having equal sets of perceptual parameters P , with $|P| = 1$. Simulations started with an initial population of size 10, distributed randomly in the environment (generally this rapidly increased to around 100). Each agent was initialised with a random initial style, each element of which was chosen uniformly from the range $[0,1]$. Each simulation was repeated 10 times with different random numbers.

$thresh_{die}$	5
P_{die}	0.1
$thresh_{rep}$	10
P_{rep}	0.1
W	30
p_{move}	0.1
$R_{foraging}$	3
$e_{extract}$	0.1
e_{add}	0.05
e_{max}	1
$e_{deplete}$	0.1
a_{max}	100
$R_{interact}$	3
D	10
s_{decay}	0.99

Table 6.1: Default settings for model parameters.

6.2 Experiment 1: Initial study of $|P|$ and S evolution

The aim of these experiments was to search the remaining parameter space for conditions under which $|P|$ and S increased. The remaining set of simulation parameters consisted of x , $R_{interact}$, N , m , D , p_{move} , L , s_{decay} and p_{mut} . A complete search of this parameter space was not attempted. I report results from a comparison of values for L , having randomly sampled the parameter space to find a region in which both $|P|$ and S rise and fall. The values for the other parameters are shown in table 6.1. Initially, the vertical cultural transmission model was used, meaning that new agents inherited their parent's styles. To reiterate, L is the parameter discussed in section 5.4.5 which determines the extent to which agents converge with other agents in their neighbourhood (recall from Section 5.5.1 that this convergence is only along the stylistic axes that each agent is aware of). I compared final average S and $|P|$ values across 10 runs for models with learning strengths of 0, 0.1, 0.2, 0.3, 0.4 and 0.5, the results of which are shown in figures 6.1 and 6.2.

$|P|$, which has a maximum possible value of 10, is strongly affected by the value of L (Figure 6.1). Without any learning the final average value of $|P|$ remains very low. However,

as L increases the final average value of $|P|$ for the population tends towards the maximum value of 10, showing a clear positive effect of learning on the evolutionary pressure for $|P|$ to increase. Remember that learning in this context simply refers to the tendency for neighbours to converge on what each agent perceives to be the centroid of their neighbours' respective styles. This supports the argument that the effect predicted in section 5.5.1 takes place, as long as the force of stylistic convergence is sufficiently strong.




Figure 6.1: Final average $|P|$ values for several runs of the model with different learning strengths.

The effect of L on the final average value of S tends in the opposite direction. With lower L , S reaches a high value, whilst for higher L , this value falls. The increase in S appears to have a maximum around $L = 0.1$. Although the graph shows average S results for this value of L as being tightly clustered towards the upper limit, we cannot be sure that the positive results for increasing S are not merely a consequence of a reduced selection pressure, and subsequent random drift. This is possible because, if agent interactions turn out to have a more erratic, random nature, then the value of S may have little effect on the actual success of individual agents, and average S would only embark on a random walk between its maximum and minimum values.

However, in that case we would expect the range of final average S values to vary quite considerably across different runs of the same simulation with different random number seeds. The simulation was run 100 times as far as 100,000 time steps and the distribution of final average S values for each run was compared using a student's t-test to the distribution of final average values of a dummy genetic variable (following the same rules of inheritance and mutation as S but not having any effect on behaviour). Final average S values were significantly higher than the dummy values ($t = 21.8607$, $d.f = 199$, $P < 2.2e^{-16}$).

Figure 6.3 shows average S values over the first 100,000 time steps for all of the runs of this simulation, also showing that average S is always steadily increasing for all runs, without exhibiting random drift. Figure 6.4 shows the complete set of S values held by one population as it varies over the first 100,000 time steps for one of these runs, also showing the variation in the population as S increases. An occasional pattern is visible in which the population forks into one large cluster of increasing S and a smaller cluster of decreasing S , usually with the decreasing S population soon dying out. From these graphs it is not possible to see the local distribution of the sub-populations which survive and those which die out, but it is certain from the rules of genetic inheritance and mutation that dissimilarity in S between agents indicates genetic distance. A possible process underlying this evolution is that sub-populations with higher S became interspersed with sub-populations with lower S after achieving high status. Thus, the random movement in physical space may drive a continual transition from a state where close kin are bunched together driving high status amongst each other in high S groups, to a state where these higher S agents can use their increased status to gain greater food access from the environment than their non-related neighbours. Figure 6.5 shows a detail of the style space in one dimension, with different agents drawn as lines and differentiated by colour. The graph shows that agents form into clusters in style space. Clustering appears to rest largely on new agents staying in the cluster that they are born into, but sudden shifts from one cluster to another can also be observed (*e.g.* around $t = 19,200$). I suggest that this is likely to be because an agent has randomly wandered into a new region of the physical space and rapidly goes about adapting to the style of its new neighbours, following the rules for style learning.




Figure 6.2: Final average S values for several runs of the model with different learning strengths.




Figure 6.3: Average S values for several runs with a learning strength of 0.1.




Figure 6.4: S values for one single run with a learning strength of 0.1.

6.2.1 Further exploration of increasing S

The continually increasing trends shown in Figures 6.3 and 6.4, and the results of the t-test comparing S with a dummy variable suggest that, in this particular case, increasing S is actually the result of a selection pressure favouring higher S individuals and not the result of a process of random drift. With higher L , S performs less well, possibly only comparably to random drift. Initially, it was expected that all else being equal, higher S individuals are more likely to end up with lower status and so be less successful than lower S individuals, which is inherent in the design of the interaction game and the status game. However, many aspects of the agents' contexts are not equal, since each agent evaluates other agents with a form of cultural subjectivity. Thus, it is conceivable that high S individuals are successful on the condition that subgroups of high S individuals can mutually reward high status to each other without having to reward it to lower S individuals. This can be done if higher S individuals can create distinct, physically local style groups within which status is rewarded. We may therefore be able to attribute this increasing S to the fact that individuals inheriting their parents' styles (the vertical transmission




Figure 6.5: Detail of one dimension of the agent style space during a run with $L = 0.1$ and vertical transmission between $t = 19,000$ and $t = 20,000$. Different individuals are drawn as lines with different colours for clarity. The graph shows the grouping of agents in the style space and occasional sudden leaps to different parts of the style space.

of style) provides a way for local kin groups to share group-specific styles that are separate from other individuals in their neighbourhood. As L increases, the convergence of styles across whole neighbourhoods is potentially greater, and local kin groups would no longer be able to isolate their styles in this way.

Thus, once a high S individual has one child, it is possible that the two individuals may gain status through mutual interactions, causing them to win status games with other individuals in the neighbourhood, as long as their styles remain distinct from others in the neighbourhood. It may be the case that within this subgroup lower S individuals actually do better, but if the subgroup has a higher S than those around it, S could still tend to increase. Furthermore, if agents' styles are diverging rather than converging, then it is plausible that the majority of such status interaction occurred during the earliest interactions of any newly created individual, when their style is still similar to their parents.

With cultural transmission turned off, rather than children directly inheriting their parents' style they started life with all style values set to a default start state of zero. This had a negative effect on the rise of S compared to the vertical transmission model, but S still showed some increase for a learning strength of 0.1 (see Figure 6.6). The failure for S to increase in the majority of these cases supports the notion that agents were only giving status rewards to close kin to whom they were stylistically close in the vertical-transmission model. However, the fact that S still increased with $L = 0.1$ suggests that it is still possible for local groups to converge in such a way that they mutually reward each other whilst playing status games with more distant neighbours. Note that agents don't need to explicitly know who their kin are in order for a kin selection process to take place. As long as conditions are in place for interactions between kin to have different properties from interactions between nonkin, an evolutionary effect can occur. Physical proximity in a geographically dispersed population can be a reliable, although by no means definite, indicator, as discussed by Dawkins (1976).




Figure 6.6: Final average S values for several runs of the model with no vertical transmission over different learning strengths.

I suggested above that kin selection may be occurring for low L in the vertical transmission

model because at the moment offspring were produced they were stylistically similar to their parents, but not too similar to other members of their neighbourhood. This hypothesis can be supported by comparing rewards paid between kin and between non-kin. I firstly compared payments only to parents with payments to all individuals. Since agents are definitely in the neighbourhood of their parents at the beginning of their lives, I also looked at how payments were distributed over age. I looked at their total status payments averaged for each age group, as well as total status payments *only* towards parents. I normalised these values by dividing by S for each agent in order to compare low and high S agents on the same scale. In both the non-vertical transmission model and the vertical transmission model, agents paid out greater status during their first few time steps than later in life. However, in the vertical transmission model, a large proportion of this was to the agent's parent, whilst in the non-vertical transmission model, the payment to the parent was not a significant proportion of the total pay-out. In figure 6.7 the average total pay-out is shown for each age group across all individuals in all runs. The two dotted lines show the pay-out in the vertical transmission case ($L = 0.1$), the upper line showing the average of each individual's total pay-out, and the lower line showing the average of the total paid only to parents. The two solid lines show the non-vertical transmission case. The sharp peak in the lower dotted line confirms that in the vertical transmission model agents were instantly rewarding kin due to inherited stylistic proximity. These results do not tell us whether other kin interactions are taking place in the non-vertical transmission model.

Figure 6.8 shows payouts from agents to various categories of kin in ten runs of the simulation with vertical transmission and $L = 0.1$ for 100,000 time steps, with the whole population sampled every 1000 time steps. The vertical axis shows average payment strengths from individuals to each category. For each agent, absolute payments to each category were divided by the payer's S value and by the number of payments made to each category). Using t-tests to compare groups, the data shows that payments to parents were significantly higher than payments to the whole population ($t = 19.8606$, $df = 13.592$, $p = 1.964e^{-11}$), as were payments to children ($t = 21.2318$, $df = 13.197$, $p = 1.372e^{-11}$), grandparents ($t = 5.9856$, $df = 17.906$, $p = 1.186e^{-05}$) and grandchildren ($t = 6.8949$, $df = 14.266$, $p = 6.679e^{-06}$), but not to siblings ($t = -0.0205$, $df = 16.131$, $p = 0.984$). This demonstrates that payments between

immediately related individuals, parents and children, constitute the dominant form of status payment, followed by payments between grandparents and grandchildren, and then more distantly related kin and other individuals, which are not significantly different from each other. Again, this supports the argument that kin payments are happening at the immediate moment of the birth of new individuals. The fact that grandparents and grandchildren share higher rewards than siblings is interesting. If we assume that the highest payments are made at the beginning of agents' lives, then this would imply that new children are immediately having children, in the following time-step (establishing a grandparent-grandchild pair with similar location and style), but that parents are not having children in successive time-steps (establishing a sibling pair with similar location and style). This observation has yet to be explained, but it may also be that the relationship between grandparent and grandchild is shorter lived, exaggerating the effect of this window of opportunity for high payments (but not necessarily contributing to increased reproductive fitness).

Finally, a new simulation with vertical transmission and $L = 0.1$ was run with payments to kin (all of the categories listed in Figure 6.8) suppressed, and a significant increase in S was no longer observed ($t = -0.268$, $df = 17.979$, $p = 0.7917$). The average number of musical interactions for each agent at each time step was roughly halved (from 6.7783 to 4.1096), so the change in result could be attributed to the effects of this general loss of interaction. However, this constitutes further evidence that kin selection is responsible for the increase in S .

6.2.2 Discussion

These models demonstrate ways in which the social determination of fitness can drive the evolution of aspects of cultural interaction. By producing effects that we were not entirely expecting, but that can be understood in terms of standard evolutionary theory, they also point to the use of models to contribute to the formulation of evolutionary theories in domains such as the evolution of human musical behaviour.

An increase in S is argued to suggest that genetic evolution can drive the consolidation of a cultural behaviour. This phenomenon was strongest in the case in which children learned their parents' styles through vertical transmission, in which case a possible explanation is that closely




Figure 6.7: Average total pay-outs made by agents mapped against age in the vertical transmission case (dotted lines) and non-vertical transmission case (solid lines). The lower of each pair of lines shows the average of the pay-outs made only to parents.




Figure 6.8: Averaged pay-outs made by agents to various different categories of kin during the vertical transmission model. The results were gathered by sampling payments made by all members of the population each 1000 timesteps. The payment values are calculated for each agent and each category by dividing the average reward paid to that category, divided by the agent's S value. This gives the payout to that category as a fraction of the agent's maximum payment ability. These payments are then averaged across the population for each category. The graphs shows that payments between parents and children are significantly higher than between other categories of kin and non-kin and that payments between grandparents and grandchildren are significantly higher than the average.

related, high- S groups could reward each other in interaction games whilst maintaining a safe stylistic distance from others with whom they play status games. The suggestion that kin pay

greater rewards to each other is supported by the evidence here, but the suggestion that they maintain stylistic groups is not supported. Instead, it was observed that kin make these high payments during the first time steps of an agent's life, at which time agent styles are still similar due to the vertical transmission of style, as shown in Figure 6.7. When agents did not inherit their parents' styles, there was a negative effect on increasing S .

To this extent, this seems like an unnecessarily elaborate way to arrive at an existing kin selection model. However, the aim was to consider style dynamics not because they would necessarily produce the required effect of increasing S , but because social learning and creativity are sensible starting points for modelling the evolution of human musical behaviour. In searching for situations in which S increases in these models, kin selection provides one possible driver for that increase. If musical behaviour did emerge from Boyd-and-Richerson style social learning processes, then this supports the intuitive notion that kin selection would have had a powerful role to play in driving the emergence of S .

One criticism of this model is that it requires us to accept quite an extreme assumption: that whilst these very specific evolutionary processes take place, all other aspects of the model remain fixed. In particular, an obvious question is: Why *would* agents be such victims of their own cultural behaviour that status games based on arbitrary cultural interactions could dominate the determination of their fitness? This scenario seems to contradict familiar trends in the evolution of social behaviour.

More specifically, if kin selection does drive increasing S , why should status games be a valid way for individuals to distribute resources beyond close kin? Surely if kin groups were exploiting a well established system of mutual trust, signalling or exchange, then defences against this exploitation would evolve rapidly. This may, of course, be the case. However, the critical questions are what this system is, how stable it is, and how it operates. By analogy with sexual selection, and following Krebs and Dawkins (1984), some behavioural adaptations in individuals exploit biases in other individuals' perceptual preferences. We would expect defences against this exploitation to emerge, but they may not eradicate it entirely. In a cultural context, the behaviour in question exploits existing rules of interaction. Defences may emerge against this exploitation, but if the original rules had sufficient advantages, then the exploitation may

also not be eradicated.

Even so, what *is* the quality that I have called status in this model, how can individuals generate it for free, and how do other individuals know it's there? An important feature of status is that it costs nothing to give (instead there are indirect costs in the gains in food that other individuals make), but confers real fitness advantages to other individuals. This makes for a surprising phenomenon in evolutionary theory, but my assertion is that real evolutionary systems do not need to exhibit such clear cut and extreme behaviour to take slight advantage of a slightly maladaptive scenario. A literal interpretation of our model's behaviour would be that enchanted individuals display excitement about the qualities of the singer, which others naïvely accept (the most fitting analogy would be a 'clap-o-meter' in a children's game show). However, a more elaborate interpretation would be that an individual's access to resources within a neighbourhood is a function of the number of supporters that individual has as well as the strength of their support. We could therefore interpret status in terms of cohesion, where an agent's S represents its level of cohesion with those who are stylistically similar (with vertical transmission being one driver for this similarity), and status represents the cohesive support available to a given individual, which comes into effect when they compete for resources with random opponents. In this case, the model offers a fairly well-developed explanation of how cohesive behaviour can emerge through what are essentially socially competitive interactions between individuals engaged in an arbitrary game. The principle of status and its potential association with a maladaptive behaviour may be a precondition for this evolutionary processes, but the evolutionary processes itself may be able to give rise to a cohesive system which emerges directly from competitive interaction in a social domain.

6.3 Experiment 2: The coevolution of S , P and p_{sing}

6.3.1 Rationale

One of the questions to emerge from these initial experiments is whether or not the system that drives the increase of S could also actually sustain the conditions under which S did actually increase. What would happen if the status game was not enforced as a fixed rule of the model, but was also considered an aspect of the model that could vary in response to genetically determined

aspects of individual behaviour?

For a second experiment, I therefore considered the effects of a new parameter, referred to as the singing probability (p_{sing}). p_{sing} determines the probability with which an agent chooses to engage in singing behaviour. If the agent doesn't engage in singing behaviour then it gets to feed directly. If the agent does engage in singing behaviour then it plays a status game with another individual as in the case above. In this model I assume engaging in musical behaviour to also include the processes of accumulating status from other individuals and varying one's musical style according to the same rules described above of learning and creativity. If an agent does not engage in musical creativity it will not engage in any of these behaviours either.

In the previous set of experiments p_{sing} (expressed as a percentage) can be considered as having been fixed at 100 and not allowed to vary. In this set of experiments, p_{sing} was initialised at zero and varied genetically in exactly the same way as S (see Section 5.5.1). I continued to allow S and $|P|$ to vary in these simulations.

I varied the ratio between the amount of feeding an agent could do following a status game interaction, and during a normal feeding session (*i.e.*, when the agent didn't choose to engage in musical behaviour). If the ratio was left at 1:1 then an agent would either feed, or chose a behaviour that resulted in possibly feeding, with equal probability.

6.3.2 Predictions

Under these circumstances it is reasonable to assume that the best behaviour for any agent would be never to engage in musical behaviour, and so p_{sing} would always evolve towards zero. However, since the cultural dynamical activities of learning and creativity are also only active if an agent takes the route of engaging in musical behaviour, then it is not completely clear that this will be the case. Although a completely non-musical agent will have maximum fitness, in the case of two agents with similar non-zero values of p_{sing} , it is possible that the agent with higher p_{sing} does more learning and gains more status. Since lower p_{sing} agents are still susceptible to enchantment, they are capable of boosting the statuses of high p_{sing} agents. It may be possible, therefore, that the correct circumstances can be established in which agents with slightly higher p_{sing} are more successful than agents with slightly lower p_{sing} , because their status is so much

higher that they have a very good chance of gaining food at every time step. Of course, at the same time there may be higher p_{sing} agents who do very badly, but as long as some higher p_{sing} agents do very well then the failure of these others is insignificant, or could even serve their success according to the lottery principle discussed earlier. In cases where a feeding advantage was given to agents winning status games, rather than agents feeding normally, this would additionally favour high p_{sing} individuals, to the direct detriment of all individuals feeding from the same area, and I expected that an increase in p_{sing} would be more feasible.

In this case I varied the learning strength, L , in stages of 0.1 between 0 and 0.5, as before, and also considered the vertical transmission model, the non-vertical transmission model, and neighbourhood ranges, $R_{interact}$, of 3 and 5 (see Table 5.1). In addition, I altered the ratio between access to feeding after winning a status game, and feeding normally. This is determined by a variable, α , which is the increase in food (expressed as a multiple) gained from winning a status game, as compared to eating normally. I compared α values of 1 and 3. P_{sing} was initialised at zero.

6.3.3 Results

I ran these simulations for 1,000,000 time steps, as before. In all cases with $\alpha = 1$, the value of p_{sing} failed to increase. In all cases with $\alpha = 3$, the value of p_{sing} increased significantly. In the latter case, other evolving variables, S and P , then followed a similar course to the results described in experiment 1. Provisionally, therefore, the hypothetical process discussed above, in which higher p_{sing} agents would gain higher status and be able to gain more food than lower p_{sing} agents, did not occur. However, since agents are now engaging in musical behaviour, including searching and learning, at a lower rate than before, we cannot be sure that the cultural dynamics discovered in previous models are still operating in the same way. For example, whilst the rate of agents' movement is determined by p_{move} in the same way as before, agents are performing fewer musical interactions during periods of interaction with other agents as they move around.

This general result, that increasing p_{sing} occurs with $\alpha = 3$ but not with $\alpha = 1$, could be tested by looking at whether this was true for other initial values of p_{sing} . Although p_{sing} was

unable to increase with $\alpha = 1$ here, it may be that higher initial values of p_{sing} would still be sustainable because they were above a threshold value.

In general, this is not a surprising result and supports the hypothesis that the structure of the social system is not capable of creating a context in which p_{sing} remains high.

6.4 Experiment 3: The coevolution of S , P , C and L

6.4.1 Rationale

Having looked at situations in which S and p_{sing} could increase, a number of other fixed parameters determining agent behaviours in the above models could also be set as genetic variables to see how they would vary over time, and how they would affect the basic outcome of increasing S and P . As with experiment 2, the aim of this experiment is to explore the stability of the results from experiment 1 whilst more parameters in the model are allowed to vary. Two additional parameters, L and C , which may affect the outcome of S and $|P|$ were made variable. L is the learning strength discussed previously. C is a new variable describing the extent to which agents possess a conservative preference in their evaluation function. That is, considering the simplified Wundt curve in Figure 5.1, the further the point (x, y) , the peak of the evaluation function, is towards the left of the graph, the greater C is. I defined C as having a value of 100 when $x = 0$ and having a value of zero when $x = 2$. Similarly, the genetic variable determining L was scaled to the range $[0,100]$, with 100 equating to the learning strength of 0.5. Both C and L were treated as genetic variables with the same inheritance properties as S and with a fixed phenotypic effect, *i.e.*, unable to vary during an individual's life-time. In these experiments, p_{sing} was fixed again at 100%, to begin with, with $\alpha = 1$. These simulations determine the extent to which the conditions for the results in the previous simulations are maintained. With all of these values varying, is it possible that S and $|P|$ still increase? Also, will selective pressures drive C and L ?

6.4.2 Predictions

These predictions are based on the assumption that the mutual boosting of status between kin, discussed in experiment 1, was the reason for increases in S . Since interaction games between

close kin can result in the mutual boosting of status, it would presumably be preferable for all individuals in kin groups, regardless of their S values, to reduce the likelihood of styles wandering away from each other, so that the chances of exchanging mutual status rewards is higher. This is most likely when agents are more conservative. But this is true as long as higher C actually does benefit high S groups who are mutually reinforcing each other. It could equally be the case that high C benefits lower S individuals who exploit those groups and cause S to decrease.

Similarly, under such circumstances, increased learning strengths, L , could benefit mutually reinforcing groups, for whom increased convergence is an advantage. Increased convergence could mean greater pressure for the evolution of increased C , as long as this increasingly conservative behaviour favoured close kin more than more distant kin. On the other hand, increased L could benefit those individuals with lower S on the periphery of mutually reinforcing high S groups, thus driving lower S . However, none of these hypothetical effects actually tells us how the combination of S , C , L and P will coevolve.

In the case of both C and L , the problem of explaining the evolution of agent behaviours is that, in this model, we cannot really specify clearly what an agent's goal is. Agents need to maximise their status with respect to other agents, but as we have seen, this is an unachievable goal given the rules of the game; there is no fixed way of securing higher status over one's neighbours. Having low S is superficially a good strategy. Having high L is a good way to gatecrash other agents' kin-based status-enhancing cliques, but that only works if the gatecrashing agents are not missing any of their neighbours' perceptual parameters. Having low L is advantageous for kin groups which form status-enhancing cliques, and which benefit from the non-convergence of styles with other agents. Having more perceptual parameters is generally better, but is associated with possible over-convergence with other agents. Having high C generally increases the chances of rewarding agents that are more similar to oneself, but may also reduce the range of cultural variability in general, which is what drives kin groups apart from other kin groups, and allows the segregation of status-enhancing cliques.

Given these various rationales for the evolution of different forms of behaviour, we cannot sensibly anticipate what the outcome of this coevolving set of variables will be, since the context

in which an agent gains fitness is determined by its social environment, and there are many possible such environments. In particular, it is tempting to talk, as above, in terms of what is good for kin groups. That is, if one expects the hypothesised kin selection process to take effect and to drive the selection of higher S agents, then one might also expect other aspects of behaviour to emerge that reinforce this process. This may happen, but it may be the case that whilst S increases under the pressure of kin selection, other factors are selected for, due to the competition between close kin within the same group that is mutually reinforcing itself.

6.4.3 Simulations

I ran simulations with vertical and non-vertical transmission rules and different initial values of L (0, 0.1, 0.2, 0.3, 0.4 and 0.5) for 1,000,000 time steps. Since C was now being treated as a genetic variable, it is possible that different rates of creative search by agents in their style space may also affect the evolutionary outcomes of the model. Therefore, I also varied a parameter which I refer to as the search range (β). This is the mean of the exponential distribution used to determine the distance between agents' current styles and the new randomly generated styles that they use when looking for new styles as part of the cultural dynamics. In previous experiments, this had been set to 3, given that the length of the range of the Wundt curve is fixed at 3. However, since agents' evaluations are carried out in the space of their perceptual systems (a $|P|$ -dimensional space), but the random generation of styles is carried out in the complete space of styles (a D -dimensional space), then it is possible that an agent's search is in a range that is generally far smaller than its evaluation range, when projected onto the space of its perceptual parameters. The reason for varying β , then, is to explore the effect of different levels of creative exploration (although note that this creative exploration is constrained by agents' self-evaluation).

6.4.4 Results

The results for both the vertical and non-vertical transmission cases (Figures 6.9 and 6.10) show that an increase in S is fragile. The only significant increase in S occurs when $\beta = 0.1$ in the vertical transmission case. In this case, L appears to be less strongly selected for, given that final L values do not vary greatly from initial L values (graph of final L on the left of Figure 6.9), and

the distribution of results for S in this case (graph of final S on the left of Figure 6.9) resembles the distribution of Figure 6.2. That is, high S depends on low L , as before, but low L does not appear to be an evolutionary outcome of these processes, except in the case where L simply isn't affected. The high L outcomes for $\beta = 1$ and $\beta = 10$ suggest that with higher inherent divergence in agents styles (a direct consequence of higher β) comes a higher evolutionary pressure for learning, and therefore convergence (*i.e.*, we can hypothesise that convergence with others is advantageous, but this would still need to be tested). Although, in this study, I do not attempt to prove that this is the case, the outcome flags up the fragility of increasing S , when other evolutionary pressures are brought into the equation. Similarly, $|P|$ increased for higher β , also reflecting a possible dependence on higher L , as in the simulations in experiment 1.

Final values of C varied widely throughout all simulations, demonstrating little selection pressure, but with lower deviation in the case of $\beta = 0.1$. In this case, with vertical transmission, final values of C vary inversely proportionally to final values of L . Thus in the case where $\beta = 0.1$ and the initial value of L was zero, with vertical transmission, C takes a high value, suggesting a freezing up of cultural change (low β , low L and high C all imply smaller changes in style for agents, according to the rules of the model), whereas as the initial value of L increases, C decreases. This may point to the emergence of higher-level qualities of the biocultural system, although a search for a more formal relationship between these values was inconclusive.

It may also be that if the kin-selection hypothesis in experiment 1 is correct and applies in this case, then in the simulation in which $\beta = 0.1$ and the initial value of L was zero, with vertical transmission, the freezing up of cultural change could potentially be interpreted as reducing to a standard kin-selection model where style accurately indicates relatedness. Being a less culturally dynamic system, it may also be that the kin selection process is more reliable and the selection pressure driving an increase in S is stronger.

I do not attempt to prove these suggestions here. This experiment demonstrates that an increase in S or in $|P|$ is not robust in light of adaptive changes in L and C with the current design of the model. This does not undermine the value of the results from equation 1. It would be interesting to search for variations of this model, however, that do allow increases in these

parameters whilst other behavioural values are free to vary, and this would lend support to the central hypothesis of this thesis. These experiments also do not exhaustively explore the dynamical behaviour of the system, which would require thousands of combinations of simulation parameters.

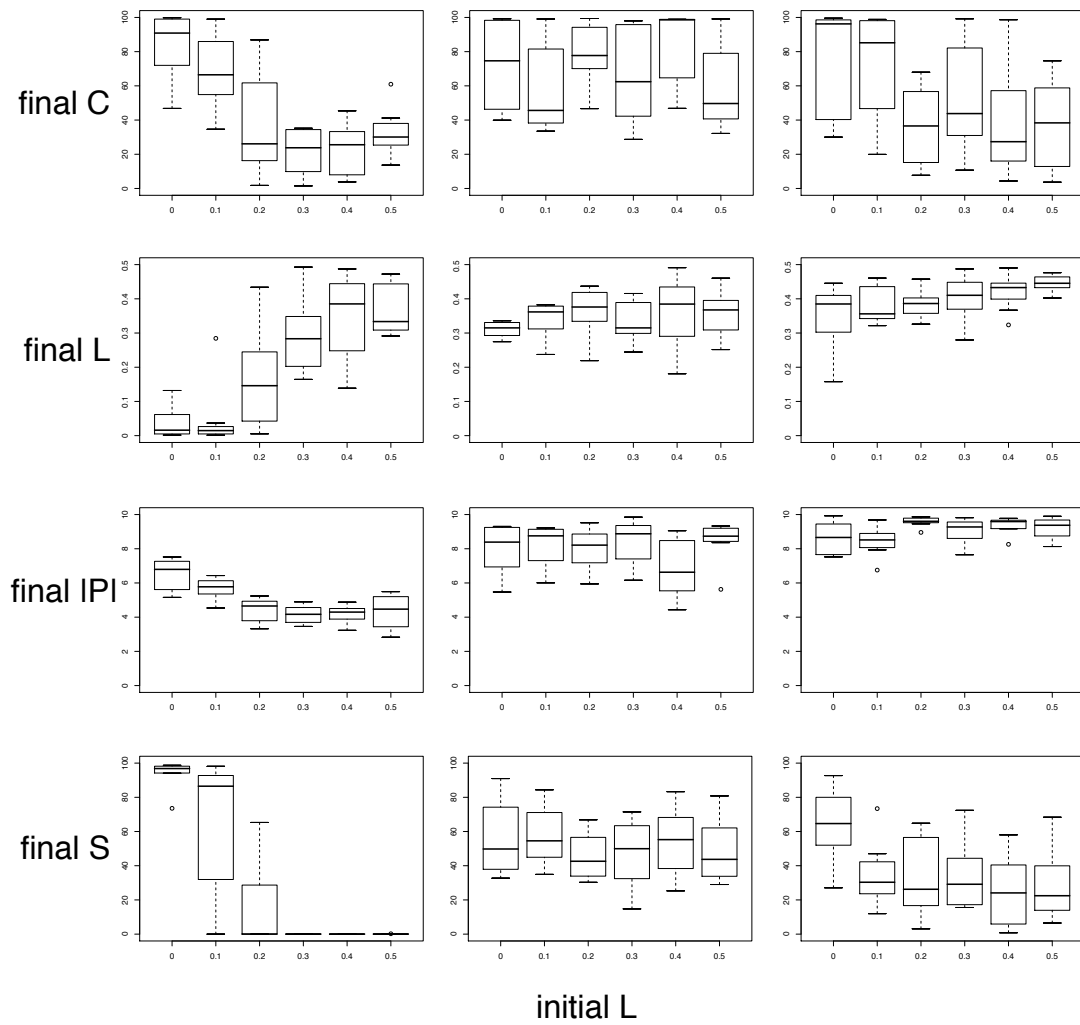


Figure 6.9: Final average values of C , L , $|P|$ and S for sets of ten simulations each with different initial values of L (x-axis) and $\beta = 0.1$ (left), $\beta = 1$ (middle) and $\beta = 10$ (right). These simulations used vertical transmission.

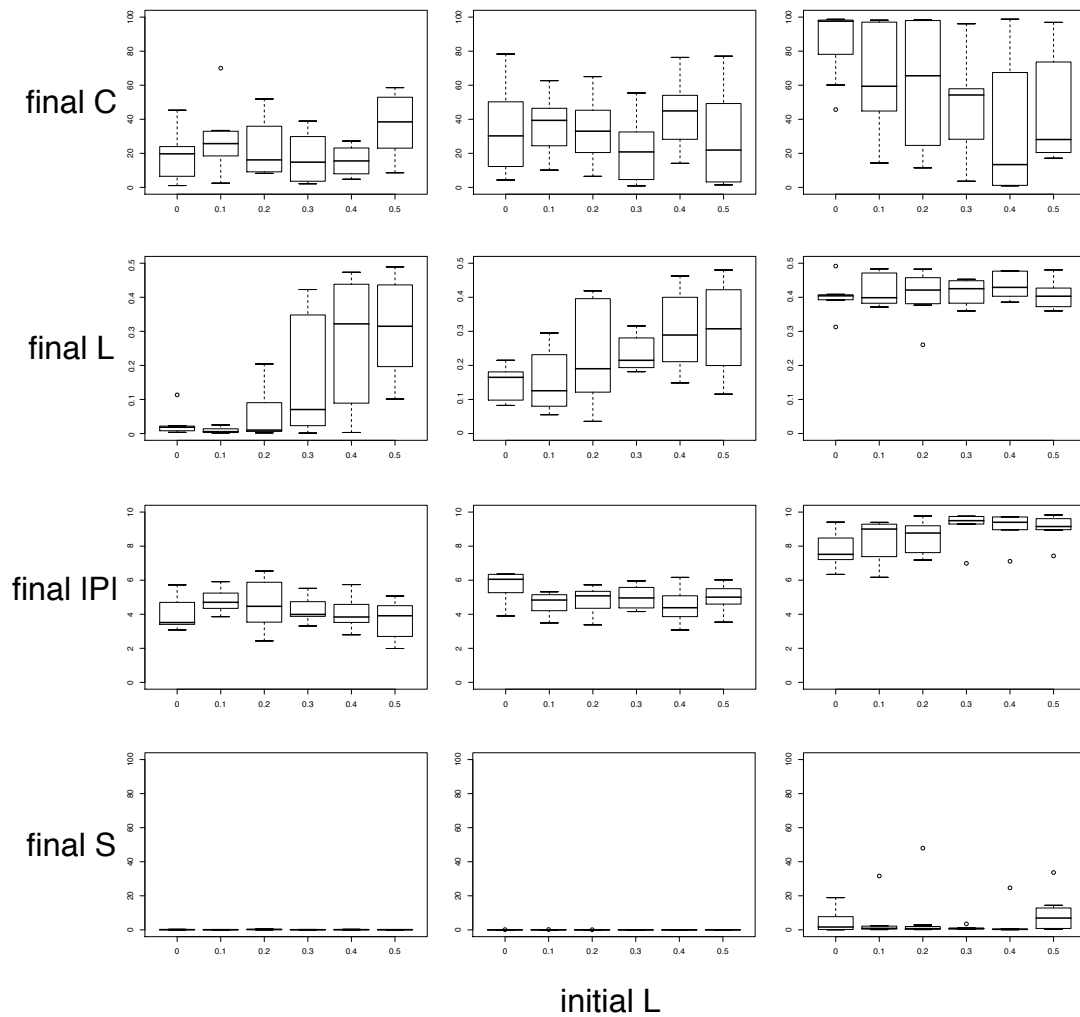


Figure 6.10: Final average values of C , L , $|P|$ and S for sets of ten simulations each with different initial values of L (x-axis) and $\beta = 0.1$ (left), $\beta = 1$ (middle) and $\beta = 10$ (right). These simulations used non-vertical transmission.

6.4.5 Summary of results

These simulations demonstrate an artificial scenario in which, arguably, aspects of musical behaviour can be seen to emerge in a population of interacting agents. Perceptual capacities were shown to increase over time, as long as learning strengths were high enough. The most compelling reason for such behaviour, although it has not been proven in the analysis of the results, follows the process described in Section 5.5.1, upon which the simulation was based; whilst

the less sophisticated perceivers continued to perceive their more sophisticated neighbours as similar and reward them, the more sophisticated perceivers perceived a difference in style which was invisible to the less sophisticated perceivers and so did not reciprocate these rewards. The results suggest that it was necessary for a force driving sufficient convergence of the agents' styles to exist in order for this effect to take place. In these models this force was provided by a tendency to 'learn' the styles of one's neighbours. In reality this learning force was a simple operator that moved each agent's style some fraction towards the centre of their neighbours' styles. Since agents moved their styles only along the axes of parameters they could perceive, agents with more developed perceptual capacities properly converged along all of these axes, whilst agents with less developed perceptual capacities converged less 'successfully', only along the axes that they perceived. This learning process consolidated the greater fitness of agents with higher perceptual capacities, and in simulations with sufficiently high learning strengths, populations would consistently converge on the maximum number of perceptual parameters. This fits with the results shown in Figure 6.1, where increased learning strength increases the effect of rising $|P|$.

The susceptibility to enchantment (S) also increased, or evolved to a non-zero equilibrium, under certain conditions. I expected S to decrease because it was superficially detrimental for individuals to reward high status to others (thus decreasing their own relative status). With the singing probability (P_{sing}) fixed at 100% (experiment 1), S increased for low learning strengths. In the case of vertical transmission of styles from parents to offspring, there was a strong selection pressure for increasing S for learning strengths from zero up to about 0.2. In the case of non-vertical transmission, the increase in S was only moderate. In other cases S converged to zero. Increases in S were argued to be the result of kin selection occurring because kin were able to stylistically converge more successfully with each other than with non-kin, pay each other high status rewards, and then carry their high status to interactions with other agents (hence a long status decay was also implicated in this process). In the case of vertical transmission, this stylistic convergence would be the result of the vertical transmission itself, and indeed, according to the results, a high learning strength served to negate that effect. In general, according to this reasoning, S was able to increase because kin groups could exploit the rules of status

games by rewarding each other with high status only. In this sense status could be interpreted as the level of support one has from other individuals in the population, and status games could be interpreted as situations in which the individual calls on his status (his group of supporters) to exert pressure on others to give him access to resources. Paying high status rewards does not directly diminish an agent's own fitness, but it does increase other agents' statuses, with the result that these other individuals are more likely to win status games, have access to food, and therefore deplete the amount of available energy in the region at each time step. Whether or not this is actually detrimental to individuals depends on other aspects of their organisation.

Evidence for the role of kin selection was provided by showing that the most significant rewards were those paid between parents and children and between grandparents and grandchildren, and a demonstration that S failed to increase significantly when these payments were blocked.

The results from experiments 2 and 3 demonstrated that the observed increase in S from experiment 1 was not stable as the evolutionary degrees of freedom of the agents was expanded. In experiment 2, a strong bias for musical behaviour was needed before agents gambled on playing interaction and status games. In experiment 3, it was shown that by allowing L and C to vary genetically, either L would tend towards a high value, in which case S would fail to increase, or S would increase ($\beta = 0.1$, vertical transmission, initial L equal to zero) but the cultural system would stagnate. Similarly, an increase in $|P|$ was dependent on sufficiently high L .

Chapter 7

Discussion of Results and Conclusions

7.1 Interpretation of simulation results

The aim of the simulations presented in the previous chapter was twofold; firstly, to test, by way of a demonstration, the hypothesis that it is possible for maladaptive cultural interactions based on style learning to drive the emergence of a self-sustaining cultural domain, along with the evolution of aspects of perception exhibiting the essential features of music, and secondly to explore how simulation modelling could be further used in developing our understanding of theories of the evolution of human musical behaviour despite the complexity of this area. In this Chapter, I discuss the results of the simulations in Chapter 6 in light of these two aims. I conclude that the simulations presented in experiment 1 in Chapter 6 do demonstrate a process of evolutionary feedback from a cultural system into biological behaviours that reinforce and redefine this system, but that the results from experiments 2 and 3 fail to support the efficacy of such a phenomenon. I also conclude that these simulations support the general argument of Chapter 4 that it is fruitful to consider complex models of cultural behaviour in order to search for alternative evolutionary processes, but at the same time that these models are problematic in that they suffer simultaneously from being overly reductionist and overly complex.

7.1.1 Summary of conclusions

The broadest conclusion drawn from these experiments is that they demonstrate a situation in which a population of agents increases its innate emotional response to musical interaction without an externally defined fitness advantage for this behaviour; the fitness of the behaviour is

established through social contexts. This is critical to the debate between sexual selection theorists and cohesion theorists because it offers a refutation to the claim that only sexual selection is capable of driving these kinds of evolutionary processes in human behaviour, and it offers a way out of the ‘freeloader’ trap of cohesion theories by building a basis for cohesion in a system of competitive social interactions, itself built on top of other existing systems of cohesion. Under the conditions of the social determination of fitness, it was demonstrated that an arbitrary interaction game can become linked to the process of determining fitness in such a way that subgroups of related individuals can bolster each other’s status by being more susceptible to the manipulation imposed by the game. Although the more ‘selfish’ individuals in the group may be fittest overall, they are less selfish (higher S) than those outside of the group, and S is able to increase.

This evolutionary processes requires the social determination of fitness to be stably established (enforced as a fixed rule in the model), and is sensitive to a number of other parameters. To reiterate the preferred conclusion stated in the previous chapter, it is possible that the increase in S is caused because closer kin can make their styles more closely converged than more distant kin, whilst maintaining the social determination of fitness. Although these are the only possibilities that have been investigated in this model, there is great potential to explore more complex interaction dynamics in this game.

These results do not themselves provide us with evidence that a changing style system should have particularly different evolutionary consequences to completely innate style under a conservative preference rule. In Chapter 3, I discussed two reasons why a changing style system could exist. The first is that a system that required individuals to learn a style provided a suitable handicap to guarantee cohesion (this is the point of view that cohesion came first, but led to competition, both in turn driving runaway evolution). The second is that a runaway cultural evolutionary process following Boyd and Richerson’s (1985) models provides the starting point for the social adoption of a competitive changing style system (this is the view that competition came first, leading to cohesion through runaway evolution). The results from Chapter 6 nevertheless suggest that cultural dynamical systems could also provide a more adaptive rate of fission between groups than would be provided by genetic identifiers, whereas an innate system

would be forced to operate at the same rate as other genetic change. Such a system may be more reliable because it could be more easily calibrated and could operate in domains more easily established by learning than by genetic determination. In complex social environments with multiple layers of group organisation, something beyond our models but possibly not beyond the prehistoric societies we wish to model, culturally dynamic modes of determining group bonds may have provided the flexibility to fit in with other forces determining social organisation, which would have been operating at sub-genetic time scales. Nevertheless, genetic evolution may have naturally taken over some of our stylistic preferences.

In general, it was clear that when allowing other aspects of behaviour to vary, such as p_{sing} , L and C , the conditions for S to increase were less likely to occur. In keeping with Boyd and Richerson's (1985) approach to the problem of biocultural coevolution, this points to the fact that, in order for the required process to occur, we need to posit certain behavioural biases that are externally justified: that are held in place by other factors. According to the current model, we need to regulate the learning strength of individuals in order for S to increase, and we need to enforce the rules of the interaction and status games. These can be treated as biases that needs to be explained in order to support the model.

The model also suggests one way in which aspects of music perception could evolve. In general, once a game is established and sustained in some way, competition between individuals through that game is maintained, and any number of cascading evolutionary effects could take place. As with sexual selection theory, this has the potential to be a creative process, rather than one constrained by an environmental niche, because individual behaviours are shaped by other individual behaviours. A perceptual mechanism could evolve that has no use outside of the game that it evolved to play, and even its own evolutionary history could be obscured behind the dramatic effects that new mode of perception had on the game itself. I believe that rhythmic perception exhibits the requisite characteristics for this kind of adaptation, although if so, it may have undergone this evolutionary development continually exapting existing perceptual mechanisms and developmental strategies. Rhythmic perception is a way of structuring auditory stimuli (particularly significant if the stimulus is indeed of the kind that we call rhythmic) so as to extract particular categories of information. We could therefore argue that musical rhythm is,

from an adaptive point of view, an arbitrary frame of reference within which musical games are played, even if certain aspects of rhythmic perception may have other adaptive origins (London, 2004). As long as the players extract this information in roughly the same way then the game has some basis for evaluation. I do not mean that rhythm should be understood as entirely arbitrary, but that its adoption may have come with a specific evolutionary history which is determined by a mix of neurological, acoustical and mathematical factors. Also, an important criterion for establishing a suitably arbitrary game for social interaction may be that it can be evaluated along lines that are sufficiently different from other sounds in the environment.

7.2 Comparison of simulations with existing modelling approaches

The models presented here propose one major development over the models of Werner and Todd (1997), which is the study of a dual inheritance system. Werner and Todd's model considers musical styles that are innate to the individual, which is clearly not the case in human musical behaviour. Werner and Todd aim to study mechanisms for the emergence of diversity in human behaviour. However, their model generates this diversity using amongst other things, a preference for novelty. In this sense the diversity that emerges in the model can be seen as a direct result of this element of the design. If, instead, the model had aimed to study *how* a cultural preference for novelty had emerged then it would have been more closely focused on an important question in human musical behaviour. The study of biocultural systems is more fruitful primarily because it is simply more appropriate to the problem. Likewise, we cannot expect to learn anything about the evolution of human musical behaviour from models which demonstrate the emergence of innate songs in a system in which the rules of engagement for singing do not vary.

Miranda's (2003) models of the emergence of musical vocabularies and of the relationship between musical signals and emotional states rest on the other side of the gene-culture division and also do not, as such, tell us anything about the role of musical behaviour or its potential evolutionary origins, except by exploring the potential of a Pinkerian explanation (Pinker, 1998). However, iterative cultural processes that demonstrate the fixation of structure in a system of interaction may tell us a lot about certain mechanisms that could be relevant to the evolution of

human musical behaviour, not only for the period following the emergence of innate cognitive capacities, but also before and during this emergence. One of the important phenomena that Miranda demonstrates is that individuals can have quite different internal representations of the same thing, but still communicate successfully. This shows another effect of a changing style system, and suggests alternative approaches to questions of creativity. Individuals with new internal representations of existing phenomena may be inherently creative. Other models by Miranda and Kirby show how musical grammars can grow in expressivity and complexity through iterative learning. Miranda's (2003) models are less problematic than both Werner and Todd's (1997) and the models presented here, because they do not pose questions about how aspects of musical behaviour affect individual survival.

The model in this thesis attempts to ground a far simpler version of the kinds of dynamics present in Miranda's models in an evolutionary situation in which musical behaviour determines fitness. Miranda, Kirby and Todd (2003) propose merging their existing models. This would certainly result in a biocultural coevolutionary system and the question remains as to whether the respective goals of each of these models can be brought together into a single model with clear behaviour. This is never a straightforward task, since as soon as agents' systems of representation and interaction are changed other aspects of the models are inevitably altered. Few modelling paradigms are so flexible that they can be merged without problems. All the same, I support the attempt, since it is necessary that a broad set of possible biocultural models of the evolution of human musical behaviour is explored. In order to understand the evolution of human musical behaviour these models must strive to search for conditions in which musical behaviour becomes an established activity in an environmental niche, in addition to demonstrating ways in which complexity, diversity and a shared repertoire can emerge from the repetition of simple interactions. The models presented in this thesis explicitly attempt to do this, and for this reason are arguably more explicit in addressing critical questions about the evolution of human musical behaviour.

7.2.1 Criticism of these results

Ultimately, the models presented in this thesis do not stand up on a number of grounds as fully developed models, and must be treated as highly exploratory. Implementation decisions such as the use of a Wundt-style novelty evaluation function can be seen partly as placeholders used to complete the construction of the model and not ultimately as committed models of human behaviour (although the notion itself is, of course, based on psychological research (Berlyne, 1971)). Likewise, the implementation of status as a quality that decays exponentially over time has a tenuous basis in any aspect of human behaviour. Similarly, assumptions about how musical behaviour is dispersed throughout the population, and what constitutes musical interaction, are both simplifying and focused on specific theoretical interests. A thorough interpretation of the model can only result in absurd statements if taken literally. It is likely that the increasing S demonstrated in this model cannot be directly applied to any aspect of the real evolution of human musical behaviour. This by no means suggests that such modelling has no value. This is a standard caveat of simulation modelling which needs to be emphasised in this case.

We can begin to address these implementation details individually, and work towards a model that is less susceptible to such clear criticisms of applicability. At the same time, however, in the qualitative domain in which the model was conceived, it may be that the stretching of the model to states where it produces interesting behaviour is more informative. Having established this behaviour in a highly exploratory model, is it possible to get it to apply in other cases?

I discuss ways in which the models could be developed below, but the view presented here still needs defending against the potential criticism that modelling with this methodology will be unsuccessful. A commonly cited view is that attempting to put theoretical notions into programming code forces one to think more rigorously about theoretical details, before simulations have even been run (*e.g.*, Todd, 2000). It is possible that in the case of these models, this initial stage of implementation and its feedback into theory has been more important than the results themselves. Indeed, it should be noted that a number of the ideas presented in Chapter 3 emerged during the preparation of the models ultimately used in Chapters 5 and 6, which had not been explicitly designed from the outset to address changing style systems. Following the approach outlined by Di Paolo et al. (2000), in a new modelling domain, iterations between theory and

implementation are required before a model can approach a stage of being comparable to empirical evidence. This is, in itself, a view that needs to be tested in larger scale contexts and rather than simplifying the model, I argue, it would also be valuable to expand it massively and study its behaviour using a variety of techniques, from statistical analysis to anthropological theory.

7.2.2 Conclusions regarding methodology

At the outset of this project my initial aim was to consider the extent to which cultural behaviour could have effected biological evolution in the evolution of human musical behaviour. A huge proportion of the work in this thesis lay in trying to reconcile the behaviour of simple computer programs with the many rich and varied arguments about the evolution of human musical behaviour. A huge gulf exists between the abstraction of computer simulation models and the complexity of real world processes. Following Di Paolo et al. (2000), the main methodological aim of this thesis was to iterate towards sensible and useable models of the evolution of human musical behaviour by firstly programming an initially apparent set of elements into the model, and exploring the behaviour of this model in order to understand what kind of behavioural and structural characteristics the model possesses. In this way, one can work towards a useable model, and only then is it really possible to begin to make hypotheses and test them by running specific simulations. Finally, the results of the model can be interpreted with respect to the original phenomenon. During the first stage of this process the researcher engages in the apparently unsystematic activity of tweaking a model until it has the best balance of useability and relevance for the problem in question. However, this process can actually be understood as part of the search through a particular (very broad) space of possible models, where the interpretability of the model's behaviour is actually the most pertinent outcome that is being varied.

An artificial life approach to the evolution of human musical behaviour seems to bias exotic explanations of self-organisation and emergence over more traditional Darwinian processes. This could be seen as a mistake made by modellers to favour those processes that generate compelling models. However, following earlier convincing artificial life experiments, the exploration of new processes is indeed explicitly what these models aim to achieve, and this is extremely valuable when Darwinian explanations for any aspect of human or animal behaviour carry such

prominence. Even if it is the single, ultimate force of biological evolution, evolution by natural selection is being associated with increasingly more elaborate and complex processes, which defy simple reduction. Increasingly complex artificial life models are likely to have a long future in elaborating on such cases.

More specifically, in the exploratory phase of this research it was necessary to define behavioural rules at the same time as hunting for specific results such as increasing perceptual complexity or susceptibility to enchantment. One can find such a result without really knowing why, only to discover that during exploration, the design of the model has drifted from the original assumptions. It is possible that the new assumptions no longer make any sense, but it is also possible that one can find new interpretations for the elements in the models that actually propose new hypothetical processes. In the models considered in this thesis, assuming that kin selection is the cause of the evolutionary pressure for increasing S , this cause was not originally proposed as a solution to the evolution of the susceptibility to enchantment, but emerged as a result of vertical transmission and the proximity of kin in the physical space. Thus, computer simulation models can be treated as opaque versions of the intuition pumps called for by Fitch (2006) to generate hypotheses, although we may have to go through multiple iterations scrutinising our models and their assumptions before they are ready to be used as hypothesis generators.

7.3 Proposition of future models

Three ways in which this model could be developed are (a) by exploring other specific rules of social interaction, (b) by exploring other specific production and perception systems, and (c) by introducing aspects of sexual selection and considering comparative studies.

7.3.1 Development of the changing style system models

With respect to point (a), the exploration of other specific rules of social interaction, the most pressing goal is to develop a model based more closely on the processes described in section 3.3. The primary assumption behind this process is that if individuals consider a behaviour valuable, and reward it, and in doing so also learn it, and finally become the focus of future rewards because they possess a similar behaviour, then we have a model for how musical behaviour

emerged at the same time as the actual reward system that sustains it. In provisional work not reported in this thesis I briefly explored the design of such a system. However, the solutions to a number of design questions about how these rules of interaction and reward should work were not apparent at that stage, and the potential number of variations in the design of behaviours was so vast as to be limiting. One of the main questions is how to interpret and model the notion of *access* to an individual for the purpose of learning. We could simply imagine that a popular *tutor* (in previous sections I have used the term *model*) would require higher rewards than a less popular tutor, and that agents would make rational decisions about which tutor to choose based on evaluations of costs and benefits. Although we can begin to design behaviours based on these notions, one must necessarily commit to an algorithm for making such a decision, and fully explore the implications of using that algorithm. Do we make agents as intelligent as possible? Do we make them manageably naïve, or do we try to establish this behaviour itself as something which natural selection might be able to optimise? This choice straddles different boundaries of theory, from signalling theory and game theory, which treats agents as relatively simple, to more human-focused notions of interaction which engage with cognitive complexity. However, a minimal cognition paradigm emerging in artificial life (*e.g.*, Beer, 1996) provides justification for the notion that cognitive complexity can be evidenced in minimal systems and attempts to address this mismatch.

Another question is what drives diversity in this model. In the current models diversity is caused by a curious search through the style space of the agents. This does the job of driving diversity, but can certainly be described as a naïve behaviour on behalf of the agents, which is also not given a functional explanation. One possibility that has also been considered in provisional tests is to allow the novelty preference of agents to vary genetically, in order to see if novelty seeking behaviours are selected for over conservative behaviours. This would superficially justify their use. Another option would be to make such behaviours less naïve by giving high-status agents (in this proposed model) the tendency to innovate new behaviours in order to maintain their distinctiveness, and thus higher status. In such a model it could be possible to allow agents behaviour to vary in more complex ways depending on their own evaluations of their social standing, so that lower status agents are driven to learn, whilst higher status agents are driven

to innovate. The parameters of the mapping between self-perceived status and changing-style behaviour would therefore make interesting genetic variables, possibly producing agents that exhibit this status dependent behaviour under the right circumstances. A final question is what gets paid from learners to tutors. The notion of status introduced in the previous models is satisfying in that agents can make payments in terms of facilitating increased access to resources, rather than directly. In the current models S would have failed to increase if rewards were paid only to the detriment of the rewarder. In this new case it could be that direct payments could support the model.

Having designed and carefully studied the models presented in this thesis, a process that took far longer than expected, the potential for making slight modifications to these models towards these more complex interaction processes now seems more manageable. The simplest step would be to start to look at how the status game itself could become the focus of an evolutionary process, as proposed above. Could we find a situation in which successful high S agents were also more successful as a result of the fact that they and their kin played status games?

7.3.2 Exploration of more complex musical capacities

Point (b) provides a link to a wealth of existing work in music cognition. As soon as the perception and production systems in the present model are altered it is likely that the simulation outcomes will change. How and why will be interesting. However, it is conceivable that more complex systems can be understood in terms of indirect parametric analyses of agents' styles, and so general properties of style dynamics could be studied across a number of models. Existing models of melodic and rhythmic perception could be brought into biocultural models in order to explore the kinds of evolutionary pressures that would be felt when these systems are used in place of the simple geometric pattern dynamics used so far. It is expected that this would be very confusing at first but that through comparison of a number of cases, the generality of the model presented in this thesis could be explored, and also a number of general principles could be studied concerning the types of genetic and cultural changes that are likely to take place in such a system. In practical terms it could be possible to define a template, for example an Interface in the Java programming language, which specifies the kinds of things each system

can do (for example, produce a song, learn from a given set of songs, innovate a new style), and which parameters should be genetically evolvable. In the latter case it may be acceptable to leave a large number of parameters open to determination by genetic evolution and test for their selected-ness as compared to neutral variables using statistical analysis. Models of music perception and production could then easily be incorporated into biocultural evolutionary models in order to explore the effects both on the perception-production system, and on the biocultural system. In particular, these systems would not exhibit the uniform, featureless space that the existing geometric model has, and would therefore define a more complex 'environmental niche' in which agents had to play out their interaction and status games. This challenging task has unfortunately fallen outside of the scope of this thesis, as the concerns of point (a) are more pressing.

In addition to more complex musical capacities, agents could exhibit more complex systems of knowledge, goals, emotions and so on, including notions of shared intentionality (Tomasello, 1999; Cross, 2007a). This would arguably model more accurately the higher social cognitive abilities present in human nature, and would, at the same time break from the paradigm of a minimal domain of interaction between agents driving the evolution of a biocultural system of interaction. The present model is caught between these conflicting interests and neither extreme is particularly attractive. Again, a flexible approach to the interpretation of such models, in terms of notions such as cognition, may be enlightening. For example, the present system arguably models cultural interaction, but using cognitively minimal agents, a clear contravention of the notion that culture relies on advanced social cognition.

7.3.3 Comparative approaches to models of the evolution of human musical behaviour

With respect to the final point, it is not clear whether any ground can be gained by combining two pressures into one model simply in order to see which is the strongest, but it is plausible that with the development of theory in other areas, sensible questions begin to emerge about how these processes could interact. Miller (2000a) has already suggested that a combination of sexual selection and social selection pressures may provide some answers to the problems of sexual se-

lection theory. One of the most important ways in which models could help to address questions about this interaction is by studying the possible cause and effect relationship between different sources of value. For example, in Section 3.7.1 I discussed the example given by Zahavi (1975) of a large house as a wealth indicator. Do we value large houses because they are indicators of wealth, or do we simply value large houses? Whilst Miller's hypothesis is compelling and well thought out, the key question remains as to the strength of this process in humans, including the strength of indicators. We assume that sexually selected traits are those things that we find *instantly* attractive, leading directly to sexual arousal. However, Miller argues that this needn't be the case and that musical capacities would provoke more long term attractiveness. This point should surely have a questionable effect on the importance of indicators. Thus, when Miller suggests that stamina is a valuable quality indicated in displays by a male's ability to dance energetically for a long period of time, one might ask whether it would not be more appropriate for a male to prove his ability by his success at hunting, or in other directly beneficial ways. If the female has time to witness a number of males in action doing something really useful, what good is the indicator? Given enough time to make decisions, it would make more sense for females to look at real phenotypic quality rather than indicators of phenotypic quality. However, if individuals who were good at singing and dancing *actually* gained rewards from their neighbours, and were *actually* directly more capable of providing for their partner and children as a result, then of course the better singers and dancers would be more attractive, but only in the same way that a large house might be attractive without indicating anything. Furthermore, assuming that our ancestors were intelligent, had good memories, and were capable of inferring qualities from behaviour, then could it be that the power of sexual selection for indicator traits would be heavily weakened? Presumably it would be possible to develop models for exploring this kind of relationship and for comparing the case in which sexual selection leads to the social determination of fitness with the case in which the social determination of fitness leads to sexual selection.

7.4 Relevance to theories of the evolution of human musical behaviour

The evolution of human musical behaviour lies shrouded in mystery, and dramatic archaeological finds or radical new techniques for understanding the innate components of musical ability are the most likely ways in which any certain knowledge can be brought to bear on the subject. However, our favourite theories are heavily swayed by suggestive evidence of the efficacy of processes. At one point sexual selection was not taken particularly seriously by evolutionary biologists, and neither was Zahavi's (1975) handicap principle or the Baldwin effect (Hinton and Nowlan, 1987). Even Darwinism itself has had low points since its inception. Alongside good solid evidence, models have helped draw attention to how these processes can work in theory, and have driven the developments of experiments that demonstrate them working in practice. The relationship between aspects of social musical interaction and individual fitness can be tested in modern human behaviour and we can attempt to distinguish cause from effect, including the relationship between sexual selection, runaway cultural evolution, and other contributors to status and sexual attraction. This thesis sets up an extensible modelling environment for looking at biocultural coevolutionary processes in the evolution of human musical behaviour, concluding that runaway biocultural coevolutionary processes are a viable explanation for the existence of music, and points to further directions in which experimental research could go, although the testability of specific hypotheses in sociological or anthropological research ultimately depends on successful communication with practitioners in those disciplines.

7.4.1 Towards alternative runaway selection processes

This thesis investigates the possibility that runaway evolutionary processes are not limited to sexual selection, and that a cohesive role for music can actually be understood in terms of kin selection operating in a particular competitive domain, built on a set of existing social rules. A sexual selection view of the evolution of human musical behaviour has traditionally been set against a cohesion view as a battle between a trusty evolutionary force with little cultural support and a widely culturally observed phenomenon with poor support from established evolutionary theory. I argue that the power of sexual selection needs to be shared amongst other theories.

To make the point bluntly, sexual selection can be viewed simply as a very specific case of the intra-specific determination of the fitness of species-specific behaviours. With the existence of cultural learning, the kinds of runaway evolutionary process exhibited by sexual selection are possible in general social forms.

It was argued in section 3.5 that a cascade of biocultural feedback processes could also have contributed to the evolution of human musical behaviour, with competitive style learning leading to the genetic take-over of some musical behaviours, as well as new, more powerful learning behaviours, pushing the domain of that competitive style learning into new territory. In general, I have proposed that the stable establishment of one domain of cohesive behaviour sets the context for new forms of competitive interaction, defining the context for a runaway evolutionary process applicable equally to evidently competitive behaviours like theory of mind, and to more mysterious modes of interaction such as music. The present models could be argued to show the evolution of one stage of this iteration in the case of musical behaviour. This suggests a basis for the emergence of a role for musical behaviour in establishing cohesion in social systems in which there was a general fluidity of social bonds between individuals of varying genetic relatedness. It also posits a role for vertical transmission in this process and could be brought to bear on the significance of mother-infant interactions in understanding music. To what extent does vertical transmission, as opposed to horizontal transmission, determine musical preferences in modern human societies, and how kin-based are the kind of music-invoked bonds proposed by cohesion theorists?

It may also be possible to seek the kinds of changing-style behaviours discussed here in other primate species, although if human behaviour had already adapted sufficiently for this runaway process to kick in then it would be reasonable not to expect to find it elsewhere. For example, as Knight (1991) argues, humans may have evolved towards new forms of sexual organisation in which alpha male dominance was no longer a successful strategy. One question then is to what extent our understanding of early hominin societies could fit with the assumptions made in these models.

7.4.2 Cultural behaviour as a major transition in evolution

Boyd and Richerson (1985) provide an argument for how culture sets up the basis for group selection, the theory that whole groups of individuals are selected for based on their collective behaviours, rather than individually on the basis of individual differences between behaviours within that group, and others (*e.g.*, Sober and Wilson, 1998) have presented additional theories for how group selection could work and be particularly effective in human cultural contexts. Boyd and Richerson argue that social learning encourages behavioural homogeneity within the group, reducing the effect of inter-individual competition. This makes for an analogy with the emergence of multicellular organisms, in which the division of cells into genetically identical clones diminished the inter-individual competition between these cells and enhanced competition at the level of the multi-cellular organism. Others, such as Dawkins (1976) have argued that group selection is unlikely because selection at the individual level is so much more powerful than selection at the group level, and should always undermine it. The model proposed here is one in which individuals gain success by leveraging the organisation of the group to their benefit. Those individuals foster group unity in the process of benefiting from it, and their offspring ultimately constitute both the winners and losers of future generations, who subsequently inhabit a slightly different behavioural niche. The point is that individual selection and group selection work together, roughly, rather than in opposition to each other. The individuals who selfishly gain from the cohesion of the group are more similar to the individuals who suffer exploitation than they are to the genuine freeloaders. Thus, selfish behaviour not only emerges in the context of a socially cohesive system but is the driving force for future group cohesion, either in the sense suggested by Sober and Wilson (1998), that groups and individuals are units evolving simultaneously in a hierarchical and coupled manner, or in the absence of a notion of group selection, but also in the absence of any notion of a static evolutionary environment.

7.4.3 Relevance to modern understanding of musical and artistic behaviour in society

For many non-scientists, the advance of Darwinism into explanations of the most personal, cherished and lived aspects of human behaviour is threatening, promising to demystify the mystical,

and in doing so to erode its everyday value. I concur that since starting to study this subject, I have taken increasingly to looking at many aspects of human behaviour as ultimately being reducible to some kind of evolutionary explanation, and this is sometimes accompanied by a certain self-awareness with respect to my personal appreciation of art. Darwinism is an addictive mode of thought, but I wish to suggest that as it becomes increasingly accepted by other academic disciplines, it is also capable of becoming increasingly *accepting* of theories of culture that earlier Darwinists had failed to appreciate. Certainly Darwinian explanations should provoke us to think about why we engage in certain behaviours. Ironically, for many it is more reassuring that these behaviours don't depend on a purpose that is defined by some other aspect of human behaviour, but are emergent, or autopoietic. Thus, if an artist claims to be doing art for art's sake, this can be genuinely supported by the argument that certain human behaviours may well be irreducible to biological or social *needs*.

Evolutionary explanations also tend to come hand in hand with generally stark statements about what constitutes human nature, although possibly against the wishes of the theorists themselves. We can be naturally aggressive, naturally peaceful, naturally polygamous, monogamous, egalitarian, hierarchical, and so on. Social anthropologists, on the other hand, have traditionally tended to treat human nature as a blank slate, moulded by cultural experience. According to such a view of the evolution of human musical behaviour, human culture maintains the creative autonomy that has been ascribed to it by social anthropologists, but in a longer term perspective, taking its share of the responsibility for producing anatomically modern humans, including human musical behaviour.

However, arguments that explain the emergence of musical and artistic behaviour as serving a good-for-the-group purpose are not of this category. In this thesis a cohesion system has been observed emerging from competitive behaviour in a social domain. I argue that the domain of interaction is still fundamentally competitive, even if it results in stronger social bonds over larger social groups.

There is always a limit to the notion that music brings people together.

Bibliography

- R. Aiello and J. A. Sloboda, editors. *Musical Perceptions*. Oxford University Press, 1994.
- E. O. Altenmüller. How many music centers are in the brain? In *The Biological Foundations of Music*, Annals of the New York Academy of Sciences, pages 273–280. 2001.
- B. Anderson. *Imagined Communities*. Verso, 1983.
- M. B. Andersson. *Sexual Selection*. Princeton University Press, Princeton, USA, 1994.
- M. A. Arbib. From monkey-like action recognition to human language: An evolutionary framework for neurolinguistics. *Behavioral and Brain Sciences*, 28(2):105–124, 2005.
- M. A. Arbib, editor. *Action to Language Via the Mirror Neuron System*. Cambridge University Press, 2006.
- Aristotle. *Physics*. Creative Commons, 2007. URL <http://etext.library.adelaide.edu.au/a/aristotle/physics/>.
- R. Aunger. *Darwinizing Culture*. Oxford University Press, New York, 2000.
- R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.
- R. Axelrod. The evolution of strategies in the iterated prisoner’s dilemma. In L. Davis, editor, *Genetic Algorithms and Simulated Annealing*, chapter 1, pages 32–41. Morgan Kaufman, San Francisco, CA, USA, 1987.
- R. Axelrod. *The Complexity of Cooperation*. Princeton University Press, Princeton, USA, 1997.

- M. J. Baldwin. A new factor in evolution. *The American Naturalist*, 30(354):441–451, June 1896.
- J. H. Barkow, L. Cosmides, and J. Tooby. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. OUP, New York, 1992.
- A. Barnard. Rules and prohibitions: the form and content of human kinship. In T. Ingold, editor, *Companion Encyclopedia of Anthropology*, pages 783–813. Routledge, 1994.
- A. J. Bateman. *Heredity*, 2:349–368, 1948.
- R. Beer. Toward the evolution of dynamical neural networks for minimally cognitive behavior. In *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 421–429. MIT Press, 1996.
- M. Begon, C. R. Townsend, and J. L. Harper. *Ecology: From Individuals to Ecosystems*. Blackwell Publishers Ltd., 4th edition, 2006.
- D. E. Berlyne. *Aesthetics and Psychobiology*. Appleton-Century-Crofts, 1971.
- D. Bickerton. Can biomusicology learn from language evolution studies??. In N. L. Wallin, B. Merker, and S. Brown, editors, *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2001.
- J. Bispham. Rhythm in music: What is it? who has it? and why? *Music Perception*, 24(2): 123–124, December 2006a.
- J. Bispham. Music means nothing if we don't know what it means - lead review of 'the singing neanderthals' by s. mithen. *Journal of Human Evolution*, 2006b.
- J. Blacking. *Music, Culture and Experience: Selected Papers of John Blacking*. University of Chicago Press, Chicago, 1995.
- S. Blackmore. Review of not by genes alone: How culture transformed human evolution. by peter j. richerson and robert boyd, chicago, university of chicago press, 2005, 332 pp, index, isbn 0-226-71284-2. *Bioscience*, 56:74–75, 2006.

- S. J. Blackmore. *The Meme Machine*. OUP, New York, 1999.
- M. Bloch. A well-disposed social anthropologist's problem with memes. In R. Aunger, editor, *Darwinizing Culture*. Oxford University Press, New York, 2000.
- M. Boden. *The Creative Mind*. George Weidenfeld and Nicholson Ltd, 1990.
- T. L. Bolton. Rhythm. *American Journal of Psychology*, 6(2):145–238, 1894.
- R. Boyd and P. J. Richerson. *Culture and the Evolutionary Process*. University of Chicago Press, Chicago, IL, US, 1985.
- D. Bratton and T. Blackwell. Understanding particle swarms through simplification: a study of recombinant pso. In *Proceedings of the 2007 GECCO conference on Genetic and evolutionary computation*, pages 2621–2628, 2007.
- S. Brown. The “musilanguage” model of music evolution. In N. L. Wallin, B. Merker, and S. Brown, editors, *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2000.
- S. Brown. Contagious heterophony: A new theory about the origins of music. *MusicaeScientiae*, 11(1):3–26, Spring 2007.
- R. Burling. *The Talking Ape*. Oxford University Press, 2005.
- D. Cannadine. The context, performance and meaning of ritual: The british monarchy and the 'invention of tradition', c. 1820-1977. In T. O. R. Eric J. Hobsbawm, editor, *The Invention of Tradition*, pages 101–165. Cambridge University Press, 1992.
- E. Clarke. Categorical rhythm perception: an ecological perspective. In A. Gabrielsson, editor, *Action and Perception in Rhythm and Music: Papers given at a symposium in the Third International Conference on Event Perception and Action*. Royal Swedish Academy of Music, 1987.
- M. Clayton, R. Sager, and U. Will. In time with the music: the concept of entrainment and its significance for ethnomusicology. *European Meetings in Ethnomusicology*, 11:3–75, 2005.

- D. Cliff and G. F. Miller. Co-evolution of pursuit and evasion ii: Simulation methods and results. In J.-A. M. J. P. P. Maes, M. Mataric and S. W. Wilson, editors, *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB96)*, pages 506—515. MIT Press Bradford Books, 1996.
- J. Clifford and G. E. Marcus, editors. *Writing Culture: The Poetics and Politics of Ethnography*. University of California Press, 1986.
- R. Conte. Memes through (social) minds. In R. Aunger, editor, *Darwinizing Culture*. OUP, New York, 2000.
- N. Cook and M. Everist, editors. *Rethinking Music*. Oxford University Press, 1999.
- G. Cooper and L. B. Meyer. *The Rhythmic Structure of Music*. University of Chicago Press, Chicago, 1960.
- T. F. Cooper and C. Ofria. Evolution of stable ecosystems in populations of digital organisms. In Standish, Abbass, and Bedau, editors, *Artificial Life VIII*, pages 227–232. MIT Press, 2002.
- L. Cosmides and J. Tooby. Cognitive adaptations for social exchange. In J. H. Barkow, L. Cosmides, and J. Tooby, editors, *The Adapted Mind*, pages 163–228. OUP, New York, 1992.
- I. Cross. Music, cognition, culture and evolution. *Annals of the New York Academy of Sciences*, 930:28–42, 2001.
- I. Cross. Music as a biocultural phenomenon. *Annals of the New York Academy of Sciences (The Neurosciences and Music)*, 999:106–111, 2003a.
- I. Cross. Music and biocultural evolution. In R. M. M. Clayton, T. Herbert, editor, *The Cultural Study of Music*. Taylor and Francis Books Inc., Oxford, UK, 2003b.
- I. Cross. Entrainment, evolution, language and music. October 2004.
- I. Cross. The origins of music: Some stipulations on theory. *Music Perception*, 24(1):79–82, September 2006.

- I. Cross. Musicality and the human capacity for culture. *MusicaScientiæ*, 2007a.
- I. Cross. Music and cognitive evolution. In L. Barrett and R. Dunbar, editors, *OUP Handbook of Evolutionary Psychology*. Oxford University Press, 2007b.
- I. Cross. Is music the most important thing we ever did? music, development and evolution. In S. W. Yi, editor, *Music Mind and Science*. Seoul National University Press, Seoul, 1999.
- M. Csikszentmihalyi. Implications of a systems perspective for the study of creativity. In R. J. Sternberg, editor, *The Handbook of Creativity*. Cambridge University Press, New York, 1999.
- C. Darwin. *On the origin of species by means of natural selection, or The preservation of favoured races in the struggle for life*. D. Appleton and company, 1860.
- C. Darwin. *The Variation of Animals and Plants Under Domestication*. D. Appleton and company, New York, USA, 1868.
- C. Darwin. *The Descent of Man and Selection in Relation to Sex*. Appleton and Co., New York, USA, 1883.
- R. Dawkins. *The Selfish Gene*. OUP, 1976.
- R. Dawkins. *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe Without Design*. Penguin, London, 1986.
- R. Dawkins and J. R. Krebs. Animal signals: Information or manipulation? In J. R. Krebs and N. B. Davies, editors, *Behavioural Ecology: An Evolutionary Approach*, pages 282–309. Sinauer Associates, 1978.
- R. T. Dean and F. Bailes. Towards a sociobiology of music. *Music Perception*, 24(1):83–84, September 2006.
- D. C. Dennett. *Consciousness Explained*. Penguin, 1993.
- D. C. Dennett. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. Penguin, London, 1996.

- P. Desain and H. Honing. The formation of rhythmic categories and metric priming. *Perception*, 32(3):341–365, 2003.
- E. Di Paolo, J. Noble, and S. Bullock. Simulation models as opaque thought experiments, 2000.
- E. A. Di Paolo. Behaviour coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behaviour*, 8(1):27–48, 2000.
- E. A. Di Paolo. Social coordination and spatial organization: Steps towards the evolution of communication. In P. Husbands and I. Harvey, editors, *Fourth European Conference on Artificial Life*, pages 464–473, Cambridge, MA, USA, 1997. MIT Press.
- J. M. Diamond. *Guns, Germs and Steel: the fates of human societies*. W. W. Norton, London, 1997.
- E. Dissanayake. Antecedents of the temporal arts in early mother-infant interaction. In N. L. Wallin, B. Merker, and S. Brown, editors, *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2000a.
- E. Dissanayake. *Art and Intimacy: How the Arts Began*. University of Washington Press, Seattle, USA, 2000b.
- E. Dissanayake. *What is Art For?* University of Washington Press, Seattle, USA, 1988.
- E. Dissanayake. *Homo Aestheticus: Where Art Comes From and Why*. University of Washington Press, Seattle, USA, 1995.
- S. Dixon. A lightweight multi-agent musical beat tracking system. In *Pacific Rim International Conference on Artificial Intelligence*, pages 778–788, 2000.
- M. Donald. *Origins of the Modern Mind*. Harvard University Press, Cambridge, MA, 1991.
- L. Dumont. *Homo Hierarchicus: The Caste System and Its Implications*. University of Chicago Press, 1980.
- R. Dunbar. *Grooming, Gossip and the Evolution of Language*. Faber and Faber, London, 2004.

- R. Dunbar. Social brain and its implications for human social behaviour, 2006.
- R. Dunbar. Evolution of the social brain. In *The Evolution of Mind: Fundamental Questions and Controversies*, pages 280–292. Guilford Press, 2007.
- R. Dunbar. Sociality amongst human and non-human animals. In T. Ingold, editor, *Companion Encyclopedia of Anthropology*. Routledge, Oxford, UK, 1994.
- R. Dunbar. The social brain hypothesis. *Evolutionary Anthropology*, 6:178–190, 1998.
- W. H. Durham. *Coevolution: Genes, Culture and Human Diversity*. Stanford Univ Press, 1991.
- E. Durkheim. *Selected Writings*. Cambridge University Press, New York, 1972.
- D. S. Eck. Real-time musical beat induction with spiking neural networks, 2002.
- N. Eldridge and S. J. Gould. Punctuated equilibria: an alternative to phyletic gradualism. In T. J. M. Schopf, editor, *Models in paleobiology*, pages 82–115. Freeman, Cooper, San Francisco, 1972.
- J. Elman et al., editors. *Rethinking Innateness*. MIT Press, 1996.
- S. Feld. *Sound and Sentiment: Birds, Weeping, Poetics and Song in Kaluli Expression*. University of Pennsylvania Press, 1982.
- R. A. Fisher. *The Genetical Theory of Natural Selection*. Dover, London, 1958.
- W. T. Fitch. On the biology and evolution of music. *Music Perception*, 24(1):85–88, September 2006.
- M. M. Flood. Some experimental games. Research memorandum RM-789. RAND Corporation, Santa Monica, CA., 1952.
- J. A. Fodor. *Representations : philosophical essays on the foundations of cognitive science*. Harvester, 1981.
- C. Geertz. *The Interpretation of Cultures*. Basic Books, New York, 1973.

- E. Gellner. *Nations and Nationalism*. Blackwell Publishers Ltd., 1983.
- E. Gellner. *Culture, Identity, and Politics*. Cambridge University Press, New York, 1987.
- N. Gilbert. Emergence in social simulation. In N. Gilbert and R. Conte, editors, *Artificial Societies*, chapter 8, pages 144–156. UCL Press, 1995.
- D. E. Goldberg. *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley, 1989.
- S. J. Gould. *The Structure of Evolutionary Theory*. Harvard University Press, Cambridge, MA, 2002.
- S. J. Gould. Evolution's erratic pace. *Natural History*, 85:24–30, October 1977.
- S. J. Gould. *The Flamingo's Smile*. W. W. Norton, London, 1985.
- S. J. Gould and R. C. Lewontin. The spandrels of san marco and the panglossian paradigm: a critique of the adaptationist programme. In *Proceedings of the Royal Society of London*, volume 205 of *B*, pages 581–598, 1979.
- S. J. Gould and E. S. Vrba. Exaptation - a missing term in the science of form. *Paleobiology*, 8 (1):4–15, 1982.
- J. Hadamard. *The Psychology of Invention in the Mathematical Field*. Dover, 1954.
- E. H. Hagen and G. A. Bryant. Music and dance as a coalition signaling system. *Human Nature*, 14(1):21–51, 2003.
- W. D. Hamilton. The evolution of altruistic behaviour. *American Naturalist*, 97:354–356, 1963.
- I. Harvey and A. Thompson. Through the labyrinth evolution finds a way: A silicon ridge in evolvable systems: From biology to hardware. In T. Higuchi, M. Iwata, and L. Weixin, editors, *Proc. of The First International Conference on Evolvable Systems: From Biology to Hardware*, 1996.

- J. Henrich and F. J. Gil-White. The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behaviour*, 22:165–196, 2001.
- R. Highfield. *Frontiers of Complexity: The Search for Order in a Chaotic World*. Ballantine Books, 1996.
- G. E. Hinton and S. J. Nowlan. How learning can guide evolution. *Complex Systems*, 1:495–502, 1987.
- E. J. Hobsbawm and T. O. Ranger, editors. *The Invention of Tradition*. Cambridge University Press, 1992.
- J. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, USA, 1975.
- H. Honing. Structure and interpretation of rhythm and timing. *Dutch Journal of Music Theory (Tijdschrift voor Muziektheorie)*, 7(3):227–232, 2002.
- D. Huron. Is music an evolutionary adaptation. In *The Biological Foundations of Music*, Annals of the New York Academy of Sciences, pages 43–61. 2001.
- J. Huxley. *Evolution: The Modern Synthesis, 3rd Edition*. Macmillan, Basingstoke, UK, 1975.
- T. Ingold. Humanity and animality. In T. Ingold, editor, *Companion Encyclopedia of Anthropology: Humanity, Culture and Social Life*. Routledge, Oxford, UK, 1994.
- T. Ingold, editor. *Key Debates in Anthropology*. Routledge, 1996.
- E. Jablonka and M. J. Lamb. *Epigenetic Inheritance and Evolution: The Lamarckian Dimension*. OUP, Oxford, UK, 1995.
- T. Justus and J. J. Hutsler. Fundamental issues in the evolutionary psychology of music: Assessing innateness and domain specificity. *Music Perception*, 23:1–27, 2005.
- A. Karmiloff-Smith. Constraints on representational change: evidence from children's drawing. *Cognition*, 34:57–83, 1990.

- S. Kauffman. *Investigations*. OUP, New York, 2000.
- S. Kauffman. *The Origins of Order : Self-Organization and Selection in Evolution*. OUP, New York, 1993.
- S. Kirby. Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners, 1998.
- S. Kirby and J. Hurford. Learning, culture and evolution in the origin of linguistic constraints, 1997.
- C. Knight. Language and revolutionary consciousness. In A. Wray, editor, *The Transition to Language*. OUP, New York, 2002.
- C. Knight. *Blood Relations: Menstruation and the Origins of Culture*. Yale University Press, London, 1991.
- S. Koelsch, E. Kasper, D. Sammier, K. Schulze, and T. Gunter. Music, language and meaning: brain signatures of semantic processing. *Nature Neuroscience*, 7:302–307, 2004.
- A. Koestler. *The Ghost in the Machine*. Hutchinson and Co., London, 1967.
- M. Kohn and S. Mithen. Handaxes: Products of sexual selection? *Antiquity*, 73:518–526, 1999.
- J. R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, 1993.
- J. R. Krebs and R. Dawkins. Animal signals: Mind reading and manipulation. In J. R. Krebs and N. B. Davies, editors, *Behavioural Ecology: An Evolutionary Approach (2nd edition)*, pages 380–402. Blackwell, Oxford, UK, 1984.
- T. S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, 3 edition, 1996.
- A. Kuper. If memes are the answer, what is the question? In R. Aunger, editor, *Darwinizing Culture*. Oxford University Press, New York, 2000.
- A. Kuper. *Anthropology and Anthropologists*. Routledge, 3 edition, 1996.

- I. Lakatos. Falsification and the methodology of scientific research programmes. In I. Lakatos and A. Musgrave, editors, *Criticism and the growth of knowledge*. Cambridge University Press, 1970.
- K. Laland and J. Odling-Smee. The evolution of the meme. In R. Aunger, editor, *Darwinizing Culture*. OUP, New York, 2000.
- K. N. Laland, J. Odling-Smee, and M. W. Feldman. Niche construction, biological evolution and cultural change. *Behavioral and Brain Sciences*, 21(1), 1999.
- C. Langton. Toward artificial life. *Whole Earth Review*, 58, 1988.
- E. W. Large and M. R. Jones. The dynamics of attending: How people track time-varying events. *Psychological Review*, 10(1):119–159, 1999.
- E. W. Large and J. F. Kolen. Resonance and the perception of musical meter. *Connection Science*, 6(2-3):177–208, 1994.
- E. W. Large and C. Palmer. Perceiving temporal regularity in music. *Cognitive Science*, 26: 1–37, 2002.
- E. W. Large, C. Palmer, and J. B. Pollack. Reduced memory representations for music. *Cognitive Science*, 19(1):53–96, 1995.
- E. W. Large, P. Fink, and J. A. S. Kelso. Tracking simple and complex sequences. *Psychological Research*, 66:3–17, 2002.
- B. Latour. *We Have Never Been Modern*. Harvard University Press, Cambridge, Mass., USA, 1993.
- F. Lerdahl and R. Jackendoff. *A Generative Theory of Tonal Music*. MIT Press, 1983.
- C. Lévi-Strauss. *The Savage Mind*. The University of Chicago Press, 1966.
- C. Lévi-Strauss. *The Elementary Structures of Kinship*. Beacon Press, 1971.
- R. C. Lewontin. Sociobiology as an adaptationist program. *Behavioral Science*, 24:5–14, 1979.

- D. Lieberman. Causal explanations of human behavior: From culture to psychology or from psychology to culture? *Psychological Inquiry*, 17:109–115, 2006.
- S. R. Livingstone and W. F. Thompson. Multimodal affective interaction: A comment on musical origins. *Music Perception*, 24(1):89–94, September 2006.
- J. London. *Hearing in Time*. Oxford University Press, New York, 2004.
- J. Lovelock. *Gaia. A New Look at Life on Earth*. OUP, New York, 1979.
- C. J. Lumsden. Evolving creative minds: stories and mechanisms. In R. J. Sternberg, editor, *The Handbook of Creativity*. Cambridge University Press, New York, 1999.
- B. Malinowski. *Argonauts Of The Western Pacific*. George Routledge And Sons, Limited, 1932.
- K. Marx. *Capital: The Process of Production of Capital (online publication)*. www.marxists.org, 1995.
- H. Maturana and F. Varela. *Autopoiesis and Cognition: The realization of the living*. D.Reidel, Boston, 1980.
- H. R. Maturana and F. J. Varela. *The Tree of Knowledge: The Biological Roots of Human Understanding*. Shambhala Publications, Inc., Boston, MA, 1987.
- J. Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, New York, 1982.
- J. Maynard Smith and E. Szathmáry. *The Major Transitions in Evolution*. Oxford University Press, New York, 1995.
- J. McCormack. Artificial ecosystems for creative discovery. In T. D. et. al., editor, *Genetic and Evolutionary Computation Conference*, pages 301–307. ACM, New York, 2007.
- J. McDermott and M. Hauser. The origins of music: Innateness, uniqueness, and evolution. *Music Perception*, 23(1):29–59, September 2005.

- W. H. McNeill. *Keeping Together In Time: Dance and Drill in Human History*. Harvard University Press, 1995.
- P. Mellars. Going east: New genetic and archaeological perspectives on the modern human colonization of eurasia. *Science*, 313(5788):796–800, August 2006.
- B. Merker. Synchronous chorusing and human origins. In N. L. Wallin, B. Merker, and S. Brown, editors, *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2001.
- B. Merker. The uneven interface between culture and biology in human music. *Music Perception*, 24(1):95–98, September 2006.
- A. Merriam. *The Anthropology of Music*. Northwestern University Press., Chicago, 1964.
- G. Miller. Evolution of human music through sexual selection. In N. L. Wallin, B. Merker, and S. Brown, editors, *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2000a.
- G. Miller. *The Mating Mind*. Random House, New York, 2000b.
- E. R. Miranda and P. M. Todd. A-life and musical composition: A brief survey. In *IX Brazilian Symposium on Computer Music: Music as Emergent Behaviour*, pages 59–65, 2003.
- E. R. Miranda, S. Kirby, and P. M. Todd. On computational models of the evolution of music: From the origins of musical taste to the emergence of grammars. *Contemporary Music Review*, 22(3):91–111, September 2003.
- M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, 1998.
- S. Mithen. *The Singing Neanderthals*. Weidenfeld and Nicholson, London, 2005.
- S. Mithen. *The Prehistory of the Mind*. Thames and Hudson, New York, 1996.
- B. Nettl. An ethnomusicologist contemplates universals in musical sound and musical culture. In N. L. Wallin, B. Merker, and S. Brown, editors, *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2000.

- B. Nettl. *The study of ethnomusicology: thirty-one issues and concepts*. University of Illinois Press, Urbana and Chicago., 2005.
- D. Nettle. Language variation and the evolution of societies. In R. Dunbar, C. Knight, and C. Power, editors, *The Evolution of Culture*, chapter 11, pages 214–227. Edinburgh University Press, Edinburgh, 1999.
- J. Noble. Talk is cheap: Evolved strategies for communication and action in asymmetrical animal contests. In J.-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S. Wilson, editors, *SAB00*, pages 481–490. MIT Press, 2000.
- F. J. Odling-Smee. Niche construction, evolution and culture. In T. Ingold, editor, *Companion Encyclopedia of Anthropology: Humanity, Culture and Social Life*. Routledge, Oxford, UK, 1994.
- F. J. Odling-Smee, K. N. Laland, and M. W. Feldman. *Niche Construction: The Neglected Process in Evolution*. Number 37 in Monographs in Population Biology. Princeton University Press, Princeton, USA, 2003.
- H. T. Odum. *Systems Ecology*. Wiley-Interscience, 1983.
- S. H. Orzack and E. R. Sober. *Adaptationism and Optimality*. Cambridge University Press, Cambridge, UK, 2001.
- D. H. Owings and E. S. Morton. *Animal Vocal Communication: A New Approach*. Cambridge University Press, New York, 1998.
- H. Papousek. Musicality in infancy research: biological and cultural origins of early musicality. In I. Deliège and J. A. Sloboda, editors, *Musical Beginnings*, pages 37–55. Oxford University Press, 1996.
- A. D. Patel. Musical rhythm, linguistic rhythm, and human evolution. *Music Perception*, 24(1): 99–104, September 2006.

- I. Peretz. Brain specialisation for music: New evidence from congenital amusia. In *The Biological Foundations of Music*, Annals of the New York Academy of Sciences, pages 153–165. 2001.
- K. Pike. *Language in relation to a unified theory of structure of human behavior*. Mouton, The Hague, 1967.
- S. Pinker. *How the Mind Works*. Allen Lane The Penguin Press, London, UK, 1998.
- S. Pinker and R. Jackendoff. The faculty of language: What's special about it? *Cognition*, 95 (2):201–236, 2005.
- H. Plotkin. *The Imagined World Made Real: Towards a Natural Science of Culture*. Penguin, 2002.
- H. Plotkin. *Evolution in Mind*. Penguin, 1997.
- G. Polya. *How To Solve It*. Princeton University Press, 1957.
- K. R. Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge*. Routledge and Kegan Paul, 1972.
- K. R. Popper. Natural selection and the emergence of mind. *Dialectica*, 32(3–4):339–355, 1978.
- T. Quick, K. Dautenhahn, C. Nehaniv, and G. Roberts. The essence of embodiment: A framework for understanding and exploiting structural coupling between system and environment. In D. M. Dubois, editor, *Third International Conference on Computing Anticipatory Systems, Symposium 4 on Anticipatory, Control and Robotic Systems*, pages 649–660, Liège, Belgium, 1999.
- M. Quinn. Evolving communication without dedicated communication channels. In *Advances in Artificial Life : 6th European Conference, ECAL*, 2001.
- A. Rehding. The quest for the origins of music in Germany circa, 1900. *Journal of the American Musicological Society*, 53:345–385, 2000.

- C. W. Reynolds. Flocks, herds and schools: A distributed behavioural model. In *Computer Graphics (SIGGRAPH '87 Conference Proceedings)*, volume 21, pages 25–34, 1987.
- R. T. P. Richard E. Lenski, Charles Ofria and C. Adami. The evolutionary origin of complex features. *Nature*, 423:139–144, May 2003.
- P. J. Richerson and R. Boyd. *Not by Genes Alone: How Cultural Transformed Human Evolution*. University of Chicago Press, Chicago, IL, US, 2005.
- B. Richman. How music fixed “nonsense” into significant formulas: On rhythm, repetition and meaning. In N. L. Wallin, B. Merker, and S. Brown, editors, *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2001.
- M. Ridley. *The Red Queen*. Penguin, 1994.
- R. Saunders. *Curious Design Agents and Artificial Creativity*. PhD thesis, Faculty of Architecture, The University of Sydney, 2001.
- R. Saunders and J. S. Gero. The digital clockwork muse: A computational model of aesthetic evolution. In G. A. Wiggins, editor, *Proceedings of the AISB'01 Symposium on AI and Creativity in Arts and Science, SSAISB*, pages 12–21, University of York, York, UK, 2001a.
- R. Saunders and J. S. Gero. Artificial creativity: Emergent notions of creativity in artificial societies of curious agents. In *Proceedings of Second Iteration*, 2001b.
- G. Schlaug. The brain of musicians: A model for functional and structural adaptation. In *The Biological Foundations of Music*, Annals of the New York Academy of Sciences, pages 281–299. 2001.
- P. Schuster, W. Fontana, P. F. Stadler, and I. L. Hofacker. From sequences to shapes and back: A case study in rna secondary structures. *Proc. Roy. Soc. Lond.*, B 255:279–284, 1994.
- W. A. Searcy and S. Nowicki. *The evolution of animal communication; reliability and deception in signaling systems*. Princetown Universtiy press, Oxford, 2005.
- T. C. Shelling. *Micromotives and Macrobehavior*. Norton, New York, 1978.

- S. Shennan. *Genes, Memes and Human History*. Thames and Hudson, London, 2002.
- B. Shore. *Culture in mind: cognition, culture, and the problem of meaning*. Oxford University Press, 1996.
- K. Sims. Evolving 3d morphology and behaviour by competition. In *Artificial Life IV Proceedings*. MIT Press, 1994.
- J. A. Sloboda, editor. *Generative Processes in Music*. Oxford University Press, 1988.
- K. Smith. The cultural evolution of communication in a population of neural networks, 2002.
- C. P. Snow. *The Two Cultures*. CUP, 1993.
- E. Sober and D. S. Wilson. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Harvard University Press, 1998.
- H. Spencer. *First Principles*. Williams and Norgate, London, 2nd edition, 1867.
- D. Sperber. An objection to the memetic approach to culture. In *Darwinizing Culture*. Oxford University Press, 2000.
- K. Sterelny. *Dawkins vs Gould*. Icon Books Ltd., Cambridge, UK, 2001.
- K. Sterelny. Review: Genes, memes and human history, by stephen shennan. *Mind and Language*, 19(2):249–257, 2004.
- R. J. Sternberg and T. I. Lubart. The concept of creativity: prospects and paradigms. In R. J. Sternberg, editor, *The Handbook of Creativity*. Cambridge University Press, New York, 1999.
- M. Strathern, editor. *Shifting Contexts: Transformations in Anthropological Knowledge*. Routledge, 1995.
- E. Szathmáry. Origins of life – the first two billion years. *Nature*, 387(662–663), 1997.
- T. Taylor. Niche construction and the evolution of complexity. In P. H. T. I. J. Pollack, M. Bedau and R. Watson, editors, *Artificial Life IX, Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, pages 375–380, 2004.

- D. W. Thompson. *On Growth and Form*. Dover, 1992.
- N. Todd, D. O'Boyle, and C. Lee. A sensory-motor theory of rhythm, time perception and beat induction. 1999.
- N. P. M. A. Todd. The auditory "primal sketch": A multiscale model of rhythmic grouping. *Journal of New Music Research*, 1994.
- N. P. M. A. Todd. Motion in music: A neurobiological perspective. *Music Perception*, 17(1), 1999.
- N. P. M. A. Todd and G. J. Brown. Visualisation of rhythm, time and metre. *AI Review*, 10: 253–273, 1996.
- P. M. Todd. Simulating the evolution of musical behaviour. In N. L. Wallin, B. Merker, and S. Brown, editors, *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2000.
- E. Tolbert. Music, meaning, and narratives of human evolution. Conference on Music and Evolutionary Thought, Institute of Advanced Study, University of Durham, July 2007.
- M. Tomasello. *Constructing a Language: A Usage-based Theory of Language Acquisition*. Harvard University Press, 2003.
- M. Tomasello. *The cultural origins of human cognition*. Harvard University Press, 1999.
- M. Tomasello, M. Carpenter, J. Call, T. Behne, and H. Moll. Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences*, 28(5):675–691, 2005.
- L. J. Trainor. Innateness, learning and the difficulty of determining whether music is an evolutionary adaptation: A commentary on justus and hustler (2005) and mcdermott and hauser (2005). *Music Perception*, 24(1):105–110, September 2006.
- S. Trehub. Musical predispositions in infancy. In *The Biological Foundations of Music*, Annals of the New York Academy of Sciences. 2001.

- C. Trevarthen. Form, significance and psychological potential of hand gestures in infants. In J. Nespoulous, editor, *Biological Foundations of Gestures*. Lawrence Erlbaum Associates, Inc, 1986.
- C. Trevarthen. The self born in intersubjectivity: The psychology of an infant communicating. In U. Neisser, editor, *The Perceived Self*. Cambridge University Press, New York, 1993.
- P. W. Turke. Concealed ovulation, menstrual synchrony and paternal investment. In *Biosocial Perspectives on the Family*, pages 119–136. Sage, Newbury Park, CA, 1988.
- E. van den Broek and P. M. Todd. Piep piep piep, ich hab' dich lieb: Rhythm as an indicator of mate quality. In W. Banzhaf, T. Christaller, P. Dittrich, J. Kim, and J. Ziegler, editors, *Advances in Artificial Life: 7th European Conference Proceedings (ECAL 2003)*, pages 425–433, Berlin, 2003. Springer Verlag.
- N. L. Wallin. *Biomusicology: Neurophysiological, Neuropsychological, and Evolutionary Perspectives on the Origins and Purposes of Music*. Pendragon Press, New York, 1991.
- N. L. Wallin, B. Merker, and S. Brown, editors. *The Origins of Music*. MIT Press, Cambridge, MA, USA, 2000.
- P. Wason. Reasoning. In B. M. Foss, editor, *New Horizons in Psychology*. Harmondsworth: Penguin, 1966.
- G. Werner and M. Dyer. Evolution of communication in artificial organisms. In C. Langton, C. Taylor, D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 659–687. Addison-Wesley Pub., 1992.
- G. Werner and P. M. Todd. Too many love songs: sexual selection and the evolution of communication. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life*, pages 434–443. Cambridge, MA: MIT Press/Bradford Books, 1997.
- C. Whitehead. Music, social mirrors and brain evolution. Conference on Music and Evolutionary Thought, Institute of Advanced Study, University of Durham, July 2007.

- A. Whiten and R. W. Byrne. *Machiavellian Intelligence II: Extensions and Evaluations*. CUP, Cambridge, UK, 1997.
- G. A. Wiggins. Characterising creative systems. In C. Bento, A. Cardoso, and J. Gero, editors, *Proceedings of the IJCAI'03 Workshop on Creative Systems*, 2003.
- A. C. Wilson. The molecular basis of evolution. *Scientific American*, 253(4):148–157, 1985.
- A. Wray, editor. *The Transition To Language*. Oxford University Press, Oxford, UK, 2002.
- A. Zahavi. Mate selection – a selection for a handicap. *Journal of Theoretical Biology*, 53, 1975.
- A. Zahavi and A. Zahavi. *The Handicap Principle: A Missing Piece of Darwin's Puzzle*. Oxford University Press, New York, 1997.