

Durham E-Theses

The Nurture of Nature: Biology, Psychology and Culture

ELIZABETH MARY HANNON

How to cite:

HANNON, ELIZABETH MARY (2010) *The Nurture of Nature: Biology, Psychology and Culture*.
Doctoral thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/337/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

The Nurture of Nature:

Biology, Psychology & Culture

Elizabeth Mary Hannon

Ph.D. Thesis

Department of Philosophy

Durham University

2009

The Nurture of Nature: Biology, Psychology and Culture

Beth Hannon

Abstract

In this thesis I explore what consequences taking development seriously in evolutionary considerations will have for how we understand the evolution of psychology and culture. I first explicate the relationship between development and evolution that informs a number of approaches to evolution, including neo-Darwinian evolutionary biology and evolutionary developmental biology. I argue that, to a greater or lesser extent, developmental processes have been misconstrued in these accounts and that the full role of development, from an evolutionary point of view, has not always been acknowledged. Instead, I suggest that a better model of the relationship between development and evolution can be found in developmental systems theory.

I explore the neo-Darwinian underpinnings of a number of accounts of the evolution of culture and psychology, including the branch of evolutionary psychology associated with the work of, among others, John Tooby and Leda Cosmides, and the gene-culture co-evolutionary account of Peter Richerson and Robert Boyd. I argue that as well as being vulnerable to the same sorts of problems that plague neo-Darwinian evolutionary biology, they face other difficulties. These accounts suppose an internalist model of the mind, and this model is neither justified nor useful. The extended mind hypothesis offers a different model of the mind whereby cognitive processes can be partially constituted by structures in the environment. I sketch an alternative account of what the evolution of human psychology and culture by combining a developmental systems approach to evolution and development with the extended mind hypothesis. This will result in a very different understanding of the relationship between biology, psychology and culture.

Table of Contents

Introduction	11
---------------------------	-----------

Part One: Development and Evolution

Chapter One: Neo-Darwinism	18
---	-----------

1. Introduction
2. Natural Selection
 - 2.1 Adaptation
 - 2.2 Types of Adaptationism
3. Genes
 - 3.1 Genetic Determinism
 - 3.2 The Gene as the Unit of Inheritance and Selection
 - 3.3 Genetic Information
 - 3.3.1 The Causal Theory of Information
 - 3.3.2 The Teleosemantic Theory of Information
4. Development
5. Conclusion

Chapter Two: Development and Natural Selection	50
---	-----------

1. Introduction
2. Developmental Constraints
 - 2.1 The Adaptationist Programme

2.2 Process Structuralism	
2.3 The False Dichotomy	
3. Organisms and Environments	
3.1 The Extended Phenotype	
3.2 Niche Construction	
3.3 Causal and Constitutive Construction	
4. Conclusion	
Chapter Three: Evolutionary Developmental Biology	78
1. Introduction	
2. Evolutionary Developmental Biology	
3. Synthetic Evo-Devo	
4. Development in Evolutionary Developmental Biology	
4.1 Modularity	
4.2 Genes in Evo-Devo	
5. Conclusion	
Chapter Four: Developmental Systems Theory	106
1. Introduction	
2. Developmental Systems Theory	
2.1 Extended Inheritance	
2.2 Development	
2.3 Evolution as Construction	
3. Holism	
4. Evo-Devo and DST	
5. Conclusion	

Part Two: Culture and Cognition

Chapter Five: Evolutionary Psychology 134

1. Introduction
2. Evolutionary Psychology
 - 2.1 The Massive Modularity Hypothesis
 - 2.2 The “Psychic Unity of Humankind”
 - 2.3 Stone Age Minds
 - 2.4 Uncovering Modules
3. Culture
 - 3.1 The Standard Social Science Model
 - 3.2 Darwinian Social Science
 - 3.3 Mismatch
 - 3.3.1 Generation of Culture
 - 3.3.2 Responding to Mismatched Culture
 - 3.3.3 The Upshot for Evolutionary Psychology
4. Conclusion

Chapter Six: Biology and Culture 180

1. Introduction
2. Gene-Culture Co-Evolution
 - 2.1 Memes
 - 2.2 Population Thinking
 - 2.3 Culture, Genes and Development
3. The Limits of Theories of Cultural Evolution

- 3.1 The Looping Effects of Human Kinds
- 3.2 Obesity
- 3.3 Objections
- 4. Looping Effects, DST and Theories of Cultural Inheritance

Chapter Seven: The Evolution of Cognition;... 206

- 1. Introduction
- 2. Cognition
 - 2.1 The Extended Mind
 - 2.1.1 Tetris
 - 2.1.2 Otto’s Notebook
 - 2.1.3 Constituting Cognition
 - 2.2 Objections
 - 2.2.1 Clark and Chalmers’ Objections
 - 2.2.2 The Internalist Understanding of Tetris
 - 2.2.3 The Causal-Constitution Fallacy
- 3. The Evolution of Extended Cognitive Processes
 - 3.1 The Cost-Benefit Analysis
 - 3.1.1 The Hidden Costs Argument
 - 3.1.2 The Hidden Benefits Argument
 - 3.1.3 The Strong Hidden Benefits Argument
 - 3.1.4 Changing Environments
 - 3.1.5 Plasticity
 - 3.1.6 Development
 - 3.2 The Manipulation Thesis and the Extended Phenotype Hypothesis

3.2.1 The Parity Principle

3.2.2 Privileging the Internal

3.3 Developmental Systems Theory

Conclusion	254
References	259

The copyright of this thesis rests with the author. No quotation from it should be published without the prior written consent and information derived from it should be acknowledged

Acknowledgements

I would first and foremost like to thank my supervisors for their help throughout. Matthew Ratcliffe provided insights, criticisms and suggestions that have been invaluable. I am particularly grateful for his engaging philosophy of biology lectures which inspired me to embark upon this thesis. Soran Reader was a constant source of motivation and help throughout, as well as a good friend. Their support was essential to my completion of this thesis.

I would also like to thank everyone else in the Department of Philosophy at Durham University. Jill Lockey was always helpful, even in the face of daft requests, and a constant source of good humour. Liat Netzer, Vicky Blake, Claire Graham, Donnchadh O’Conaill, Liz McKinnell, Olley Pearson, Owen Earnshaw and Stephanie Eichberg provided useful feedback, proof-reading, moral support and cake, for which I am very grateful and promise to return the favour. Thanks also go to Fern Elsdon-Baker and Ali Brown for their help throughout. Any mistakes that remain in this thesis are, of course, my own. I am particularly grateful to Matthew Conduct, my parents, Mary and Brendan Hannon, and my grandparents, Kevin and Stella Mahon for their support throughout; their patience is boundless.

Finally, I am indebted to the British Society for the Philosophy of Science for their financial support, without which I could not have produced this thesis.

Part One

Development and Evolution

Introduction

My aim in this thesis is to explicate the relationship between development and evolution in a number of different evolutionary accounts, and to explore what a more developmentally-informed account of the evolution of psychology and culture would look like. In particular, I want to explore what an account of the evolution of psychology and culture would look like if we combine the extended mind hypothesis with a developmental systems perspective.

The question of how biology, psychology and culture relate to one another has a long history. A good deal of work on this issue can be understood as sitting somewhere on a spectrum with biological determinism at one extreme and environmental determinism at the other. Although Darwinian evolutionary theory has been applied to this issue since Darwin first published *On the Origin of Species* (1859), no consensus has yet been reached on what light evolutionary considerations can shed on this matter. In this thesis, I take it that there is some insight to be gained from asking evolutionary questions about these matters. This is not, by any means, the claim that a complete account of human psychology and culture can be had by understanding the evolution of human psychology and culture; rather, it is the much smaller claim that an evolutionary perspective can make some contribution to an understanding of human psychology and culture. A theme that reappears throughout this thesis is the value of adopting a variety of approaches to any particular research question, and I will be critical of views that tend to obscure or *a priori* rule out alternative lines of inquiry.

Granting that an evolutionary perspective can help us understand human psychology and culture, much will then hang on the particular account of evolution we adopt. By far the most widely accepted evolutionary account is neo-Darwinian evolutionary biology, and a variety of accounts of the evolution of psychology and culture have used this as a starting point for their own analyses. I will focus on one neo-Darwinian approach to evolutionary psychology associated with the work of, among others, John Tooby and Leda

Cosmides (e.g. Tooby & Cosmides 1992). Buller (2005) distinguishes Evolutionary Psychology from other evolutionary accounts of psychology by the use of the upper case, and I will follow this convention here. I will also examine gene-culture co-evolutionary accounts including meme theory and particularly the work of Richerson and Boyd (2005). Not all of these co-evolutionary accounts adopt a strict neo-Darwinian approach, but they are all part of a class of theories sometimes termed “interactionist” (e.g. Buss et al. 1999; Oyama 2000b, 2000c). The interactionist maintains that biological and psychological traits emerge from the interaction between our biology and our environment. Proponents of this approach suggest that interactionism resolves the age old “nature versus nurture” debate. Rather than couching the issue in terms of nature versus nurture, the interactionist position adopts a “nature *and* nurture” stance (Robert 2003). However, interactionism does not mark any significant conceptual break from the older nature/nurture debates. Nature and nurture continue to be conceived of as two distinct causes; the interactionist’s job is to detail how they interact with one another (Oyama 2000b, 2000c).

Neo-Darwinian evolutionary biology has tended to marginalise development in evolutionary theory. Indeed, neo-Darwinian biologists will often assent to this, tending to view development as beside the point from the perspective of their own research interests. However, there has been a growing interest in exploring the role of development in an evolutionary context. The emergence of developmental systems theory and evolutionary developmental biology – two accounts dealt with in this thesis – testify to this. In particular, I am interested in developmental systems theory here; I will argue that other approaches to a more developmental evolutionary theory suffer from serious problems. Developmental systems theory rejects the claim that nature and nurture are two distinct causes of phenotypes. Rather, the developmental systems theorist maintains that nature is the outcome of nurture. As Oyama puts it, “in [developmental systems theory] ‘nature’ is not a phantom reality standing behind the phenotype: the phenotype in its surrounds *is all the nature there is*, and this is plenty” (2000c: S341). Nothing about phenotypes is already given; there is no essential nature lurking beneath the noise generated by developmental processes. Rather, phenotypes are just the result of the complex interactions of often vast arrays of developmental resources.

The developmental systems perspective alters our conception of both the explanans and the explanandum of evolutionary theory. The explanans is now thought of as including developmental processes and outcomes. The explanandum shifts as well so that we are no longer explaining the evolution of a skin-bound organism, or even an extended phenotype; rather, we must attend to developmental systems and the life cycles they produce. The developmental system can extend far beyond the skin of the organism, and can affect how we individuate life forms. Structures in the environment can partially constitute developmental processes and can form part of what secures heritable similarity. When such an approach is used to analyse the evolution of human psychology, it finds a good deal of resonance with an idea from the philosophy of mind termed the extended mind hypothesis. This hypothesis also suggests that boundaries created by our skin are less important than might be supposed, and argues that cognitive processes can be partially constituted by structures in the environment. I will sketch a picture of what an account of the evolution of psychology and culture might look like when we take developmental systems theory and the extended mind hypothesis together.

This thesis is organised into two parts; the first deals with evolutionary biology, while the second focuses on the evolution of psychology and culture. In part one my aim is to explicate the relationship between development and evolution in a number of different evolutionary theories, and argue that developmental systems theory avoids the pitfalls inherent in other accounts of evolution.

Chapter one will be largely expository, and thus only concerned with outlining neo-Darwinian evolutionary biology; criticisms of this position will arise in later chapters. I identify two key commitments of neo-Darwinism that lie behind its understanding of the relationship between development and evolution by natural selection: the adaptationist programme and the gene as the unit of inheritance. I also detail what evolution and natural selection require and entail in order that these ideas can be well understood as they arise throughout the rest of this thesis.

In chapter two I will focus specifically on the relationship between development and evolution by natural selection in a number of different approaches to thinking about evolution. I will begin by returning to one of the

two key commitments of neo-Darwinism, the adaptationist programme, and discuss the role this commitment plays in marginalising the contribution of developmental processes in evolutionary processes. I will contrast two responses to the adaptationist programme: process structuralism and niche construction. Although process structuralism is in many ways quite different to neo-Darwinism, I argue that both share common ground in that they have adopted similar models of the relationship between development and evolution. The niche construction approach, on the other hand, takes a different view of natural selection than that implied by the adaptationist programme, and argues that organisms are not merely moulded to suit their environments as the adaptationist programme assumes, but that organisms can also mould their environments to suit their own needs. I will argue that the niche construction approach can easily be extended to include interactions between developmental processes and the wider environment. I will look at some criticisms of this approach that stem from Dawkins' (1999a) extended phenotype approach, and also from the work of Godfrey-Smith (1996). However, I will argue that these criticisms miss the mark and that the niche construction model offers a more realistic account of the relationship between development and evolution by natural selection than that offered by either process structuralism or neo-Darwinian evolutionary biology.

Chapter three turns to evolutionary developmental biology and will put pressure on the second neo-Darwinian commitment that maintains the gene is the unit of inheritance. Because evolutionary developmental biology is a relatively new field, there are a wide variety of opinions that comprise it, so I will spend some time identifying different trends which fall under the evolutionary developmental biology umbrella. Evolutionary developmental biology makes a strong case for the importance of taking development seriously when trying to understand evolution. However, while the gene is not viewed as the sole unit of inheritance in evolutionary developmental biology, it is taken to underpin hereditary transmission. This is then taken to justify viewing the gene as distinct from other developmental resources in some important way. I will argue that the privileging of the gene in evolutionary developmental biology, though a less extreme version than is often found in neo-Darwinian evolutionary biology, is based on a circular argument, and given this, we can no longer accept the claim that it is genes that ultimately explain heritable similarity.

Having, in chapter two and three, undermined both the adaptationist programme and the idea that the gene is the only unit of inheritance, I will turn in chapter four to developmental systems theory. While neo-Darwinism begins with a picture of evolution, and builds a picture of development from that, developmental systems theory begins with development, and only then tries to understand how evolution might occur. The result is that developmental systems theory offers a very different kind of evolutionary account. I will argue that developmental systems theory does not fall prey to either of two styles of argument presented against it; it is neither threatened by the emergence of evolutionary developmental biology, nor is it unworkably holistic. Both evolutionary developmental biology and developmental systems theory have the ability to enrich one another and developmental systems theorists are already addressing the shortcomings in evolutionary developmental biology being lamented by evolutionary developmental biologists themselves. This, in turn, demonstrates that developmental systems theory can be put to work by scientists as well as philosophers, and deflates concerns about an unworkable holism in the developmental systems approach.

Part two turns to the evolution of psychology and culture and will deal with a variety of approaches to this issue. As well as demonstrating short-comings in a number of these approaches, my aim here is to sketch an account of what a developmental systems perspective can bring to the study of the evolution of culture and cognition.

Chapter five deals with Evolutionary Psychology. Evolutionary Psychology is based on neo-Darwinian biology and thus is vulnerable to all the same criticisms levelled at neo-Darwinian evolutionary theory in part one. Nonetheless, I discuss Evolutionary Psychology here for two reasons. First, it makes stronger claims about evolutionary biology than can be supported by neo-Darwinian evolutionary theory, so that even if neo-Darwinism can be defended from the criticisms made in part one of this thesis, Evolutionary Psychology will continue to face difficulties. Second, even accepting Evolutionary Psychology's interpretation of neo-Darwinian biology, it faces problems specific to its research focus – human psychology and culture. The sorts of problems faced by Evolutionary Psychology here may need to be dealt with by other attempts to give evolutionary accounts of psychology and culture. Evolutionary Psychology

is currently caught in a dilemma between two different claims. On the one hand, Evolutionary Psychologists suggest that culture is largely generated by our evolved cognitive capacities. On the other hand, they claim modern humans living in industrialised cities are often maladapted. I will argue that these two claims cannot be made compatible, and Evolutionary Psychologists are forced to either accept that they have little to say about modern culture, or they must abandon their particular model of human psychology.

In chapter six I will examine gene-culture co-evolutionary accounts, such as meme theory (Dawkins 1989; Blackmore 1999) and that developed by Richerson and Boyd (2005). One important criticism of these sorts of approaches concerns the nature of inheritance. Natural selection is often taken to require almost complete vertical transmission. Theories of cultural evolution, as well as developmental systems theory, allow both vertical and horizontal transmission, and this has been taken by critics to rule out the possibility of natural selection operating on cultural variants or on developmental systems. Boyd and Richardson have developed one way of dealing with this criticism, but it relies on a model of the mind that will be rejected in chapter seven. I will outline another way this criticism can be handled using ideas developed in chapter four on developmental systems theory. However, I will argue that evolutionary accounts that focus on humans will need to contend with a particular difficulty. Following Hacking (1996, 1999), I will argue that human kinds differ from natural kinds in their instability as a result of looping effects. This places important limitations on evolutionary accounts that focus on humans, limitations that are not a feature of accounts of the evolution of non-human animals.

Finally, in chapter seven I turn to the evolution of cognition. The extended mind hypothesis – understood here as the idea that our cognitive processes may be constituted by resources beyond our brains – has been viewed as somewhat radical. However, if one accepts the account of development and evolution put forward by developmental systems theorists, this model of cognition seems a good deal more plausible. I will begin by outlining the extended mind hypothesis, and will counter some recent criticisms levelled at it. Both Rowlands (1999) and Menary (2007) have offered evolutionary justifications for the extended mind hypothesis but I will argue that neither offer convincing accounts of the evolution of extended cognitive processes; on the contrary, they undermine the extended

mind hypothesis when they adopt Dawkins' (1999a) extended phenotype approach. Instead, I will suggest that a developmental systems perspective provides an account of development and evolution that is sympathetic to the goals and methodology of those who argue for the extended mind hypothesis and more effectively allows us to deal with common worries about the extended mind hypothesis. Combining a developmental systems perspective with this extended mind hypothesis will result in a very different understanding of the relationship between biology, psychology and culture.

Much of the science surveyed here is in its early stages, and specific revisions may have to be made. But what I hope to achieve here is to make the case for there being a role for development in the study of the evolution of cognition and culture, and that the developmental systems approach offers the best way to do this. My aim here is to indicate what such a developmental systems approach to human culture and psychology might look like, and the sorts of issues it might handle; how such a framework is put to work will require more investigation.

Chapter One

Neo-Darwinism

1. Introduction

The Modern Evolutionary Synthesis that led to neo-Darwinism was, broadly, the coming together of large-scale evolutionary work – such as work on speciation, biogeography, and palaeontology – with the more experimental Mendelian genetics. Given the disparate disciplines involved, it should be no surprise that neo-Darwinism is a broad church. The specifics of the research programmes that are essential to the work of a palaeontologist may be less important – often irrelevant – to the day-to-day work of the population geneticist. This is, of course, unproblematic. Productive research can result from prioritising some theoretical issues while bracketing others. But given these different approaches, a straightforward and comprehensive description of neo-Darwinism is not an easy thing to arrive at. There are certain commitments which everyone working within the neo-Darwinist camp will assent to, but just how important any particular issue will be very much depends on what research question is being pursued. It is not always clear that there is agreement on the terms and concepts being used. For example, although the gene seems fundamental to the neo-Darwinian project, there appears to be no clear-cut agreement on what a gene is (Stotz et al. 2004). As a result, I cannot hope to give a comprehensive survey of neo-Darwinism and the range of opinions within this framework here. Instead I want to focus on the two key commitments of neo-Darwinism that motivate its characterisation of development in relation to evolution by natural selection. The first of these commitments is to the adaptationist programme which maintains that traits – or at least an important category of traits – are best explained in terms of problems posed by ancestral environments (Lewontin 1978, 2001b). What “best explanation” amounts to here will be discussed, with particular reference to Godfrey-Smith’s (2001a) distinction between different forms of adaptationism. The second commitment is to the gene as the unit of inheritance. This is taken as the starting point for a justification of the privileged role of the gene in development and the subsequent marginalisation of development in evolutionary

considerations. My aim here is merely to detail these aspects of neo-Darwinism; I will deal with each commitment more critically in later chapters.

One of the aims of the first part of this thesis is to explicate the relationship between development and evolution by natural selection as envisaged by a number of different approaches to evolution. The two commitments of neo-Darwinism that I will explore here underpin the neo-Darwinian understanding of this relationship. This is uncontroversially the case with regards to the commitment to the gene as the unit of inheritance; neo-Darwinists justify their neglect of developmental mechanisms by taking this commitment as a starting point. Given this, I will explore the relationship between this commitment and the relationship between development and natural selection in detail in this chapter. The manner in which the commitment to the adaptationist programme leads to a neglect of development in evolutionary considerations is less straightforward, and results from a problem shared by other approaches to evolution. Thus, while I will outline what the adaptationist programme amounts to here, the way in which the adaptationist programme leads to a neglect of the full role of development in evolution by natural selection will be dealt with in chapter two.

2. Natural Selection

Although natural selection is not the only evolutionary process, within neo-Darwinism it is taken as the most important. There are different ways of unpacking the claim that natural selection is the “most important” process, and I will discuss these in more detail below. For now I want to focus on natural selection itself. In its “classical” articulation, evolution by natural selection is thought to follow from three key points: variation, differential fitness, and heritability (Godfrey-Smith 2009). Within any population there is some degree of variation. Some differences that exist between individuals will lead to a tendency for some organisms to survive and reproduce with success than other members of the population. Because offspring more closely resemble their parents than strangers, those advantageous traits possessed by the parents that ensured their greater chances of survival and ability to reproduce may be inherited by the offspring so that they too possess this advantage. Taken together, these three points can lead to radical changes in the composition of the population over a

large number of generations as a result of natural selection. In each generation, those individuals without the advantageous trait will fail to reproduce, or at least reproduce at a reduced rate when compared to those individuals with the advantageous trait. The percentage of individuals in the population with the advantageous trait will increase as the percentage of individuals with the disadvantageous trait decreases. After sufficient time, the population may no longer include individuals with the disadvantageous trait and instead may be solely composed of those individuals who possess the advantageous trait.

Godfrey-Smith (2007, 2009) has noted that it is sometimes assumed that evolution by natural selection must follow from variation, differential survival and heritability. However, this is mistaken; while these conditions will lead to natural selection, *evolution* by natural selection does not necessarily follow. Natural selection does not entail evolution, or change, in a population. Stabilizing selection, for example, involves both selection and an absence of change or evolution. In a hypothetical population, individuals either possess trait *A* or trait *B*. Trait *A* confers an advantage in some situations so that *A*-individuals are favoured over *B*-individuals. In other situations, however, it is *B*-individuals who are favoured over *A*-individuals. So there are two selection processes acting on the population. In this example, the processes are such that the composition of the population remains in a dynamic equilibrium around a certain point and thus there is no (significant) change in gene frequency. Nonetheless, there are selection processes at work favouring either trait *A* or *B*, and these traits continue to exist in the population because one or other of the selection processes favours them.

Natural selection may also act to preserve the composition of a population against deleterious variants and drift. Drift, another cause of evolutionary change, results from “differences in survival and reproduction that merely reflect the operation of chance” (Godfrey-Smith 2009: 27). Some organisms will be less successful at escaping predators, finding food, attracting mates, and so on than competitors because they possess traits that make them less competent in these tasks than conspecifics. Those individuals who are able to avoid predators due to some trait will tend to survive, while those that lack this trait will tend to be eliminated. However, not all “eliminations” will occur because the organism possessed traits that were poorly suited to the environment, and not all organisms

will survive and reproduce at higher rates than conspecifics just because they possess traits more suited to their environment. Sometimes success or failure will occur because of random events. Drift is a problem felt more strongly when examining small populations. Chance events that remove particular organisms from the population will have a larger effect on the proportion of a given phenotype in a small population. Across larger populations, chance events tend to even out so that chance eliminations of one trait are compensated by chance eliminations of other traits, leaving the overall composition of the population largely unaffected.

Evolution has sometimes been equated with change in gene frequencies.¹ In chapter four I will argue that not all cases of evolution need consist in changes in gene frequencies. However, this takes us beyond neo-Darwinism, the focus of this chapter. Here it is worth noting some points about defining evolution as a change in gene frequencies. Evolution occurs as a result of a number of different mechanisms, natural selection being only one. Drift, for instance, will account for some changes in frequency. So, not all cases of gene frequency change are cases of natural selection. Similarly, not all cases of natural selection are cases of gene frequency change. Stabilizing selection acts to resist change in gene frequency. This need not be a problem if our focus is on *evolution* by natural selection, but it is important to note that gene frequency change and natural selection cannot be equated with one another. Gene frequencies change for reasons other than natural selection, and natural selection does not necessitate gene frequency change.

A further complication arises out of the neutral theory of molecular evolution. This suggests that a good deal of change at the genetic level is the result of mutations that have no phenotypic effects and thus are neutral with respect to natural selection (Kimura 1968). If evolution is defined as change in gene frequencies, changes at the genetic level that have no effect on the phenotype will nonetheless be considered evolutionary changes. This does not apply in the opposite direction, however; changes in the phenotype that are not correlated with changes at the genetic level are not typically considered to be

¹ See for example Merrell (1994: 16) and Silvertown and Charlesworth (2001: 53).

evolutionary changes.² Thus there is an asymmetry here between the genetic and the phenotypic level.

2.1 Adaptation

An important concept in evolutionary biology is that of adaptation. I will explain shortly the role this concept can play in neo-Darwinian evolutionary biology, but first I want to discuss the concept itself. Any trait thought to be favoured by selection – a trait that ensures the individual it belongs to survives while competitors are eliminated, and which can be inherited by offspring – is called an adaptation. Within evolutionary biology, there tends to be a distinction drawn between a trait that is an adaptation and a trait that is adaptive. An adaptation need not be adaptive, and an adaptive trait need not be an adaptation (Sober 1984). An adaptive trait is one that will tend to help the organism survive and reproduce. Calling a trait adaptive tells us nothing about whether or not that trait has been favoured by selection in the past; the term adaptive refers only to present utility. To become an adaptation, a trait must have been adaptive at one point. If adaptive traits are heritable, and the environmental conditions do not change over subsequent generations so drastically as to render the trait non-adaptive, the adaptive trait may become an adaptation. An adaptation is only rightly so called when that trait was selected for in the past. However, at some later point, this adaptation may no longer aid survival and reproduction, perhaps due to changes in the environment of the lineage. The adaptation may now be neutral with respect to natural selection, or may even become maladaptive. But so long as the trait is not fatal to the organism, it may take some time for it to disappear from the population once it ceases to be adaptive.

² As an example of a phenotypic change that does not correspond to a genetic change, take female butterflies. Whatever leaves the female butterfly hatches on and eats in her very early life will tend to be chosen as the site for laying her own eggs (Jablonka & Lamb 2005: 240). Her female offspring will repeat the cycle. However, if this plant becomes rare and is no longer locally available, female butterflies will have to find another plant for their eggs. The female butterflies that hatch on these new leaves will now lay their eggs on this new plant, and their female offspring will follow suit. There has been a significant change in the behavior of these butterflies and a change that will be inherited by their female offspring, but it was not the result of any genetic mutation. Nonetheless, this would not count as an evolutionary change if evolutionary change must involve change in gene frequencies. This sort of phenomenon will be discussed in more detail in chapter four.

Williams (1966) suggested that when we talk about a trait that has been favoured by selection – that is, an adaptation – we use the term “function” to designate the role played by the trait that ensured it was favoured by selection, and we use the term “effect” to talk about any other roles that trait may have. So in the case of the heart, its function is to pump blood, while one of its effects is to make a distinctive noise as it pumps blood.

A related distinction developed by Sober (1984) pertains to what it is that is “visible” to selection. When animal breeders select one animal over another for breeding, they do so because that animal possesses desirable properties, such as a passive manner, high milk yield or a fast running speed. In these cases, the breeders are *selecting for* properties. Similarly, natural selection selects for properties, not the bearers of those properties. However, while we can say that properties are selected for, there is *selection of* the bearers of those properties. “Selection for” pertains to those properties that ensured the organism survived and out-reproduced its conspecifics, whereas “selection of” refers to any trait that hitch-hiked along with the target of selection.³ So for example, the heart was selected *for* its ability to pump blood and so is an adaptation for this function. On the other hand, although the sound a heart makes may be adaptive (for instance, allowing for medical diagnosis), it is not considered to be an adaptation. In this case, there was selection *of* the sound of a heart beat:

To say that there is selection for a given property means that having that property causes success in survival and reproduction. But to say that a given sort of object was selected is merely to say that the result of the selection process was to increase the representation of that kind of object. (Sober 1984: 100)

So Williams’ distinction picks out a special case of Sober’s distinction. That is, while Williams’ distinction was between different sorts of properties, Sober’s distinction is broader. Williams’ account distinguishes only between properties, namely, the effects and functions of traits. Sober’s distinction, however, applies

³ For example, pleiotropy is the name given to the phenomenon whereby traits are developmentally linked. If selection favours the effects of one trait, the other trait will be propagated throughout the population too due to pleiotropy.

to the bearers of properties too; although there can only be selection *for* properties, there can be selection *of* both properties and bearers of properties. Taking Sober's and Williams' distinctions together we can say that there can only be selection for properties (functions), but that there can be selection of both properties (effects) and the bearers of properties.

Gould and Vrba (1982) have highlighted another distinction they believe is important. A trait can have numerous effects that had hitherto been neither a detriment nor an advantage to the organism; due to a new environmental context, however, such effects may become adaptive. They argue that an adaptation that has an effect (rather than function) that becomes adaptive ought to be known as an exaptation for this effect: "We suggest that such characters, evolved for other usages (or for no function at all), and later 'co-opted' for their current role, be called exaptations" (1998: 55). This approach appears to divide adaptive traits into two sorts: those that are adaptations (i.e. have a history of selection) and those that are exaptations (i.e. have no history of selection). To highlight this distinction, Gould and Vrba use the example of birds' feathers. Feathers are thought to have initially conferred an advantage due to the role they played in regulating body temperature; that is, they are thought to be adaptations *for* thermoregulation. Later they had an adaptive role to play in flight and so, in Gould and Vrba's terminology, became exaptations for flight. Once feathers became exaptations for flight, new selection pressures would have been felt that favoured certain properties of feathers over others, and these are considered to be secondary adaptations for flight:

The evolutionary history of any complex feature will probably include a sequential mixture of adaptations, primary exaptations, and secondary adaptations... and co-opted structure (an exaptation) will probably not arise perfected for its new effect. It will therefore develop secondary adaptations for the new role. (Gould & Vrba 1998: 65)

What this means is that in order to designate a trait as being an adaptation for a particular task, it is not enough to know that performing that task is adaptive. We must be able to show that it was the performance of this task that led to the propagation of the trait when it first appeared. We cannot talk of feathers having

the *function* of aiding flight; rather, they have the *effect* of aiding flight. Once feathers are used in flight, a host of secondary adaptations may arise due to the new selection pressures flight brings. New kinds of musculature and more aerodynamic shapes may be selected for the role they play in aiding flight.

Buss et al. (1998) question the usefulness of the concept of exaptation. Exaptations are adaptive traits without a history of selection and, they argue, it is unclear what relevance current utility has to evolutionary considerations:

All evolutionary explanations of the existence of species-wide mechanisms are to this extent explanations in terms of the past fitness effects of that kind of mechanism that led to the current existence of the mechanism in the species. The fact that a mechanism currently enhances fitness, by itself, cannot explain why the mechanism exists or how it is structured ... There are good reasons to think that it is not scientifically illuminating to demonstrate a feature's current correlation with fitness ... unless such correlations reveal longer term, past selective pressures. (1998: 540)

In one sense this is correct. Designating a trait an exaptation tells us only that it is currently adaptive. It is an exaptation precisely because it is adaptive without being an adaptation, and so does not require an explanation in terms of natural selection. But the lack of a selection history for these traits does not make the concept irrelevant. There is an important distinction between not needing to explain the prevalence of a trait in terms of natural selection and that trait being irrelevant to evolutionary considerations. It is important for the evolutionary biologist to be able to discriminate those traits that are in need of an explanation in terms of natural selection from those that do not. Having some concept to designate traits that are not in need of such an explanation helps us to avoid the mistake of assuming it was an adaptation for this use. That is, it gives a name to the class of traits that are adaptive, but not adaptations. Given this, the concept is not useless. It may not, however, be very useful in practice. Separating out exaptations from secondary adaptations will be very difficult. If there exists heritable adaptive variation in feathers, presumably the feathers that aid better flight will be selected, thus giving them a history of selection, and making them

adaptations, not exaptations. So do feathers cease to be exaptations for flight once there is selection for their effect on flight? And if so, how do we identify the point where a trait stops being an exaptation and starts being an adaptation? Gould might argue that the feather as a whole remains an exaptation, while specific modifications of the feather might be taken as secondary adaptations. It is not clear that biological structures can be so easily carved up into adaptations and exaptations. If we do not have clear criteria for identifying exaptations, it is hard to see how it can be practically applied.

2.2 Types of Adaptationism

Neo-Darwinism is committed to an adaptationist programme. This entails thinking about anatomy, physiology and behaviour in terms of natural selection. The adaptationist tries to identify adaptations and find explanations for why the functions of these traits were selected for. Godfrey-Smith (2001a) identifies three ways one might motivate this adaptationist programme. The first, called empirical adaptationism, rests on the claim that all or most traits possessed by an organism are in fact adaptations. That is, the majority of biological form can be explained with reference to natural selection and this situates natural selection as the most powerful force in evolution:

Natural selection is a powerful and ubiquitous force, and there are few constraints on the biological variation that fuels it. To a large degree, it is possible to predict and explain the outcome of evolutionary processes by attending only to the role played by selection. No other evolutionary factor has this degree of causal importance. (2001a: 336)

This is the empirical adaptationist position stated in its strongest terms. It is possible to be slightly more moderate by introducing some constraints on variation and allowing some other evolutionary mechanisms. For instance, Mayr weakens this claim slightly when he says that “*almost* any feature of an organism can be and has been shown to be of selective significance” (2001: 172, italics added). Natural selection provides the explanation for the vast majority of biological form. This claim is one that can only be settled by empirical research.

The second form of adaptationism identified by Godfrey-Smith is explanatory adaptationism. This position makes no claims about what is responsible for the evolution of all, or even most, of an organism's traits. It is concerned primarily with a subsection of traits – complex adaptations – and maintains that natural selection is responsible for their evolution. Although the claim that complex adaptations are primarily the product of natural selection is uncontroversial, where explanatory adaptationism makes its unique claim is in its suggestion that complex adaptations are the most important set of traits and thus ought to be the main focus of any evolutionary theory:

The apparent design of organisms, and the relations of adaptedness between organisms and their environments, are the *big questions*, the amazing facts in biology. Explaining these phenomena is the core intellectual mission of evolutionary theory. Natural selection is the key to solving these problems – selection is the *big answer*. Because it answers the biggest questions, selection has unique explanatory importance among evolutionary factors. (Godfrey-Smith 2001a: 336)

The natural question to ask about this position is why the apparent design of the organism should be the most important issue in biology. It is certainly interesting and we might well be interested enough to seek an explanation, but as Godfrey-Smith notes, this seems to be "just a fact about *us*" (2001a: 347). From this philosophical and historical standpoint, it is easy to see why natural selection should be viewed as important. Arguments from design played an important role in theological, philosophical and scientific thought. This context provided a spring board for scientific enquiry; the kinds of questions raised against this backdrop allowed natural philosophers, naturalists and scientists to ask fruitful questions and structure their research. However, this only justifies the focus on adaptations as a methodological approach (more on which shortly). Explanatory adaptationism is not a claim about what is useful for doing research, rather it is making a claim about the world: namely, adaptations are the most important features of the biological world and would be even if we were not here to witness them. Adaptations, then, demand explanation in a way other features of the biological world do not. But it is difficult to see what, besides our own aims and

interests, makes parts of the natural world more important than others in this way. Complex adaptations may be a distinct subset of traits that require an evolutionary explanation, but it is unclear why – beyond the historical reasons already mentioned, as well as perhaps current worries concerning creationism – complex adaptations should be prioritised so highly above other traits. Godfrey-Smith argues that, unlike the empirical adaptationist, the explanatory adaptationist is not making a claim that can be empirically tested; it is as much a philosophical as it is a scientific position.

The third and final form of adaptationism, according to Godfrey-Smith, is methodological adaptationism:

The best way for scientists to approach biological systems is to look for features of adaptation and good design. Adaptation is a good “organizing concept” for evolutionary research. (2001: 337)

The kinds of questions that are asked when we adopt the adaptationist programme, the methodological adaptationist argues, are productive – they lead to new knowledge about the biological world and they raise further issues for investigation. For example, Amundson (1994) points out that adaptationists have argued that thinking in terms of adaptations not only throws light on adaptations, but where we find less than optimal design, we may be able to identify developmental constraints. That is, if we find an organism that does not appear to be optimally adapted to its environment, one possible explanation for this is that there is something about the details of that organism’s developmental processes that means it cannot develop in such a way so as to be optimally adapted. We can then investigate the development of that trait and establish the nature of the developmental constraint. Thus the methodological adaptationism approach is held to be illuminating not only with respect to evolutionary biology, but also with respect to developmental biology.⁴ Thinking about the biological world in terms of adaptations has been productive, but as Godfrey-Smith points out, the

⁴ However, Amundson (1994) argues that methodological adaptationists are mistaken. They do not help developmental biologists identify developmental constraints because adaptationists and developmental biologists mean different things when they speak of developmental constraints, and given this, both are interested in identifying two distinct things.

case for methodological adaptationism rests on an inductive argument; the adaptationist programme may have been useful in the past, but that does not mean that it always will be. If it should turn out that a great many important issues are uncovered when we abandon the adaptationist programme, this would provide good evidence against it being considered the *best* way to organise research (though we may still retain a commitment to a more moderate adaptationism that views adaptation as just one of a selection of good concepts to organise research around, depending on our research interests).

These different versions of adaptationism are by no means mutually exclusive, and it is possible to hold any combination of these views. One might maintain that the adaptationist programme is the best way to structure research because most traits are adaptations, for instance. I will return to the varieties of adaptationism later, in relation to some specific claims of gene selectionists such as Dawkins (1989, 1999a). Before I do so, I will first outline the role of the gene in neo-Darwinian evolutionary theory.

3. Genes

Neo-Darwinism is committed to the idea that the gene is the only biological unit of inheritance.⁵ It is the gene that ensures offspring more closely resemble parents than strangers in ways that can be visible to natural selection. A parent and its offspring might come to possess a similar scar by coincidence, but this sort of similarity will not be inherited by future generations and so cannot have effects that selection can act on. The neo-Darwinist maintains that only traits underpinned by genes can persist in a lineage long enough for complex adaptations to arise. This legitimises the claim that genes play a uniquely important role in neo-Darwinian evolutionary biology.

In this section I want to examine in detail the role genes are thought to play in evolution and the consequences this has for thinking about development. I will first look at the issue of genetic determinism. Neo-Darwinists have been keen to stress that they are not committed to genetic determinism, despite criticisms to the contrary. Examining what the neo-Darwinist is denying in this debate will

⁵ Cultural inheritance systems have also been proposed that are thought to work in parallel with genetic inheritance, and have their own units of inheritance (e.g. Dawkins 1989). I will discuss cultural inheritance in more detail in chapter six.

enable a clearer understanding of the role the genes are thought to play. Although neo-Darwinists are committed to the gene as the unit of inheritance, there exists disagreement over whether genes are the only, or one among many, units of selection. I will discuss what is at stake here and argue that whatever position the neo-Darwinist adopts, this does not affect the way in which development is mischaracterised within neo-Darwinism. Finally, I will discuss why a commitment to the gene as the only unit of biological inheritance allows the neo-Darwinist to neglect development.

3.1 Genetic Determinism

Neo-Darwinists have been accused of genetic determinism. The following quotations certainly suggest critics have reason for their claims:

But always, without exception, living things are designed to do things that enhance the chances of their genes or copies of their genes surviving and replicating... (Ridley 1997: 18)

[Organisms'] activities are governed by genetic programs containing historically acquired information... (Mayr 1997: 21)

Now [replicators] swarm in huge colonies, safe inside gigantic lumbering robots, sealed off from the outside world, communicating with it by tortuous indirect routes, manipulating it by remote control. They are in you and me; they created us, body and mind; and their preservation is the ultimate rationale for our existence. They have come a long way, those replicators. Now they go by the name of genes, and we are their survival machines. (Dawkins 1989: 19-20)

However, neo-Darwinists have vehemently denied any genetic determinism (e.g. Dawkins 1989: 270-271). Kaplan (2000) has identified three varieties of genetic determinism that will help us to assess what is going on here.

The strongest form of genetic determinism has it that genes contain the “complete information” for the organism (Kaplan 2000: 11). If we have the DNA of the organism, then we have enough to build a perfect duplicate of that organism (or at least, we will have once the requisite technological breakthroughs have been made). This is more the stuff of science fiction than modern science and is not a claim supported by neo-Darwinists. A less stringent form of genetic determinism, what Kaplan calls “intervention is useless” determinism, acknowledges that the genes do not do all the work in building organisms. However, if a trait is discovered to be underpinned by a gene (if a gene is found “for” a trait) then the idea is that the trait will inevitably develop.⁶ No amount of intervention will alter the appearance of this trait. Again, this position is rejected by most. For instance, Evolutionary Psychologists, who adopt a neo-Darwinian approach to evolution (see chapter five), cite as motivation for their research the desire to find ways to prevent bad, or at least undesirable, behaviour (e.g. Daly & Wilson 1988). Dawkins, too, argues that “anybody can see that, as a matter of fact, genes do not control their creations” (1989: 271). Note here that although he argues genes do not *control* organisms, organisms are the *creation* of the genes. I will return to this point shortly.

The weakest and most subtle form of determinism Kaplan admits is one that denies that genes are sufficient for the development of a trait but maintains:

- (a) the genetic is the natural place to look when attempting to explain, predict, and control traits with even partial genetic etiologies, and that (b) traits with partial genetic etiologies are best understood as primarily genetic, and it is only through directed intervention that the expression

⁶ The locution “gene for X” here is taken to pick out the function of the gene (the property of the gene in virtue of which it was favoured by selection). There is a broader sense in which the “gene for” locution is also used. Knock-out experiments involve researchers mutating or otherwise dampening the causal role a gene is supposed to play in development. When this affects the development of some trait, the researchers may then talk about having found the gene for that trait. Clearly, they have identified a part of DNA that plays a causal role in the development of that trait, but they have not established whether there was selection for that gene in virtue of its having that effect. The sense of “gene for” discussed above is a subclass of this use of the phrase. In both instances, we might expect the removal of the gene in question to have some phenotypic effect, though note that there seems to be a high degree of redundancy in the genotype (Sterelny & Griffiths 1999: 128). But only where we have also provided adequate grounds to suppose that there has been selection for the gene can we talk of the gene embodying instructions for the development of that trait.

of these genes with partial genetic etiologies can be avoided or controlled. (2000: 12)

This characterizes the neo-Darwinian approach to genes best. Genes are not considered to be sufficient for the appearance of any trait; other resources must also be present. An organism may require food, oxygen, perhaps some parental care or socialisation to develop in a species-typical way. The path from gene to trait is not straightforward. Development is highly complex and can be sensitive to the slightest changes in the timings of events. Further, the appearance of a trait can be prevented by means other than the mutation of the gene(s) in question. Thalidomide can severely disrupt the development of arms and legs despite the fact the individual may have copies of all the genes involved in normal arm and leg development. All of this is readily accepted by the neo-Darwinist.

Nonetheless, there is something important about the contribution made in development by the genes that sets them apart from other developmental resources. Unlike those other resources, genes are thought to create bodies in virtue of the fact that they contain or embody programmes for development. Whether or not we wish to consider this a form of determinism, it at least constitutes a commitment to the distinct ontological status of genes within development. Although development requires many non-genetic resources, and though genetically underpinned traits may be altered by directed intervention, genes are thought to play a more important role in development than any other resource. I will now turn to the justification for this commitment.

3.2 The Gene as the Unit of Inheritance and Selection

Variation and elimination will ensure that some individuals survive and reproduce, while others do not. However, without heritability whatever trait has been selected for in one generation may not appear in the next. Heritability increases the chances that whatever advantage in survival and reproduction an organism's parents had, that organism might also possess.⁷ The neo-Darwinist is

⁷ There is never a guarantee that such an advantage will be inherited by offspring. First, offspring only resemble their parents, they are not exact duplicates. Second, whether or not a trait is advantageous may depend on the state of the environment. If the environment changes in some way from one generation to the next, what was an advantageous trait may become either neutral

committed to the idea that the only biological entity of note to be inherited from one generation by the next is the gene.⁸ More correctly, it is not any physical entity that persists between the generations, but rather copies of that entity that persist. Although errors in copying sometimes occur – introducing mutations – this copying procedure is considered to have a good degree of fidelity. This line of thought has led Dawkins to speak of the “potential near-immortality” of genes (Dawkins 1989: 35). Given this, the gene is considered to be the unit of inheritance.

Genes can be copied with a good degree of fidelity for many generations. Crick (1970) suggested that there is a one way line of influence, from the genes to the phenotype; this became known as the Central Dogma of molecular biology. That is, genes can impart information to proteins and thus bodies, but not vice versa. Genes, or at least the information they are thought to embody, are thought to be unaffected by events in the life of the organism they reside in, beyond whether or not that organism survives and reproduces. Whatever environmental interactions an organism engages in throughout its life, this will not change what its offspring can potentially inherit; there is no inheritance of characteristics acquired by an organism during its lifetime.⁹ Genes in one generation have two causal arrows pointing away from them: one points to the phenotype in that generation, the other points to the genes’ copies in the next generation. This suggests genes are the unmoved mover in development. At the very least, they play a causal role in the development of the phenotype, but they are not thought to be altered by this fact.

A minimum requirement for a trait to propagate throughout the population as a result of natural selection is that the trait should reappear in subsequent generations. And as genes are the only biological entities that are thought to be

or disadvantageous with respect to selection. That is, the offspring may inherit the trait, but not the advantage associated with it.

⁸ Because it is the genes that are inherited, neo-Darwinian evolutionary theory is committed to the process of inheritance taking place only between organisms and their descendants. However, models of natural selection do not require the link to be between parents and offspring. Both developmental systems theory, the topic of chapter four, and theories of cultural evolution, the topic of chapter six, do not assume heritable similarity applies only to parent-offspring relationships.

⁹ In chapter four, evidence which challenges this claim will be addressed, as well as the consequences of this for evolution and development.

capable of persisting over many generations, it is only those traits that are somehow underpinned by genes that can be subject to natural selection. Consequently, any trait thought to have been favoured by selection – an adaptation – must also be a trait underpinned by genes. One might also assume, on this line of thought, that most traits that showed heritable similarity might be underpinned by genes too.¹⁰ Recall that it is not only adaptive traits that can be propagated throughout a population as a result of selection. “Hitchhiker” traits, such as those developmentally linked to adaptations, may also spread throughout the population, although there will not have been selection for them. Stochastic events may also account for the spread of some traits. Given the commitment to the gene as the unit of inheritance, hitchhiker traits and traits that persist as a result of stochastic events must also be underpinned by genes. Thus, the class of traits that can be considered to be underpinned by genes is broader than just those traits that are adaptations.

It is one thing to say that the development of a trait involves genes – this much is uncontroversial – but quite another to suggest as Dawkins does that genes “manipulate” organisms for their own ends (1989: 20). Dawkins’ claim here stems from his particular version of neo-Darwinism, gene selectionism. Gene selectionists maintain that the gene is the unit of selection and arrive at this conclusion by asking the question “what benefits from selection?” They conclude that what benefits is whatever persists through the numerous rounds of selection:

Natural selection in its most general form means the differential survival of entities... and that some of the entities must be potentially capable of surviving – in the form of copies – for a significant period of evolutionary time. Small genetic units have these properties: individuals, groups and species do not. (Dawkins 1989: 33)

The thought here is that for natural selection to produce adaptations, especially complex adaptations, there must be something that is capable of persisting from one generation to the next for a large number of generations. The gene

¹⁰ Some exceptions might be made here. If the neo-Darwinist accepts a cultural inheritance system, such as memetics, then heritable similarity in behavioural traits may not have to rely exclusively on genes. See chapter six for more on this.

selectionist argues that what persists is the gene because the gene is the only biological entity that is replicated in each generation. Organisms do not replicate on this view, they merely reproduce. Offspring of sexually reproducing organisms are a mixture of both parents' traits (as well as unique developmental occurrences brought about by interactions with the environment). Asexually reproducing species are not considered to replicate either, on the grounds that their offspring do not inherit the characteristics acquired by their parents as a result of interactions with their environment.¹¹ Dawkins argues that the replicator in natural selection is the largest section of DNA not to be “shuffled into oblivion” by crossing-over in sexual reproduction (1989: 35). Although organisms only last a generation, these sections of DNA – the genes – may persist through copies; they are the entities that differentially survive. They differentially survive as a result of the effects they have in the world, or at least this is one reason why they may persist; drift may result in the persistence of some genes too. Genes that play a role in the development of an organism such that the organism is better able to survive and reproduce will be more likely to be replicated. Those genes that have a deleterious effect on the development of the organism will not be replicated.

This is complicated somewhat by neo-Darwinian accounts of kin altruism. Relatives share a certain percentage of genes, and this means that there are two ways a gene may come to be replicated. First, a gene may have an effect such that it increases the chances of reproduction for the organism in which it resides, or second, it might increase the chances for reproduction in siblings, or other genetic relatives. This second scenario may involve some level of reproductive sacrifice on the part of the organism in question, but ultimately may better ensure the replication of the gene if the relative also possesses a copy.¹² Thus, a gene which plays a role in development such that the organism or its kin are better able to hide from a predator, source food, attract mates, or any of the other things

¹¹ Note that this move is problematic. Although some acquired characteristics are not inherited, this is not always the case. Further, just as alterations to the phenotype are not all passed on, neither are alterations to the genetic material. Many mutations are removed by “proof-reading” processes which are thought to have evolved to minimize the effects of harmful mutations (Sterelny & Griffiths 1999: 63-64). Thus, neither genetic material nor phenotypes are perfectly replicated.

¹² See chapter five for further discussion of kin altruism.

an organism can do to ensure survival and reproduction for itself or its kin, will increase its chances of being replicated. That is, genes may be selected as a result of the effects they have in the world. And once such selection has taken place, these effects become, following Williams' (1966) terminology, the function of those genes. On this account, where we find an adaptation, we find a means by which genes have ensured they were replicated into future generations. And the more we suppose that organisms are the product of natural selection, the more we will be inclined to see organisms as a collection of traits that help to ensure the replication of the genes.

As well as coining “replicator” for genes and any other entity that displays the “longevity, fecundity and copying-fidelity” essential for a replicator, Dawkins suggests that organisms are “vehicles” for their genes (1989: 33). Vehicles are those things that interact with the environment with a greater or lesser rate of survival and reproduction. The organism is most commonly thought of as such a vehicle, although higher levels of biological organization might also be considered.¹³ The organism is, on this view, a way for genes to copy themselves into future generations. Adaptations are what successful genes created and which ensured they were replicated. The genes, as replicators, sit inside bodies or vehicles (or “survival machines” as Dawkins puts it) and it is the interactions of the vehicles with their broader environment that determines whether the genes are replicated or not. Hull (1980) clarified Dawkins’ distinction between replicators and vehicles by invoking a slightly different distinction between replicators and interactors. Interactors, as the name suggests, are those things that interact with the environment and influence survival and reproduction in so doing. Hull argues that the replicator/vehicle distinction suggests that genes, as the paradigm example of replicators, only replicate. But on Dawkins' account the genes have two roles; they “produce copies of themselves and... influence their own survival and the survival of their copies” (Hull 1980: 318). That is, genes are both replicators and interactors. The gene, with its dual role as replicator and interactor, can help to ensure its own replication independently, as, for example, in the case of meiotic drive (Dawkins

¹³ Dawkins (1999a) argues that the vehicle may not coincide with the skin-bound organism. His extended phenotype hypothesis will be discussed in chapter two.

1989: 235). Other interactors include organisms, and may include groups of organisms. Non-genic interactors are a means to an end for replicators; they help fix the chances of replication for genes.

Not all neo-Darwinists are gene selectionists, and there have been a number of different positions adopted on the issue of the unit of selection; however, the hierarchical model of selection appears to be the main alternative to gene selectionism in the current debate, and is a position that will feature in later chapters.¹⁴ The claim of the gene selectionist is that we can adequately explain evolution by natural selection by adopting a gene's eye view; adaptations are best explained by thinking about how those adaptations help ensure gene replication. Neo-Darwinists who reject gene selectionism do so because they maintain that the gene's eye view alone will not suffice to generate adequate explanations for evolution by natural selection, and instead they argue for a hierarchical model of selection. The idea here is that we must also attend to higher levels of biological organisation in order to understand differential survival and reproduction; selection does not just act at the level of the gene, but also at the level of the organism, populations and species (Sterelny & Griffiths 1999: 38-43). I will discuss in more detail the idea of selection acting at different levels of organisation in chapter three and six, but here I want to clarify what is *not* at stake. All sides of this debate accept that the gene is the only biological replicator.¹⁵ Further, all sides maintain that adaptations are underpinned by genes in some sense that makes the genes more significant than other developmental resources. Brandon's (1996: 58-65) distinction between units and levels of selection clarifies matters here. Those that maintain a hierarchical model of selection are concerned with the issue of the *level or levels* at which selection can work, rather than with what constitutes the unit of selection. That is, the debate concerns whether attending to the gene as interactor is sufficient, or whether interactors at higher levels of biological organisation must also be attended to in order to generate good explanations for differential reproduction and selection. Within neo-Darwinism at least, the current debate involves the level of selection.

¹⁴ See Lloyd (2001) for a more detailed analysis of the units of selection debate and the various positions adopted within it.

¹⁵ Developmental systems theory, the topic of chapter four, rejects this position, but I am concerned here only with different points of view within the neo-Darwinian framework.

The unit of selection – at least as defined by Dawkins as that which persists through replication – is not in contention.

To return briefly to the discussion above concerning the varieties of adaptationism, Godfrey-Smith (2001a) argues that Dawkins is an explanatory, rather than empirical, adaptationist. Of course, as Godfrey-Smith points out, the different forms of adaptationism are not mutually exclusive. One can believe that adaptations are the most in need of explanation (explanatory adaptationism) and also believe that the organism is comprised primarily of adaptations (empirical adaptationism). So I am not claiming that Dawkins is not an explanatory adaptationist, but rather that he is also an empirical adaptationist of sorts. Godfrey-Smith bases his claim that Dawkins is not an empirical adaptationist on Dawkins' acceptance of the neutral theory of molecular evolution, discussed above, which maintains that a good deal of evolutionary change involves mutations at the level of DNA but which are neutral with respect to the phenotype (Dawkins 1986: 302-303). In accepting the neutral theory of molecular evolution, one accepts that a good deal of evolutionary change is not to be explained by selection, and so, *prima facie*, Dawkins does not appear to adopt empirical adaptationism. However, while Dawkins might happily accept that not all evolutionary change is the result of natural selection, there is a more limited case of evolutionary change where he does adopt something like an empirical adaptationist perspective. At the level of organisms especially (or perhaps extended phenotypes, to be discussed in chapter two), Dawkins' approach maintains that natural selection is the main driver of evolutionary change. The claim here is that interactors above the gene, in particular multi-cellular organisms, are the result of natural selection. Bodies are, in fact, an adaptation for gene survival:

An individual body is a large vehicle or “survival machine” built by a gene cooperative. They cooperate because they all stand to gain from the same outcome – the survival and reproduction of the communal body – and because they constitute an important part of the environment in which natural selection works on each other. (Dawkins 1986: 192)

This form of empirical adaptationism maintains that evolutionary change at higher levels of biological organisation is primarily the result of natural selection. Empirical adaptationism of this sort may seem to be inevitable if one accepts that complex organic structures can be produced only by generations of cumulative selection. Certainly, multi-cellular organisms are highly complex entities. For instance, the degree of co-ordination required between the different organs to maintain life is difficult to overestimate. Dawkins suggests that the organism itself seems to be a complex adaptation, and complex adaptations cry out for explanation more than anything else: “whatever we chose to call the quality of being statistically-improbable-in-a-direction-specified-without-hindsight, it is an important quality that needs a special effort of explanation” (1986: 15). For Dawkins, this explanation is in terms of natural selection. Thus, Dawkins has combined explanatory adaptationism with empirical adaptationism at higher levels of biological organisation.¹⁶

Even granting that the gene is the sole unit of inheritance, this is not enough to establish why the gene is thought to play a qualitatively different role in development. That is, we can acknowledge that the gene has a special ontological status making it distinct from other developmental resources in virtue of the fact it is the unit of inheritance or the replicator, but what is it about this distinct ontological status that means that the gene plays a radically different role in development? Conceivably, the gene could possess this distinction marking it out from other developmental resources, and yet this might not affect its role in development. The following section will discuss the justification for moving from a commitment to the gene as unit of inheritance to the idea of genetic programmes.

3.3 Genetic Information

Genes are often spoken about in informational terms within neo-Darwinism.

Genes are thought to contain or embody information:

¹⁶ Note, however, that Dawkins’ is not the strongest form of empirical adaptationism, even at these higher levels of biological organisation. Dawkins is willing to accept that drift may play a role in phenotypic evolutionary change too. I will discuss the role drift is thought to play in this context in more detail in chapter two.

It is raining DNA outside. On the bank of the Oxford canal at the bottom of my garden is a large willow tree, and it is pumping downy seeds into the air... Not just any DNA, but DNA whose coded characters spell out specific instructions for building willow trees that will shed a new generation of downy seeds. Those fluffy specks are, literally, spreading instructions for making themselves... It is raining instructions out there; it is raining tree-growing, fluff-spreading algorithms. That is not a metaphor, it is the plain truth. It couldn't be any plainer if it were raining floppy discs... What lies at the heart of every living thing is not a fire, not warm breath, not a "spark of life." It is information, words, instructions. (Dawkins 1986: 111-112)

Developmental biology can be seen as the study of how information in the genome is translated into adult structure, and evolutionary biology of how the information came to be there in the first place. (Szathmáry & Maynard Smith 1995: 231)

Genes are not mere chemicals, they are instructions for building new vehicles to ensure their own replication. Genes form programmes, recipes, or plans for organisms. They may not strongly determine development, in the sense discussed above, but they play a role in development unlike any other developmental resource: "[the genome] is a set of instructions which, if faithfully obeyed in the right order and under the right conditions, will result in a body" (Dawkins 1999a: 175). But what exactly do Maynard Smith and Dawkins mean when they speak of "information" and "instructions" in this context? There are at least two ways in which we might understand this. The first relies on the causal theory of information, and the second relies on the teleosemantic theory of information. I will discuss each in turn.

3.3.1 The Causal Theory of Information

Dretske (1981) develops one way to think about the concept of information. He first distinguishes information from meaning. Information is "an objective

commodity, something whose generation, transmission, and reception do not require or in any way presuppose interpretive processes” (1981: vii). That is, information exists in the world regardless of the existence of intelligent agents. Meaning, on the other hand, is wholly dependent on intelligent agents to interpret, understand and so on. Dretske illustrates the difference between meaning and information with the example “I have a toothache” (1981: 43-44). This always has the same meaning whether or not I have a toothache. However, the utterance of the sentence only conveys information about the condition of my tooth when I do in fact have a toothache. When I do not, the utterance of the sentence contains no information whatsoever (that is, it contains no information about my tooth, though it may contain information about the fact that I am lying and so on). Often, meaning and information coincide, for the obvious reason that this allows us to communicate effectively; however, this is a contingent fact. When I deceive you, as when I say I have a toothache when I do not, my aim is for meaning and information not to coincide. Information is closely associated with truth. Something can only be information if it reliably reflects the world as it actually is. Dretske argues that the terms “false information” and “misinformation” are part of a colloquial understanding, but his “nuclear sense” of information would disallow such usage: “information is what is capable of yielding knowledge, and since knowledge requires truth, information requires it also” (1981: 45).

Dretske uses much of the terminology of communication theory, a mathematical theory of information developed by C. E. Shannon. This approach to information is concerned with the average quantities of transmitted information, rather than information content. The event or state of affairs to which the information pertains is known as the source. Information is conveyed via the signal. The person or thing that intercepts the information is known as the receiver. The channel of communication is defined as being the set of background conditions that generate only irrelevant or redundant information from the receiver’s point of view (1981: 114-115). So, for example, the temperature may be the source and a thermometer may be the receiver. A drop in the temperature will cause a drop in the level of mercury in the thermometer. The channel conditions for this will be such things as the chemical properties of mercury, the conductivity of the thermometer’s glass bulb, the atmospheric pressure, and so on. However, it is easy to see how the channel conditions, the

source and the receiver may all play each others' roles. Just as the thermometer gives us information about the temperature, the temperature can give us information about the thermometer. Similarly, the channel conditions might become the object of interest to us. The glass bulb of the thermometer may be subjected to various temperatures in order to determine the properties of the mercury inside – for example, to determine the rate of expansion of mercury as a function of temperature. Thus the relationships in this account are symmetrical and there is no privileged source of information. It is important to note the difference between an information-link and a direct causal one. If *A* contains information about *B*, it is not necessary for *B* to have caused *A*, but there must be some form of systematic dependency between the two. For example, the images on my television screen give me information about the images on your television screen when we are tuned to the same channel, even though the images on my television screen did not cause the images on your screen, but rather they are effects of a common cause.

Is this, then, the sense of information that biologists have in mind when they discuss genes? On Dretske's account, the genotype would be the source, development would be the signal, the phenotype the receiver, and all other non-genetic resources involved in development would be the channel conditions. If the channel conditions were held constant and the genotype varied, then the receiver – the phenotype – would co-vary with it. But in Dretske's account the roles of source, receiver and channel conditions may all be reversed. Anything that has a systematic dependency (co-varies reliably) with anything else may be considered a source of information about that thing. So, on this view of things, both the genotype and the environment are sources of information about the phenotype.

This does not seem to match up with how neo-Darwinists use the concept of information. To begin with, they do not speak of the environment as a source of information; the environment is treated quite differently from genes. As Maynard Smith states:

Biologists draw a distinction between two types of causal chain, genetic and environmental... the nature-nurture divide has become fundamental in biology (2000: 189).

Dretske's account of information creates a parity in terms of information between all factors involved in development. But this is exactly what neo-Darwinists deny. Speaking of the gene as an "informational unit" was intended to call attention to its distinct and privileged position in development.

A further problem for this account is that it cannot ground the idea that genetic information can be "misread" or "misrepresented" (Godfrey-Smith 1989, 2008). Neo-Darwinians talk about the phenotype misinterpreting messages from the genes, for example, when an individual born without fully developed limbs as a result of exposure to thalidomide is nonetheless said still to possess the genes for the full development of limbs. The environmental perturbation (thalidomide) is said to interfere with the genetic information being realised in the phenotype. This would mean that the information exists even when the receiver does not pick it up. But on the causal theory, information is just reliable covariation; if there is no covariation, there is no information. Thus, Dretske's causal theory of information fails to ground the sense of information employed by neo-Darwinians in two ways: first, it does not provide any justification for privileging the gene as a source of information, and second, it cannot account for misrepresentation. If a concept of information applicable only to the genetic (at least in the biological realm) is to be made respectable, it requires a theory that can do both these things.

3.3.2 The Teleosemantic Theory of Information

Maynard Smith (2000) has argued that the teleosemantic theory of information provides justification for the particular sense of information employed by neo-Darwinians in relation to genes. This theory of information relies on the concept of intentionality. Intentional mental states are a common example used to illustrate this. Such states are aimed at, or about, something else – to believe *that*, to hope *that*, etc. Intentional states do not need to represent the world as it actually is; I can believe that the sky is green. This belief is mistaken, but it is nonetheless an intentional state directed at something. Intentional states carry intentional or semantic information; my belief that the sky is green carries this sort of information even though it is not true of the world (Sterelny & Griffiths 2001: 104). Maps similarly contain this sort of information. A map of London

contains information about something, even if it happens to contain some spurious street names. Indeed, a map of an entirely fictional city has semantic information. In the biological context, intentionality is grounded in the idea that a biological entity might come to have a particular effect as a result of natural selection. We can say a gene contains information about a particular developmental outcome if ancestral copies of that gene were replicated in virtue of the fact that they ensured that developmental outcome. That is, genes contain intentional information about particular developmental outcomes. The function of the gene is what the gene contains semantic information about.

Maynard Smith (2000) asks us to contrast two scenarios. In the first, a programmer designs a programme based on the game “Fox and Geese” in which four geese try to corner the fox, while the fox tries to evade capture. In this instance, we have no difficulty assigning intentional information or instructions to those bits of code that determine the behaviour of the geese. In the second scenario, different random codes determine the behaviour of the geese. Those bits of code that lead to the geese cornering the fox are preserved, while those that fail to do this are discarded. The successful bits of code are mutated and the process repeats itself so that, through trial and error, the geese become adept at cornering the fox. Maynard Smith argues that the code directly programmed by a programmer ought to be considered equivalent, in terms of its semantic information content, as the code arrived at by this trial and error process. Armed with this new definition of information, Maynard Smith argues he can now justify talk of the gene carrying instructions for the development of the phenotype:

The DNA and proteins carry instructions, or a program, for the development of the organism; that natural selection of organisms alters the information in the genome; and finally, that genomic information is “meaningful” in that it generates an organism able to survive in the environment in which selection has acted. (2000: 190)

To talk of a gene's role in this way is to make a normative claim – it is what the organism ought to be doing. When it is failing to perform this role, something has gone wrong. It can now be said, for instance, that the purpose of a particular

set of genes is to ensure the full development of the limbs; however, environmental influences may force the phenotype to misinterpret these genes. This allows us to resolve the two issues that arose with the causal theory of information: the concept of information needed to be applicable only to genes, and the concept of information needed to allow us to talk of “misinformation” or “misrepresentation.” First, because the gene is the sole biological unit of inheritance for neo-Darwinists, it is the only developmental resource that can be selected for, and thus the only resource that can be a source of information or instruction for development. This concept of information justifies the privileged position of genes in development. Second, semantic information need not map onto anything in the world to contain information about something. This allows us to talk about developmental outcomes that never occur. So we can say an organism contains genes with instructions to build some adult trait, even if that trait is never actually realised, as for instance when an organism does not survive into adulthood. The teleosemantic theory of information, then, is taken to justify the use of terms such as information, instructions and programmes in biology in such a way that it only applies to genes and, thus, sets them apart from other developmental resources in development.

As Griffiths (2001) notes, this account hinges crucially on the claim that genes are the sole (biological) unit of inheritance. Should it be the case that more than genes are inherited, this account can no longer be used to justify the privileging of the gene. In chapter four I will argue, in line with Griffiths and others, that a great deal more than genes are inherited.

4. Development

Viewing genes as containing or embodying instructions for development has the consequence that the specific details of development can be ignored when thinking in terms of evolution. There are two sets of entities to attend to: replicators and interactors. Interactors engage with the world in such a way that there is differential survival and replication of replicators. Genes, organisms, and perhaps groups or species, have been considered interactors, while genes alone are the biological replicators. The details of how the instructions in the genes are realised in non-genetic interactors are irrelevant for thinking about selection. We do not need to know *how* the genetic instructions are followed; it is enough to

know that they are followed. All we need focus on is the instructions themselves (what is replicated), and the final product of these instructions (what interacts), usually the organism.

This allows us to think of evolution by natural selection as a change in gene frequency (with an acknowledgment that natural selection is not the only process that may cause such a change in gene frequencies). Neutral mutations aside, to track gene changes is to track phenotypic changes. The gene is three things: replicator; interactor in its own right; and, in virtue of the fact that genes represent at least the evolutionarily relevant aspects of phenotypes, a good proxy measure of interactors at higher levels of biological organisation.

Of course, there may still be work to be done at these higher levels of organisation in terms of, for instance, establishing the specific selection pressures that a population might be exposed to. But for more abstract disciplines such as population genetics, the gene is understood as a very powerful entity that allows the researcher to ignore both development and even the phenotype itself. And for those who do work at the level of the whole organism and attempt to identify adaptations and selection pressures, the details of development can also be neglected. It is enough to think about selection pressures and how they might impact on genetic replication.

Bracketing developmental concerns, or indeed any particular concerns, can be an entirely legitimate way for a scientist (or anyone else) to make a problem more tractable. However, the issue here is not merely one of neglect; as I shall be arguing in coming chapters, it is also one of misrepresenting the nature of development and, as a result, misrepresenting the relationship between development and evolution. West-Eberhard has suggested that “among the consequences of neglect of [developmental] mechanisms in modern evolutionary biology are the problems that arise when the black box of mechanism is filled with imaginary devices” (2003: 11). Genetic programmes are the imaginary devices West-Eberhard has in mind.

Despite objections to the contrary (e.g. Dawkins 1989: 271), there is a sense of genetic control being invoked whenever we talk of genetic programmes or instructions for development. It may not be the strongest sense of control which suggests that the outcome of development is inevitable, but it does suggest that genes play a more important role in development than other resources. And

because of the normative status of genetic information, even when genetic instructions are not followed, we are justified in saying those instructions were for an outcome that ought to have occurred. Genes are “in charge” in a way that nothing else is. We may say a teacher is in charge of a class, even if the children do not always do as they are told. It is the role of the teacher to be in control, just as it is the role of the gene to be in control of a certain developmental outcome. This is not genetic determinism in its strongest sense, but it does imbue genes with some guiding force in development that marks them out as qualitatively distinct from other developmental resources.

There are a number of ways of handling the fact that, on the one hand, genes have this instructive role in development, whilst on the other, development is not insensitive to environmental conditions. The genetic programme is sometimes understood as containing conditional rules for development: if in environment *A*, then develop trait *X*; if in environment *B*, then develop trait *Y*... and so on. Environments are not static, and an organism that could not respond to, for instance, changes in temperature or light conditions would not survive. Homeostatic mechanisms in the body can often compensate for such short-term fluctuations (sweating in response to high temperatures, the expansion of the pupil in response to decreased visibility, etc.). Environments can fluctuate over longer periods of time too. Food might be plentiful in one generation but scarce in the next. An organism that could develop in one way in a food-rich environment and another in a food-poor environment, such that it was adaptive in both, would have an advantage over an organism that was adaptive in only one of these environments. Thus we might suppose that we would, where they were to occur, see selection for genes that were sensitive to environmental conditions during development. Note that even though disjunctive genetic programmes allow for environmental input, there remains a sharp difference between environmental inputs and genes. The genes continue to embody instructions and they determine which inputs play a role in normal development, and what role they should play. To extend the programming metaphor, environmental inputs act as raw data for the programme to operate on. Dawkins has suggested that the idea of a genetic programme is less useful than that of a genetic recipe (e.g. Dawkins 1981: 567). This move was motivated by his desire to emphasise the fact that the gene selectionist is not advocating a return to the one gene-one trait

concept. Rather, genes are thought to have to work together to produce any complex adaptation or to create whole organisms. The relationship between a cake and a recipe is thought to be akin to the relationship between the phenotype and the genes. There is no one-to-one mapping of the final product to the set of instructions for its creation. The phenotype, like the cake, is characterised by emergent properties. This metaphor nicely illustrates the difference between the genes and other developmental resources. While the genes act as the recipe, other developmental resources act as the raw ingredients for the final product.

But although genes-as-recipes allows for the possibility of emergent effects, Dawkins does not abandon the genes-as-programmes idea either. Rather, he suggests that these programmes are not metaphorical. He suggests that his assertion that we are “survival machines – robot machines blindly programmed to preserve the selfish molecules known as genes” (1989: xxi) should be taken as literally true (1981: 572). So although the gene-centric neo-Darwinist does not appear to argue for a strong sense of genetic determinism, and allows both environmental inputs to development and emergent properties to appear throughout development, nonetheless genes are given a distinct ontological status among developmental resources. Non-genetic developmental resources are the raw material of development, while genes set the ground rules for how such raw material should contribute. The genetic recipe in no way contradicts the idea that “genes do indirectly control the manufacture of bodies” (Dawkins 1989: 23).

Genes are the “master molecules,” the unmoved mover directing development. Development is merely the realisation of the instructions embodied in the genes. Certainly, this does not have to lead to one and only one outcome and can be sensitive to the organism's environment; the phenotype will not be predictable from the genotype. But development, from an evolutionary point of view, is considered nothing more than the realisation of the genetic instructions such that an interactor can be created. Development is taken as the output of evolution and it does not tend to feedback into the evolutionary process. The result of this is that development can be neglected or black boxed.

5. Conclusion

In this chapter I identified two commitments of neo-Darwinian evolutionary theory; the first was to the adaptationist programme, while the second was to the

gene as the unit of inheritance. I discussed why this latter commitment leads to development being viewed, from an evolutionary point of view at least, as little more than the realisation of genetic instructions and, given this, why the details of developmental processes can be neglected. This commitment, and the view of development from an evolutionary perspective that is thought to follow from it, will be challenged in chapters three and four.

The neo-Darwinist might argue that development is not entirely ignored in the neo-Darwinian approach. Developmental constraints can be thought to constrain evolution by natural selection such that optimal adaptations do not emerge, and these can be of interest to the neo-Darwinian pursuing the adaptationist programme. Despite the apparent attention to development in this context, I will argue in the following chapter that the adaptationist programme is based on a flawed model of the relationship between development and natural selection and, as a result, it leads to a neglect of the full role of development in natural selection.

Chapter Two

Development and Natural Selection

1. Introduction

This chapter will focus on various characterisations of the relationship between development and evolution by natural selection. In the first part of this chapter I will outline the relationship between natural selection and developmental constraints advanced by proponents of two different approaches: neo-Darwinism and process structuralism. In chapter one I detailed two key commitments of neo-Darwinism. The first is to the adaptationist programme and the second is to the idea of the gene as “master molecule” in development. I will deal with the latter commitment in chapters three and four, but here I want to focus on the adaptationist programme. Although development tends to be neglected in neo-Darwinism, where it is likely to feature at all is in terms of developmental constraints on natural selection. The process structuralists also focus on developmental constraints. They argue that organisms are tightly interwoven systems that cannot be understood primarily as the products of natural selection. Instead, we must attend to the details of development to discover what regulates organic form. I will argue that despite many differences, neo-Darwinism and process structuralism share in common a model of the relationship between development and natural selection. Further, I will argue that this model is based on a false dichotomy whereby development and natural selection are viewed as opposing processes.

In the second part of this chapter I will discuss one way in which natural selection and development can be seen to work in tandem. I will do this by first discussing the role of the organism in its environment. Both Dawkins’ (1999a) extended phenotype hypothesis and Lewontin’s (1978, 2001) niche constructivism share in common an acknowledgement of the fact that organisms alter their environments, but these approaches make very different uses of this fact. The extended phenotype approach extends the “reach” of the gene so that it not only programmes for traits of the skin-bound organism, but also programmes for better dams, nests and so on. The basic adaptationist model remains intact on this account. The niche constructivist account, on the other hand, allows the

organism to play a role in the generation of selection pressures so that the adaptationist model no longer holds. I will argue that Lewontin offers the better account and will demonstrate how this approach can be extended such that developmental processes can also be seen to generate selection pressures. This allows development to play a contributing role to natural selection rather than being either the output of, or an opposing force to, natural selection.

2. Developmental Constraints

2.1 The Adaptationist Programme

Although neo-Darwinism is not particularly concerned with how genotypes are supposed to be realised as phenotypes, there is one context in which development does tend to be mentioned, and that is in its role as a limiting factor in evolution by natural selection. It is only the extreme empirical adaptationist (if such a person exists) who believes that there are no constraints on variation. Less extreme views acknowledge a more circumscribed space within which natural selection can work.

Dawkins (1999a) outlines the role he sees development playing in evolution and outlines the various ways in which natural selection may be thought to have been prevented from producing organisms that are perfectly suited to their environment. An important factor constraining natural selection for maximum fitness is that of historical constraints. Natural selection cannot just put a well-designed wing where an arm once was. Instead, that arm must be slowly modified over time until it becomes a wing. Natural selection does not begin with a blank slate every time a new adaptation is required, but must work with the organism as it currently exists. And importantly, every step along the way should tend to leave the organism better off than before. This will severely circumscribe the path than can be taken (if such a path can be taken at all) from arm to wing. Dawkins remarks that with such constraints, “far from expecting animals to be perfect we may wonder that anything about them works at all” (1999a: 39).

However, although Dawkins acknowledges such historical constraints, he argues they do not necessarily establish the boundaries of natural selection. Instead, he invokes drift as a means of escaping local optima, an idea first

developed by Wright (1932). Wright imagined an “adaptive landscape” in which the peaks represent phenotypes that are adapted to their environments. The higher the peak, the greater the fit between organism and environment. The valleys represent highly unsuitable, and even unviable, phenotypes. Wright suggested that organisms could end up stranded on smaller peaks – local optima – unable to “cross” this adaptive landscape through the valleys to the higher peaks. This is because every step in the cumulative selection process must produce a more adaptive phenotype than the one which came before. A lineage may evolve a sub-optimal trait, but be unable to evolve a more adaptive trait without having to first sacrifice some degree of adaptiveness. However, Wright argued that this fate could be avoided if drift was able to play a big enough role.¹⁷ Changes brought about by drift do not need to be adaptive and so this would allow lineages to break free from local optima and move towards more adaptive peaks. Dawkins suggests drift may explain why we tend to find organisms that appear to be well-designed despite historical constraints:

... animals ought to be risible monstrosities of lashed-up improvisation, top-heavy with grotesque relics of patched-over antiquity. How can we reconcile this reasonable expectation with the formidable grace of the hunting cheetah, the aerodynamic beauty of the swift, the scrupulous attention to deceptive detail of the leaf insect?... I have emphasized drift in this role. (1999a: 40)

Dawkins argues that drift may resolve this “real paradox” of apparent adaptiveness with the limitations imposed by historical constraints. On this account, although historical, developmental constraints apply, natural selection – with a little help from genetic drift – will tend to ensure adaptive traits emerge and will ameliorate any oddities that arise as a result of the details of evolutionary trajectories. Developmental constraints are understood as working in opposition to natural selection; where natural selection tends towards innovation and the production of optimal phenotypes, development acts to conserve already established phenotypes.

¹⁷ The effects of drift are thought to be negligible in larger populations, but it can have a more noticeable effect on smaller populations. See chapter one for a more detailed discussion of drift.

Another factor that may limit a lineage achieving maximum fitness concerns cost constraints. This may reduce the fit between organism and environment because although one adaptation might ease a given problem faced by an organism, it may create new problems such that the organism is worse off overall. For instance, a greater wing span might be more effective at scaring predators and competitors; however, the cost of maintaining such a wing span in terms of energy consumption might make such an innovation ultimately detrimental. For example, whatever extra food resources are gained by such a wing span may not adequately compensate for the extra energy required to develop and maintain such wings. Dawkins argues that “any view of biological optimization that denies the existence of costs and trade-offs is doomed” (1999a: 47). We cannot assume when we see a trait that it is the optimal design for the problem it was selected for. An organism has a finite supply of energy to devote to developing and maintaining traits and it may be that it is just the best solution that could be managed given the other traits that need to be maintained. Thus, the costs of other traits place a limit on the optimal design of a given trait.

Note that although the best trait for a given environment is not assumed to evolve on this account, the approach is nonetheless highly adaptationist in orientation. The presumption is that some traits may be sub-optimal because the overall adaptedness of the organism is the result of natural selection. Evolution by natural selection is thought to produce highly adapted organisms for a given environmental context. Gould and Lewontin (1979) raised this point in their critique of the adaptationist programme: where a *trait* appears less than optimally adaptive, the assumption is that this is the result of natural selection for an adaptive *organism*. That is, where one adaptationist explanation fails (a given trait is not as adaptive as we might expect), another adaptationist explanation is offered in its place (the organism must be as adaptive as we would expect). They argue that where a trait does not appear optimally adaptive, the possibility that some evolutionary processes other than natural selection might be invoked is not entertained on the adaptationist programme.¹⁸

¹⁸ Dawkins' discussion of escaping from local optima does appear to allow an evolutionary process other than natural selection to operate (drift). Though, arguably, while drift is invoked in this instance to escape a problem for natural selection from development, it is not taken as seriously in other contexts. The assertion that complex adaptations can *only* arise through

In sum, the previous state of the lineage and the various costs of traits act to constrain natural selection to a greater or lesser extent. Of course, development is essential for the realisation of the genotype; however, this aspect of development is generally taken for granted. Where development features in evolutionary considerations, it is understood as a conservative force in evolution constraining natural selection.

2.2 Process Structuralism

If all extant phenotypes were mapped onto a multi-dimensional graph representing the different traits of organisms – morphospace – we would find that organisms tend to cluster around particular points, while other areas of the graph stand quite empty. There are (at least) two questions we can ask about such a graph, and these two questions exemplify the difference between the process structuralist and the adaptationist. The adaptationist asks why phenotypes are pushed into these clusters. The adaptationist programme suggests that some phenotypes will be unsuccessful in given environments, and thus will be eliminated by natural selection. According to this account, the empty spaces on the graph represent those phenotypes unsuited to their local environment. The adaptationist programme is concerned with the push factors that explain this clustering pattern – the reasons why phenotypes tend away from these empty spaces.

The process structuralist asks a different question, namely, why are phenotypes pulled to these points on the graph? According to this approach, not all phenotypic forms we can imagine will in fact be possible largely due to facts about development. These “pull” factors make it difficult, if not impossible, for actual phenotypes to occupy those empty regions of the graph. It is these pull factors that are of interest to the process structuralists. The adaptationist and the process structuralist make different bets about what would happen to the spread of phenotypes in morphospace in the absence of selection (Amundson 1994). The adaptationist assumes the spread of phenotypes will tend to even out across morphospace. Without selection, phenotypes in what was the empty region will

cumulative natural selection is surely mistaken if drift is, at least some of the time, required to escape from local optima. Complex adaptations may, given this, be the result of both natural selection and drift.

not be eliminated; there will be nothing to sieve out all but the most lethal of mutations. Process structuralists, on the other hand, assume that in the absence of selection, the spread of phenotypes in morphospace will remain largely unchanged. Organic form is so tightly constrained by facts about development that selection has little variation to work with.¹⁹

Motivation for the process structuralist's position, as Griffiths (1996b) points out, stems from considerations of phylogenetic inertia; that is, the persistence of certain traits in a lineage long after selection pressures are considered to have been removed. The pentadactyl limb of tetrapods is considered an example of phylogenetic inertia (Griffiths 1996b: S2). This is a homologous structure (i.e. inherited from a common ancestor) in which the basic relationship between the bones that comprise it has remained largely the same. This limb is found in all amphibians and mammals. Because this limb features in such different contexts (a bat's wing and a frog's leg), it seems to be the case that the original selection pressures that produced the limb must no longer be present: “highly conserved traits are not obviously conserved by some universal selective advantage which they confer and which leads to stabilizing selection” (Griffiths 1996b: S2). This sort of consideration would seem to support the process structuralist's claim that, should selection pressures cease, the distribution of phenotypes in morphospace will not be significantly altered.

Given this then, the process structuralist maintains that, before we concern ourselves with evolution by natural selection, we ought to attend to the details of generation of form:

The fundamental assumption of the new paradigm is that form and variation are neither random nor arbitrary, and that much of evolution can be understood, not in terms of the maximisation of fitness, but

¹⁹ Depending on how big a role developmental constraints are thought to play, the adaptationist might not expect a completely even spread of phenotypes in morphospace in the absence of natural selection. Similarly, the process structuralist may allow more or less of a role for natural selection. Those that allow a bigger role for natural selection will expect to see a slightly broader spread of phenotypes in morphospace in the absence of selection. Nonetheless, even with these more moderate positions, the spread of phenotypes as envisaged by the adaptationist will look quite different to that envisaged by the process structuralist.

through the processes which generate form and variation at every level before natural selection could be said to act. (Ho 1988: 13)

Development is thus prior to natural selection. Development is not an infinitely plastic process, allowing any kind of variation to arise and be subjected to natural selection. Rather, development proceeds by quite stringent guidelines (what Webster and Goodwin (1982) call “generative laws”) such that natural selection will be seriously circumscribed by them. Understanding how organic form is generated becomes the pressing concern and natural selection the afterthought.

According to Webster and Goodwin, “the general aim of structuralist theory is to make the order of a unified system intelligible” (1982: 41). This order arises, they argue, from the combination of three ideas first discussed by Piaget: wholeness, transformation and self-regulation. The principle of wholeness is designed to refute the reductionism and atomism believed by Webster and Goodwin to characterise what they term the “evolutionary paradigm” (neo-Darwinian biology). Wholeness refers to two aspects of organisms. First, structures (organisms) are wholes that, while their component elements might change, nonetheless maintain themselves. Given this, if we want to understand the whole, it is not enough to understand its parts. Second, the structure is thought to “control” the parts insofar as certain properties of those parts arise as a result of their relation to the whole structure. That is, the parts have important relational properties. This too means that we cannot understand the structure, or organism, by examining its parts in isolation. These considerations are thought to forbid the reduction of a structure to its component parts. However, parts of structures are thought to have intrinsic as well as relational properties, thus “the structure is not 'all-powerful'” (1982: 40). Webster and Goodwin argue that, given this relationship between biological structures and their parts, structures “are not reducible to the sum of their parts,” and “cannot be understood atomistically” (1982: 40).

Any structure is thought to be a member of a set of transformations. This concept is similar to that of morphospace; there exists a range of morphologies that are physically possible, and also a range that are physically impossible for a given structure. Not all structures will share the same transformation set. Equally, however, the transformation sets may overlap at certain points so that lineages

may sometimes move into a new area of morphospace at a crucial point. The clusters in morphospace represent transformation sets, while the empty spaces represent structures that are not “coherent” (Webster & Goodwin 1982: 41). Evolutionary change is then understood as a change in structure in line with those allowed by the system of transformations to which the structure belongs. This means that evolutionary change cannot be changed from any one form to any other form. That is, there are strict limitations on possible forms:

A “random” change in any of these [developmental] factors will not result in a “random” change of structure, but in an orderly change to another possibility, another member of the system of transformations, and typical form will be conserved. (Webster & Goodwin 1982: 43)

The final principle outlined by Webster and Goodwin is that of self-regulation. This refers to the fact that structures can often maintain their integrity despite exposure to perturbations.²⁰ This is in keeping with their claim that “there is not a unique relationship between composition and form” (1982: 34). The idea here is that a similar trait can be developed by organisms with different genotypes, for example. Taking these three principles together should lead us to conclude that organisms, as structures, are “law-governed, self-organizing totalities in which the parts are in some way mutually constitutive” (1982: 40).

If we want to understand why organisms have the form they do, we need to understand what it is that creates the available forms. According to the process structuralist, the answer to this question will be found in the details of development. Once the laws governing development are known, it is argued, it will become clear what forms are allowed and what forms are forbidden, and importantly, why this is the case:

The biological domain is, therefore, conceivable as a domain which creates itself and within which general and systematic generative processes are at work. It is in terms of these law-governed processes,

²⁰ This is similar to the idea of canalization, though canalization is usually taken to be the result of natural selection.

supposing that they exist and can be theoretically formulated, that the production of the diversity of forms must be understood. (Webster & Goodwin 1982: 26)

The process structuralist is interested in explaining the generation of form and how this constrains future transformations. Nothing in this account rejects natural selection. After all, all natural selection requires is a population with heritable variations. How variations come about is entirely irrelevant so long as some of these variations can continue to appear in later generations. Nonetheless, the process structuralist account does place very strong limitations on the power of natural selection in evolution. Process structuralism, if correct, undermines the explanatory relevance of natural selection. The adaptationist assumes that thinking in terms of natural selection offers the best explanation for the anatomy, physiology and behaviour of organisms, or at least the best explanation for the most important aspects of anatomy, physiology and behaviour.²¹ Against this, the process structuralist would argue that given the starting point (the initial structure of the organism), there are only some transformations open to it, governed by “generative laws.” Out of that set of transformations, natural selection may be the reason one form was adopted over another. However, it is the laws governing structure that create the space in which natural selection can work, and these laws significantly limit the scope for change through natural selection. If we want to explain biological form, we must first map out what structures are possible and explain why. Only then can we concern ourselves with the details that thinking about natural selection can add. On this view, the “push” factor of natural selection is less explanatorily useful than the “pull” factor of developmental constraints:

... an analysis of [the] dynamic stability of life cycles can never be complete without an understanding of the generative dynamics that produces organisms of particular forms, because their intrinsic stability may play a *dominant* role in determining their abundance and their persistence. (Goodwin 2001: 132)

²¹ See chapter one for a more detailed discussion of the different types of adaptationism.

Natural selection might have a role to play in microevolutionary processes, explaining small shifts within populations, but it is not thought to be the cause of macroevolutionary trends in the history of life. The mistake adaptationists have made, according to the process structuralist is, as Sterelny puts it, to assume “macroevolution is microevolution repeatedly re-summed” (2000b: S372).

We are presented with a choice then. If we suspect that development plays an important causal role in evolutionary trends and that it does a good deal of work in explaining organic form, we may be forced to downsize the explanatory relevance of natural selection.²² If, on the other hand, we suspect that natural selection is responsible for most details of organic form, then the explanatory significance of developmental concerns may be undermined. A moderate stance between these two poles leaves us with a situation where some traits are marked off as best explained in terms of developmental constraint, while others require explanation in terms of natural selection, creating non-overlapping magisteria of development and natural selection.

2.3 The False Dichotomy

The debate between adaptationists and process structuralist is one that pitches development against natural selection. What natural selection attempts to change, development seeks to constrain. Although there may be instances of such a direct conflict between development and natural selection, this basic characterisation of development and evolution creates a false dichotomy (Griffiths 1996a, 1996b; Oyama 2000b; Schwenk & Wagner 2004).

The dichotomy is false because development need not be a conservative force, and natural selection need not be a force for change. As discussed in chapter one, selection can account for stability of form. Natural selection will tend to prevent the spread, for instance, of a deleterious mutation that arises in the population. Stabilising selection can act so as to tend a population towards a certain distribution of different phenotypic traits, or to preserve just one

²² Though note that this will not undermine the explanatory adaptationist who is not concerned with the extent to which organic form generally can be explained by natural selection, but rather views complex traits, however few there may be, as the most in need of explanation. See chapter one for a more detailed discussion.

phenotypic trait. Natural selection does not entail change and, given this, phylogenetic inertia alone does not demonstrate the existence of developmental constraints on natural selection. Although pentadactyl limbs are put to different uses in bats and frogs, for instance, this does not rule out the possibility that selection is acting to conserve them. It may be the case that new selection pressures are acting in each case. As Griffiths (1996a, 1996b) notes, there is just insufficient evidence for the claim that phylogenetic inertia is the result of developmental constraints.

Further, this spread of phenotypes in morphospace – actual phenotypic variation – may not represent all developmentally possible phenotypes. Selection acting early enough in development will ensure developmentally possible phenotypic variants never actually develop. Schwenk and Wagner worry that this means “development devolves to another case of selection,” and this means that developmentalist accounts such as the process structuralist's have “little to contribute to our understanding of phenotypic evolution” (2004: 392). However, the fact that selection *can* operate at all stages in development does not establish that it actually *does* significantly affect the range of phenotypic variation available. Rather, the fact that we cannot tell whether the phenotypic variation available is the effect of developmental constraints also implies that we cannot tell whether the variation available is the effect of selection. The adaptationist suffers from the same paucity of evidence.

Schwenk and Wagner (2004) argue that accounts of developmental constraints work on the assumption that developmental constraints are given first, and only after this does selection operate. This would imply that selection only acts on adult phenotypes, which is clearly not the case:

... selection acts on the phenotype seamlessly throughout its ontogeny. It is therefore not possible to separate mechanistically the generation of phenotypic variation during development from the action of selection because the very failure of a developmental system to produce variant phenotypes might itself be due to selection. (Schwenk & Wagner 2004: 392)

And developmental structures that may, in a particular instance, prevent the appearance of an optimally adapted trait, might themselves be the result of natural selection.

Just as natural selection can act to promote stability, development can enable evolutionary change. I will say a good more about this point below, but the basic idea here is that development does not just limit evolutionary change, it also makes it possible. A constraint is both enabling and limiting: “if the developmental regularities manifest in an organism... reflect 'constraints' then the main thing the organism has been constrained from is chaos” (Griffiths & Stotz 2000: 32). Development and evolution can act in tandem to bring about evolutionary change, or to preserve a given form. Developmental constraints and natural selection are not so easily separable.

Both process structuralism and neo-Darwinism assume a common model of the relationship between development and natural selection which envisages them as opposing forces in the history of life. We are presented with the predicament of deciding whether it is developmental factors or natural selection that best explains the form of organisms. This dichotomy is a false one. Development and natural selection are neither antagonistic forces nor are they separate domains that can be considered in isolation. The remainder of this chapter, as well as chapters three and four, will explore the ways in which development and natural selection may be integrated.

3. Organisms and Environments

In order to present another way in which the relationship between development and evolution by natural selection can be understood, I want to first explore attempts to understand the evolutionary consequences of the activity of organisms in their environments. Dawkins (1999a) and Lewontin (1978, 2001) have developed quite different approaches to incorporating the observation that organisms are not in fact passive in their environments, but often alter the environment in important ways. I will argue here that Lewontin's approach is the better of the two. Dawkins (1999a) extended phenotype hypothesis, while acknowledging the impact of the organism on the environment, fails to incorporate this point into broader considerations about natural selection. Lewontin, on the other hand, follows through on these ideas and develops a

different understanding of the process of natural selection. Although the niche constructivism approach originated by Lewontin has tended to focus on the behaviour of organisms, I will demonstrate why developmental considerations are also important here. Such an account changes the relationship between development and natural selection so that, instead of being wholly opposed processes, development may enable evolution by natural selection. Finally, I will turn to some criticisms of Lewontin developed by Godfrey-Smith (1996). While Godfrey-Smith makes clearer some distinctions in Lewontin's work, I will argue that this does not undermine Lewontin's central point.

3.1 The Extended Phenotype

Lewontin argues that the adaptationist programme is based on a model in which the organism is shaped to better fit its environment: “the organisms themselves being nothing but the passive medium through which we see the shape of the world” (Lewontin 2000: 44). In this analysis, neo-Darwinism assumes an autonomous environment against which the organism is judged (Lewontin 1978; 2001). If the organism better fits its environment than any of its conspecifics, it will tend to outlive and out-reproduce those conspecifics. On the other hand, if the organism is found wanting, it (and its genes) will die out. The organism is thought to be “honed” or “moulded” to match its environment by the process of natural selection. Those particular traits that ensure a competitive advantage for an organism over its conspecifics, and which can be inherited by offspring, are selected for and become adaptations. The environment, on this line of thought, is considered to pose problems that adaptations solve. Lewontin uses the metaphor of a lock and key – the environment is the lock and the organism, as key, must be filed to fit this lock. Ultimately, the organism is the product of its environment:

... information from the environment is present in the organism. In a few cases this is vividly literal – a frog carries a picture of its environment around on its back. Such information is usually carried by an animal in the less literal sense that a trained observer, dissecting a new animal, can reconstruct many details of its natural environment. (Dawkins 1998: 21)

There is a one-way line of influence; the autonomous environment exerts pressure onto the organism. If the organism is to survive, it must evolve to possess the correct set of adaptations. The anatomy, physiology and behaviour of the organism (or at least those aspects of anatomy, physiology and behaviour considered important by the explanatory adaptationist) are then to be explained primarily in terms of solving problems posed by ancestral environments. However, it is trivially true that organisms affect their environments. Beavers build dams that dramatically alter the landscape. Termites construct highly complex mounds that can alter the landscape, sometimes for thousands of years (Turner 2004: 345). This appears to pose a problem for anyone who relies on the adaptationist programme and its assumption of an autonomous environment.

One way the neo-Darwinist may respond to at least some of these issues is to employ Dawkins' notion of the extended phenotype. If we accept that organisms are just the genes' method of ensuring their own replication, Dawkins asks why we should assume that the effects of genes stop at the skin of the organism. Perhaps structures in the environment too might be considered part of what genes create to further their chances of replication. He imagines two beavers, each of which has genes which underpin their abilities to build and maintain a dam (Dawkins 1999a: 233-234). One beaver has a slightly different set of alleles which enable the building of a slightly better dam. Perhaps that beaver is stronger and can gather more wood, or perhaps that beaver is biased towards building dams at more advantageous points along a river. However it happens, one beaver is more adaptive than its conspecifics in virtue of building a more effective dam, and the genes which ensure this is the case have a greater chance of being replicated. The traditional neo-Darwinian way of explaining this would be to see the beaver's phenotype – its greater strength, say – as that which ensured the replication of the genes. Instead, the extended phenotype approach claims that it is the superior dam, rather than the strength of the beaver that built the dam *per se*, that ensures the genes are replicated. The beaver's strength, and the building activities that result in the better dam, are viewed as being akin to developmental processes. On the traditional neo-Darwinian view, development is usually seen as the means to the end of adaptation. Building the dam is just a continuation of the developmental trajectory on the extended phenotype

approach. It is not the greater strength of the beaver that is the adaptation, it is the dam:

Nobody has any trouble understanding the idea of genetic control of morphological differences. Nowadays few people have trouble understanding that there is, in principle, no difference between genetic control of morphology and genetic control of behaviour... The extra step from behaviour to extended phenotype... is as conceptually negligible as the step from morphology to behaviour... an animal artefact like any other phenotypic product whose variation is influenced by the gene, can be regarded as a phenotypic tool by which the gene could potentially lever itself into the next generation. (Dawkins 1999a: 199)

Dawkins argues that the only obvious objection to the extended phenotype approach would seem to be an arbitrary commitment to the skin of the organism as a barrier to the effects of genes. If we can say that possession of gene *A* fixes the chance that an organism will develop a particular physiological trait – that is, *A* is the gene for this trait – then we should also be able to talk about a gene *B* that fixes the chance that an organism will have a certain impact on its environment – that is, *B* is the gene for a particular environmental effect. These alterations to the environment of the organism by the organism do not threaten the adaptationist programme; they are merely further examples of adaptations. Just as selection sorts genes that affect physiology, it can also sort genes that affect the organism's environment. The organism has been moulded by natural selection, via its genes, to manipulate its environment just as the organism has been moulded by natural selection, via its genes, to develop a particular bodily trait. So, rather than weakening the adaptationist programme, the extended phenotype in fact strengthens the adaptationist programme by extending its explanatory reach.²³ However, while this approach acknowledges the important impact organisms have on the world around them, it only tells half the story.

²³ A consequence of adopting an extended phenotype view might be, if we take Dawkins' replicator/vehicle ontology seriously, a further diminution of the status of the organism. The vehicle, perhaps commonly thought to coincide with the boundaries laid down by the skin of the organism, might in fact include such things as beaver dams and bowers. This would further

3.2 Niche Construction

Lewontin argues that a basic error in the neo-Darwinian model of evolution by natural selection arises in that it must assume the pre-existence of niches. A niche consists of that subset of the environment that has some impact on an organism – food, predators, habitats, mates, and so on:

The ecological niche is a multidimensional description of the total environment and way of life of an organism. Its description includes physical factors, such as temperature and moisture; biological factors, such as the nature and quantity of food sources and of predators, and factors of the behaviour of the organism itself, such as its social organisation, its pattern of movement and its daily and seasonal activity cycles. (Lewontin 1978: 159)

Not every aspect of the external world exerts selective pressure on all organisms. Land-based mammals, for instance, do not typically experience selection pressure for the ability to breathe under water, and no one supposes otherwise. The adaptationist model acknowledges that different organisms, or groups of organisms, experience different sorts of selection pressures. The selection pressures a population is exposed to are created by the niche those organisms occupy. But why should some selection pressures apply to one group of organisms and not another which may be living in close proximity? The adaptationist, according to Lewontin, offers no explanation for why a population experiences a certain set of selection pressures and instead assumes that such pressures are just given: “the history of life is then the history of the coming into being of new forms that fit more and more closely into these preexistent niches” (2001: 63). But, argues Lewontin, the environment is not divided up into sets of selection pressures. The niches that exist for land-based organisms did not exist

undermine the kind of approach taken by Lewontin (see below) who argues for a greater role for the organism as actor in evolution. In chapter four, the status of the organism as a result of this view, as well as the kind of account developed by developmental systems theory, will be addressed.

prior to the evolution of land-based organisms. Niches are not given, rather they are created by the organisms that occupy them.

As mentioned above, organisms continuously interact with their environment and alter the environment in the process. It is in these activities that the organism creates its niche. Often, though not necessarily, such alterations make the environment more suitable for the organism. This suggests that when we see an organism well-suited to its environment, we are not at liberty to suppose that this is because the organism has been moulded to the environment. Rather, it may have been the environment that was moulded to the organism, or more likely, some complex two-way accommodation was reached between organism and environment.

Lewontin (2001) distinguishes four ways organisms can engage in niche construction. First, organisms make certain aspects of the environment relevant and others irrelevant. Wind currents are not particularly relevant to an organism that lives on a river bed. Cliff faces may be of little relevance to birds that nest in trees but provide vital nesting sites to other birds. Gravity may be an important aspect of the world for macro-organisms (particularly land-dwelling larger macro-organisms), but Brownian motion has far more relevance for bacteria (Lewontin 2001: 65). Different organisms can live in close proximity but have very different relationships with the world around them: “it is the life activities... that determine which parts of the world, physically accessible to all of them, are actually parts of their environments [niche]” (2001: 64).

Second, organisms alter the world around them as they make it part of their niche. Organisms can create their own resources. They may disperse the seeds of the fruit they consume, ensuring an ongoing supply of food for themselves and their offspring. *Macrotermittines*, a species of termite, gain the required nutrition from the wood they consume by cultivating fungi which partially digest the wood for them (Turner 2004). More mundanely, organisms use up resources around them and produce waste. Turner (2004) anticipates an objection to this point that he terms the “fly-in-the-soup” problem.²⁴ The objection rests on the idea that although an organism may, for instance, use oxygen and produce carbon dioxide,

²⁴ The name of this objection is inspired by the joke:
Diner: “Waiter, waiter, there's a fly in my soup!”
Waiter: “Don't worry sir, he won't eat very much” (Turner 2004: 350).

such effects will be tiny and can be safely ignored: “the capacities of environmental sources and sinks are typically so vast that any physiological draw-downs or build-ups by organisms will affect these environments only negligibly” (2004: 331). But, argues Turner, this line of argument will not work. It is not always the case that the environment in question really can swamp the effects of the activity of an organism. To use Turner’s example, the burrow of a rodent will be affected by that rodent's respiratory processes in a significant way. Further, even where the effects of a single organism on an environment may seem small, populations of such organisms can have much more significant effects. The contribution of a single tree to the percentage of oxygen in the atmosphere may be small, but the contribution of a rainforest is significant.

Third, physical signals from the world are often converted by the organism into another form. Short day plants, such as rice, can convert sunlight into chemical signals. If there is an insufficient amount of this chemical in the plant's system, this triggers the flower to open up. Sound waves hitting our ear drums are converted into different types of signals. Signals from the environment are made relevant to the organism in ways dependent on their biology.

Fourth, organisms “create a statistical pattern of environment different from the pattern in the external world” (Lewontin 2001: 64). A particular food source may fluctuate between being in plentiful supply to being scarce, but the organisms that consume this food stuff may occupy wide enough territories, seasonally migrate, or store food such that they do not experience fluctuations in food supply. Importantly, it is not the case that fluctuations in food supply have *no* effect on the organisms in question. Rather, the point here is that the organism is not subject to the *same* patterns of abundance and scarcity that characterises the external world.

We can add a fifth way that an organism can make some aspects of the world relevant. This is in terms of development. Both developmental plasticity and canalization are important here. Developmental plasticity refers to the “modifiability of morphology during development... and environmentally sensitive behavior” (West-Eberhard 2003: 35). This phenomenon is commonly discussed in terms of the environment providing “cues” or “inputs” to development processes which helps determine the course of future

development.²⁵ Plasticity is a “universal property of living matter” (West-Eberhard 2003: 34). However, the environmental cue in no way wholly determines the final developmental outcome. What counts as an environmental cue, for instance, varies from one lineage to another, and may even vary within a population. Timing is a crucial factor in development. A particular developmental resource present at one time can have massive developmental repercussions, while at another time be largely irrelevant as seen, for example, in the differing effects of thalidomide on adults and embryos. Thus, environmental cues that play a role in developmental outcomes do not adequately explain those outcomes. The rest of the developmental system makes those cues more or less relevant.

Canalization refers to the degree to which the development of a trait is buffered from environmental perturbations and is related to developmental plasticity, in that this plasticity that may underpin canalization. Ariew suggests that “the degree to which a trait is canalized is the degree to which the developmental process is bound to produce a particular endstate despite environmental fluctuations both in the development's initial state and during the course of development” (1999: 117). Canalization can be achieved in two ways. First, a different developmental trajectory may be taken to the same phenotypic end which can involve utilising different developmental resources towards the same developmental end. In this way there is a degree of redundancy in developmental systems. Here it is developmental plasticity that ensures a stably reappearing adult phenotypic trait. Because the developmental trajectory is altered and involves a different aspect of the environment, new selection pressures may be experienced at this level. However, because the adult phenotypic trait is preserved, there may be continuity of selection pressures at this level. The details of an organism's development then make some aspects of the environment relevant and others less so.

The second way a trait may be said to show (a degree of) canalization is if a particular aspect of the environment is entirely irrelevant to the organism. That feather development in birds is unaffected by the clothes on my washing line is

²⁵ See chapter four for an alternative way of framing how environmental factors interact with developing systems.

not particularly biologically interesting, but serves to make the point that developmental systems play a role in determining what parts of the external environment are relevant and thus might constitute a source of selection pressure.

Lewontin argues that because the adaptationist model of natural selection understands natural selection as the change in organisms (or, more correctly, populations of organisms) as a result of changes in environments, it misses out on an important causal relationship – the organism's effects on its niche. Instead, evolution by natural selection should be understood as the co-evolution of organisms and their environments.²⁶ This is not to undermine the role environments may play in natural selection, but rather to augment the role of the environment with the activities of the organism in a two-way interaction. Organisms do not have a free-hand:

The error is to suppose that because organisms construct their environments they can construct them arbitrarily in the manner of a science fiction writer constructing an imaginary world. The coupled equations of coevolution of organism and environment are not unconstrained. Some pathways through the organism-environment are more probable than others, precisely because there are real physical relations in the external world that constrain change. (Lewontin 2001: 65)

Lewontin's approach to niche construction demonstrates a way in which the idea of an active organism can be adopted while retaining natural selection. Selection pressures are still present in this approach, but rather than viewing them as autonomous pressures, they arise in part out of the activities and development of the organisms. This creates a coupled system and a different evolutionary dynamic than that envisaged by neo-Darwinism.

²⁶ Lewontin illustrates this point in the following way: the adaptationist model represents evolution by natural selection as $dO/dt = f(O,E)$ and $dE/dt = g(E)$, where O is the organism, t is time, and E is the environment. The organism or rather, as Olding-Smee et al. (2003) point out, a population of organisms changes over time as a function of the organism and the environment, while the environment changes over time as a function that acts solely on the environment. On Lewontin's model, the first differential equation remains the same, however the second becomes $dE/dt = g(O,E)$.

While Dawkins' extended phenotype account does make more space for the kind of organism-environment interactions discussed here, it does not ameliorate the larger point at the heart of Lewontin's critique of the adaptationist programme. The point is this: the traditional adaptationist approach tends to assume an environment in which the selective pressures generated by the environment stand distinct from the activity of the organism. While allowing the organism to alter its environment, the extended phenotype hypothesis maintains a very similar model of selection. Rather than the organism being tested against the autonomous environment, it is the extended phenotype that is tested. The boundary has been moved beyond the skin of the organism so that what was once considered the environment is now part of the extended phenotype, but this new boundary functions like the last. These altered aspects of the world – dams, nests, and so on – are part of the extended phenotype, just as cells, tissues, bones and so on are. Selection pressures now result from everything outside of the extended phenotype. This newly demarcated environment is largely independent of the extended phenotype, at least in evolutionary considerations. The boundary of the extended phenotype separates that which is moulded from that which provides the mould.

However, this does not adequately deal with the problem of presuming the pre-existence of niches. Even if we accept the extended phenotype hypothesis, we are still left with a confusion concerning why some aspects of the world external to the extended phenotype are relevant and so create selection pressures, while other aspects are not. That is, the extended phenotype will have its own niche, but the extended phenotype approach does not resolve the problem of the origin of niches; it merely redraws the boundary between phenotype and niche.

What Lewontin's account demonstrates is that the activities (and we can include developmental processes) of organisms – or even extended phenotypes – make certain features of the world relevant and, as a result, those organisms are implicated in the selective forces which act on them. The neo-Darwinian conception of natural selection, even on the kind of view extolled in *The Extended Phenotype*, maintains a two-stage process whereby (1) the environment hones a population via natural selection, after which (2) the population may go on to interact with the environment, forming extended phenotypes. The niche construction approach makes a third move: (3) those interactions with the

environment shape the selective forces acting on the population. This creates a feedback loop between the activity of the organism or extended phenotype and the selective forces the organism or extended phenotype is exposed to. The organism, or extended phenotype, becomes implicated in its own evolution.

If we take Lewontin's point then it seems as though the adaptationist programme not only overlooks important aspects of the relationship between organisms and environments, but is based on some conceptual confusion. Natural selection cannot be viewed as something that sits outside of the activity and development of organisms and their interactions with the environment. The environment cannot be thought to pose problems for the organism, or behave like a lock to which the organism must be filed to fit. The shape of the lock and key are determined by one another.

3.3 Causal and Constitutive Construction

Godfrey-Smith (1996) raises two arguments against Lewontin. The first is that Lewontin is too strong in his criticism of the adaptationist programme. The second concerns a worry for niche construction.

Godfrey-Smith distinguishes a number of different explanatory strategies that might be adopted to account for the relationship between organisms and environments. Asymmetric externalism is an explanatory strategy that explains the properties of organisms with reference to the properties of environments, but "explicitly or implicitly denies that these properties of the environment are to be explained in terms of other properties of the organic system" (1996: 132). This position is contrasted with externalism which focuses on the causal impact of the environment on the organism, but does not deny that there is a two-way causal interaction:

I understood externalists as giving positive theories about how various organic properties depend on environment; externalists were not viewed as needing to take any particular stand on whether the state of the environment depends on the activities or properties of the organic system. (1996: 132)

Asymmetric externalism will shade into externalism. An explanatory strategy can very explicitly reject the idea that organisms can influence environment in any evolutionarily significant way; however, strategies that neglect the role of the organism and adopt an implicit asymmetrical externalism can look like those that acknowledge the role of the organism, but choose to focus attention elsewhere. Lewontin suggests that the adaptationist programme, through neglect of the effects of the organism on the environment, adopts an asymmetric externalist stance. Godfrey-Smith (1996) acknowledges that many paradigm cases of adaptationist explanations demonstrate complete neglect of the role of the organism, and suggest implicit asymmetrical externalism.²⁷ However, he argues, there are other cases often dealt with by the adaptationist programme that explicitly accept an important role for the idea of co-evolution. These cases revolve around the idea of evolutionary arms races. Predators and their prey, as well as parasites and their hosts, are thought to be embroiled in such arms races. As the predator evolves to become more skilled at catching its prey, this creates selection pressures for the prey to out-manoeuvre the predator. If the prey does out-manoeuvre the predator, the prey has now created new selection pressures that will act on the predator, and so on. Given this, Godfrey-Smith argues that “for the most part, adaptationists can only be accused of neglecting rather than denying the influence of organic systems in their environments” (1996: 132). And although Lewontin argues that global neglect equates to something like Godfrey-Smith's implicit asymmetric externalism, the arms race examples demonstrate that in fact the neglect is not systematic. Rather, the adaptationist programme seems closer to externalism than asymmetric externalism.

One response to this might be to acknowledge the co-evolutionary model adopted in arms races, but maintain that this externalism applies only to a sub-section of organism-environment interactions. While the interactions of predators and prey, or hosts and parasites, will certainly account for some of the selection

²⁷ The example Godfrey-Smith cites here is the effect of pollution on the colouration of moths (1996: 136). Pollution due to the Industrial Revolution in England led to the darkening of tree trunks. Peppered moths, which had tended to have a light colouration and had been well camouflaged on the undamaged tree trunks, became easy prey. Peppered moths that had a darker colouration were suddenly at an advantage – they were now far less easily spotted by predators. The darker moths became more common in the population as a result. Since the decline of industrial pollution of this sort, the numbers of lighter coloured peppered moths have risen again.

pressures felt by the organism, they will not be the only ones. The physical, non-organic world will also exert pressures. A systematic neglect of how the organism affects the physical environment seems to characterise adaptationist explanations, thus it might be argued that in this domain – the non-organic environment – the adaptationist programme is implicitly asymmetrically externalist.

Godfrey-Smith also argues that not all the examples of activities by organisms count as genuine cases of construction, but rather that some of the cases that Lewontin picks out are in fact examples of externalism. Godfrey-Smith distinguishes between two senses of construction. The first is causal construction and this is exemplified by beavers building dams, termites their nests, and so on: “this is a matter of physical, causal intervention in the world, intervention which effects a change in external affairs” (1996: 145). The second involves the organism changing its relationship to the environment in some way, without altering any intrinsic features of the environment. This includes examples such as organisms which make aspects of the external environment relevant to them, which alter the physical signals they receive from their environment, and which transform the statistical patterns in their environment. Flowering plants that convert sunlight into chemicals, or birds that migrate do not, by these activities alone, alter any intrinsic property of the environment.²⁸ Such activities alter the relationship between the organism and its environment, but these are not instances of construction according to Godfrey-Smith’s approach. Rather, they are instances of the organism modifying itself to meet the demands of its environment, and this is an instance of externalism:

There is a difference between organic actions that make changes to intrinsic properties of external things and organic actions that do not, and a useful theoretical framework should not obscure this difference. (Godfrey-Smith 1996: 147)

²⁸ A bird that migrates may go on to alter intrinsic properties of its new environment, through building nests, predation, and so on. However, it is not migration *per se* that has effected these changes, but rather the nest building, hunting, and so on.

Godfrey-Smith has identified a genuine distinction between the sorts of cases Lewontin discusses. However, he has not undermined Lewontin's argument. While it is the case that some of the instances Lewontin has cited (and the further developmental cases I raised) involve internal accommodation of external phenomena, this in itself is not problematic. Lewontin's claim was not that environmental properties have no role to play in explaining organic properties, but rather that there is a two-way interaction between organisms and environments. Internal accommodations are to be expected on this, more aptly titled, *co-constructivist* account. To refute the niche constructivist we would need to demonstrate not that internal accommodations do not happen, but that organisms do not, in their own behaviour and development, make some aspects of the world relevant and others irrelevant. In other words, we would need to demonstrate that niches are given rather than constructed.

Further, Godfrey-Smith's argument against niche construction hinges on a confusion between two senses of "environment" at play here. Brandon (1990) uses the term "external environment" to capture the broadest sense of "environment" – all aspects of the physical and organic world. This is the sense of environment employed by Godfrey-Smith. Lewontin, on the other hand, is concerned with a much narrower conception of the environment, the niche: "the sum of all the selection pressures to which the population is exposed" (Odling-Smee et al. 2003: 419). There is a genuine difference between causal construction and constitutive construction but, from the point of view of Lewontin's account, Godfrey-Smith's conclusions are the wrong way around. For niche construction, it is not causal construction that should ultimately count as niche construction, but rather constitutive construction. This is not to say that causal construction is irrelevant, but rather causal construction matters in virtue of any relational changes that result from it. For example, an organism will emit infrared radiation, and this amounts to a causal change in its environment. Some property of the environment (the amount of infra-red radiation present) has been altered. However, it is conceivable that this has no impact on the kinds of selection pressures the organism faces; none of its predators or prey may be sensitive to radiation of this wavelength, and so on. On the other hand, a nocturnal creature may not causally construct its environment just in virtue of being active during the night; nonetheless it has a different sort of relationship

with its surroundings than a diurnal organism living in close proximity. The nocturnal organism is exposed to a different set of selection pressures than the diurnal organism as a result of its activities; they occupy different niches. For a niche construction account, causal construction matters only insofar as it alters selection pressures, and this depends on a change in the *relationship* between the organism and the environment rather than on a change in an intrinsic property of the environment *per se*. Niches are relational – they pick out only those aspects of the external environment that are relevant to the organism. Godfrey-Smith argues:

There is a difference between dealing with the world by intervention with it, and dealing with the world by effecting an internal change; a difference between adapting to an irregular terrain by acquiring better balance and more nimble feet, and adapting to an irregular terrain by laying an acre of concrete on it. (1996: 146)

We can agree that there is a difference between these cases, but argue that this is not a difference that matters for Lewontin's point. Developing more nimble feet, or laying concrete, will have the effect of altering selection pressures. Lewontin's argument is not undone by acknowledging that populations adapt to their environments, his is a *co*-evolutionary model. The point is to also acknowledge that as the population adapts to its environment, it alters its niche, and in doing this alters the selection pressures to which it is exposed. As the population changes in response to selection pressures, we have environmental properties explaining organic ones. But as the population evolves so too does the niche, and the new properties of this niche are explained in terms of the properties of the organisms inhabiting it.

Godfrey-Smith's distinction, though real enough, does not undermine Lewontin's argument. In essence, Lewontin's aim is to demonstrate that the niche – the source of selection pressures – cannot be understood in isolation from the organism and its activities. Both relational and causal construction (where the latter causes changes in the relationship between the organism and its environment) alter selection pressures, and thus both demonstrate the inadequacy of a purely externalist model of natural selection.

4. Conclusion

Process structuralism and niche constructivism offer two quite different responses to the adaptationist programme. The process structuralist finds the adaptationist programme to be incompatible with developmental concerns and suggests that natural selection does not have a role to play in explaining much of organic form. But process structuralism and the adaptationist programme share a model of the relationship between development and natural selection such that developmental constraints and natural selection are considered to act in opposition to one another. Where natural selection tends towards change, developmental constraints tend towards conservatism. What is at stake in the argument between process structuralism and neo-Darwinism is whether it is natural selection or developmental constraints that are responsible for given traits, not the nature of the relationship between development and natural selection. However, I argued here that this model of the relationship between development and natural selection underlying both process structuralism and the adaptationist programme rests on a false dichotomy, and that natural selection and development need not act in opposition to one another.

The niche constructivist also finds the adaptationist programme to be unrepresentative of the relationship between organisms and environments, notwithstanding the extended phenotype hypothesis, and rejects the adaptationist programme. However, niche construction distinguishes between the adaptationist programme and natural selection, so can retain the latter while abandoning the former. This means that, unlike the process structuralist, the niche constructivist does not need to view development and natural selection as opposing forces in evolution. Niche constructivism allows development and evolution to be viewed as deeply intertwined. Development does not constrain evolution, rather it makes evolution possible. Development, and the phenotypic traits it leads to, make some aspects of the environment relevant and act as sources of selection pressures. These selection pressures may then alter development and phenotypes such that new selection pressures are generated, and so on. Development and natural selection, far from being opposed to one another, are mutually reinforcing processes. Niche constructivism allows us to attend to development from an

evolutionary point of view without necessitating the abandonment of natural selection.

Lewontin's approach points to one way development and natural selection may be understood as contributing to the larger process of evolution. Evolutionary developmental biologists are also interested in the integration of evolution and development, but approach the issue from a different angle. This will be the subject of the next chapter.

Chapter Three

Evolutionary Developmental Biology

1. Introduction

In chapter one, I identified the two key commitments of gene-centric neo-Darwinism: (1) that the adaptationist programme offers the best way to explain (at least the most interesting aspects of) organic form and change in organic form, and (2) that the gene is the unit of inheritance. In the previous chapter I discussed problems with the predominantly externalist explanatory strategy of the adaptationist programme. Here I want to put pressure on two related ideas: first, that it is enough to track trends in gene frequencies in order to identify evolutionary changes, and second, that development can be understood as the realisation of a genetic programme. Population genetics, which concerns itself with tracking changing gene frequencies, consciously neglects development and, as I will discuss later, this may be a mistake. However the problems run deeper than just neglect. If development is not investigated, it becomes easy to mischaracterise: “among the consequences of neglect of mechanisms in modern evolutionary biology are the problems that arise when the black box of mechanism is filled with imaginary devices” (West-Eberhard 2003: 11). In this chapter and the next the genetic programme will be seen to be such an imaginary device.

Thinking in terms of genetic programmes is a way to bracket off questions about development in order to make other questions more tractable. It was assumed that if the black box of development was opened, the results would not impinge on evolutionary biology in any significant way. A complete theory would of course include these developmental details, but sidelining development was acceptable due to the assumption that phenotypes are importantly underpinned by genes. That is, developmental biology might flesh out how it is that the genotype is realised in the phenotype, but will not alter the fact that we can adequately understand evolution as a change in gene frequencies in a population and that development is the realisation of some representation of the

phenotype in the genotype.²⁹ We might ask whether gene-centric neo-Darwinism is justified in assuming that, once the black box of development is opened, it will not disrupt the picture of evolution it presents. As the relationship between development and natural selection comes under closer scrutiny in various research disciplines, this assumption looks increasingly problematic.

Evolutionary developmental biology (evo-devo) is one such research discipline which attempts to open the black box of development and this has been to the detriment of some of the assumptions about development made by the neo-Darwinist. Evo-devo aims to create a new synthesis between evolutionary theory and developmental biology. Evo-devo features considerations brought to bear by both the process structuralist and the neo-Darwinist but unlike either of these, both natural selection and development are thought to be necessary to explain biological systems and they are not construed as necessarily opposing processes. Indeed, evo-devo research attempts to understand how development makes evolution possible, and how evolution-enabling development has itself evolved.

This chapter aims to do two things. The first is to continue a line of argument from the previous chapter. Namely, development and evolution are neither irrelevant to one another, nor are they in opposition. By detailing some of evo-devo's central research questions and the results they have generated, we will further undermine the false dichotomy of the process structuralists and the adaptationists, as well as making clear the relevance of development in evolutionary considerations.

The second aim of this chapter is to lay the ground for an argument I will make in chapter four. Gilbert (2003: 349) has suggested that developmental systems theory (the position I wish to defend in chapter four) faces its biggest threat from evo-devo. In this chapter I want to make clear the commitments of evo-devo, and point to some conceptual difficulties inherent in the theoretical underpinning of this discipline. This will allow me to show in the following chapter that evo-devo and developmental systems theory do not share the same theoretical commitments and that developmental systems theory avoids the conceptual problems identified with evo-devo. The role of the gene is at the

²⁹ See chapter one for a more detailed discussion of these points.

centre of this theoretical dispute; although evo-devo rejects the idea of genetic programmes, it retains a privileged position for the gene in development and evolution. Here I will argue that evolutionary developmental biologists have not provided adequate grounds for granting genes this role in development.

2. Evolutionary Developmental Biology

Evo-devo is a relatively recent discipline. Although interest in the relationship between development and evolution is as old as interest in evolution, evo-devo has largely been made possible by advances in technology that have allowed scientists greater access to biological phenomenon at the molecular level. Wagner, for instance, suggests that “the molecular genetic revolution in developmental biology is in some respects like the invention of the electron microscope,” because “a new level of biological organization has come within the grasp of science” (2000: 95). Understanding the phenomenon witnessed at this and other levels, and relating them to evolution, is at the core of evo-devo. However, perhaps due to its recent arrival, there is no consensus on what exactly counts as evo-devo.³⁰ Different research programmes exist which relate development to evolution, but there is no agreement as to whether they are all research programmes within evo-devo, or whether some are better classified as something else (e.g. developmental evolution).

There are two general questions that can be asked about the relationship between development and evolution. The first concerns how evolution affects development. The second focuses on how development affects evolution. Some researchers work exclusively on answering one or other of these questions, though others view the two as interlinked (Robert 2002: 592). We can broadly distinguish between a number of forms of evo-devo. The first is conservative evo-devo and it is primarily concerned with the first question: providing evolutionary explanations for the new developmental processes and mechanisms that molecular genetics has identified. These mechanisms and processes “would not be different than any other character specific study of variation, like the evolution of DNA sequences or the evolution of morphological characters”

³⁰ For various taxonomies of the research programmes concerned with the relationship between evolution and development and that may or may not be classed as “evo-devo,” see Hall (2000), Gilbert (2003), and Sarkar & Robert (2003).

(Wagner 2000: 95). Offering evolutionary explanations for developmental mechanisms would extend the explanatory reach of neo-Darwinism, rather than challenge it in any way. However, even this approach forces development further into the picture than before. If we want to explain why one mechanism was favoured over another, we must at the very least attend to the role these mechanisms are playing in development.

A different version of evo-devo goes further than the last, and seeks to establish a new synthesis between developmental biology and neo-Darwinian evolutionary theory. Synthetic evo-devo is at least committed to the second question: how does development help explain evolution? This approach may also ask how evolution explains development. The aim of synthetic evo-devo is neither to abandon neo-Darwinian biology, nor to have development subsumed by it; rather, it seeks to form a genuine union between the two, creating something altogether different: “evo-devo is a synthesis of evolution and development with emergent properties not found from analysis of development or evolution alone” (Hall 2000: 177). Synthetic evo-devo is the moderate account of the relationship between evolution and development. Unlike conservative evo-devo, it seeks to place development on a par with evolution as an explanans.

More radical versions of evo-devo also exist. One such approach challenges the distinct position the gene holds in development and evolution. This position shades into developmental systems theory which will be the topic of chapter four. Here I am largely interested in the moderate, synthetic version of evo-devo as this probably ought to be considered the mainstream stance within the field.³¹ I will attempt to lay out its key commitments, with a particular eye to highlighting the differences between synthetic evo-devo and developmental systems theory to better adjudicate on the claim by Gilbert that what I am calling synthetic evo-devo is a genuine threat to developmental systems theory.

3. Synthetic Evo-Devo

Synthetic evo-devo attempts to affect a synthesis between evolution and developmental biology. That is, rather than subsuming one discipline into the

³¹ For example, two of the most well known evolutionary developmental biologists, Brian Hall (1992, 1999, 2000) and Scott Gilbert (2003), both defend this synthetic version.

other, evo-devo hopes that the two disciplines can merge to produce something new (e.g. Hall 2000). This is no easy task. Developmental biology and evolutionary biology have, for the most part, gone their separate ways since the Modern Synthesis. Each has their own distinctive research programme and language, and focuses on different aspects of the organic world. As well as establishing a coherent synthesis between such disparate disciplines, evo-devo must also make clear why paying attention to development is important for evolutionary considerations. To clarify, there are two ways evolution and development might come together. The first is merely additive. Evolutionary and developmental biology are made compatible such that the language and concepts used are consistent and both share the same broad research interests. One discipline picks up just where the other left off so that a coherent account of both evolution and development is created. The relationship between evolution and development imagined by the neo-Darwinist is of this sort. Evolutionary theory would provide all the ultimate explanations for organic form and change of form, while developmental biology would provide the proximate explanations. This scenario is not what evolutionary developmental biologists have in mind. Rather, when evolutionary theory and developmental biology combine, the hope is that the product is something unlike either of the two contributing disciplines. Neo-Darwinian evolutionary theory is not held to be wrong, on this approach, but incomplete. And in completing it, the research questions are altered such that evolutionary considerations are also developmental considerations. Godfrey-Smith (2001a) and Sterelny (2000) note the change in focus evo-devo motivates. For neo-Darwinian evolutionary biology the “\$64k question” was “why are organisms adapted?” (Sterelny 2000: 376). For evo-devo, that question becomes “how is adaptation possible?” And though Sterelny argues that developmental considerations do not pose a threat to the neo-Darwinian project, like Robert (2002) I will argue that such a change in the focus of evolutionary research constitutes a significant break with neo-Darwinian evolutionary theory.

A key focus of evo-devo research is the generation of variation. Neo-Darwinian evolutionary biology assumes variation is the result of genetic mutation. For cumulative selection leading to complex adaptations, such genetic mutations need to be translated into phenotypic variation. Although this process may be complex, neo-Darwinian theory nevertheless assumes it can be

adequately understood, from an evolutionary point of view, as the result of this mutation (perhaps in conjunction with other genes). That is, if a phenotypic variant arises in the population, it is enough to track any genetic variants that coincide with it. These genetic variants are assumed to underpin the trait in question and, from an evolutionary point of view at least, adequately account for its appearance. The picture of evolution emerging from evo-devo, however, challenges this assumption. That is, phenotypic novelty is not adequately understood as the result of genetic mutation. Indeed, phenotypic novelty may precede genetic novelty. Thus, the variation required for selection is not, on the evo-devo account, primarily the result of genetic mutation (though this will be involved) but rather the result of the reorganisation of developmental processes.

Phenotypic and developmental plasticity are phenomena that are thought to produce phenotypic variation. These forms of plasticity may be the result of selection. Selection can sometimes favour developmental mechanisms or traits that produce a range of responses appropriate to the organism's context. This, for example, is an idea often relied on by Evolutionary Psychologists. Putative psychological modules receive an input from the environment, and produce an adaptive response.³² Also, coral reef fish can undergo a female-to-male sex change in certain environments (Lutnesky 1994), and numerous organisms hibernate or migrate as seasons change and may undergo physiological changes such as extra coat growth. Canalization, discussed in chapter two, may ensure a degree of robustness in a particular trait as a result of the selection for plastic developmental mechanisms. These responses are thought to be adaptations; selection favoured mechanisms that produced plastic responses over mechanisms that produced “one size fits all” responses.

The environmentally contingent production of such adaptive responses, however, need not be an adaptation: “phenotypic accommodation is adaptive adjustment, without genetic change, of variable aspects of the phenotype following a *novel* input during development” (West-Eberhard 2005: 610, italics added). In such cases, because the input is taken to be novel, there cannot have been selection for the particular developmental outcome that occurs. That

³² See chapter five for a more detailed discussion of Evolutionary Psychology and psychological modules.

adaptive phenotypic novelty does not require genetic mutation is most famously illustrated in the case of Slijper's two-legged goat. Slijper, a Dutch veterinarian, performed a post mortem on a goat born with a congenital condition that resulted in the paralysis of its front legs. Nonetheless, the goat achieved mobility from an early age by walking upright on its hind legs:

... it had developed several behavioral and morphological specializations... including the ability to hop rapidly when disturbed, enlarged hind legs, a curved spine, and an unusually large neck. (West-Eberhard 2003: 51)

The phenotype was altered in numerous different ways such that the goat could move about with relative ease. A coherent, integrated phenotype developed, despite the fact that the normal developmental trajectory had been severely perturbed.

Coyne (2009) has suggested that such phenomena have no relevance for evolution. The novel input that disrupted the goat's development was environmental rather than genetic, and so it is assumed the novel phenotypic arrangement would not have been shared by any offspring the goat might have had. Further, Coyne argues that this example does not demonstrate “some inherent self-regulatory property of development” but rather examples like this “reflect an evolved phenomenon: natural selection has given bones and muscles the adaptive property of developing in response to stresses they experience” (2009: 383). However, this is to miss the point of the example. Dealing first with Coyne's second assertion – that the goat's phenotype merely reflected the evolved response of its bones and muscles to respond to stresses – this significantly underplays the degree to which the goat's morphology and physiology differed from its conspecifics. It is not just the case that in Slijper's goat certain muscles were better developed than in other goats through greater use. There was a significant rearrangement of parts such that aspects of the goat's novel anatomy closely resembled that of other bipedal animals, such as humans, orang-utans and kangaroos (West-Eberhard 2005: 612). Beyond that, the goat developed a novel trait, tendons that attached the thickened and elongated pelvic muscles to the newly shaped pelvic bone (West-Eberhard 2005: 611). The arrangement of the

goat's anatomy was drastically altered, and novel traits were introduced. This appears to constitute more than just the evolved capacity of bones and muscles to respond to stresses. Certainly, the evolved capacity for bones and muscles to respond to stresses will have been important, and will feature as part of the explanation, but alone this is insufficient.

This undermines Coyne's assertion that there is no evolutionary significance in this example. If the environmental inducer was no longer present, we can assume the goat's offspring would not have been bipedal, and in this sense the goat's novel morphology would not have had any evolutionary consequences. But granting this, Slijper's goat still has repercussions for how we think of development and evolution. This example appears to suggest that complex, adaptive phenotypes do not require a genetic programme to guide their development. Recall, genes are thought to contain instructions for development as a result of being selected for in virtue of the effect they have on phenotypes. Given the novelty of Slijper's goat, there can have been no selection for any genes involved in the development of this phenotype (or at least, no selection for them in virtue of their producing this sort of phenotype). A functional, integrated phenotype developed without generations of selection for genes that contained instructions for this developmental outcome. This undermines a key assumption of neo-Darwinian biology that adaptive, complex phenotypes must be adaptations. This is not to say that if such novel phenotypes arise, and if they prove to be heritable, they may not be selected for. Indeed, I will shortly discuss examples of this. Rather, the point here is just that developmental processes appear to be such that rather than requiring very small, incremental changes in order to produce complex, adaptive traits, quite big results can be had by small alterations in aspects of the developmental system.³³ Genetic programmes do not

³³ Dawkins (1998) distinguishes between two kinds of large, sudden evolutionary change: "Boeing 747" saltation and "Stretched DC-8" saltation. The DC-8 airplane was modified by elongating its fuselage to create the Stretched DC-8. The "Stretched DC-8" saltation involves something similar; an existing structure is enlarged (or perhaps made smaller) or the structure itself might be repeated: "it refers to large and sudden changes in magnitude of some biological measure, without an accompanying large increase in adaptive information" (1998: 26). This kind of saltation is considered possible by Dawkins. The Boeing 747 airplane, unlike the stretched DC-8, involved a completely new design. Thus the saltation named after it involves "a big increase in information content or complexity" (1998: 26). Dawkins argues that such an increase in complexity or information content could only be achieved gradually, as a result of cumulative selection and so rules out the possibility of "Boeing 747" saltation. Although "Stretched DC-8"

appear to be required to do all the work in explaining phenotypic variation. Developmental processes, not encoded for in the genes, appear to be able to account for novelty and variation too.

Robert (2004) uses an example with more straightforward evolutionary consequences. Geomyoid rodents (pocket gophers and kangaroo rats) possess external, fur-lined cheek pouches in which they store food. This is in contrast to other rodents which possess internal pouches that do not have a fur lining. The external pouches tend to be larger, and more efficient at storing body water, and they are thought to be a more recent evolutionary innovation (2004: 100). The standard adaptationist account might lead us to expect the external pouch – given it is both a complex adaptation and an improvement on the older internal pouch – to have arisen as the result of cumulative selection. An examination of the developmental details, however, reveals something different. Rather than the external cheek pouch appearing in small steps between the original internal pouch and the current external pouch, it looks as though the external pouch appeared in more or less its current form from the start. The mechanism that leads to the development of the cheek pouch, epithelial evagination, occurs in a slightly different place for external pouches such that the evagination process becomes entangled with the development of the lips, which in turn is connected to the development of the snout. As a result, the pouch develops with an external opening. Further, the fur lining of the external pouch develops as a consequence of the evagination process in this new location interacting with the developing facial epithelia (2004: 101). The very small change in the location of the evagination process leads to an innovative and adaptive novel phenotypic trait.

This small change in the location of the initiation of the evagination process may involve a genetic mutation. The neo-Darwinist might be tempted to explain

saltations are considered possible in neo-Darwinism, it is assumed they must be underpinned by genes in order to be evolutionarily significant. Further, they must be underpinned by genes in order that the different aspects of the organism adjust in their development and positioning: “we know that a single mutation can orchestrate changes in many diverse parts of organs” (Dawkins 1998: 27). Slijper’s goat and further examples to be discussed below do not fit easily into Dawkins’ categorisation of acceptable and unacceptable saltations. With the exception of some novel tendons, the same structures are present in Slijper’s goat. However, the difference between Slijper’s goat and a normally developed goat is not merely one of magnitude; Slijper’s goat adopted quite a different sort of body plan, not unlike that of bipedal animals. Further, such a qualitatively different, yet adaptive, phenotype was arrived at without a genetic mutation to “orchestrate” the new timing and placing of developmental processes.

this in the usual way: a genetic mutation leads to a novel phenotypic trait that selection can act on. This genetic mutation is then understood as a change in the genetic programme or recipe which issues instructions guiding development. But recall that in the case of the goat above, the phenotype that emerged could not be understood as resulting from instructions in the genetic programme. If there is no need for instructions to guide development in that case, there is no reason to suppose they are required here.

West-Eberhard (1998, 2003, 2005) argues that evolutionary change should not be assumed to begin with genetic mutation. Normal developmental trajectories can be disrupted by environmental factors. Phenotypic accommodation allows for these perturbations to alter the normal path of development, but nonetheless produce adaptive phenotypes, as in the case of Slijper's goat. If an environmental perturbation is the cause of the developmental disruption, many, if not all, organisms in the population in that generation may now possess significantly different phenotypes from previous generations, and these novel phenotypes may be adaptive. This will not always be the case, of course, and an environmental perturbation may be seriously detrimental to the organisms involved, or may be buffered against as in cases of canalization, but in at least some instances such adaptive novelties may arise. West-Eberhard suggests that this flexibility helps determine evolutionary trajectories through a process of genetic accommodation. That is, phenotypic accommodation may enable genetic accommodation. Genetic accommodation occurs where there exists individual variation in responsiveness to environmental stimuli. Some individuals may need a high degree of exposure to an environmental stimulus before this results in the development of a new phenotype, while others require far less, depending on variations in their individual biology. For instance, the environmental stimulus may be adding to a chemical already produced by the organism. Some organisms will already produce a large amount of the chemical and so will only require a small amount from the environment in order to produce the adaptive response. Other organisms in the population, due to low levels of the chemical in their system, will require more of that chemical from their environment in order to produce the same adaptive response. If the environmental stimulus is, or becomes, scarce or just tends to fluctuate, those organisms that require a smaller exposure will have a better chance of developing

the adaptive phenotype. Any organism that, through some genetic mutation, or because of some previously latent genetic capacity, requires less exposure than its conspecifics will be favoured by selection.³⁴ Generations of selection pressures acting in this direction may result in the population no longer needing very much, or any, exposure to the stimulus. A phenotype that first appeared as the result of an environmental perturbation can in later generations, through this process of genetic accommodation, develop in the absence of that environmental factor.

West-Eberhard argues that because behaviour is usually more flexible than morphology or physiology, it will tend to be what changes in an altered environment:

The produced response subjects other attributes of the phenotype to an altered selective regime (e.g., a particular behavior may produce new physiological or morphological demands). This means that given sufficient genetic variation in morphology, a recurrent behavioral response to the environment can affect the evolution of the structures affected or employed as a result. Thus, behavior being especially plastic, behavior must often take the lead in evolution. (West-Eberhard 2003: 180)

Because behaviour tends to be particularly variable, we might expect to see behavioural changes when we see environmental changes. And because organisms play a role in generating the selection pressures they are subject to, when an organism's behaviour changes, so too will the selection pressures.³⁵ West-Eberhard is proposing a two-step model of evolution by natural selection here: "developmentally mediated variation and then selection resulting in gene frequency change" (1998: 8419). How does phenotypic accommodation followed by genetic accommodation affect the relationship between development and

³⁴ This will not always be the case. If the cost of internally reproducing the effects of the environmental stimulus outweighs the benefits of the new phenotype, then we cannot assume selection will favour genetic accommodation. See chapter seven for a more detailed discussion of this point.

³⁵ See chapter two for discussion of the organism's role in generating selection pressures.

natural selection? West-Eberhard's approach may seem to marginalise development from an evolutionary perspective, leaving us with the standard neo-Darwinian picture. For example, Sterelny (2000) has suggested that developmental considerations have not yet been demonstrated to be anything other than beside the point as far as population geneticists are concerned. He argues that population genetics requires two things: that there exists heritable (i.e. genetic) variation in a population which leads to differential success in survival and reproduction, and that this success is a function of the environment the organism finds itself in. Given this, he argues that a more detailed understanding of the mechanisms that generate phenotypic variation certainly broadens our knowledge of biological phenomena, but in terms of the primary target of population genetics – tracking gene frequencies as they respond to environmental demands – knowledge of how the variation came about is superfluous to demands:

... the role of variation is acknowledged, but there is a tacit assumption that the mechanisms that generate variation will not bias or block evolutionary response to selective pressure... in this adaptationist perspective, phenotypic variability is not the cause of a particular adaptive shift, even though it is a necessary condition of any evolutionary change. (2000: S373)

However phenotypic variation comes about, what matters from the neo-Darwinian perspective is that adaptive phenotypes are propagated throughout the population, and it is selection rather than development that explains this spread. Only when the variation is underpinned by genes does it become relevant to selection in this account. If we want to understand evolution by natural selection, we can skip past the phenotypic accommodation and the developmental processes that enable this, and continue to focus on genes.

It is not clear, however, that Sterelny has succeeded in undermining the attempts of evolutionary developmental biologists to make development relevant to understanding evolution. Sterelny assumes that if developmental considerations have not been demonstrated to be essential for evolutionary considerations, then those who seek to integrate development and evolution have

failed. This is not the assumption of evolutionary developmental biologists. Indeed, evolutionary developmental biologists maintain that sometimes the population genetics approach is the most appropriate.

Wagner (2000) demonstrates when a population genetics approach suffices, and when it is inappropriate. In the case of stable sex ratios in populations, the population genetics approach is appropriate. Higher mammal populations have a sex ratio of 1:1, held in a dynamic equilibrium. Events will tend to push the population away from this ratio, but equally mechanisms exist that return the population to this point. Whenever one sex becomes rare, those remaining members of that sex have an increased fitness compared with the opposite sex. The rarer sex is more likely to meet a member of the opposite sex with which it can mate. Thus it is more likely to have offspring than the members of the abundant sex. Selection will favour any mechanism that the rarer sex possess that biases the sex of offspring in the favour of that rarer sex. This will lead to the sex ratio edging back towards a ratio of 1:1. To explain this evolutionary phenomenon, Wagner argues that the standard population genetics approach is best. This is not to deny that there are no developmental issues at stake here. Indeed, there are a great many. Sex determination involves many different, complex, and not always fully understood developmental mechanisms (Wagner 2000: 96). In fact, it is this variability in developmental mechanism which renders stabilising selection the more relevant factor in this evolutionary explanation. Different species rely on different developmental mechanisms and, even within a single species, different mechanisms play a role in sex determination. We cannot explain the general tendency towards 1:1 sex ratios with any single mechanism. On the other hand, stabilising selection is common to all scenarios:

Hence, given the variety of molecular mechanisms involved in sex determination and the many levels at which sex determination can be influenced, there is no particular developmental mechanism that specifically can account for the phenomenon of a 1:1 sex ratio. The decisive mechanisms that explain the 1:1 ratio are all realised at the population dynamic level and involve frequency-dependent fitness in favor of the rarer sex. (Wagner 2000: 96)

Wagner contrasts this case with eyespots on butterfly wings (2000: 96). The eyespot organiser, a group of cells, is thought to be central to the development of eyespots. This organiser induces surrounding cells to produce a pigment, creating the distinctive wing pattern. The eyespot organiser utilises pre-existing genes involved in the development of the anterior-posterior compartment boundary (in both butterflies and *Drosophila* wing development) to produce this phenotypic novelty. The group of cells called the eyespot organiser are thought to have endowed the genes in these butterflies with a new regulatory function. Some genetic mutations may have been involved in the appearance of this wing pattern, but to understand the evolutionary innovation we need to attend to the rearrangement of the genetic and developmental architecture underpinning this trait (Wagner et al. 2000: 821). Here, the kind of explanation that cites changes in gene frequencies is less informative: “without prior knowledge of the regulatory relationships among these genes in the a-p compartment boundary it would have been impossible to understand which genetic changes were sufficient to establish an eyespot organizer, i.e., the evolutionary novelty” (Wagner 2000: 97).

In both of these cases, selection and development are involved and important. However, in explaining these phenomena, selection and development do not play equal roles:

It is obvious in the case of the origin of eye spots patterns population genetics is much less informative than in the case of sex ratio evolution. It does not help much to say that there were one or two mutations that created eyespots and that these mutant alleles were selected. There is not much we can learn from such a statement. (Wagner 2000: 97)

Wagner suggests that developmental explanations will have more “explanatory force” than explanations in terms of population genetics in certain instances. Although multiple factors are always involved in any evolutionary phenomenon, some may be more relevant than others in terms of satisfying our explanatory goals. Robert points out that this is more than the claim that population genetics offers incomplete explanations but also that “the best explanation of evolutionary

change is not always made exclusively in terms of changes in gene frequency in a population” (2004: 98). Wagner argues for the adoption of a “shifting pluralism” which recognises “the idea that there are multiple causes and mechanisms involved in every evolutionary process and that their relative importance for the outcome of evolution shifts from situation to situation” (2000: 97).

Sterelny's suggestion that population genetics does not need developmental considerations is correct, in the sense that it can continue to ask its central research question as it has not been demonstrated to be false. But the fact that the population genetics model can ignore developmental considerations without being wrong is entirely consistent with the evo-devo argument. Sterelny misses the mark in his discussions of the relationship between development and population genetics. Evolutionary developmental biologists do not claim population genetics approaches are mistaken, or that they will always be inappropriate and ought to be replaced, but that they will be incomplete and will sometimes be inappropriate.

Population genetics, and neo-Darwinian evolutionary theory more broadly, might concern itself with the changing gene frequencies in the population, and explaining at least some of these changes as responses to environmental pressures. It tries to answer the question “why are organisms adapted?” Developmental biology tries to answer the question “how do organisms develop?” Evo-devo tries to answer the question “what is it about development that makes adapted organisms possible?” Both of the first two questions may continue to be addressed, but they will not be answering the third question. Evolutionary developmental biologists argue that the neo-Darwinian approach is incomplete, that in some situations these sorts of explanations will lack “explanatory force,” and that more productive lines of enquiry can be pursued through addressing their central research question. Evo-devo does not seek to displace the population genetics approach outright but, where appropriate, to augment it with more informative explanations and to bring a new research question to the fore.

4. Development in Evolutionary Developmental Biology

While evo-devo does not deny a role for population genetics in evolutionary theory, it diverges from neo-Darwinian biology in terms of how development is understood. In the remainder of this chapter, I will explore the evo-devo characterisation of development. I discussed above the idea that development does not appear to require a guiding programme in anything like the sense implied by Dawkins “lumbering robots” manipulated by genes (Dawkins 1976: 20). But if genetic programmes do not accurately characterise development, what does? The shifting pluralism that Wagner (2000) argues for is, as he notes, characteristic of the entire evo-devo enterprise. Organisms are understood as hierarchically organised systems that have emergent properties at each level of organisation. Thus evo-devo rejects the reductionism of the neo-Darwinian approach that assumes phenotypes are adequately explained, from an evolutionary point of view, in terms of genes. This hierarchical organisation is combined with the idea that organic systems are modular. This concept, and how it informs the evo-devo understanding of evolution, will be explored below.

Although evolutionary developmental biologists reject the idea that genes strongly control development, nonetheless the gene is held to play a more important role in development than other developmental resources. I will also explore the evo-devo justification for this claim and argue that it is insufficient.

4.1 Modularity

The concept of modularity does a lot of work in evo-devo, but it is a difficult concept to give a rigorous definition for, and it has tended to be used operationally (Robert 2004; Bolker 2000). At the very least it is taken to pick out aspects of an organism that show a high degree of internal integration while being, to some extent, dissociable from the rest of the system. Modules have been defined in terms of functions, processes and structures, and are thought to exist at all levels of organisation, from the genome to entire organs and limbs, and even to whole organisms and populations (Hall 2003; Bolker 2000: 771; West-Eberhard 2003: 62). Such modularity could account for the ability of organisms to develop integrated phenotypes despite perturbations during development. For instance, modular processes might change location but continue, more or less, to develop in their normal way as in the case of the developmental process that led to the external cheek pouches of the geomyoid

rodents. Similarly, transplanting key ectoderm and mesoderm cells from the region of a frog embryo where a forearm normally develops to some other part of another embryo results in the growth of a forearm in this new location: “other regions, when grafted, produced hind limbs, tails, hearts, kidneys, depending on their original positions in the donor embryos” (Hall 2003: 228).

Alternatively, if the development of one module is affected, the semi-independence of modules may ensure the effects of the perturbation are localised to one module, rather than being directly experienced by the entire phenotype. Other modules may have to be reorganised to compensate for the damaged module or modules, but modularity allows for such reorganisation. So for example we could imagine that in the case of Slijper's goat some event disrupted the normal development of the goat's front legs, but because biological systems are constituted by modular structures and processes, other structures and processes could be reorganised to compensate and ensure phenotypic integrity.

That organisms can be thought of as being comprised of modules should not be taken as a restatement of the atomism Gould and Lewontin (1979) claimed at work in the adaptationist programme – that is, organisms can be understood as merely a collection of adaptations and do not need to be examined at the level of the organism as a whole. Modular traits are not entirely independent from other modules at the same level:

Subindividual phenotypic components [modular traits] are only *semi-independent* because they share traits, or overlap, with other subunits [modular traits], because they have physical connections with others, and because they cannot function, survive, and reproduce on their own. (West-Eberhard 2003: 82)

Indeed, it is largely a matter of context as to what constitutes a module – modules are identified where there appears a good deal of integration in a particular structure or process, compared with the structures and processes surrounding it. This context-sensitivity may, in part, contribute to the difficulties had in attempting to define modularity compared with recognising it.

Modularity differs from the atomistic approach of neo-Darwinism in at least one other way. Biological systems are often thought of as nested hierarchies

of modules. This, argues West-Eberhard (2003), allows us to understand some apparently odd phenomena concerning homologous structures and processes. Two structures (or processes) are said to be homologous if the structures (or processes) appear to be very similar.³⁶ They are usually taken as evidence of common ancestry when discovered in two separate lineages.³⁷ Structures identified as homologous are sometimes found to be underpinned by distinct developmental mechanisms. Hall (1992) discusses such an example observed in an experiment on highly inbred mice. The mice were subjected to selection for longer tails and, after seven generations, several strains were produced with tails that had increased in length to the same extent. But this extra length had been achieved in different ways in the different strains. Some strains had a similar number of vertebrae as the first generation of mice, but those vertebrae had grown longer. In other strains, the vertebrae had remained the same length, but their number had increased:

The developmental processes producing increased tail length are profoundly different in these two lines; early respecification of basic segmentation in the former, expanded growth of elements already present in the basic body plan of the latter... Are the tails and/or vertebrae in these selected lines homologous, either with one another, or with the tails or vertebrae in the unselected parental line? (Hall 1992: 184)

West-Eberhard's suggestion is that homologous structures of this sort are composed of a nested hierarchy of modular subunits. Some of these subunits can change without it affecting the form of the homologue (West-Eberhard 2003: 60).

³⁶ There are two broad cases of homology discussed in relation to morphology. The first is serial homology. This occurs within a single organism when a structure or process is repeated at different points in the developmental system (e.g. body segments, feathers, specific cell types). I will focus here on homologies present in (at least) two distinct lineages, or in a single lineage at different times, sometimes referred to as “special homology.”

³⁷ Homologies are often defined as similarities in processes or structures *as a result of* common ancestry; however, Griffiths (2007) suggests this is misleading. Homology as a concept preceded Darwinian evolution, and has continued to be used independently of any commitment to evolutionary biology as a descriptive category within comparative anatomy.

Thus the tails of the seventh generation mice may be considered homologous, and perhaps some (though not all) of the vertebrae in their tails may be homologous too, but the developmental processes (or at least aspects thereof) that produce the tails are not homologous.

This hierarchical organisation means that when we discuss a module, we must be clear about the level of organisation we are focusing on (Bolker 2000). That final structure (adult phenotype) can be preserved while the developmental processes that lead to it are modified indicates the important role of development in evolutionary considerations. The dichotomy invoked by both neo-Darwinism and process structuralism that supposed that development conserved while selection altered is shown to be incoherent in these cases. Development both conserves and changes. The adult structure remains in place, but the developmental processes that cause the appearance of these structures are altered. The evolutionary developmental biologist maintains that it is not just genes or whole organisms that are subject to selection, but also modular developmental processes. Developmental processes are subject to evolutionary change, and thus ought not to be neglected in evolutionary theory as has been the tendency in neo-Darwinism.

A module is not to be understood as merely the sum of its parts, on the evo-devo account, but instead is characterised by emergent effects: “the emergent qualities of different levels of organization are one of the reasons why biology needs to be studied at different levels, and why molecular biology, cell biology, or genetics alone cannot solve all of the important questions of evolutionary biology” (West-Eberhard 2003: 61). Modules at all levels act to mediate the effects of genes on phenotypes (Gilbert 2003), and are understood as the “building blocks of evolution” (Schlosser 2002: 2). Modules play two roles. They may constitute units of evolution (e.g. Hall 2000; Schlosser 2002; Gilbert 2003), and they also underpin the evolvability of biological systems. That is, they are both the product of evolution, form part of the explanation for how evolution is possible, and can act as a level of selection. Identifying such modules and determining how it is they interact with one another will be crucial to understanding how organisms can come to have adaptations.

4.2 Genes in Evolutionary Developmental Biology

In the introduction to this chapter I mentioned Gilbert's claim that evo-devo was the biggest threat to developmental systems theory, the subject of the next chapter (Gilbert 2003: 349). In the following chapter I will assess this claim, but here I want to discuss the role of the gene in development and evolution which is understood in very different ways by evolutionary developmental biologists and developmental systems theorists. I will outline the evo-devo position on these issues here, and in the next chapter I will deal with the developmental systems approach and discuss how it differs.

Evolutionary developmental biologists, like developmental systems theorists, reject the strong claims made by gene-centric neo-Darwinism that suggests that, developmental constraints aside, all that is evolutionarily relevant in development can be understood in terms of genes. But unlike the stance adopted by developmental systems theorists, evolutionary developmental biologists are committed to the idea that genes are ontologically distinct from other developmental resources.³⁸ Robert (2004) calls this the gene-in-context approach. Genes are not considered able to have any causal effects on development in the absence of other developmental resources. This much is uncontroversial, though perhaps some of the rhetoric of gene-centric neo-Darwinism suggests otherwise. Nonetheless, the genes in this context play a more important role than anything else. The gene-centric neo-Darwinian accommodation of context sensitivity in development is sometimes put in terms of genetic programmes involving the conditional rules discussed in chapter one (i.e. if *X*, follow developmental pathway *A*, if *Y* follow pathway *B*). The language of genetic programme seems to suggest that development can be understood, by and large, by attending to the genes alone; a position evo-devo rejects.

The idea of the gene as the unmoved mover in development has been abandoned in the evo-devo approach. Taking the genes-in-context approach more seriously than the neo-Darwinist, the evolutionary developmental biologist

³⁸ As discussed above, evo-devo is not univocal on these matters. Some advocate an interpretation of genes in development that is closer to that of developmental systems theory. Robert et al. (2001) outline the different positions adopted within evo-devo on the issue of genetic causation in development, as well as the positions adopted on inheritance that I will discuss shortly. However, those stances within evo-devo that more closely resemble developmental systems theory are not the views offered in the mainstream presentations of the discipline (e.g. Hall 1992, Gilbert 2003, Robert et al. 2001) and, given this, I have focussed on these latter claims here.

investigates how, for example, modular organisation impacts on the expression of genes. Further, the idea that genes are regulated by other genes, as well as other elements in the cellular milieu, is taken as a live topic for research. For example, evolutionary developmental biologists point to the fact that there is little contribution from the zygote nucleus early in development. Indeed, the zygote nucleus can be removed without disrupting the initial stages of development in some species (Hall 1992: 88). The egg cytoplasm is the site of most activity in these early stages:

... it is highly misleading to view the female gamete as a large DNA molecule. The egg cytoplasm is full of all sorts of agents, both nutritive and informational, which have an important role to play in the initial development of the new soma. Indeed, key processes of early development may be initiated by cytoplasmic agents rather than nuclear ones... (Arthur 1987: 107)

But although this attention to the role of other developmental factors marks a shift from the neo-Darwinian approach, non-genetic factors in development continue to be cast in the supporting role for gene action, albeit with a stronger supporting role than before. They are a repository of information for development unlike any other developmental resource.

This qualitatively different role for genes over other developmental resources, even with the genes-in-context approach, can be seen in the treatment of homologous genes. As well as homologous structures at a higher level being (partially) realised by non-homologous processes and structures at a lower level, higher level non-homologous structures may be underpinned by lower level homologous processes and structures. For example, experimenters identified a gene called *eyeless* in *Drosophila* that, when mutated, resulted in flies with no eyes. Similar knock-out experiments, targeting a gene called *aniridia* were performed with mice, and resulted in mice without any irises. Finally, a mutation in a gene called *small eye* in humans was found to result in the development of small eyes. Since their initial identification, these genes have been renamed *Pax6* as all three genes are now thought to be homologous (Burian 2005: 221). Although the eyes of mammals differ quite considerably from the compound

eyes of *Drosophila*, and thus the eyes themselves are not considered homologous, key genetic resources for their development are nonetheless thought to be homologous:

... homologous processes can trigger the formation of organs that perform similar functions, but – because the structures produced have many features that are not shared and not derived from a common source – are not closely homologous. (Burian 2005: 252)

As well as being homologous across a surprising range of species, *Pax6* in *Drosophila* has some other surprising properties. When the *Pax6* gene is activated in other tissues of *Drosophila*, for instance in its leg or antenna, an eye will develop (Burian 2005: 221). *Pax6* is an example of a homeobox gene. Homeobox genes are involved in the establishment of the basic body plan of a wide range of organisms, including animals, plants and fungi, and lead to strange developmental outcomes when experimentally manipulated. Mutations in the homeobox genes in *Drosophila* have led to legs developing where we would normally expect an antenna, and a rearrangement of the thoracic segments (Robert 2004: 27). This has led some to consider *Pax6* specifically, and the homeobox genes more generally, “master control genes” (Robert 2004: 27; Burian 2005: 224).

This privileging of the gene can also be seen when non-genetic factors and epigenetic processes are defined in terms of genes, and the activity of genes they enable. Hall's definition of epigenetics is illustrative of this genes-in-context approach:

Epigenetics or epigenetic control is the sum of the genetic and non-genetic factors acting upon cells to selectively control the gene expression that produces increasing phenotypic complexity during development. (1992: 89)

Primarily it is their role as enablers of gene expression that renders epigenetic processes important to development.³⁹ Genes may rely crucially on other aspects of the developmental system. They are not just turned on or off by this context; their effect on development is also influenced by the context. That is, context matters both up and downstream from gene expression. Nonetheless, it is gene expression that is credited with producing “increasing phenotypic complexity during development.”

One reason to suppose that evolutionary developmental biologists such as Hall maintain a qualitatively distinct role for genes in development might be due to the acceptance of a degree of preformationism in his account of development:

In one sense, epigenesis has triumphed for embryonic structures are not all preformed in the egg. Yet, in another sense, preformation “explains” some aspects of development. The genetic basis for development lies preformed in the DNA of the egg and subsequently in the zygote. The basic raw material for protein synthesis is preformed in the ribosomes and endoplasmic reticula of the egg. (Hall 1992: 86)

... fundamental developmental processes... are controlled by epigenetic expression of preformed information. (1999: 115)

Preformationism is an old idea in the history of biology that maintains that development involves the growth of existing structures and processes. Epigenesis, on the other hand, maintains that development involves the emergence of structures and processes. Gene-centric neo-Darwinism has been characterised as “neo-preformationist” (e.g. Griffiths & Knight 1998: 225; Oyama et al. 2001b: 4). Although structures and processes found in adult phenotypes are not thought to be contained in the genes, the instructions for their realisation are; there is some “representation” of the adult phenotype already present in the genotype.⁴⁰

³⁹ Waddington (1942) was the first to coin the term “epigenetics.” In common with Hall's interpretation, he emphasised the role of the genes as “primary determinants in development” (Robert et al. 2001: 956).

⁴⁰ See chapter one for an explanation for why genes are thought to represent the phenotype.

If genes are understood as embodying preformed instructions for development, it is clear why they are held as a class apart from other developmental resources. But does this claim undermine the evo-devo commitment to development involving more than just genes? Unlike the preformationism of gene-centric neo-Darwinism, preformed information is not sufficient for explaining development: “early development is a mixture of preformed and epigenetic information” (Hall 1999: 115). Development remains characterised by emergence at different levels of organisation, and cannot be reduced to this preformed information. So whatever sense of information is being employed in Hall's account, it is not complete information, or even near-complete information, for the realisation of the adult phenotype.

Perhaps Hall and like-minded evolutionary developmental biologists intend genes to embody preformed information about much more immediate developmental outcomes than adult phenotypes:

Genes have homes: nucleus, cytoplasm, cell, and so forth. Just as our individual influences and interactions change the further we venture from our homes (street, community, suburb, town, municipality, province/state, country), so the gene's influence varies with distance from its locus in the nucleus. The gene lives, and has a home, and plays an active role as homemaker, albeit with household help. (Hall 2001:228)

The idea here may be that genes contain or embody semantic information about very local phenomena. At higher levels of organisation, however, emergent properties may be considered to develop from the complex causal interactions between all developmental resources. That is, in virtue of their role as bearers of semantic information, genes continue to have a distinct ontological status from other developmental resources. Nonetheless, because developmental outcomes often involve more than genes and have emergent properties, genes cannot be considered to control or programme for these outcomes. Indeed, Dawkins' recipe analogy seems more appropriate here than “programme,” but note that this might suggest that the evolutionary developmental biologist is actually quite close to the gene-centric neo-Darwinian position after all. We might respond to this by arguing that although Dawkins has claimed that genes are better understood as

“recipes” than “programmes,” this has not had any impact on how Dawkins views the relationship between natural selection and development. Evolutionary developmental biologists, on the other hand, have taken the fact that there is more to development than genes seriously from an evolutionary perspective. Thus, although there might be some agreement on a metaphor for development, this does not translate into agreement on a good deal else.

What reasons might an evolutionary developmental biologist have for privileging the gene in development? This may be explained in part by the fact that a good deal of time and energy has already been spent attempting to understand what genes do, meaning that their effects are better understood than a good deal else in the cellular milieu and giving an inflated sense of their importance in development:

Given the ever-growing diversity of molecular tools and techniques, as well as the veritable explosion of information they produce, it is small wonder that [evo-devo] focuses so heavily on gene regulation and changes in gene expression... (Robert et al. 2001: 960)

Another factor responsible for viewing genes as set apart from other developmental resources and may explain their distinctive informational role in development is that genes are considered to be the sole unit of heredity transmission for many evolutionary developmental biologists. Evolutionary developmental biologists acknowledge that a good deal more is inherited by the zygote than naked DNA: “organisms do not start life as naked DNA... the zygote is not a blank slate on which zygotic genes alone can write their instructions” (Hall 1998: 202-203). However, beyond DNA, what is inherited is characterised by Hall as either epigenetic potential, or products of maternal gene activity. Epigenetic potential refers to the ability of genes and their products to interact with their surrounding milieu in order to produce phenotypes:

Epigenetic *potential* is heritable, epigenetic *processes* are not. There is an important difference between the ability to respond to environmental and genetic cues (epigenetic potential) and the actual responses to those cues (epigenetic processes). (Robert et al. 2001: 960).

Products of maternal gene activity include cell membranes, organelles, cytoplasmic polarities and all those other structures that comprise the zygote:

... maternal cytoplasmic control, preformed organelles, and their role in spatial segregation at the initiation of development, are viewed as epigenetically inherited only if we forget that parental (usually maternal) genomes produced them. (Hall 1998: 203)⁴¹

Although the evolutionary developmental biologist grants that more than DNA is passed to the offspring, only the genes constitute units of hereditary transmission. This is not to say that those other entities are causally inert and so do not play a role in ensuring heritable similarity. Rather, the claim here seems to be that these other entities are ultimately the products of (maternal) genes. So, on the one hand, evolutionary developmental biologists such as Hall argue that developmental outcomes are emergent, the result of complex causal interactions of developmental resources, and specifically, that genes require the activity of other developmental resources in order to have any effects. On the other hand, evolutionary developmental biologists allow that there are some cases where genes play the dominant role, such as in the production of organelles inherited by offspring. This suggests that cellular phenomena can still be understood primarily in terms of genes, an approach rejected for extra-cellular phenomena. This accords with the quotation from Hall above concerning the ever-decreasing power of the gene as we move further away from its “home.”

The justification for this presumably stems from the idea that genes are instructional in a way that other developmental resources are not. As argued by

⁴¹ Hall here mentions “preformed organelles.” There are perhaps two senses in which this is meant. The first is that if we accept that the zygote contains some structures from the beginning of its existence, then before development begins, there will be structures with form, i.e. development will begin with already formed structures. This is a very trivial sense of preformationism. The point is that, according to those that reject neo-preformationism, from this starting point development involves the emergence of new structures and processes, rather than merely the growth of pre-existent structures. The second sense in which these organelles might be considered preformed is that, as products of the maternal genome, they may be acting as bearers of maternal genetic information, and thus embody preformed semantic information in the same way genes are thought to.

Maynard-Smith (2000a) and outlined in chapter one, genes are construed as informational because they are considered the sole unit of inheritance and so are the only entities that can have a history of selection. This history of selection allowed neo-Darwinists to talk about genes representing phenotypic traits. But evolutionary developmental biologists acknowledge that genes are not the only unit of inheritance. Rather, they maintain that genes are the only unit of hereditary transmission; while numerous entities are contained within the zygote, it is only genes that underpin heritable similarity. Those other inherited entities are not causally inert in development, but instead they can be construed as gene products because genes play this instructional role in development. Whatever role these other entities play, they are realising genetic instructions of some sort. But this argument seems to be circular. The evolutionary developmental biologist's argument seems to be that because genes have a special role to play in development, this means other inherited entities can be construed as gene products or messengers carrying genetic information, thus allowing us to draw a distinction between units of inheritance and units of hereditary transmission. But the justification for the instructional role of genes stemmed from the fact they were considered to have a unique role as units of inheritance. We can ignore other inherited entities as distinct units of hereditary transmission because genes are privileged in development, but genes are supposed to be privileged because only they are inherited. In the absence of any other argument for why genes alone should be considered to have an instructional role in development, the distinction between those things that are merely units of inheritance, and those things that are also units of hereditary transmission seems unwarranted.

If more than DNA underpins heritable similarity, we undermine the special status of the gene as the bearer of semantic information and thus the gene's status as either the ultimate driver of, or as playing a privileged role in, development. If an entity becomes a bearer of semantic information because it was selected for, then either everything that has a history of being selected for becomes a bearer of semantic information, or we ought to abandon the concept altogether and attempt to justify prioritising genes in some other way (Sterelny et al. 1996; Sterelny & Griffiths 1999; Griffiths 2001; Godfrey-Smith 2008).

6. Conclusion

Evo-devo has been responsible for some remarkable discoveries and, in eschewing the dichotomy outlined in chapter two, it has highlighted the important work developmental considerations can do for evolutionary theory. However, it faces difficulties in justifying the role it grants genes in development. I mentioned here that Gilbert (2003) argued that evo-devo had successfully occupied the territory developmental systems theorists had established for themselves. In the following chapter I will outline the developmental systems approach and demonstrate the ways in which it differs from evo-devo, and in doing so will deny Gilbert's assertion. In particular, evo-devo and evolutionary developmental biology differ in how the gene is understood, and given the difficulties discussed here with the evo-devo characterisation of the role of genes in development, this will stand in developmental systems theory's favour. Nonetheless, evo-devo has had some remarkable successes and I will argue that developmental systems theory can be enriched by attending to some of the work of evolutionary developmental biologists.

Chapter Four

Developmental Systems Theory

1. Introduction

Developmental systems theory (DST), like evo-devo, attempts to bring developmental considerations back into evolutionary theory. However, as I will show here, it does so in a way quite different to evo-devo. There are two broad strands to DST. The first strand, what Dupré terms its “negative phase” (2008: 174), casts a critical eye over the neo-Darwinian approaches to evolution and development. Those writing in the DST tradition have developed detailed critiques of many of the implicit and explicit assumptions in mainstream evolutionary biology, particularly those that have resulted in the neglect or mischaracterisation of development. The second, constructive strand is concerned with developing an alternative way to frame the relationship between development and evolution that avoids the problematic issues identified in the first strand. This second strand is often either misunderstood or overlooked. It is misunderstood when critics accuse DST of being unworkably holistic such that it is of no use to science. It is also misunderstood when critics accuse developmental systems theorists of arguing that genes are unimportant. Both of these misunderstandings will be addressed here. The constructive strand is also neglected, for instance, when Gilbert (2003) suggests that evo-devo has rendered DST irrelevant because evo-devo corrects the mistakes in neo-Darwinian biology and also does what DST calls for; namely, evo-devo is taking development seriously. Certainly, evo-devo goes some way to addressing the neglect of development but, as I argued in chapter three, evo-devo is not without its own problems. Rather than pitting evo-devo against DST, however, I will argue here that not only do evo-devo and DST have something to contribute to one another, but also that productive dialogue between the two is already taking place.

2. Developmental Systems Theory

Developmental systems theory offers a very different approach to development and evolution than that proposed by neo-Darwinists. Both DST and neo-Darwinism are ultimately interested in answering the same question: how do we

explain the morphology, physiology and behaviour of organisms? However, they approach this question from very different angles. Chapters one and two detailed one way in which neo-Darwinism allowed development to be neglected in evolutionary theory; outside of their role as a limiting factor in evolution, the details of development are irrelevant to the project of identifying and explaining adaptations and tracking gene frequencies on this account. This view stems from the two commitments of neo-Darwinism: the adaptationist programme and the gene as the unit of inheritance. If evolution is what is to be explained, and development is an *effect* of evolution, development can be viewed as lying outside of the research concerns of neo-Darwinism. Evo-devo, discussed in chapter three, emphasises the role of development as a causal factor in evolution, but remains tied to the idea that the gene is the unit of hereditary transmission. DST, on the other hand, rejects both commitments of neo-Darwinism. It adopts the co-constitutive approach developed by Lewontin (1978, 2001) discussed in chapter two, and extends the notion of inheritance far beyond the gene. This leads to a very different model of development and evolution. In particular, by extending the notion of inheritance, there is no longer any justification for the privileging of the gene and the subsequent marginalisation of development. Tracking genes is no longer viewed as a short-hand way of tracking development and organisms. Instead, we must attend to development itself. In this spirit, DST first asks how developmental outcomes are achieved, and then asks how these systems have evolved. By beginning with development, the developmental systems approach jettisons any commitment to the gene as the unique unit of inheritance or hereditary transmission. This allows developmental systems theorists to re-evaluate how we view developmental interactions. Finally, we can incorporate this new approach to development into a co-constructionist approach to evolution by natural selection.

2.1 Extended Inheritance

In chapter three I discussed the fact that though evolutionary developmental biologists do not assume that the organism inherits naked DNA, this is not taken to challenge the idea that genes are the sole unit of hereditary transmission, since other cellular entities are construed as products of the maternal genes. However, I argued that this argument is circular; non-genetic components of the zygote are

considered genetic products because genes have a privileged role in development, and genes have this privileged role in virtue of the fact that they are the sole unit of hereditary transmission. In the absence of any reason to view the gene as ontologically distinct from other developmental resources, the zygote must be viewed as containing many different entities that have been inherited and that are essential for development. For example, membranes surrounding the nucleus and cell are essential for development; membranes can only develop if there are older membranes to act as templates. Also, cell differentiation depends upon the cytoplasmic polarities in the zygote, and DNA transcription requires the chromatin marking system (Griffiths & Gray 2001: 195). The study of epigenetic inheritance – the inheritance of cellular entities beyond DNA – is in its early stages, but the results so far suggest that there is far more to inheritance than DNA. The developmental systems approach does not stop at the cell; many extracellular and even environmental resources for development can be considered part of the expanded inheritance of an organism. Many mammals require certain gut bacteria to allow them to digest plant cellulose and they inherit these bacteria through eating their mother's faeces. Rabbits not only receive essential gut bacteria, but also have their food preferences influenced this way. In humans, the mother's diet during breastfeeding appears to influence the food preferences of the infant (Jablonka & Lamb 2005: 163-164). Beyond this, organisms can inherit features of their environment too: nests and/or nesting sites, dams, social groups, educational systems, and so on. All of these things can be passed from parent to offspring and they will play a role in the development of the offspring. Indeed, these things ensure heritable similarity. A child can inherit its mother's love of carrot juice because this is what she drank while breastfeeding the child.

Differences in the details of what is inherited can be acted on by selection. Some species of butterfly tend to lay their eggs on the same plant that their mother laid her eggs on (Jablonka & Lamb 2005: 190). If, in one generation, a butterfly lays her eggs on a different plant, her female offspring will also lay their eggs on this new plant. Thus it is possible to have variation within a population. Sometimes this variation may be adaptive. Perhaps this new plant offers better protection from predators – the eggs may be better camouflaged – or is a more nutritious foodstuff for the larvae.

Neo-Darwinists may argue that, for instance, the imprinting for food preferences or sites for future egg laying that occurs in offspring are underpinned by adaptations for such environmental responsiveness and can be understood as disjunctive genetic programmes responding to environmental inputs. DST does not suggest that, for instance, imprinting is only determined by the environmental stimulus in question. Indeed, this suggestion would run counter to the view of development inherent in DST. But what DST suggests is that these aspects of the environment, and the cellular entities mentioned above, may play a similar role in explaining heritable similarity as genetic material. Arguing that imprinting is ultimately a matter of a genetically underpinned adaptation begs the question. It presumes the priority of the genes. However, the special status of genes is dependent upon them being the sole unit of inheritance, the very issue at stake here.

Sterelny et al. (1996) have attempted to combine the notion of an extended inheritance with neo-Darwinian biology. Their extended replicator approach, while broadening the scope of what is inherited (though not quite as broadly as DST), retains the basic structure of neo-Darwinian evolutionary biology. The class of replicators is broader than neo-Darwinian accounts typically acknowledge, but very little else has changed. Development remains merely the result of evolution and is thus largely irrelevant to evolution itself, so that tracking the changes in replicator frequencies will remain a good way to track evolutionary change. If such an account can be made to work, this would mean that neo-Darwinism could take on board the idea of DST's expanded inheritance without it threatening the coherence of the overall theory.

However, there are limits to how well the extended replicator theory can capture the numerous forms of extended inheritance (Jablonka & Lamb 2005: 375-376; Griffiths & Gray 2001: 196-197). The theory of extended replicators suggests we understand inheritance in terms of different "channels" or "systems" of inheritance. With this approach we would have a genetic channel of inheritance, an epigenetic channel of inheritance, a cultural channel of inheritance, and so on. Creating distinct channels of inheritance in this way fails to recognise several key features of development. I will outline these features first, and then discuss the problem they pose for any attempt to interpret the notion of extended inheritance from within the neo-Darwinian framework.

2.2 Development

Oyama et al. (2001: 2) describe four themes that they take to be characteristic of development and that they suggest need to be reflected in any theory of evolution. These themes are: (1) the joint determination of every trait by multiple causes; (2) the context sensitivity and contingency of development; (3) development as a constructive process; (4) the distributed control in developmental systems. I will deal with each of these in turn.

It is, in one sense, trivially true that every trait is the outcome of multiple causes. In chapter one I outlined several varieties of genetic determinism, the strongest of which maintained that nothing but genes were required for development. Such a position is held by no one in this debate. But although no one denies the fact that developmental outcomes have multiple causes, this can be largely neglected in evolutionary considerations while genes hold a privileged position in development. So while everyone acknowledges that a trait has multiple causes, genetic causes are more salient in development than non-genetic causes if genes are understood as embodying instructions for development. Because DST begins by trying to establish the nature of developing systems, and only then attempting to understand how developing systems could evolve, these aspects of development, previously considered true but irrelevant, instead become central.

The context-sensitivity and contingency of development, DST's second theme, is also not controversial in a general sense; however, as I will discuss shortly, those approaches that focus on the gene as issuing instructions in development tend to recast this context sensitivity in terms of some inherent potential in genes. The developmental systems approach, rejecting this construal of the gene, suggests that in any given situation a developmental resource may play an important role in the development of a trait. However, the role of the developmental resource is determined by the context in which it appears. In another context the same resource may have a different effect or none at all. The overall state of the developmental system determines what developmental factors are important.

Context sensitivity may, the neo-Darwinist argues, be handled by disjunctive genetic programmes. These disjunctive programmes contain

conditional rules: if *X* is present, follow developmental path *A*; if *Y* is present, follow developmental path *B*.⁴² There are two responses to this. The first is that although neo-Darwinism may, to some extent, be able to handle context-sensitivity, it may not always be best placed to notice its occurrence. This is because it is often in attending to developmental details that we notice this context sensitivity. DST, by beginning with development, is primed to notice these cases. Neo-Darwinism, on the other hand, need only attend to developmental details when they obviously conflict with specific claims. But given that the neo-Darwinian approach does not provide any motivation to ask questions about development, such conflicting evidence may not be discovered. So, both neo-Darwinism and DST may agree that context sensitivity is a genuine feature of biological systems. But while DST has this built into its fabric, the neo-Darwinian approach either acknowledges context sensitivity on a more *ad hoc* basis, or simply overlooks it. Unlike the issue of the joint determination of developmental outcomes by multiple causes which neo-Darwinian theory sees as irrelevant, context sensitivity is not irrelevant to its project and thus the ease with which it may be overlooked is problematic.

There are further problems with the neo-Darwinian response to context sensitivity. The disjunctive programme suggests that the different developmental outcomes the programme is associated with are determined by the programme responding in different ways to different inputs from its environment. That is, it suggests that in different contexts, the gene or genes behave differently (i.e. issue different instructions). However, this does not reflect the facts of the situation. A gene can behave identically at two different times – that is, it can continue to lead to the production of the same type of protein – but the effects of this gene activity have very different developmental outcomes. Morange (2000) uses the *Notch* gene to make this point. This gene is involved in two distinct processes at two different times. During early development, it is involved in lateral inhibition – an intercellular process whereby cell differentiation is influenced by the path taken by its neighbouring cells. Later on in development, the same gene is involved in the control of neurite growth in nerve cells. The *Notch* gene is involved in the production of the very same protein in both cases,

⁴² See chapter one for a more detailed discussion of these disjunctive programmes.

and this protein is involved in the same basic cellular interactions in both cases. At this level, nothing has changed. And yet we have two different developmental outcomes at two different times. It is the state of the system in which this gene and its product are located that has changed. The *role* played by the *Notch* gene and its products in the two different processes – lateral inhibition and neurite growth – is not determined by anything intrinsic to the gene or its product, but rather is determined by the state of the system in which they are situated.⁴³ The disjunctive programme takes our attention away from the biological reality of the situation by allowing us to assume that the different processes the *Notch* gene is involved in can be explained in terms of the gene receiving certain inputs and producing certain outputs. However, the change is not at this level of biological organisation. The disjunctive programme approach is inadequate for handling this kind of context sensitivity. The neo-Darwinian neglect of development means that issues of context sensitivity, although permitted through disjunctive programmes, may not come to light. Worse, disjunctive genetic programmes appear to be an inadequate explanation for context sensitivity.

The third and fourth themes of DST – development as construction and distributed control – are closely related and I will discuss these ideas together. Central to the idea of development as construction and distributed control is the rejection of any form of preformationism. Preformationism was a theory developed in the seventeenth century and was committed to the idea that development involves a quantitative but not qualitative change (Pinto-Correia 1997). Nicolaas von Hartsoeker's image of a tiny man curled up inside a sperm most famously illustrates the point, though this is a strong version of preformationism; others believed that the essential form of the organism was already present and just required some "assembly." Whether this image was taken literally or not, it captures the basic commitment of preformationism;

⁴³ Recent work has gone so far as to question whether there really is any such thing as a gene, and a number of different concepts appear to have been conflated in this one term. At the very least, a good deal of developmental work appears to be required before anything like a protein can be produced, and this involves more than just DNA. See, for example, Moss (2004), Dupré (2005), Stotz et al. (2004) for more detailed discussion of these points. I will not address the complications that arise from this as they do not damage my argument here. On the contrary, it further undermines the claim that genes can be uniquely informational and that they precede development.

development involves growth but not any kind of qualitative change. The essential form of the organism was taken to be present in either the sperm if one was a spermist, or the egg if one was an ovist. On such a view, there was no need to explain how biological order came to be; it was always there. Preformationist theories do not tell us about development; rather, they deny the need for development.

Authors in the DST tradition have argued that gene-centric theories such as neo-Darwinism are a new form of preformationism (e.g. Oyama 2000a; Oyama et al. 2001). What comes preformed on this account is genetic information. Certainly, there may be scope for some indeterminacy in outcome; given different environmental contexts, development may proceed in different ways. However, the routes development must take are already broadly determined by the genetic instructions. Biological order arises from the genetic instructions:

Today we think of preformationism as an archaic relic of outmoded thought, and we snicker at the absurd idea that there are little people curled up in sperm or egg cells. But replacing curled-up people with curled-up blueprints or programs for people is not so different. That is, there is not much conceptual distance between *aggression* in the genes, on the one hand, and *coded instructions for aggression* in the genes, on the other. What is central to preformationist thought is not the literal presence of fully formed creatures in germ cells, but rather a way of thinking about development – development as *revelation* of preformed essence rather than as contingent series of constructive interactions, transformations, and emergences. It is a way of thinking that makes real development irrelevant because the basic “information,” or form, is there from the beginning, a legacy from our ancestors. (Oyama 2000b: 136)

DST rejects the idea that genes should be construed as containing programmes or instructions for development and does not assume one, or some set of, resources control or issue instructions for development: “the claim that development occurs because it is programmed to occur or because it has been selected by evolution is merely a promissory note redeemable against future developmental biology”

(Oyama et al. 2001: 4). Preformationist approaches suggest there is nothing of interest to investigate in development (at least from an evolutionary point of view) and so discourage lines of research.

In fact, the context sensitivity of developmental resources suggests that trying to locate the source of control of development in terms of developmental resources is a mistake. Each resource plays a role in the developmental outcome, and the role each resource plays is determined by the other resources involved in, and by the current state of, the system. Thus, the correct level of description to talk about control is at the level of the system itself rather than in anything like genetic instructions. The notion of the system controlling development is different in several respects to the idea that genes control development. Unlike the control that genes have over development in the neo-Darwinian picture, where the causal arrow moves in one direction – from genes to the rest of the developmental system – on the DST account the arrow points in multiple directions. The system exerts control over its component parts, but the system itself emerges from the interactions of its component parts. If control of development is most usefully located at the level of the system, and that system itself is to be understood as being composed of a “vast and heterogeneous assembly of interactants” that are “system-dependent and change over time” (Oyama et al. 2001: 5), the route development will is not “represented” in any resource prior to its realisation. What this means is that there is not a predetermined path (or set of paths) available for development to be pushed down. Rather, developmental systems theorists argue, development occurs because developmental resources interact with one another in ways typical given their context and do not require instructions. This leads to the conclusion that development is a process of construction or an “ad hoc” process (Moss 2004).

2.3 Evolution as Construction

DST also includes a commitment to evolution as a process of construction. DST pays due heed to the kinds of phenomena raised by Lewontin (1978, 2001) that I discussed in chapter two. The niche construction account developed by Lewontin, and more recently expanded upon by Odling-Smee et al. (2003), drew attention to the fact that the particular selection pressures faced by an organism are not just given, but are in fact the result of the interaction between the organism and its

environment. The organism constructs its niche; it makes certain aspects of the physical environment relevant to it, and in doing so generates the selection pressures it is subject to. Natural selection then involves looping effects as a result of the activity of organisms in a population turning some parts of the physical environment into their niche, while at the same time the niches provide selection pressures altering the population. But although there is much that is right in the niche construction account, DST does not incorporate it wholesale. The first reason for this is that Odling-Smee et al.'s approach relies on a dichotomous account of inheritance which shares similarities with Sterelny et al.'s (1996) extended replicator approach discussed above (Griffiths & Gray 2001). Both accounts are based on the idea of distinct streams of inheritance. In Odling-Smee et al.'s account there are two streams of inheritance: the genetic and the environmental. Bearing in mind the context sensitivity of development, any inherited entity will have its effects determined by the other entities being inherited, and by the larger collection of developmental resources available for any individual developmental system. This means that it might not always be best to think of inheritance in terms of individual entities or distinct channels of inheritance, as this will tend to obscure the interactions between them. As an example, Sterelny (2001; Sterelny et al. 1996) suggests we conceptualise inheritance such that we have a genetic system and an epigenetic system. However, the chromatin marking system, one part of the epigenetic inheritance channel or system, operates by modifying the pattern of gene expression (Sterelny & Griffiths 1999; Jablonka & Lamb 2005). The effects of each of these "channels" cannot be fully understood separately from one another. For example, the role the genes play and the role the chromatin marking system plays can only be determined by the overall system in which they are located. Separate channels of inheritance may be taken to imply that they each contribute to developmental outcomes additively and that each channel has some degree of independence from the others. This is something that cannot be assumed *a priori*. DST rejects any formulation which may make it easier to downplay this context sensitivity: "a central theme of the DST research tradition has been that distinctions between classes of developmental resource should be fluid and justified by particular research interests, rather than built into the basic framework of biological thought" (Griffiths & Gray 2001: 206). Because development is context sensitive,

whatever is inherited in one such channel will have its affects in part determined by what is inherited via other channels. Inheritance itself is developmental. It is the causal interactions between the elements of the system that underpin heritable similarity; intrinsic properties of inherited entities alone will not be enough to explain heritable similarity. Once we recognise this, the replicator/interactor distinction becomes less convincing. Recall that interactors were those things involved in causal interactions with the world, while replicators were merely copied.⁴⁴ If what is “replicated” includes developmental processes, this blurs the distinction between replicator and interactor. Further, what persists from one generation to the next must first be constructed. Rather than the “replicator” preceding development, it is the product of development. And because it is the entire system that is implicated in any developmental outcome, DST views the entire developmental system as the replicator in evolution, thus integrating interactor with replicator.

Like neo-Darwinian evolutionary theory more generally, Sterelny et al.’s (1996) extended replicator approach relies on a distinction between replicators and interactors. But once an extended inheritance is acknowledged, as well as the fact that any developmental resource has its affects determined in part by the context it finds itself in, the replicator/interactor distinction breaks down and the extended replicator approach seems less convincing. Note that this should not be confused with Hull’s argument, outlined in chapter one, that genes are both replicators and interactors. Genes act as replicators and interactors in different contexts on Hull’s account. During meiosis and mitosis, the gene acts as a replicator; however, during protein production (and the programming of development more generally) genes act as interactors. Although these processes are causally related (programming development is supposed to increase the chances of replication), replication and interaction are two distinct processes on the neo-Darwinian approach. Proponents of DST, on the other hand, suggest that the replicator is the product of development and that the “replicator” is the entire developmental system. Rather than the replicator enabling interaction (via programming development), it is interaction that constitutes the replicator. Thus the dichotomy between interactor/replicator is rejected.

⁴⁴ See chapter one for a more detailed discuss of the replicator/interactor distinction.

The second distinction that proponents of DST draw between themselves and niche constructionists such as Odling-Smee et al. (2003) concerns another dichotomy inherent in this approach; namely, evolution by natural selection proceeds as a result of two causal processes. In the first, the organism alters its environment, and in the second, the environment alters the organism (Griffiths & Gray 2001: 206). While Lewontin argues that while neither organism nor niche can be understood without the other on this niche constructionist account, there are nonetheless two distinct things being considered; a boundary is maintained between organism and environment. DST changes the focus from organisms to developmental systems, which may or may not coincide with one another: “the developmental system is not two things, but one, albeit one that can be divided up in many ways for different theoretical purposes” (Griffiths & Gray 2001: 206). There may be a particular research question that necessitates taking the skin of the organism as a boundary in order to make the question tractable, and this is perfectly compatible with DST. Further, in a specific context, some proposed boundary might in fact be important. This does not conflict with DST either. However, what proponents of DST reject is having firm boundaries picked out in advance of any investigation into a particular biological phenomenon. The boundary between organism and environment can be less important than is often thought, and being committed to this boundary from the start can allow us to overlook important instances where this boundary is irrelevant, such as in extended inheritance. Given this, talking in terms of developmental systems, which can extend into the environment, engenders no preconceptions about what does and does not constitute them. In not specifying a boundary to the developmental system in anything other than general terms, this means that questions about how resources interact in development cannot be assumed to be answered, but instead must be investigated. Nonetheless, Griffiths and Gray argue that the niche construction model is essentially correct in its emphasis on the fact that environments are not given, but are made relevant by developmental systems and, further, they argue that the models designed by Odling-Smee et al. ought to be used, but used “tactically” (2001: 206). That is, we may draw a boundary at the skin of the organism (or indeed somewhere else) if this allows us to address our research question, and in some contexts this boundary may indeed

be important, but we should not grant this boundary any weight in broader theoretical considerations.⁴⁵

3. Holism

The fact that DST allows, in principle, a great many things to be inherited and constitute a developmental system has led to worries that it is an unworkable approach for biologists to adopt. The concern is that it advocates a holism that cannot be made compatible with empirical investigation. However, such a concern is unfounded. DST is holistic, but not methodologically so. Proponents of DST do not suggest a change in the way experiments are performed, for instance; rather, they suggest further experiments. As already noted, it is entirely compatible with DST to bracket certain issues for the purposes of designing experiments. For instance, one can treat certain features of the developmental system as background conditions in order to establish a given causal relation. Knockout experiments are consonant with DST, so long as the fundamental principles of DST are kept in mind. How a given developmental outcome was achieved has not been fully understood just because a relevant gene has been identified. How that gene interacts with those “background conditions” is equally pertinent. Further experiments could be performed by making the genes part of the background conditions, and varying some other developmental resource. Methodologically then, fears of holism are misplaced.⁴⁶

A different worry about holism might stem from the fact that theories of natural selection seem to require discrete entities to form the basic units of the theory: “the sorts of things that can be counted, that... have clear boundaries and that... do not overlap so much that they cannot be distinguished from one another” (Griffiths & Gray 2001: 209). DST blurs the line between organism and environment. The developmental system will not be identical to our standard ideas of an organism; it will frequently extend beyond the skin of the organism to incorporate developmental resources traditionally considered “environmental.”

⁴⁵ This point suggests that Keller’s (2001) worry that DST may lead to an elision of the body may be unwarranted here. DST does not preclude the possibility that the skin of the organism might be an important boundary in some contexts. Rather, the issue here is not to decide in advance what those contexts will be.

⁴⁶ A good many of the contributors in Oyama et al. (2001) discuss DST-inspired scientific research demonstrating that this framework does not preclude practical research.

Further, the developmental system includes many more things than have typically been considered relevant for evolution. Many resources will play a role in the reconstruction of a life cycle and some of these developmental resources and interactions will be entirely unique to the individual developmental system, such as new languages, novel food stuff, medicines and so on. Such resources and interactions may be so entirely unique to an individual life cycle that they would seem to have no bearing on evolutionary questions. Given all this, how can we make principled distinctions in order to pick out fundamental units for DST?

Developmental systems theorists respond to this worry by making a distinction between a developmental system and an *evolutionary* developmental system. When thinking in terms of evolution we can abstract away from the unique details of individual developmental systems to construct evolutionary developmental systems (Griffiths & Gray 2001). These systems are constituted by all those resources, and the interactions between them, that reliably recur in each generation of a lineage and that ensure that the normal life cycle is reliably reconstructed.

Given that many populations are polymorphic – that is, have more than one form typical of the population (for example, sexual dimorphism and seasonal colour morphs in some butterfly species) – the evolutionary developmental system must account for the multiple phenotypes that constitute the lineage. This is an issue common to all evolutionary theories; variation within the population is an essential component for natural selection and thus descriptions of the lineages should reflect this. DST does this by incorporating everything required to reconstruct a normal life cycle, or the normal life cycles, into the evolutionary developmental system for that lineage. The evolutionary developmental system is, then, a description of the lineage at a particular time.

DST refers to the organism-like entities in its theory as “life cycles.” The life cycle is the product of the developmental system. If developmental systems extend beyond the skin, how do we individuate life cycles? There are many biological processes that appear cyclical and involve the “repeated assemblies” of developmental resources; which of these repeated assemblies should be designated as a life cycle and which are merely components of a life cycle? A cell has its own “life cycle,” as do the leaves of trees (Griffiths and Gray 2001:

209). Symbiotic relationships may be particularly problematic; on a DST account, are there two life cycles in a symbiotic relationship, or just one? Griffiths and Gray suggest that if individual cycles can “give rise to new cycles of itself that are not coupled to the other member of the symbiosis in the characteristic way” then such cycles constitute distinct life cycles. Facultative symbiotic relationships are of this sort. For example, if a bird population commonly eats a fruit from a tree and so disperses its seeds, both bird and tree might continue relatively unharmed without one another if, for instance, there were other species that dispersed the tree’s seeds, and other sources of food for the bird population. On the other hand, some cycles may no longer be able to replicate themselves in isolation from their symbiotic partners, in which case the cycles would be sub-cycles of a larger life cycle. These are obligate symbiotic relationships. For instance, the eukaryotic cell is thought to have evolved from a symbiotic relationship between what became the cell nucleus and the cell organelles that became strongly obligate. While at some point each lineage may have been able to extract itself from the symbiosis, it is now contingently irreversible. “Contingent irreversibility” is a term coined by Maynard-Smith and Szathmáry (1995: 9) and utilised by Griffiths and Gray (2001) to capture the sense in which, while technically possible, there are certain evolutionary paths that are so unlikely and difficult, we may disregard them. Thus, two lineages that become so intricately intertwined in a symbiotic relationship that they are contingently irreversible can be regarded as having become a single lineage. In each generation of this lineage, the symbionts constitute a single life cycle or developmental system. Life cycles are then those things that can reconstruct themselves in each generation. This allows us to pick out those things that are part of the evolutionary developmental system and those things that are not. Any resource which contributes to the developmental system that produces a life cycle, and which persists because of its contribution to that lineage of life cycles, is a part of the evolutionary developmental system.

Not all developmental resources will have the same life span as the life cycle. A developmental resource may have a different periodicity to the life cycle, as we might expect if they are to be inherited by offspring. Some developmental resources may appear to outlive any individual life cycle, for example, a beaver dam may persist long after the beavers that originally constructed it have died.

Other developmental resources, such as sun light and gravity, will not depend on the life cycle for their persistence. But a niche that is “transmitted” vertically (parent-to-offspring) will still have to be maintained by each generation. The offspring will depend on the niche, and the niche will depend on the offspring. Indeed, it is this high degree of mutual dependence that makes the organism-niche pairing better considered as a single developmental system. In at least some cases, it may be misleading to talk about niches being transmitted. Rather, each generation inherits certain resources which the next generation must use reconstruct, or at least actively maintain, the niche. A dam, for instance, may only persist because the next generation continue to maintain it. Offspring life cycles may inherit just enough to reconstruct the complex relations between the skin-bound organism and its ecological niche. This reconstruction is the construction or the development of the life cycle. So although some aspect of what we traditionally call the environment may seem to outlast the organism, the persistence of the resource is not really persistence at all. Rather the resource is reconstructed, or continually maintained, by each generation.

Sunlight and gravity, unlike dams, do not have to be maintained. But recall Lewontin’s (1978, 2001) point in chapter two that a niche is not just given but is constructed; the organism or developmental system makes certain aspects of the environment relevant. So although, for instance, sunlight is not itself the result of what Godfrey-Smith (1996) calls “causal construction,” its role in the developmental system comes about as a result of “constitutive construction.” And in each generation the relationship between any persistent resource such as sunlight or gravity must be constructed anew.

Natural selection, on the adaptationist approach, seems to require something to which organisms can be more or less suited. The worry then is that given that developmental systems extend out into the environment, developmental systems subsume anything that might have been considered to generate selection pressures. That is, anything that causally interacts with the organism, and does the same for the offspring of that organism, and indeed for many generations in that lineage, seems to be a good candidate for being considered part of the developmental system of each generation:

Because the focus is on how the complete life cycle is achieved, everything needed for that life cycle is assumed to be present. So anything that impinges on the process is an element of the system itself. It is this that creates the impression that all change in the system itself must be endogenously driven and creates the apparent puzzle about the source of selective pressures. (Griffiths & Gray 2001: 207)

DST does not wish to abandon the mechanism of natural selection, so does this imply that selection pressures are generated from within developmental systems? Griffiths and Gray (2001) suggest that this worry about endogenous selection pressures can be overcome. This first move is to shift focus away from the evolutionary developmental system, which is a description of the lineage, and onto populations of individual developmental systems. There will be some variation between individuals within a population. Resources normally inherited by offspring may be missing or scarce because parents did not generate the resources, the wider population did not generate the resources, or because resources independent of the activities of the population (persistent resources) are rare. Also, variation may occur because developmental systems incorporate new resources which result in new developmental outcomes. All these situations will lead to different developmental outcomes and thus variation within the population. This explains variation; the next step is differential reproduction. While neo-Darwinism explains differential variation as the result of selection pressures from an immovable environment, DST incorporates the idea of differential reproduction based on better or worse interactions between developmental resources within any individual developmental system:

One variant does better than another, not because of a correspondence between it and some preexisting environmental feature, but because the life cycle that includes interaction with that feature has a greater capacity to replicate itself than the life cycle that lacks that interaction. (Griffiths & Gray 1994: 300)

At the most extreme end, developmental systems that lack essential resources die out as they will be incapable of successfully reconstructing their own life cycle.

Moving up from this point, variation in developmental resources will lead to developmental systems replicating with more or less success. The least successful developmental systems die out because they do not replicate or because they replicate at rates so small compared to other developmental systems that they are soon swamped by the more successful systems. New developmental resources may also appear which allow the developmental system to cope better when other resources are rarer and thus to replicate at higher rates. Developmental systems may compete for the same resources. Because of facts about the organisation of the developmental systems involved, some may be more efficient at securing the resource, or require less of it to survive, and thus will be more or less capable of replicating themselves. These resources do not provide external selection pressures; they are part of the developmental systems involved.

A final worry that Sterelny (2001) raises concerns expanded inheritance. Although a good deal of what is inherited by offspring comes from their parents, the wide sense of inheritance advocated by developmental systems theorists also includes developmental resources that can be inherited from other, non-parental sources. Since beaver dams are the work of an entire lodge, they appear to fall into this category. Parents do not pass beaver dams to their offspring. At best, one generation of beavers transmits dams to the next, but this picture is complicated by the fact that generations commonly overlap. This kind of horizontal transmission is problematic and Sterelny has suggested that this undermines any attempt to define evolutionary developmental systems or the life cycles they produce: “if this is transmission at all, it is diffuse and development is holistic” (2001: 344).⁴⁷ The objection here is that diffuse inheritance rules out the possibility of cumulative selection based on differences between individual life cycles; there will be no heritable variation that can be explained by inheriting the shared resource. Griffiths and Gray have suggested two possible responses to this. The first is that, although diffuse transmission will not provide a heritable basis for cumulative selection processes, “changes in ecological and cultural

⁴⁷ Sterelny et al.’s (1996) extended replicator approach adopts a more limited form of extended inheritance than that employed by developmental systems theorists. The extended replicator approach restricts what can be inherited to only those things that are inherited vertically.

inheritance... could play important evolutionary roles both in opening up new sets of adaptive possibilities and by facilitating the dynamics of evolutionary change” (2001: 202). Another point can be added here. Although the shared resource may be available to all developmental systems in the population, not all developmental systems will interact with this resource in the same way. The role of a developmental resource is determined by the overall state of the developmental system, and there will be variation in the developmental systems within a population. A developmental system may make better or worse use of a resource from the point of view of selection. While the resource itself is not the difference maker which explains the differential fitness within the population, the entire system, including the shared resource, explains why one developmental system survives and reproduces with more success than another.

Griffiths and Gray’s (2001) second response to this worry invokes the idea of trait group selection as developed by Sober and Wilson (1998). A trait group is one in which all members possess a similar trait. Such groups may be long lasting, fleeting, or anything in between. Trait group selection occurs when there is a strong correlation between possessing a given trait, and belonging to a group of individuals with the same trait. Beavers will tend to form groups with other dam-building beavers. Organisms that engage in reciprocal and kin altruism tend to group together. The population is structured in such a way that interactions between organisms are biased; some organisms are likely to spend a good deal of time together, others are unlikely to do so. This gives groups a degree of coherence. Organisms in one trait group can have a selective advantage over conspecifics in other trait groups in virtue of the fact that they are in the same trait group. Sober and Wilson (1998) are particularly concerned with the evolution of altruism in their work. They suggest that in a population containing both altruists and free loaders, altruism need not be undermined by selection for free loaders as is often assumed, if the population is structured so that it contains smaller, interacting groups.⁴⁸ Altruists will do badly in groups that have either a

⁴⁸ The worry here is that altruists will tend to be worse off than free loaders as altruists will share their resources with others, while free loaders will take resources from altruists, while keeping all their resources to themselves. In this way, the argument goes, a population of altruists will be eventually displaced by free loaders. This has led to what is known as “the problem of altruism.” Despite the fact that altruism does not look like it should exist, nonetheless acts of altruism are witnessed in the natural world. Examples include the alarm call of the vervet monkey

majority or equal number of free loaders, and the free loaders will be favoured by selection. However, although the free loaders will do well relative to the altruists where they comprise at least half of the group, in majority altruist groups the situation is different. In these groups, both altruists and free loaders do better than the free loaders in the free loader groups. But because altruists make up the majority of the most successful groups, they survive and reproduce in such a way so as to dominate within the larger population. So although some altruists do worse than some free loaders in the population, ultimately altruists are favoured by selection in this sort of scenario. It is, to use Hull's (1980) terminology, the group that is the interactor in these sorts of situations. Selection at the individual level may often favour the freeloaders, but groups composed primarily of altruists will do much better than groups composed primarily of free loaders. Thus, if there is some correlation between possessing a particular trait and belonging to a certain group (so that altruists tend to group with other altruists), there could be cumulative selection for altruistic traits. Dam building and other forms of diffuse inheritance are forms of altruistic behaviour on the part of one generation to the advantage of the next, or even to the advantage of other members of the same generation. Every time a beaver contributes towards building the dam, all the members of the lodge are better off. Similarly, organisms that engage in kin and reciprocal altruism ensure organisms other than themselves benefit from their behaviour. This can allow the evolution of traits in organisms or developmental systems as a result of selection acting at the level of the group. Thus diffuse inheritance need not undermine evolution by natural selection.

What does all this mean for the developmental system and the life cycle? Sterelny (2001) worried that the extended inheritance promoted by DST meant that we would not be able to distinguish the fundamental units (lineages, individuals, and so on) for an evolutionary theory. This problem does not arise

which alerts conspecifics to danger, but puts the individual who raises the alarm at increased risk, or the worker bees who forego their own chances of reproduction in order to further the chances of the queen reproducing successfully. Gene selectionists provide one response to this apparent dilemma (e.g. Dawkins 1989: 88-93). Genes, understood as acting to ensure their own replication, will programme development such that the organism will sacrifice itself if this will increase the chance that copies of that gene in relatives will be replicated. See chapter five for further discussion.

from all instances of extended inheritance, just to those that do not meet the vertical transmission criteria. However, this is only a problem if we are committed to a model of natural selection which allows it to act at only one level. Gene selectionist accounts are often like this; genes are the units of selection, and group selection, if acknowledged at all, plays a very minor role in evolution and cannot be the basis for complex adaptations (Lloyd 2005). But DST adopts a hierarchical model of natural selection whereby selection can act on many different levels: at the level of trait groups, at the level of life cycles, at the level of the gene, and at many other levels in between (Griffiths & Gray 2001: 210).⁴⁹ The possibility of trait group selection means that there can be natural selection where inheritance is more diffuse. And even if what is inherited cannot form the basis for cumulative selection, it can nonetheless play a role in an evolutionary trajectory. In chapter six I will discuss further ways in which, even when there is a high degree of horizontal transmission, this need not preclude the possibility of natural selection.

4. Evo-Devo and DST

In chapter three I mentioned Gilbert's (2003: 349) claim that DST faces its biggest threat from evo-devo. He argues that while DST has offered some useful insights into biological phenomena, and while many of these insights have found their way into evo-devo research, evo-devo looks set to take over the niche that developmental systems theory has carved out for itself. Prior to evo-devo, DST played an important role in calling attention to the otherwise neglected role of development in evolutionary biology. However, Gilbert argues that given evo-devo has put development back in the picture, and done so in the lab rather than from "the armchairs of philosophy" (Robert et al. 2001: 958), evo-devo stands to make DST largely irrelevant. Essentially, Gilbert appears to be arguing that evo-devo does what DST does, but better; it has brought developmental considerations into work in the lab and this has led to productive research programmes. However, Gilbert's argument largely hinges on the extent to which evo-devo and DST resemble one another in their theoretical commitments, and it is not clear that they do. If DST represents a distinct view of evolution and

⁴⁹ See chapter one for a discussion of the hierarchical model of selection.

development to that on offer from evo-devo, then evo-devo does not threaten to render DST irrelevant. Here I will argue that evo-devo neither takes over the conceptual space occupied by DST, nor offers a better way to understand evolution and biology.

There are a variety of grounds on which evo-devo and DST differ. Evo-devo, as discussed in more detail in chapter three, is committed to the gene as the unit of hereditary transmission, whereas DST adopts a very broad understanding of what secures heritable similarity. As I argued in chapter three, the claim that the gene is the sole unit of hereditary transmission appears circular once it is acknowledged that a good deal more than genes are required for any developmental outcome, an uncontroversial claim in evo-devo. DST rejects any privileged role for genes, whereas, at least on mainstream versions of evo-devo, the gene remains importantly distinct from other resources. In Gilbert's account, this is because the gene provides specificity to developmental interactions:

DST has, generally, made the error of not assigning instructive or permissive influences in the interactions. To most of the developmental systems theorists, all the participants are on the same informational level. In this way, the genome is just one other participant, just as are cells or the environment... the specificity of the reaction (that it is a jaw that forms and not an arm; that it is a salamander jaw that forms and not a frog jaw) has to come from somewhere, and that is often a property of the genome... Instructive partners provide specificity to the reaction, whereas permissive partners are necessary, but do not provide specificity... The gene has to be given its proper respect. (Gilbert 2003: 349-350).

Gilbert acknowledges, however, that it is not always the gene that provides specificity – sometimes it may be an environmental factor. He cites the sex determination of turtles based on nest temperature as an example of a non-genetic resource providing specificity (2003: 350). The resources that provide specificity here seem to be the difference makers in development against a background that is taken to be invariant. DST does not deny the importance of these resources, it merely cautions against assuming that there is something

inherent in these resources which means that, should the context change, the resources will play the same role. Nest temperature, genes, and any other developmental resource may be construed as difference makers only in the context of those “permissive” resources. Elsewhere they may have a different effect or none at all. DST does not entail viewing all developmental resources as contributing in exactly the same way to every developmental interaction. Different resources play different roles in specific developmental processes. What the DST proponent argues is that we should not prioritise a particular developmental resource in a general sense. As Gilbert notes, genes may be “instructive” in some scenarios (sex determination in mammals), and “permissive” in others (sex determination in turtles and some other reptiles). What the developmental systems perspective seeks to prevent is a neglect of the importance of the “permissive” factors in enabling the “instructive” resource to have the kind of effect it does. The contingent nature of the role played by any individual developmental resource – the degree to which it depends on its context within the developmental system – should always be remembered. Further, the acknowledgement in Gilbert’s own account of the degree to which context matters – sometimes it is the gene that specifies, sometimes it is the environment – suggests that the developmental systems approach is merely foregrounding an uncontentious aspect of development. DST, by denying the privileging of the gene, does not amount to the claim that all developmental resources are the same. Many resources will have an entirely unique role to play. But a unique role is not the same as a privileged role. The developmental systems perspective attempts to ensure that this conflation of unique and ontologically distinct does not occur.

Evo-devo has done a good deal of important work in identifying semi-autonomous aspects of the developmental system that can act as levels of selection. This work will benefit DST by drawing attention to some of the levels at which selection might work, allowing the development of a more fine-grained hierarchical model of selection. Equally, DST can bring something to the study of modularity. Griffiths and Gray note that, by incorporating developmental systems considerations, we might identify modules in places that might otherwise have been neglected because they breached the skin boundary: “there is no reason to think that extended forms of inheritance such as symbiont transmission or cultural traditions will be any less modular in their

developmental consequences than genetic factors” and, given this, “the range of phenomena that can and should be given selectionist explanations is considerably increased” (2001: 203). Further, the attention to context sensitivity intrinsic to the developmental systems approach will be important here. That certain developmental processes or structures can appear semi-autonomous is in fact the result of context dependence. When we identify modular structures or processes, what we have identified is a process or structure that is not significantly impinged upon by other structures or processes in that particular type of developmental system. We cannot assume that the same developmental resources in a different developmental system will result in the same quasi-independent processes occurring or the same structure developing. In chapter three I discussed *Pax6*, a type of gene found in a wide variety of organisms, which appears to play a role in (among other things) eye development. This gene has a similar effect in a wide variety of circumstances. For example, *Pax6* from *Drosophila* was transplanted into a frog and induced an ectopic eye (Burian 2005: 252). Modularity can exist at any level of biological organisation and the section of DNA, or the genome more broadly, that we might identify as *Pax6* looks like a good example of a semi-autonomous aspect of the developmental system. However, it is important to note that *Pax6* does not always have the effect of initiating eye development. *Pax6* will induce eye development in a number of different places in a mouse. However, in the pancreas, although *Pax6* is active (it produces, or contributes to the production of, a protein), it does not lead to the development of an ectopic eye:

What switches on the cascade that makes the eye or that produces pancreatic proteins is a group of interacting signal-transduction modules hooked into the right context. The multiple modules required to initiate the different processes are composed not only of genes but also their own gene products and a series of additional proteins that must interact correctly, with each other and with the nucleotide sequences that respond to the signals. (Burian 2005: 252)

The developmental systems perspective focuses our attention on the fact that although some developmental resources may seem to have the same effect in

a wide variety of situations, this does not entail that the resource is entirely responsible for those effects. Often, the developmental resource is situated in a broader developmental system that makes possible the particular contribution of the developmental resource in question. Modularity may well be an important feature of developmental systems which has permitted evolution, but we must be cautious not to create a new set of boundaries between individual modules and the rest of the developmental system. That is, we need to bear in mind that any processes or structures that appear semi-autonomous do so in part because of the context in which they appear, rather than exclusively as the result of any inherent properties they might possess.

Although Gilbert has pitched evo-devo against DST, such a move seems unnecessary. Rather, argue Griffiths and Gray, both have something to offer the other:

... we believe that [evo-devo] and DST are essentially complementary. DST does not provide a theory of phenotypic integration and modular evolution, but rather stands in need of one, and [evo-devo] is beginning to supply such a theory (Griffiths and Gray 2004). Conversely, nothing in the fundamental inspiration of [evo-devo] precludes it embracing a wider conception of the developmental system, not as emerging from interactions between genes, but as emerging from interactions between the whole matrix of resources that are required for development. (2005: 423)

Indeed, Gilbert has written about the need for the work of ecological developmental biologists (developmental biology which attends to the important role of the environment in development) to be heeded by evolutionary developmental biologists. At times, Gilbert's message seems particularly close to DST:

Experimental isolation of the embryo from "outside" influences during analyses of developmental mechanisms has proven a useful and powerful approach; the problem is that it excludes a priori the environment itself as a contributor to or influence on development... Our

“self” turns out to be a permeable rather than a discrete entity... the processes that generated our adult form themselves arose from interactions between us and our environment... Environmentally dependent development also calls into question our epistemology as developmental biologists... Studies of ecological developmental biology demonstrate that this internalist approach, although powerful, is incomplete. (Gilbert & Bolker 2003: 6-7)

Gilbert shares a similar aim to DST; both wish to integrate ecological considerations into developmental considerations. Indeed, a special issue of *Evolution and Development* on ecological developmental biology edited by Gilbert and Bolker (2003) included a contribution from H. Fred Nijhout who works in the DST tradition (see, for example, Nijhout 2001). Evo-devo and DST do not compete for the same space, but rather they are adding to each other in productive ways. This should also serve as a further example against claims that DST is either unworkably holistic, or has nothing to say to biology as it is practiced.

5. Conclusion

Critics of DST speak of an unworkable holism implied by the theory or of its unsuitability for scientists in the lab or field. Collaboration with evolutionary developmental biologists offers just one indication of the wrong headedness of this line of thought.⁵⁰ The developmental systems perspective offers the best way to approach the integration of development and evolution in that it does not rely on an unsupported preformationism about genes. The developmental systems perspective takes some relatively uncontroversial positions – developmental outcomes require more than genes, context matters in development – and

⁵⁰ Godfrey-Smith (2001b) has argued that as well as providing a framework for empirical research, DST has a second role as a “philosophy of nature” which “comes *after* empirical science and tries to redescribe structures in the world that have already been described by the sciences” (Godfrey-Smith 2001b: 284). He is quick to point out that these two roles are not always distinct, and that each one may influence the other. A philosophy of nature is not as constricted as a scientific research programme; although closely informed by science, it can for instance embrace holism and reject some of the language and interpretations used by scientists. This sort of activity can help with negotiating issues concerning science and society more broadly: “the scientific ideas should be fed into such discussions in a philosophically processed form, not in the raw language of science” (2001b: 284).

combines these with claims about extended inheritance. The result is a model of development and evolution that differs significantly from that presented by neo-Darwinian evolutionary theory. Though both agree on the importance of natural selection, DST eschews the overwhelming focus on a one-way line of influence from environment to organism, and incorporates the idea of evolution by natural selection as a matter of co-construction. Further, while neo-Darwinism has cast development as the process of following genetic instructions, DST rejects the notion that there is something given prior to development itself. This leads to the conclusion that there can be no “replicator” prior to development, and thus the interactor/replicator dichotomy breaks down. There is no information for development preformed in the genes, instead development is a matter of construction: “nature and nurture are not alternative causes but product and process” (Oyama 2000a: 148), so that “it is always a nurtured nature” (Oyama 2000c: S341).

In the second part of this thesis I will turn to the specific issues of the evolution of culture and cognition. I will critically discuss research in this area, particularly that done under the heading of Evolutionary Psychology, inspired by neo-Darwinian evolutionary theory. In bringing developmental considerations to the fore in evolutionary theory, the developmental systems approach leads to quite a different picture of the evolution of culture and cognition than research that is based in neo-Darwinian evolutionary biology.

Part Two

Culture and Cognition

Chapter Five

Evolutionary Psychology

1. Introduction

If we can offer evolutionary explanations for physiological structures such as the heart or the eye, and if we can offer evolutionary explanations for animal behaviour, such as the bowerbird's efforts to attract mates with its ornate nest and the blackbird's efforts to defend its territory, then perhaps we can explain the human brain, human psychology and human behaviour in the same terms. Of course, evolutionary explanations do not explain *everything* about a given physiological structure. Individual differences will exist between one heart and another. Similarly, there may be some variation in exactly how one organism defends its territory when compared to its conspecifics. So we should not expect an evolutionary explanation of human psychology to explain everything about how individuals think and behave. Nonetheless, many researchers have attempted to apply the lessons of neo-Darwinian evolutionary biology more generally to the particular case of human psychology, behaviour and culture. Some have even gone so far as to claim that they are mapping out human nature – those qualities that are universal among humans (e.g. Tooby & Cosmides 1992; Pinker 2002). There are a broad range of approaches taken in the attempt to understand the human mind in terms of evolution, and they operate under names such as sociobiology, human ethnology, evolutionary anthropology and evolutionary psychology. Even under the heading of evolutionary psychology there is a divergence of opinions and approaches. One of the most well-known approaches has developed out of the work of John Tooby and Leda Cosmides (1989, 1990, 1992), Donald Symons (1992), and David Buss (1992, 1999), among others, and it is this version of evolutionary psychology I will focus on here. Buller (2005) refers to this brand of evolutionary psychology using the upper case – Evolutionary Psychology – and I will follow this convention here.

I will begin by detailing the rationale offered by Evolutionary Psychologists for their project. Evolutionary Psychology follows from an acceptance of both the adaptationist programme and the gene as unit of inheritance, two core commitments of neo-Darwinism outlined in chapter one.

Because Evolutionary Psychology relies on neo-Darwinism, any problems inherent in neo-Darwinism will also threaten Evolutionary Psychology. Given I have already argued that neo-Darwinian biology is deeply flawed, it follows so too is Evolutionary Psychology. Nonetheless, I discuss Evolutionary Psychology here for three reasons. The first is that Evolutionary Psychology faces problems particular to its own discipline, so that even if the case I have made against neo-Darwinian biology can be answered, further work needs to be done to defend Evolutionary Psychology. The second reason for discussing Evolutionary Psychology here is that, in discussing its problems, it will highlight issues any account of the evolution of culture and cognition may need to contend with so that at least some of the arguments here may have broader applicability. I am particularly concerned in this thesis with what a theory of evolution more sensitive to the details of development might tell us about the evolution of human psychology and culture. The final reason to discuss Evolutionary Psychology here is that it will help to make clear what differences such a developmentally-informed approach to the evolution of cognition and culture will make.

Once the claims of Evolutionary Psychology have been outlined, I will discuss two classes of problems particular to Evolutionary Psychology.⁵¹ The first sort stem from the fact that Evolutionary Psychology only appears to utilise a subset of the tools provided by neo-Darwinism. Many of the key claims of Evolutionary Psychology, such as the idea that no significant evolutionary change could have occurred in the last 10,000 years, are difficult to support from a neo-Darwinian point of view. The second class of problem that Evolutionary Psychology faces arises from the combination of the adaptationist strategy and their particular explanandum: our cognitive architecture. The aim of Evolutionary Psychology is to uncover our hidden mental structures and this involves determining the nature of the selection pressures that caused these structures to evolve. However, there are good theoretical reasons to think that it will be difficult, if not impossible, to identify many of these selection pressures. This, if correct, would undermine the ability of Evolutionary Psychologists to

⁵¹ This is by no means an exhaustive treatment of difficulties faced by Evolutionary Psychology. More detailed critiques can be found in, for example, Dupré (2001) and Buller (2005).

generate hypotheses about the hidden structures of our minds using this methodology.

Finally, I will argue that Evolutionary Psychologists' arguments concerning culture are incoherent. Evolutionary Psychologists reject approaches to culture and behaviour that pay little attention to evolution and biology, such as those found in parts of the social sciences (Evolutionary Psychologists call this the Standard Social Sciences Model, or the SSSM). Evolutionary Psychologists also claim that modern humans are unsuited, or maladapted, to modern environments. I will argue that there exists a tension between these two claims such that Evolutionary Psychologists cannot maintain both and, as a result, face a choice. They can continue with their commitment to their version of neo-Darwinism and accept they have little to say about modern human behaviour or culture, and thus must accept the SSSM. Alternatively, they can abandon their commitment to their version of neo-Darwinism, and lose much that is supposed to distinguish Evolutionary Psychology from other approaches to the evolution of cognition and culture.

I will not challenge the model of the mind Evolutionary Psychologists have adopted here, but I will deal with this in chapter seven. What I want to do here is, first, spell out what Evolutionary Psychology is committed to and, second, demonstrate that, even if neo-Darwinism is correct in its model of development and evolution, Evolutionary Psychology has not made its case.

2. Evolutionary Psychology

Fundamental to Evolutionary Psychology is the idea that human psychology is the result of a collection of adaptations to a past environment. Recall from chapter one that adaptations need not be adaptive. That is, while the process of natural selection is ongoing, an adaptation must be adaptive. However, the environment can change, and a trait that once helped the organism negotiate its environment can be rendered useless, or even a hindrance. Evolutionary Psychologists argue that the evolution of complex adaptations, such as the ones that they believe underpin human psychology and behaviour, take a very long time (from a human perspective at least) to evolve. Our current environments are quite different from those of our ancestors. While we need to negotiate traffic, processed food, large cities, and so on, our ancestors had to master the skills

required to forage for food, negotiate the savannah, avoid predators, live in small tribes, and so on. Given the length of time required for the evolution of complex adaptations, and the novelty of our current environments, we should not expect evolution to have rendered us well adapted to our current environments. Evolutionary Psychologists term the aspects of our ancestral environment to which we are thought to be adapted to as the “environment of evolutionary adaptedness” (EEA). The EEA is not a particular time or place but “the statistical composite of selection pressures that caused the design of an adaptation” (Cosmides & Tooby 1997). For humans, Evolutionary Psychologists believe the EEA coincides with the Pleistocene (Tooby & Cosmides 1992: 143). That is, we have evolved to suit particular features of our Pleistocene ancestors’ environment.

The key claim of Evolutionary Psychology, and what marks it out from earlier attempts to explain human behaviour in terms of evolutionary theory (such as sociobiology, for instance), is its massive modularity hypothesis. Evolutionary Psychology rejects the possibility of the mind as a general purpose problem-solver (Tooby & Cosmides 1992). Evolutionary Psychologists argue that instead of being a general-purpose problem-solving device, the mind is organised into modules. Each module has been designed to solve a particular problem our ancestors would have routinely faced. Given the large number of problems faced by our ancestors, we should expect to find a large number of modules that allowed them to solve these problems; hence, our minds are massively modular. That is, human psychology will not be characterised by a single adaptation that solved all, or even many, of the problems our ancestors faced in the EEA. Instead, there will be multiple adaptations or modules. These specialised modules are not thought to map directly onto the brain. Modules are functionally, rather than anatomically, individuated.⁵² The relevant bits may be in

⁵² Other brands of modularity, such as those proposed by Fodor and Chomsky make different claims than those of Evolutionary Psychology and should not be confused with them. Fodor (1983), for instance, claims modules are “informationally encapsulated” which is not often a recognised property of the modules Evolutionary Psychology promote. However, Buller and Hardcastle argue that, despite the claims of Evolutionary Psychologists to the contrary, Evolutionary Psychology’s modules must in fact be informationally encapsulated to make sense (2000: 309-310). Pinker has suggested that it would be more useful to think of mental modules more loosely as “mental organs” as Chomsky has done (1997: 31). Organs are less rigorously delineated – they cannot always be “encircled by a dotted line” – and they interact with many other parts of the body in complicated ways. This way of thinking allows Evolutionary Psychologists to reject information encapsulation as a property of their “modules.” The type of

various different locations in the brain – less like “the flank steak and the rump roast on the supermarket cow display,” and more like “roadkill, sprawling messily over the bulges and crevasses of the brain” (Pinker 1997: 30). Evolutionary Psychologists have adopted the massive modularity hypothesis less because of positive reasons for such a position, and more because of what they perceive to be the implausibility of a general-purpose problem-solver. In examining their arguments against general-purpose problem-solvers, we can see just what Evolutionary Psychologists have in mind when they discuss modules.

2.1 The Massive Modularity Hypothesis

Tooby and Cosmides (1992) offer a number of justifications for their massive modularity hypothesis. First they make a poverty of stimulus argument, that is, they argue that humans have cognitive capacities that exceed what they could have learnt from experience. This, they suggest, indicates that human minds are more than blank slates. The paradigmatic example cited here is language acquisition (e.g. Tooby & Cosmides 1992: 94-95; Pinker & Bloom 1992). The rules of grammar are highly complex and structured; nonetheless children routinely come to master them. A child who grows up with no exposure to a linguistic community may never fully master language in later life; however, it appears as though children do not require as much exposure to language as might be expected given its complexity (Tooby & Cosmides 1992: 45; Pinker & Bloom 1992; Pinker 1994). Children are not explicitly tutored in how to use grammar, they are not routinely corrected when they make grammatical errors, and they themselves are often exposed to grammatical errors. Indeed, children are thought capable of inventing complex, grammatically structured languages when only exposed to “pidgin” languages (Pinker 1994: 33-39). That is, there seems to be a paucity of the right sort of stimuli to account for a child’s ability to learn grammar. It appears that a child’s mastery of the structure of language surpasses what we should expect, given their limited exposure. Evolutionary Psychologists (among others) take this to suggest that when we are exposed to language, some

module suggested by Chomsky et al. are “domain-specific systems of truth-evaluable mental representations that are innate and/or subject to informational restrictions,” rather than the information-processing devices and mental representations of Evolutionary Psychology (Samuels 2000: 18).

pre-existing capacity allows us to extrapolate from the small amount of input we receive from our linguistic community such that we can come to speak a language fluently. This pre-existing capacity has been termed the “language acquisition device” and this constitutes one of the many modules thought to comprise the human mind. Evolutionary Psychologists contend that there are many other domains where the ability of humans to perform complex tasks can only be explained by similarly specific modules.

Secondly, Tooby and Cosmides also argue that an entirely “blank slate” mind, or in their terms, domain-general content-independent problem solver would be fatal to an organism. A completely domain-general mind seems liable to become paralysed by combinatorial explosion. At any time, may be a large number of potential courses of action one might take, and each one of these will lead to further possible courses of action, and so on. In order to establish the most adaptive behaviour it appears necessary to evaluate a very large number of possibilities. These numbers can quickly becomes so large as to be unmanageable, rendering the individual incapable of evaluating the options: “any design for an organism that cannot generate appropriate decisions, inferences, or perceptions because it is lost in an ocean of erroneous possibilities will not propagate, and will be removed from the population in the next generation” (Tooby & Cosmides 1992: 103). In order to avoid this, some constraints are required such that not all courses of action need to be considered in any decision making process. This is, among other things, known as the “frame problem”; without a way to frame the world such that certain courses of action are seen as available, while others are effectively rendered invisible, the individual cannot make any decisions at all. The “unframed” perspective of a blank slate mind, or a general-purpose problem solver would not just be less efficient than a collection of domain-specific systems, it would be useless:

... it is usually more than efficiency that is lost by being limited to a general-purpose method – generality may often sacrifice the very possibility of successfully solving the problem, as, for example, when the solution requires supplemental information that cannot be sensorily derived. (Tooby & Cosmides 1992: 179)

The selection pressure in these cases would be maximally strong and we should see selection for domain-specific systems.

This much, however, only establishes that an entirely “blank” mind is unfeasible. We might expect that some constraints are necessary to rule out combinatorial explosion, but allow a good deal of flexibility within those boundaries; not quite a completely content-independent, general-purpose problem solver, but not a highly domain-specific content-rich collection of modules either. However, Tooby and Cosmides present further arguments which suggest that natural selection will favour increasingly domain-specific content-rich modules.

Tooby and Cosmides argue that there cannot be a general problem solver because there is no such thing as a general problem (Tooby & Cosmides 1992: 142). Our ancestors would have faced a wide range of quite distinct problems – avoiding predators, attracting mates, finding food, and so on. What counts as a solution to one problem will not count as a solution to another. Pinker notes that a general-purpose problem solver is a jack of all trades and a master of none (Pinker 1997: 28). To successfully negotiate the environment, each problem faced by an organism requires a specific solution that cannot usually be generalised beyond that particular problem domain: “different adaptive problems are often incommensurate” (Tooby & Cosmides 1992: 111). So, in order to achieve greater success at solving any one adaptive problem, it is necessary to have very specific solutions. And given the large number of problems our ancestors would have faced in their environment, they would have required a large number of distinct solutions. There are two ways specific solutions could be found: the organism, starting with some broad constraints to rule out combinatorial explosion, could learn through trial and error different solutions to all the different problems it faced, or the organism could possess the sort of highly specialised modules Tooby and Cosmides propose. They offer three reasons why selection would favour the latter.

First, the greater the generality of the problem-solver, the more inefficient it will be. Learning takes time, and inventing a solution takes even longer. This problem is only compounded by the fact that there are thought to have been many different problems, each one taking time to solve. In the meantime,

mistakes can be made. This suggests that whenever a more specific problem-solving module was available in the population, selection would favour it:

Such mechanisms will be far more efficient than general-purpose mechanisms, which must expend time, energy, and risk learning these relationships through “trial and possibly fatal error.” (Tooby & Cosmides 1992: 111)

Second, an individual is not always in a position to determine which of the options available to it is the most adaptive. Two courses of action may be as adaptive as one another most of the time, with one only very occasionally delivering a slightly higher degree of adaptedness. Such a slight difference in adaptedness might easily go unobserved by an individual. Further, if the benefit of one course of action over another is not immediate, it may not be possible to determine in advance which is most adaptive. For example, juvenile diet might affect adult fertility but this might be a difficult pattern to discern, especially if fertility improves only marginally. The thought is then that an individual, without any guidance in these matters, would then fail to systematically opt for the most advantageous option:

Many adaptive courses of action can be neither deduced nor learned by general criteria alone because they depend on statistical relationships that are unobservable to the relevant individual. For a content-independent system to learn a relationship, all parts of the relationship must be perceptually detectable ... Natural selection, through incorporating content-specific decision rules, allows the organism to behave as if it could see and be guided by relationships that are perceptually undetectable and, hence, inherently unlearnable by any general purpose system (Tooby & Cosmides 1992: 111-112).

If a heritable cognitive variant should arise that disposes those individuals who possess this trait to opt for the fitness-enhancing option, there will be selection for this innovation and it will be pushed towards fixation in the population.

Third, even when the most adaptive course of action is observable in principle (i.e. the consequences of two courses of actions are immediately apparent to the observer), what grounds does an individual have for deciding one outcome is adaptive and the other is not? Avoiding a predator is quite a different thing to catching prey, and different again to discerning nutritious from poisonous fruit. Success in each of these domains looks quite different. How is an organism to tell which course of action to take when what counts as success is unclear? Things are complicated further by kin altruism. Kin altruism is the name given to the phenomenon whereby an organism harms its own chances of survival and reproduction in order to help kin survive and reproduce. This is explained by gene selectionists and Evolutionary Psychologists in terms of the fact that kin share a certain percentage of genes with each other, and thus it may be a worthwhile strategy for one token of a gene to sacrifice its “host” organism in order that its copies get to replicate. Hamilton’s rule is thought to establish the relationship between the cost an organism will endure for a genetic relative and the degree of relatedness between that organism and its family member (Dawkins 1989: 88-93). Siblings share half of their genetic material, and thus should endure higher survival and reproductive costs for one another than cousins would for one another, given they only share one eighth of their genetic material. So, according to this line of thought, what counts as an adaptive course of action, from gene’s eye view, is for an individual organism to risk harm and the ability to reproduce in certain situations, depending on the degree of risk and the degree of relatedness to the kin member such behaviour may help. This would suggest that even a general rule such as “adaptive success means avoiding harm” would not ensure the individual always acted adaptively. A general-purpose problem solver cannot latch onto some “general cue or criterion for success or failure that can apply across domains” (Tooby & Cosmides 1992: 111). Again, an individual with at least some “substantial built-in content-specific structure to discriminate adaptive success from failure” would be at a distinct advantage.

Because selection will always favour modules that are more domain-specific and content-rich, Tooby and Cosmides argue, selection will have pushed the organisation of human psychology towards massive modularity:

... human psychological architecture must be far more frame-rich and permeated with content-specific structure than most researchers (including ourselves) had ever expected. (1992: 112)

2.2 The “Psychic Unity of Humankind”

As well as expecting the mind to be massively modular, Evolutionary Psychologists also expect all normally developed humans to share the same sort of modular structure, and Tooby and Cosmides refer to this shared cognitive architecture as the “psychic unity of humankind” (1992: 79). There is nothing in the previous arguments that makes this so. That is, it is possible that massive modularity is a universal characteristic of normally developed humans, but that individuals may differ in terms of the sorts of modules they possess. For example, all normally developed humans have eyes, but there exists different eye colours in the human population. Similarly, humans differ in terms of their blood group. There could be completely different, perhaps competing, modules. In fact, despite claims that Evolutionary Psychology is concerned with a universal human nature, this option – that there exist different modules in different people – is a key claim of Evolutionary Psychology in one particular case. Men and women are supposed to have quite different modules in a variety of domains, though especially those concerned with mating (e.g. Silverman & Eals 1992; Buss 1992; Wilson & Daly 1992).⁵³ This exception aside, there are no more instances where humans are thought to differ in terms of the adaptations they possess: “humans must share a complex, species-typical and species-specific architecture of adaptations” (Tooby & Cosmides 1992: 38).

The justification for the claim that these adaptations constitute a universal feature of humankind stems from two arguments. First, all psychological adaptations must be underpinned by gene complexes, as this is the only way they could have been heritable, visible to selection, and thus adaptations: “the organization of our mental modules comes from our genetic program” (Pinker

⁵³ See Dupré (2001) and Buller (2005) for critiques of much of the work done by Evolutionary Psychologists in this area.

1997: 23).⁵⁴ The more complex the adaptation, the more genes will be required to underpin it:

Complex adaptations are intricate machines that require complex “blueprints” at the genetic level. This means they require coordinated gene expression, involving hundreds or thousands of genes to regulate their development. (Tooby & Cosmides 1992: 78)

That is, if the adaptation was not universal within the population, it would be lost during sexual recombination. So if such modules exist at all, they must be universal.

The second argument is based on an analogy between our physiology and our psychology. *Gray's Anatomy* is said to describe “in precise anatomical detail individual humans from around the world” and this “demonstrates the pronounced monomorphism present in complex human adaptations” (Tooby & Cosmides 1992: 38). Normally developed humans all possess two arms, legs, eyes, ears, kidneys, a heart, a brain, a stomach, and so on. If such uniformity is present in our physiology, by analogy we should expect the same of our psychology.⁵⁵

Because we share the same sort of cognitive adaptations, Evolutionary Psychologists maintain that this should lead to similarities in how we behave, and this to similarities between the cultures that such behaviour generates:

Our immensely elaborate species-typical physiology and psychological architectures not only constitute regularities in themselves but they impose within and across cultures all kinds of regularities on human life,

⁵⁴ See chapter one for a more detailed discussion of these issues.

⁵⁵ Buller (2005: 424-428) offers a detailed critique of this claim. In particular, Evolutionary Psychologists are conflating two distinct grains of analysis here. Kidneys, hearts, and limbs are at the same level of biological organisation as the brain, and at one level of description are invariant in normally developed humans (indeed, at one level of description, they are invariant across all the primates). But modules (or the physical bit of the brain that are thought to correspond with the cognitive module) are at a level of description where a great deal more individual variation is found in anatomy. Thus the comparison is unjustified.

as do the common features of the environments we inhabit. (Tooby & Cosmides 1992: 89)

The regularities imposed on cultures by our shared psychology constitute what Evolutionary Psychologists call *metaculture* and are thought to characterise a great deal of human life:

Such statistical and structural regularities concerning humans and human social life are an immensely and indefinitely large class... adults have children; humans have species typical body form; humans have characteristic emotions; humans move through a life history cued by observable body changes; humans come in two sexes; they eat food and are motivated to seek it when they lack it; humans are born and eventually die; they are related through sexual reproduction and through chains of descent; they turn their eyes towards objects and events that tend to be informative about adaptively consequential issues; they often compete, contend, or fight over limited social or subsistence resources; they express fear and avoidance of dangers; they create and participate in coalitions; they desire, plan, deceive, love, gaze, envy, get ill, have sex, play, can be injured, are satiated: and on and on. (Tooby & Cosmides 1992: 89)

But how do we account for the apparent cultural variation we observe? Or, for that matter, how do we account for the differences between individuals in any given culture? There are several ways Evolutionary Psychologists attempt to accommodate this variation. The first is in terms of “minor, superficial, nonfunctional traits” (Tooby & Cosmides 1992: 38). Evolutionary Psychologists are interested only in those cognitive traits that seem complex and they, like many others, assume complex traits must be the products of natural selection. In this sense, Evolutionary Psychology is adopting the explanatory adaptationism (and is thus vulnerable to the same criticisms).⁵⁶ Our universal human nature

⁵⁶ Explanatory adaptationism maintains that complex adaptations are the most pressing or most important issue for evolutionary biology. The difficulty with this position, as discussed in

derives from our shared adaptations. Differences that arise from structures that are not adaptations are not of interest to Evolutionary Psychology, and further, not important enough to undermine the claim that we share a similarly structured psychology.

A second way in which individual and cultural variation may be explained is in terms of *epidemiological culture*. Epidemiological culture arises when one or a few individuals develop a behaviour or idea that spreads throughout the population. This type of culture most closely resembles traditional notions of what constitutes culture. Ideas and behaviours spread like a virus (hence, “epidemiological”) to other members of the community. The ability to learn these behaviours and pick up these ideas is importantly underpinned by the modular mind and these modules bias which aspects of behaviour and which ideas are picked up: “domain-specific mechanisms influence which representations spread through a population easily and which do not” (Tooby & Cosmides 1992: 121).

The third way Evolutionary Psychologists explain individual and cultural variation is in terms of *evoked culture*. Evoked culture is thought to be the product of a mechanism or module responding to local conditions:

... manifest expressions may differ between individuals when different environmental inputs are operated on by the same procedures to produce different manifest outputs. (Tooby & Cosmides 1992: 45)

Modules are not thought to be strongly deterministic; they do not result in one type of psychology or behaviour, but rather in the right behaviour for the right context, at least in the environments to which these modules are adapted. Tooby and Cosmides use an analogy with a jukebox to highlight their point (1992: 115). Identical jukeboxes are placed at different locations across the globe. Each one has a clock and a navigational device that measures their longitude and latitude. The jukebox is programmed to play particular songs for given times and places. So, in each location, the jukebox plays a different song. Nonetheless, because the

more detail in chapter one, is that the claim about the importance of complex traits seems only to be a statement about our own interests.

jukeboxes are the same, they all have the same songs, and they all have the same programmed rules. They may play different songs, but this is because they are receiving different inputs from their environments. Similarly, the idea of evoked culture is meant to capture those aspects of cultures that may visibly differ from one another, but are underpinned by the same rules or evolved mental modules.

Tooby and Cosmides seem to imply at least two ways organisms can respond in this flexible manner to environmental input:

Thus, when we use terms such as “evolved design,” “evolved architecture,” or even “species-typical,” “species-standard,” “universal,” and “panhuman,” we are not making claims about every human phenotype all or even some of the time; instead we are referring to the existence of evolutionary organized developmental adaptations, whether they are activated or latent... For this reason, adaptations and adaptive architecture can be discussed and described at (at least) two levels: (1) the level of reliably achieved and expressed organization (as, for example, in the realized structure of the eye), and (2) at the level of the developmental programs that construct such organization. (Tooby & Cosmides 1992: 82)

The first is at the level of the adaptation itself. In this case, the same trait responds in different ways to different environmental conditions, and any change made is reversible. The expansion and contraction of the pupil in response to light conditions is an example of this sort of environmental responsiveness. Similarly, one might seek out one sort of food in one environment, and another sort of food elsewhere, or dress warmly in cold environments, and wear lighter clothing in hotter places, and so on. So one way human cognition might be responsive to environmental conditions is if a given module generates different responses given different inputs.

The second way in which such flexibility can be achieved is at the level of development. Tooby and Cosmides seem to acknowledge some form of disjunctive genetic programme:

... different coordinated designs, psychological or physiological, cannot be the direct product of suites of genetic differences. Different genetic programs (corresponding to subsets of genes) are activated in one morph or another, but are present in all individuals. In short, the conclusion from evolutionary genetics is that different species have different designs because of different genes, but within a species, different designs emerge from the same genes... (Tooby & Cosmides 1990: 45)

Genetic instructions or programmes for development do not appear to necessarily be followed, and can be triggered or activated by the right sort of inputs.⁵⁷ At some point in development, an environmental input sends development down one of a number of potential trajectories and can lead to distinct developmental outcomes. For example, many species are thought to have their sex determined by environmental factors. The temperature at which they are incubated determines the sex of turtles for instance (West-Eberhard 2003: 121). Developmental plasticity of this sort explains the existence of polyphenic traits (that is, the existence of distinct phenotypic forms) in a population for the neo-Darwinist. Unlike the previous example, this form of plasticity is usually not reversible. In this case, rather than a given trait reacting in different ways to different environmental stimuli, it is the genetic programme that generates different developmental outcomes for different environmental inputs.

The difficulty with developmental plasticity from an Evolutionary Psychologist's perspective is that it allows distinct phenotypic forms with the same genetic basis to arise in a population. That is, it allows for different modules to develop. Once the possibility of this form of developmental plasticity is permitted, we seem to be suggesting that what tends to be shared by humans are not cognitive adaptations, but rather disjunctive genetic programmes. However this is a point not acknowledged by Evolutionary Psychologists. For instance, Barkow et al. write that "there is a universal human nature, but ... this universality exists at the level of evolved psychological mechanisms, not of expressed cultural behaviors" (1992: 5). Rather than claiming the universality is

⁵⁷ See chapter one for a more detailed discussion of disjunctive genetic programmes.

due to shared developmental programmes, it is here explained by shared psychological mechanisms.

Some explanation for this may stem from Tooby and Cosmides' justification for the elision of developmental adaptations in other areas of their work. While they acknowledge that adaptations can be discussed at two levels (organised structure and development), Tooby and Cosmides state that they do not "usually bother to terminologically distinguish between successfully assembled expressed adaptive architecture from more fundamental developmental adaptations that construct them" (1992: 82). While it is perfectly legitimate to discuss biological organisation at whatever level is relevant to one's research interests, this is distinct from conflating different levels of organisation. In this instance, Tooby and Cosmides treat the developmental adaptations and cognitive adaptations as, for their purposes, equivalent. However, this conflation is problematic; an important class of developmental adaptations are excluded on this formulation. They have neglected the class of developmental adaptations that are not expressed in a given individual as a cognitive adaptation. However, their justification for speaking of a universal human nature stems only from the inclusion of this class of developmental adaptations.

Disjunctive genetic programmes are unproblematic in neo-Darwinian biology, and play an important role in explaining polyphenic traits in a population. Evolutionary Psychologists do not seem to deny the possibility of such disjunctive programmes, and rely on them to argue for the general principle that humans share a universal human nature as a result of natural selection. However, the possibility that such disjunctive programmes could allow the development of different modules in different individuals (sex differences aside) is not seriously countenanced, and no justification for this neglect is offered. Evolutionary Psychology, then, does not use the full range of tools made available to it by neo-Darwinism.

2.3 Stone Age Minds

Cosmides and Tooby (1997) have suggested "our modern skulls house a stone age mind." The time that has elapsed since the Pleistocene is considered too brief, by evolutionary standards, to allow for "the design of complex circuits" in brains to have evolved. As a result, modern humans are considered better suited to the

Pleistocene than the modern world. I will argue in part three of this chapter that this claim is incoherent, but for now I want to focus on their evolutionary justification for the claim that we possess stone age minds. Neo-Darwinian evolution theory leads us to expect that there should be a time lag between adaptations and environments. If we focus solely on genetic inheritance, as neo-Darwinian evolutionary biologists and Evolutionary Psychologists do, it is clear why this is assumed to be the case. Most genetic mutations are harmful, so the chances that an advantageous mutation will arise are small. Numerous generations will pass before the right sort of variation even exists in the population and selection can act. And once the right sort of variation is present, it can take a very long time for selection to result in the propagation of the trait throughout the population. If the adaptation is a complex one, the result of cumulative evolution, this process is extended even further in time. In this way, the speed with which environments can change far exceeds the speeds at which complex adaptations can arise. Given Evolutionary Psychology's commitment to the idea that our cognitive architecture is a complex adaptation, and that genes are the sole unit of inheritance, we should assume that our cognitive architecture took a long time to come about in the first place, and that it will take a long time for any significant alterations to this architecture to evolve.

But there are better reasons than Evolutionary Psychologists suppose for imagining at least some of our psychological mechanisms could have been altered since the Pleistocene. A relatively recent evolutionary novelty involves the ability to digest lactose. This is thought to have arisen in response to the pastoral, dairy-based lifestyle of some human groups between 8,000 and 6,000 years ago (Odling-Smee et al. 2003: 342-343). The suggestion is that dairy farming of some description came about first, and this created a selection pressure for the ability to digest lactose. This would mean that lactose tolerance has evolved far more recently than Evolutionary Psychologists allow for cognitive adaptations.

One response to this might be to argue that modifications such as lactose tolerance are very small and simple relative to the complex developmental or cognitive adaptations Evolutionary Psychologists are concerned with. However, there is an important difference between the evolution of an adaptation in its entirety, and modifications of an adaptation. A structure as complex as the eye is

highly unlikely to evolve in anything like 10,000 years for humans. However, much smaller modifications of the eye are possible in this time. Further, not all genes are considered equal in neo-Darwinism (and mainstream evo-devo); while some are directly involved in the production of proteins, others play a regulatory role. These genetic switches or “master control” genes are thought to regulate other genes. They embody the conditional rules of the disjunctive programme; depending on the input, they activate different sets of genes. Modifications of such genes could have powerful developmental effects in the sense that it could result in an alteration to the conditions under which different phenotypic forms are expressed. Tooby and Cosmides argue that such switches are less effective than organisms that display flexibility at the phenotypic level (1990: 46). That is, the reversible plasticity made possible by a single module responding in a flexible way to different inputs will be favoured by selection over a developmental plasticity that tends to require that one developmental path be taken over another early on in development and cannot be reversed. In response to this, however, Buller notes that we see in nature both forms of plasticity, and that Evolutionary Psychologists offer no reason to suppose that developmental flexibility will always be selected against (2005: 116-119). While a module (or any other trait) that can display reversible plasticity allows the organism a greater degree of responsiveness to its environment, it may be a more energy demanding trait in terms of its development and maintenance. This is just to say that there are no *a priori* reasons to suppose selection will favour one form of plasticity over another. Further, even if flexibility in adult traits would be favoured by selection over developmental flexibility, this alone is not enough to allow us to conclude that we will only possess the former kind of flexibility. If the right variation in the population never arises, the adaptively optimum solution will not evolve.⁵⁸ Given this, there does not seem to be any justification for Tooby and Cosmides ruling out the possibility that at least some of our developmental programmes/modules rely on genetic switches that could have been altered by the changing selection pressures in post-Pleistocene environments. Again,

⁵⁸ Further problems with the assumption that the most adaptive solution to a problem will evolve are discussed in Gould & Lewontin (1979).

Evolutionary Psychology appears to only attend to some aspects of neo-Darwinian evolutionary theory and not others.

2.4 Uncovering Modules

So far I have discussed the theoretical justification for what our minds should look like according to Evolutionary Psychologists. I now want to look at some issues concerning the methodology of Evolutionary Psychology. Human psychology should be structured, on this account, into a large number of modules, each dealing with a specific problem domain and equipped with some specific rules or knowledge for dealing with that domain. Given this general theoretical background then, Evolutionary Psychologists attempt to uncover these modules. This, note, is importantly different to the approach taken by evolutionary biologists. Often, the trait in question is observable by researchers; it is often the observation of the anatomical or physiological trait that will initiate and guide research. For instance, if we observe a peacock's extraordinary tail, mimicry in butterflies, or dam building in beavers, we might want to know why these organisms came to possess these traits. Even where evolutionary explanations are sought for behaviour, it is the behaviour itself that is the explanandum, rather than a proxy for some underlying cognitive adaptation. Evolutionary Psychology does not have this advantage. The structure of the mind is hidden, and it is the unearthing of this structure that Evolutionary Psychology promises. Evolutionary Psychology can "supply the missing middle: the psychological mechanisms that come between theories of selection pressures on the one hand and fully realized sociocultural behavior on the other" (Barkow et al. 1992: 6). Evolutionary Psychologists do not begin with either a module or a developmental programme. Instead, there are two possible approaches available to the Evolutionary Psychologist. The first is to look for cross-cultural regularities in behaviour and postulate underlying adaptations to explain the regularities. Identifying such regularities may be complicated by at least three issues. First, modules (or developmental programmes) do not lead to tightly regimented behaviours, so this variability will muddy the waters. Second, metaculture, evoked culture and epidemiological culture are not thought to be independent cultural streams, but rather interacting aspects of culture resulting in further potential for individual and cultural differences, thus making regularities resulting from similar modules

or developmental programs harder to detect. Finally, our adaptations were thought to be designed for the Pleistocene rather than current environments. Our radically altered modern environments may skew developmental programmes, or provide such odd inputs for psychological mechanisms that the manifest behaviours we witness may not be the behaviours our developmental programmes were selected for.

Despite this, Evolutionary Psychologists believe at least some regularities can be identified. Pinker (2002: 435-439) lists what he terms “surface universals”; that is, universal behaviours rather than universal cognitive structures. The job of the Evolutionary Psychologist is then to postulate cognitive adaptations that generate these regularities or surface universals, and the selection pressures that would explain their evolution. But there are a number of problems with this approach. For example, marriage is typically cited as a cultural universal, but as Dupré points out, this is a rather broad category:

Anthropologists describe systems of marriage that are monogamous, polygamous, occasionally polyandrous, hypergamous or hypogamous (women marrying up or down in status, though equal status is said to be the commonest case), between people of the same sex, and in some cases not involving sexual relations at all. (2001: 59)

Even granting the existence of such regularities does not, by itself, establish an adaptation for such behaviour. There are many ways to explain common cross-cultural behaviours, and Evolutionary Psychology does not have a principled means to decide which of these potential explanations is appropriate in any given scenario. A cross-cultural regularity may equally be explained by the presence of that behaviour in a common ancestral population that has been culturally inherited by modern societies. Evolutionary Psychologists cannot rule out that the common practice has been culturally, rather than genetically, inherited as a form of epidemiological culture. Perhaps only some of the regularities could be explained in this way, but the question is, which ones? How could we tell the difference between those regularities that are the result of cognitive adaptations, and those that are the result of a form of cultural inheritance? Evolutionary Psychology does not provide a solution to this.

A regularity need not only be the result of common heritage. It is possible that some regularities are the result of epidemiological culture spreading very widely. For example, the sweet potato was only found in America, until approximately 1,000 years ago when Polynesian explorers brought it to New Zealand and some Pacific Islands. Later, European explorers and colonisers came across the sweet potato in the West Indies and introduced it to Africa, India and Indonesia. It was then introduced into Papua New Guinea by Indonesian traders. By the time Europeans made contact with the communities living in the remote highlands regions of Papua New Guinea in the early twentieth century, the sweet potato had become an essential component of their agriculture and diet, as well as the agriculture and diet of numerous groups of people on many different continents (Bourke 2009). This example demonstrates the ease with which a good idea – the growing and consuming of a particular food stuff in this instance – can spread. Indeed, Levy (2004) argues that a marginal advantage of a given cultural practice over alternatives is enough to ensure the propagation of that cultural practice. Even groups of people typically conceived of as untouched by the outside world, such as the inhabitants of the highlands of Papua New Guinea, are susceptible. And such issues are only exemplified by issues surrounding globalisation: “surely the large majority do share, to a considerable extent, values shaped by exposure to the same transnational media” (Dupré 2001: 59). While growing sweet potato is unlikely to be mistaken for the result of a cognitive adaptation, other behaviours are less clearly examples of what Evolutionary Psychologists would call epidemiological culture. Putting concerns about the applicability of the term “marriage” to such a wide array of behaviours aside for a moment, we could imagine something like marriage as a particularly successful element of epidemiological culture too. But if this were the case, it would look just like a cultural regularity that sprang from a universal adaptation. Again, Evolutionary Psychology does not give us a principled way to tell apart a regularity caused by a universal cognitive adaptation, and a regularity caused by a particularly successful element of epidemiological culture.

The Evolutionary Psychologist might respond to this by arguing that their job is not done just when they identify a cross-cultural regularity. They must also posit selection pressures that would generate the kind of adaptation that could explain such a regularity. Understanding the selective pressures faced by our

ancestors, and the kinds of adaptations that might evolve in response to these pressures, might be thought to allow us to pick out those cross-cultural regularities that can be explained by these adaptations. Once we have a better idea of what kind of cognitive adaptations would have been useful, it may be easier to see if current populations actually demonstrate evidence of such adaptations. However, thinking about the selection pressures our ancestors faced generates a range of problems of its own. We cannot directly access the selection pressures our ancestors were subject to, and can only infer what they must have been from rather indirect evidence. This is because, while many of our physiological adaptations were thought to be for the physical environment, Evolutionary Psychologists believe that the selection pressures that shaped our psychology stemmed primarily from our social environment, and social environments are difficult to reconstruct. Contemporary hunter-gatherer societies are often used by Evolutionary Psychologists in lieu of access to ancestral societies. The specifics of their physical environments might differ from one population to the next, but it is social environments that we are interested in, and these should demonstrate enough regularity in structure to give us a clue about life in the Pleistocene. However, things are not so simple. Buller notes two problems:

First, it is naïve to think that the social lives of extant hunter-gatherer populations have not changed significantly in the last 10,000 years... Second, as the anthropologist Laura Betzig points out, there is considerable variation in the lifestyles of extant hunter-gatherer populations. (Buller 2005: 95)

Hunter-gatherer populations are not all alike and have developed quite different social systems. For instance, the average daily caloric intake as collected by women across hunter-gatherer populations ranges from 2% to 67% and average parental care ranges from ten minutes a day to 88% of the day (Buller 2005: 95) Further, contemporary hunter-gatherer populations do not live lives entirely independent of the modern human world. To look at the culture of a hunter-gatherer population is not to peer into our evolutionary past. Noss and Hewlett (2001), in their study of societies in the Aka forest in the Central African

Republic, document the greater participation of women than men in hunting than men and the equal participation of both parents in childcare. They have described the Aka as “probably one of the most gender egalitarian cultures in the ethnographic record” (2001: 1028). This stands in contrast, for instance, to the Baka women of Eastern Cameroon who do not hunt, but do carry spears for their husbands, brothers or fathers during a hunt. The Aka also exemplify the role played by exposure to other cultures. While the women do the majority of the hunting now, historically this activity had equal gender participation. The reason for this change in participation, suggested by Noss and Hewlett, concerns the recent employment opportunities afforded to the Aka by large businesses not indigenous to the area. The employers, for their own cultural reasons (they are Euro-Americans and Western-trained Africans), appear willing to employ only men, and so these men are not available to participate in hunting. Further, this exposure of hunter-gather societies to other cultures is not a feature of the modern phenomenon of globalisation (though globalisation appears to have intensified a pre-existing trend). The spread of the sweet potato demonstrates the degree to which it has long been possible for cultures to interact and to influence one another. Buller notes that anthropologists have documented interactions between hunter-gatherer populations and the rest of the world at least since the dawn of colonialism (2005: 95). These interactions can result in cultural changes such as the greater participation of women in hunting in the Aka, but can also result in a homogenisation of cultures, as in the case of waged labour in the men of the Aka. This means that hunter-gatherer populations will not be particularly useful in shedding light of the social milieu, and thus selection pressures, faced by our ancestors.

A more theoretical problem facing the identification of the relevant selection pressures concerns how exactly selection pressures arise. Lewontin’s (1978, 2001b) niche constructionism highlights the role an organism plays in making salient some aspects of the environment over others.⁵⁹ The life of a nocturnal animal may not impose on the life of a diurnal animal living in geographical proximity. A tree may be important to a bird which uses it to nest in, but irrelevant to a bird which nests closer to the ground. Large stones on the

⁵⁹ See chapter two for a more detailed discussion of Lewontin’s niche constructionism.

surface of the ground may be important to a bird that uses them to crack the shells of snails, but irrelevant to the earthworm underneath. The life activities of an organism determine which aspects of the environment are potential sources of selection pressures. This suggests that we would need to know something of our ancestors' cognitive structure prior to the development of our current adaptations in order to know what parts of their environments were relevant. It is difficult enough to ascertain what cognitive adaptations we might currently possess, determining the adaptations of our very distant ancestors is even more problematic. This is further complicated by the fact that there is no fossil record to consult.⁶⁰

For our ancestors' motivational states and cognitive processes would have been selectively responsive to certain features of the physical and social environments, and only those features would have affected subsequent adaptive evolution of early human psychology. At this point we again collide with our ignorance of our early ancestors. And, given that psychologies don't fossilize, this ignorance is likely intractable. (Buller 2005: 98-99)

Finally, even if we identify a selection pressure that did influence the course of our evolution, this does not give us enough information about what the adaptive solution would actually look like:

But selection never designs solutions to adaptive problems from scratch. Adaptations all emerge through modifications to preexisting structures. The form of a solution to an adaptive problem, then, will always depend heavily on the form of the preexisting structure that got modified... Consequently, we can never infer the structure of an evolved solution to an adaptive problem from the nature of the problem itself. (Buller 2005: 103-104)

⁶⁰ Even if a fossil record was possible, it would not definitively solve our problems here. It is not enough to see a structure to know how it was used. We might imagine a fossilised feathered bird faced selection pressures for flight, but would be wrong if the bird was in fact flightless.

For instance, if the problem a population faces is how to travel through its environment, there are multiple ways this might be solved: swimming, flying, walking and so on. Some knowledge of the previous state of the organisms in that population might help narrow down the possibilities, but such information is difficult to obtain in the case of our ancestors' cognitive structure.

As well as difficulties associated with the practicalities of evidence gathering, even if we could identify the selection pressures that existed in the Pleistocene, this alone would not guarantee the appropriate solution in fact evolved:

Some problems have a do-or-die character such that failing to solve them results in failure to survive or reproduce. But not every adaptive problem absolutely *must* be solved... often a population can survive and reproduce without variation that is differentially responsive to some selection pressure. Thus, even if we could identify all the adaptive problems facing our ancestral human populations, we still couldn't be assured that our ancestors evolved solutions to those problems. (Buller 2005: 103)

So, while solving the frame problem seems essential, for instance, other problems proposed by Evolutionary Psychologists may not have this "do-or-die" quality and we cannot assume *a priori* that such solutions in fact exist as part of our evolved psychology. One reason for this is because the evolution of an adaptation is dependent on the right sort of variation being present in the population; if the variation is not there, the adaptation will not evolve. Another reason concerns the fact that numerous selection pressures may have been in operation. The evolution of a new psychological mechanism is developmentally costly and, for instance, if this cost exceeds any benefit the mechanism might confer, then such a mechanism will not evolve. Even if we could identify one selection pressure, this does not inform us about other competing pressures. Again, this is a problem for evolutionary biology generally, not just Evolutionary Psychology. But unlike evolutionary biologists, Evolutionary Psychologists cannot directly examine the trait in question to see whether it appears to solve the proposed adaptive problem.

The Evolutionary Psychologist might respond by arguing that although there exists difficulties with both identifying the selection pressures faced by our ancestors and determining cross-cultural regularities that arise from cognitive adaptations, taken together the sorts of problems raised here can be avoided, or at least ameliorated. Perhaps we cannot be sure which cognitive adaptations would actually arise given various selection pressures, but we could posit an adaptation and then check for a corresponding cross-cultural regularity. If we find such a thing, then we can be reasonably sure we have identified a cognitive adaptation. However, that our ancestral environment posed a certain problem to our ancestors and that we see continued evidence of the solution to that problem does not rule out the alternative explanations. For example, I mentioned the possibility that a cross-cultural regularity could be explained by common heritage or just the spread of an aspect of epidemiological culture. Such aspects of epidemiological culture, if they are to be spread and maintained, are most likely to do so if they are found to be useful for the individuals involved. That is, epidemiological culture will often find success where it allows individuals to solve problems posed by their environments. The very same grounds that Evolutionary Psychologists cite to support their claim that a cross-cultural regularity must be explained by a shared cognitive adaptation also support the claim that the regularity is a particularly successful aspect of epidemiological culture. Evolutionary Psychology does not allow us to distinguish between these two cases.⁶¹

Beyond these worries with gathering evidence, there is a deeper problem in terms of the Evolutionary Psychologist's ability to tell us anything about modern human behaviour and culture, and their ability to use modern human behaviour and culture to support their evolutionary claims. There is a tension between, on the one hand, the claim that modern humans are maladapted to modern environments, and, on the other hand, the claim that we can explain culture in terms of our evolved psychological architecture. The remainder of this chapter will expand on this point and its consequences for Evolutionary Psychology.

3. Culture

⁶¹ Indeed, Levy (2004) argues that cross-cultural regularities in gender inequality may be better explained by something approaching the SSSM than by evolved cognitive adaptations.

Evolutionary Psychologists argue that culture is to be understood as the product of our evolved cognitive architecture, and that to gain a better understanding of culture we must understand this cognitive architecture and what it was selected for. They maintain that the social sciences have failed to adequately explain human culture and society, as a consequence of their commitment to the Standard Social Sciences Model (SSSM), and subsequent refusal to seriously countenance an evolutionary approach to psychology and culture. Evolutionary Psychologists also argue that other approaches to human behaviour and culture that incorporate evolutionary considerations (e.g. sociobiology) have failed to apply a genuinely Darwinian theory to behaviour and culture due to their focus on current adaptedness and behaviour rather than psychology. I want to argue here that Evolutionary Psychology cannot simultaneously maintain its opposition to these sorts of sociobiological approaches and oppose the SSSM. I will argue that Evolutionary Psychology's commitment to the idea that modern humans are ill-suited to their current environments means that it must either recognise a need for the SSSM in explaining modern human behaviour and culture, or it must deny that we are maladapted to modern environments, and move closer to sociobiology and its related fields. This latter move may require Evolutionary Psychology to reject, or at least seriously weaken, its commitment to the massive modularity thesis.

3.1 The Standard Social Science Model

Evolutionary Psychologists describe the social sciences as being largely underpinned by the Standard Social Sciences Model. The SSSM, according to Tooby and Cosmides (1992), combines two things, an account of how people develop, and an account of how culture operates. People are the products of their cultures and cultures are emergent phenomena such that they need to be studied at that level (rather than, say, the level of biology or psychology) according to Tooby and Cosmides' presentation of the SSSM. There are two causal arrows of interest here. One points from culture to individuals; minds are moulded by the culture or cultures they are surrounded by; they are the "blank slates" onto which

culture is written.⁶² The other points from an earlier culture to a later one: “the sociocultural level is a distinct, autonomous, and self-caused realm” (Tooby & Cosmides 1992: 28).⁶³ This leads to the idea that to study social science, one need not attend to the details of evolutionary biology or psychology:

In discussing culture, one can safely neglect a consideration of psychology as anything other than the nondescript “black box” of learning, which provides that capacity for culture... evolved, “biological,” or “innate” aspects of human behavior or psychological organization, having been superseded by the capacity for culture. The evolution of the capacity for culture has led to a flexibility in human behavior that belies any significant “instinctual” or innate component... (Tooby & Cosmides 1992: 32)

Evolutionary Psychologists are highly critical of this approach to studying culture:

After more than a century, the social sciences are still adrift, with an enormous mass of half-digested observations, a not inconsiderable body of empirical generalizations, and a contradictory stew of ungrounded, middle-level theories expressed in a babel of incommensurate technical lexicons... (Tooby & Cosmides 1992: 23)

Evolutionary Psychologists instead argue that culture is not an autonomous entity, but rather that it is the product of our cognitive adaptations and should be understood by first understanding those adaptations. We will only gain a firm

⁶² Levy (2004) argues that Evolutionary Psychology has created a straw man with their depiction of the SSSM. I will discuss this in more detail below.

⁶³ The picture of the SSSM painted by Tooby & Cosmides is structurally very similar to that painted by neo-Darwinism, with culture in the place of the genotype. That is, the causal arrows points from the genotype in one generation to the phenotype of that same generation and the genotype of the next generation. Similarly, the causal arrows point from culture at one time to the psychological and behavioural aspects of the phenotype at the same time, and culture at some later point in time.

grasp on understanding culture when we approach it from an evolutionary perspective:

... culture is the manufactured product of evolved psychological mechanisms situated in individuals living in groups. (Tooby & Cosmides 1992: 24)

...nothing the organism interacts with in the world is nonbiological to it, and so for humans cultural forces are biological, social forces are biological, physical forces are biological, and so on. The social and the cultural are not alternatives to the biological. They are aspects of evolved human biology and, hence, they are kinds of things to which evolutionary analysis can properly be applied. (Tooby & Cosmides 1992: 86)

I want to argue that Evolutionary Psychology may in fact require the SSSM. Before I do this however, I first need to discuss why Evolutionary Psychologists reject another approach to understanding modern humans using evolutionary theory.

3.2 Darwinian Social Science

While Evolutionary Psychology focuses on evolved psychological mechanisms, the group of approaches Symons calls Darwinian social science, or DSS (human behavioural ecology, sociobiology, evolutionary biological anthropology, Darwinian anthropology) rely on a “psychologically agnostic science of human behaviour” (Symons 1992: 146). Instead of focusing on psychological mechanisms, DSS concentrates on behaviour. The basic strategy of this approach is to study human populations and determine if their behaviour is adaptive (that is, whether cultural practices and systems increase the reproductive success of individuals and their kin in that culture). The basic assumption is that humans will tend to behave in ways that maximise their reproductive success, and that such behaviour is the result of an adaptation. Symons criticisms of DSS are focussed in particular on Crook and Crook (1988). Crook and Crook argue that human beings have evolved the capacity to behave adaptively in given social

contexts. Even in quite novel circumstances, the claim is that humans will behave in ways that maximise their inclusive fitness. They focus on the system of co-fraternal polyandry in Tibetan society. Their hypothesis is that this marriage system will prove to be adaptive given their particular ecological and socio-economic situation. Such communities farm land in harsh, high altitude conditions and suffer other socioeconomic difficulties, such as the high taxes demanded by local landlords. A good deal of labour is required to successfully farm this land. There is a basic tension, according to Crook and Crook, between achieving what we might call short-term reproductive success (that is, having as many children as is physiologically possible), and not over-populating the land and placing too much strain on limited resources. Optimal inclusive fitness is reached when the happy medium between these two points is reached. Crook and Crook argue that the marriage system employed in these communities maximises reproductive success given these ecological constraints. Women marry the eldest brother in a family, and subsequently his younger brothers. This means the male children in a family will continue to live together looking after the farm, thus providing the required labour for the running of the farm. If the sons married different women, and the farm land was divided up among the brothers, the land would soon become unviable for farming. In this system, the land is kept intact as the brothers live together. Presumably, also, one woman who marries several brothers produces fewer children on average than would be produced if each brother married separately, creating less demand on the limited resources available. Thus, this system is thought to be adaptive; it ensures the land remains suitable for farming and it provides labour to work the land (and perhaps dampens population growth). This marriage system is thought to be the result of an adaptation that allows humans to establish systems that maximises their inclusive fitness in different, even novel, environments:

The central prediction made in a Darwinian perspective is that humans are endeavouring, consciously or unconsciously to optimize their reproductive success... The genetic adaptation consists in the provision of a flexibility that allows reproductively optimizing behaviour to vary with context. (Crook & Crook 1988: 98-99)

What may have been demonstrated is that the value system of the individuals concerned have been shaped by adaptive learning to match the socioeconomic context in ways that are functionally effective with regard to reproduction. (Crook & Crook 1988: 98)

Symons has two points in relation to this. The first is that evidence of adaptive behaviour is not evidence of an adaptation. This, in itself, seems uncontroversial. Symons' second point concerns Crook and Crook's assertion that some general-purpose learning device could exist that would allow humans to behave adaptively in any given ecological or sociocultural environment. Evolutionary Psychologists consider the idea of such a general-purpose learning device to be incoherent. The commitment to the MMH entails that in novel environments – environments unlike our EEA – we should not expect people to behave adaptively in all respects:

Since the adaptations that underpin human behavior were designed by selection to function in specific environments, there is a principled Darwinian argument for assuming that behavior in evolutionary novel environments will often be *maladaptive*. (Symons 1992: 154)

Thus, we should not expect modern humans to display adaptive behaviour in any particular scenario. Indeed Evolutionary Psychologists suggest that we can assume the opposite, that humans are often maladapted to modern environments. The environment to which our cognitive adaptations have evolved to fit is, according to Evolutionary Psychologists, the Pleistocene. This is an environment (more correctly, a wide collection of environments) very different from modern industrialised cities. We can also assume the environment of our hunter-gatherer ancestors was quite different to that of agricultural societies. Humans, then, are not well adapted to the environments they have spent the last 10,000 years inhabiting. Thus, DSS is not, according to Symons, the correct way to employ Darwinian thinking to human behaviour and culture: “the hypothesis that human behavior is surprisingly adaptive does not derive from Darwinian theory and is *almost certainly wrong* in modern industrial environments” (1992: 155, my own

italics). Whether or not the co-fraternal polyandry of the Tibetans is adaptive, Symons argues, it is certainly not an adaptation.

3.3 Mismatch

The mismatch between our cognitive adaptations and our modern environments presents a problem for Evolutionary Psychology. Mismatches are not uncommon in evolutionary biology, and it is this that is often cited as the reason for the extinction of one species or another. But when this occurs with non-human animals, the changes that occur in the environment are not usually considered to have been caused by the species in question.⁶⁴ A new species competing for the same food or habitat, a new predator, human interference, or catastrophic geological or meteorological events are typically cited in cases of mismatch. However, the change in the human environment since the Pleistocene is, by and large, not the result of autonomous environmental change. For instance, when Symons discusses why we should expect modern humans to be maladaptive, he mentions industrial cities. Presumably, our more recent ancestors were maladapted to agricultural societies, given life in such societies would differ in numerous ways from hunter-gather lifestyles. These are clearly environments created by human activity; it is human activity that has changed the environment. One issue then is to account for such environmental change. I will argue that Evolutionary Psychology cannot do this.

There are two aspects to thinking about culture and mismatch here. The first is the generation of culture such that it becomes ill-suited to our cognitive adaptations. The second is how we respond to a culture mismatched to us. These issues may be closely related – the maladaptive behaviours that arise due to mismatching might contribute to the generation of cultures that are mismatched in new, or more extreme ways. But as both the generation of culture and the response of individuals to that culture raise various issues of their own, I will treat them separately here.

3.3.1 Generation of Culture

⁶⁴ Though see chapter two on niche construction for reasons to doubt this assumption.

Evolutionary Psychology is committed to the idea that our cognitive architecture was established in the Pleistocene and has changed very little since. It is this cognitive architecture that explains our behaviour and the generation of culture:

Evolutionary Psychologists' general methodological strategy is to explain culture through explaining *the evolved cognitive architecture that generates culture*. If we understand what our cognitive mechanisms "that actually produce behavior," were adaptations for, then we should be able to "predict behavior far more closely, even in modern cultures." (Tooby & Cosmides 1989: 37, emphasis added.)

Against this, however, Evolutionary Psychologists argue that our environments have changed quite radically. And Evolutionary Psychologists do not assume that these changes are merely superficial; recall Tooby and Cosmides', and Symons' suggestion that modern humans are *often* ill-suited to modern environments. The changes that have occurred in human environments are substantial enough to warrant Evolutionary Psychologists concluding that often our behaviour underpinned by cognitive adaptations will be maladaptive in modern environments. Evolutionary Psychology maintains that culture can be understood as the ultimate output of our cognitive adaptations, but how does this square with a culture so different from the EEA that we are maladapted? To put this tension into sharp relief, compare the following two quotations:

By directly regulating individual behaviour and learning, [innate psychological mechanisms] *directly govern cultural dynamics*; the key to understanding cultural processes must therefore lie in the discovery and subsequent mapping of the properties of these complex and specialized psychological mechanisms. (Tooby & Cosmides 1989: 30, emphasis added.)

Adaptive tracking must, of course, have characterized the psychological mechanisms governing culture during the Pleistocene, or such mechanisms could never have evolved; however, once human cultures were propelled beyond those Pleistocene conditions to which they were

adapted at high enough rates, the formerly necessary connection between adaptive tracking and cultural dynamics was broken. (Tooby & Cosmides 1989: 35)

On the one hand, it is argued that culture is the product of our evolved cognitive mechanisms, while on the other hand, culture has run away from us. How do we explain, on the one hand, culture developing beyond our capacities to deal with it, and on the other hand, the claim that significant aspects of culture are directly governed by our evolved cognitive architecture? I will explore some possible answers to this question, but none will be particularly attractive to the Evolutionary Psychologist.

Evolutionary Psychologists, as discussed above, divide culture into three categories: metaculture, evoked culture, and epidemiological culture. Metaculture refers to the deep similarities that are supposed to exist between different human cultures. Evoked culture refers to the different cultural practices that result from the same adaptations receiving different inputs from the environment. In cold environments, for instance, finding materials and making clothing might be the output of an adaptation. In a warm climate, the same adaptation will not lead to this behaviour. To count as evoked culture, it must be behaviour that has occurred in the past and that has been adaptive and thus selected for. A cognitive adaptation that produced some entirely novel behaviour in an entirely novel environment could not be thought to be contributing to evoked culture.

The third type of culture Evolutionary Psychologists discuss is epidemiological culture. This arises out of individuals making discoveries of some sort (a quicker way to make a spear, a catchy tune, a new word), others learning from this individual's behaviour, and the trait slowly spreading throughout the population. This type of culture is the most likely explanation for cultural change, especially where cultural change is not preceded by any significant changes in the non-cultural aspects of the local ecology. (Where a cultural change occurs due to ecological change, the Evolutionary Psychologist can maintain that this cultural change is the result of evoked culture so long as the ecological change brings about conditions similar to those that existed in the Pleistocene.) Evolutionary Psychologists maintain that cultural differences are

superficial compared to the deep commonalities between cultures. If epidemiological culture is what accounts for (some of) the differences between cultures,⁶⁵ it must only be superficial. It is difficult to argue about what is and is not superficial, such things can come down to explanatory goals or just personal interest.⁶⁶ However, the ways in which modern humans are mismatched do not appear to be superficial by the standards of Evolutionary Psychology. For example, Evolutionary Psychologists, like evolutionary biologists more generally, put a good deal of emphasis on adaptations relating to reproduction. A good deal of time and effort has been put into avoiding reproduction in more modern societies, from celibate monks and nuns, to the availability of contraceptives. So if mismatch arises out of epidemiological culture, Evolutionary Psychologists must accept that epidemiological culture does not just concern “superficial” matters, but can alter behaviour in all manner of ways.

Evolutionary Psychologists also maintain that epidemiological culture is constrained by our evolved cognitive architecture. First of all, our cognitive adaptations will bring about epidemiological culture:

... epidemiological culture is... shaped by the details of our evolved psychological organization. Thus, there is no radical discontinuity inherent in the evolution of “culture” that removes humans into an autonomous realm. (Tooby & Cosmides 1992: 119)

Second, our evolved cognitive architecture will determine what ideas and behaviours hit upon by one or a few individuals go on to form epidemiological culture. Someone may come up with a good idea or a clever way of doing something, but we need to be able to recreate the “representations or regulatory elements” in our own psychological architecture (Tooby & Cosmides 1992: 118). We require, according to Evolutionary Psychologists, mechanisms to allow us to infer the representations in the minds of those we observe:

⁶⁵ Evoked culture will also explain some differences between cultures.

⁶⁶ See the discussion in chapter one of a similar issue in relation to explanatory adaptationism.

... this task of reconstruction would be unsolvable if the child did not come equipped with a rich battery of domain-specific inferential mechanisms, a faculty of social cognition, a large set of frames about humans and the world drawn from the common stock of human metaculture, and other specialized psychological adaptations designed to solve the problems involved in this task... Mechanisms designed for such inferential reconstruction evolved within a pre-existing complex psychological architecture and depended on this encompassing array of content-structuring mechanisms to successfully interpret observations, reconstruct representations, modify behavior, and so on. (Tooby & Cosmides 1992: 119)

So, epidemiological culture is constrained by our evolved psychological mechanisms on both sides. Our adaptations constrain the kinds of ideas or behaviours we might come up with, and our adaptations help determine what kinds of ideas or behaviours will succeed and spread throughout the population.

However, if epidemiological culture is shaped by the details of our evolved architecture in this way, it is difficult to see how this architecture can account for the generation of modern cultures to which we are supposed to be so unsuited. On the other hand, if epidemiological culture is responsible for the mismatch between our adaptations and our sociocultural environment, then the mechanisms that are thought responsible for the generation and transmission of culture cannot be particularly tightly constrained. In this latter case, Evolutionary Psychology will have little to say about epidemiological culture. Knowing about the evolution of the mechanisms involved in the generation of epidemiological culture will shed little light on the direction epidemiological culture has taken. As an analogy, evolutionary biologists may be able to tell us a great deal about the hand as an adaptation, but this adaptationist story will not tell us very much about the ability to write, play the guitar or sow. That is, the adaptationist story may put some constraints on what the human hand can do, but these constraints are quite broad. To understand writing, for instance, at best the evolutionary story will only feature as a partial explanation, and in many cases, may not be explanatorily relevant at all. Whatever the hand was adapted for, this has only a small bearing on writing now. Similarly, if the mechanisms that generate

epidemiological culture provide only the broadest of constraints on the kind of culture that can be generated, then knowledge of why such mechanisms evolved in the first place, though a potentially interesting story in its own right, will have little bearing on attempts to understand modern behaviour and culture.

Evolutionary Psychologists might attempt to diffuse this problem by claiming that epidemiological culture may have snowballed such that our modern environments are strikingly different than our EEA, but meta- and evoked culture retain their tight grip. Given this, Evolutionary Psychology still has a role in uncovering these aspects of our cultures and explaining them in terms of natural selection. This move will not work however. Epidemiological culture overrides much of meta- and evoked culture, otherwise there would be no interesting sense in which we were maladapted from the Evolutionary Psychologist's point of view. That is, if our cognitive adaptations were still producing meta- and evoked culture, then we would remain generally suited to our environments. If we are mismatched in important ways (recall the example of contraceptives above), then the adaptations responsible for evoked and meta-culture cannot be producing culture. Modern culture is then primarily epidemiological culture. And if Evolutionary Psychologists are forced to concede that the mechanisms which generate epidemiological culture only offer the broadest of constraints, they have rendered themselves irrelevant in attempts to explain culture and behaviour. That is, if modern cultures are different enough to allow us to assume, as Symons, Tooby and Cosmides do, that we are often maladapted, then Evolutionary Psychology has some work to do to explain this. In terms of the generation of culture, Evolutionary Psychology has to decide whether it wants to maintain that epidemiological culture is highly constrained and leave the appearance of mismatched cultures mysterious, or loosen the constraints placed on epidemiological culture, in which case Evolutionary Psychology loses its ability to explain very much about modern cultures.

3.3.2 Responding to Mismatched Culture

Evolutionary Psychologists might grant that the generation of epidemiological culture is largely unconstrained by evolved architecture and not something it has a great deal to say about. However, they may maintain that Evolutionary Psychology still has a role to play in explaining maladaptive behaviour in novel

environments. It cannot explain why such novel cultures come into existence, but it might explain why we behave in the maladaptive ways in response to such environments.

There are three ways in which we might expect individuals with adaptations to one sort of environment, or set of environments, to behave maladaptively in a new environment. The first is if an individual behaved in a manner appropriate for our EEA, but not for a modern environment. For example, a dispute between two people might end in violence. This might have been an adaptive response in the harsher environment of our ancestors, but is very likely to be maladaptive now. I will refer to this sort of a mismatch as a “positive mismatch” (in that a thought process or behaviour is actually elicited). The second kind of mismatch I will call a “negative mismatch” as it concerns the failure of any thought process or behaviour to be elicited. For example, Cosmides and Tooby (1997) claim that electrical sockets pose more of a threat to most Americans than snakes, yet these same people are more frightened of snakes. They argue that while we have evolved to be wary of snakes, we lack any such adaptation for electrical sockets (or many other features of modern environments). Mismatches occur because we *fail* to respond in this sort of case. The third way mismatch may come about is that, in a novel environment, cognitive adaptations receive new inputs which generated entirely novel outputs. This unpredictable output may be adaptive, but it is more likely to be maladaptive – there are just more ways to be maladaptive – in which case I will refer to it as “unpredictable mismatch.” I will argue that behaviour generated by positive mismatch is incompatible with the claim that modern cultures have become increasingly unsuitable for our cognitive adaptations. Evolutionary Psychology initially looks on firmer ground with failure to behave appropriately as a result of negative mismatch, but even here I will argue that its role is seriously circumscribed. Finally, I will argue that Evolutionary Psychology can tell us nothing about behaviour that stems from unpredictable mismatch.

To demonstrate why a commitment to the idea that we are currently maladapted rules out the possibility that we frequently produce old ways of behaving (positive mismatch), let’s take a hypothetical scenario based loosely around a hypothesis often proposed by Evolutionary Psychologists. According to Evolutionary Psychologists, men prefer as mates younger women due to their

greater reproductive potential, while women (at least of reproductive age) tend to prefer older men, as they are likely to have greater resources to provide for them and their children. As a result, we should expect younger women to marry older men. Now let's imagine a culture which, for some reason, considers marriage between two people of similar ages ideal, and marriage between people of vast age differences abhorrent; married couples with large age gaps face being ostracised by their communities, and so on. To act according to our supposed evolved preferences would be maladaptive. There are two possibilities here. Either people tend to act in accordance with their cultural norms, or they can act in accordance with their evolved preferences. If they act in accordance with their cultural norms, then Evolutionary Psychology has little to say about the behaviour of most people and (given the previous discussion) little to say about the origin or persistence of these cultural norms. At most, it may say something about the transgressors of the cultural norms (and even then, not all of them – some instances will have entirely different causes). At best, Evolutionary Psychology is highly marginalised. The real action is going on at the level of cultural norms, not evolved psychology.⁶⁷

If people tend to act in accordance with their evolved preferences instead of cultural norms, this would suggest that cultures should never stray too far from conditions similar to the EEA. Enough people breaking the rule about marriage would soon see that marriage system deteriorate and eventually abandoned. Societies might experience occasional perturbations as the result of epidemiological culture, but would generally hover around an equilibrium position.⁶⁸ So, if mismatch leads to old behaviours in novel environments, we

⁶⁷ Evolutionary Psychologists may argue that their prime explanatory target is psychology not behaviour. Thus people may act in accordance with cultural norms but, for instance, feel conflicted in doing so, and Evolutionary Psychology can explain why. This is certainly an option available to them (though the discussion in the earlier half of this chapter casts doubt on the success for such a strategy). However, such a move would be to acknowledge that they are unable to discuss culture.

⁶⁸ It might be argued that modern cultures represent only a perturbation in a system that will eventually right itself. 10,000 years is a short period of time from an evolutionary point of view. This line of thought, however, will not work. Cultural change can happen at a much faster pace than genetic evolution, thus 10,000 years is a lot of time for the system to have corrected itself. Tooby and Cosmides also deny that cultures are likely to resolve themselves to some optimal condition: "... there is no a priori reason to suppose that any specific modern cultural or

should expect mismatch to be rare as the old behaviours regenerate older, more “optimal” culture. On the other hand, if, as Evolutionary Psychologists suggest, mismatch often occurs such that we ought to assume it, and that the previous 10,000 years have been characterised by a trend away from EEA conditions, this indicates that people do not tend to act in accordance with their evolved preferences. In this case, Evolutionary Psychology has little to say about modern human behaviour or culture.

Evolutionary Psychology also faces problems with negative mismatch (the absence of an appropriate response to modern environments due to the absence of adaptations for modern environments). If we grant that epidemiological culture can exceed the limitations of our evolved cognitive adaptations, as we must do if modern cultures are unsuitable, we must also grant that we will be capable of behaving in ways we do not have specific adaptations for. Knowing that our EEA did not have a given feature tells us nothing about what modern humans can or cannot do. At best it may tell us about some of the things we cannot do because, first, our EEA did not contain the relevant structures *and*, second, epidemiological culture has not provided us with the wherewithal to deal with the particular problem. The inability to behave adaptively in a given situation is only partially explained by evolutionary considerations of the sort Evolutionary Psychologists employ. That is, the explanation will also have to refer to why epidemiological culture does not allow us to deal effectively with the problem. Recall Cosmides and Tooby’s (1997) example of a negative mismatch: Americans are more at risk from electrical sockets than snakes, but are more wary of snakes than electrical sockets. Perhaps the absence of electrical sockets in our EEA explains why this is so. But if epidemiological culture can allow us to do things we have not been adapted for (agriculture, playing computer games, using contraception), then there is no principled reason why it might not also compensate for the lack of an adaptation to assess risk in relation to electrical sockets. So, to fully explain why we assess the risk of electrical sockets so poorly (if in fact we do), it is not enough to consider what we have been adapted for. Evolutionary Psychology alone cannot mark out the limitations of our cognitive abilities. At best, Evolutionary Psychology provides part of the

behavioral practice is 'adaptive'... or that modern cultural dynamics will necessarily return cultures to adaptive trajectories if perturbed away” (1989: 35).

answer. In any given situation, it will be an open question as to whether epidemiological culture can supplement what evolution did not equip us with.

The final option is that new environments provide input to our cognitive adaptations, and this results in some entirely novel behaviour being produced. This novel behaviour may occasionally, through blind luck, be adaptive, but we can probably assume that it usually will not be – again, there are far many more ways for something to go wrong than right. This sort of approach has potential in that it might explain why modern cultures have drifted so far from the EEA. If some small changes (perhaps due to epidemiological culture or change in the broader ecology) fed into our cognitive adaptations and produced strange and unpredictable behaviour, this potentially could have a snow-ball effect, pushing culture further and further away from its “optimal” state. However, the ability to explain why modern cultures are so far away from this optimal position comes at a high cost for Evolutionary Psychology. Evolutionary considerations will do little to explain these novel behaviours. Certainly, evolution can be thought of as creating the cognitive architecture that underpins these behaviours, but it cannot explain why a given behaviour was generated in a given environment – there is no evolutionary logic to the generation of this behaviour. The behaviour is a by-product of cognitive adaptations, not their function. This is similar to the point above about the evolution of the hand and writing as a by-product. Knowing what a trait is adapted for tells us very little about by-products of that trait. The situation is in fact worse for Evolutionary Psychology than it is for the biologist who can tell us something about the evolution of the hand. The biologist can at least point to facts about physiology that constrain and enable the holding of a pen and so on. Given Evolutionary Psychology focuses solely on function (what the cognitive mechanism was adapted for) rather than the neurological details of mechanisms, it will have even less to say about any by-products of these cognitive adaptations.

3.3.3 The Upshot for Evolutionary Psychology

Evolutionary Psychology cannot explain how modern mismatched environments came about. It also cannot explain how this mismatch manifests itself. Behaviour as a result of positive mismatch, except in a minority of cases, seems ruled out as this would push modern cultures back to more “optimal” states. At best

Evolutionary Psychology only has a partial role in explaining the failure to act as a result of negative mismatch, and it has nothing at all to say about behaviour generated by unpredictable mismatch. This leaves Evolutionary Psychology with three options: deny mismatch occurs, accept DDS, or accept the SSSM. I will outline the consequences of each of these options for Evolutionary Psychology.

Evolutionary Psychologists could reject the claim that modern humans are ill-suited to modern environments. They might argue that humans are not maladaptive in modern cultures because we replicate the conditions of the EEA (in the manner discussed above in relation to positive mismatch). The differences between conditions created by current cultures and the EEA would be merely superficial on this account. This gives Evolutionary Psychologists two options. They can argue that agriculture, industrialisation, systems of writing, birth control, states of all political hues, China's one child policy, and so on, all recreate at some deep, structural level conditions of the EEA. Alternatively, they can argue that these things are merely superficial. Neither strategy seems particularly promising. It is very difficult to see how, for instance, birth control and industrialisation recreate EEA conditions; this seems outright implausible. On the other hand, arguing that such behaviours and cultural practices are merely superficial elements of modern cultures seems difficult to accept too. For instance, given that reproduction is at the heart of natural selection, modern contraceptive use seems to amount to more than a superficial element of culture. In any case, it is evident that Evolutionary Psychologists do not currently think such cultural practices are superficial, given their assumption that we are often maladapted to our modern worlds. Although arguing that we are in fact well adapted to modern environments would allow Evolutionary Psychologists to hold on to most of their other theoretical commitments (unlike the alternatives I will shortly explore), it is not clear how such a move could be achieved. An awful lot of work would need to be done to establish that modern environments really do substantially replicate our EEA.

Alternatively, Evolutionary Psychology could move closer to DSS and claim that we are capable of more flexible responses to the environment which ensure our adaptedness. In line with the sort of argument presented by DSS, evolution would have had to endow us with the ability to respond effectively to brand new circumstances and to learn new ways of coping with them such that

we remain adaptive. Cultures may change over time, but we are capable of behaving adaptively in such new environments because we have some general capacity for culture. This would allow evolutionary considerations to continue to play a role in explaining modern human behaviour and culture. However, this move cannot be made by Evolutionary Psychologists unless they relinquish their MMH. Recall Tooby and Cosmides (1992) argument for MMH was motivated by the thought that a general-purpose problem solver – in this case a general ability to generate and acquire cultural practices – is incoherent. If Evolutionary Psychologists accept such a general-purpose capacity for culture they have undermined their argument for domain-specific, content-rich modules. It is possible to develop accounts of the mind that mix special-purpose modules with general-purpose capacities. However, if Evolutionary Psychologists wish to move in this direction, they will need to develop an entirely new set of arguments. The argument offered for the MMH was the implausibility (even impossibility) of a domain-general problem solver. If they grant that such general-purpose problem solvers are possible, they have removed the foundation of their particular argument for any modularity. New grounds would need to be sought to establish the idea that any modules are required.

Finally, if Evolutionary Psychologists wish to maintain that we are fundamentally ill-suited to modern environments, they cannot explain modern cultures or behaviours. Evolutionary Psychologists could accept a division of labour with the SSSM. Evolutionary Psychology can attempt to explain the evolution of our cognitive architecture, but leave discussions of human behaviour and cultures since the Pleistocene to those employing the SSSM. Evolutionary Psychologists might concede that evolutionary considerations will not be enough to explain everything about human behaviour and cultures. Indeed, they do admit this much, though what they expect to not be able to explain is merely superficial differences, rather than the much larger portion of human behaviours and cultures at stake here. But even if they acknowledge that they cannot fully explain modern human behaviour and culture, they might still argue that the SSSM ought to be avoided due to their belief that its basic commitment to the mind as a “blank slate” is utterly wrong-headed. However, it is not clear that the SSSM really is committed to the mind being a blank slate:

Certainly, no-one – not even Skinner and his followers – has ever believed in the blank slate of Pinker's title. Even behaviorists believe that the human mind has in-built learning mechanisms and preferences, in the absence of which schedules of reinforcement would be useless. All sides in the nature-nurture debate (at least all minimally rational sides) are interactionists. (Levy 2004: 461)

The SSSM then does not deny that the mind has some structure, and that this structure may have some evolutionary explanation, it rather maintains that this structure is flexible enough to allow for some general capacity for culture. If the point is to study culture, this general capacity will not be particularly interesting or explanatorily relevant. This is similar to the view of the mind employed by DSS, except that the SSSM does not require that the mind produces adaptive behaviours and cultural systems. Thus the SSSM model of the mind is more plastic than that assumed by DSS, but this has the benefit of allowing Evolutionary Psychology to claim modern environments are not what humans adapted to, and thus we should not expect modern humans to tend to act to increase their inclusive fitness.

Evolutionary Psychology cannot help us explain – even in quite broad terms – the kinds of cultural changes that have occurred since the Pleistocene, nor can it explain how individuals react to these changed circumstances. It cannot give us guidelines, drawn from evolutionary considerations, about what we should and should not expect people to do. If cultural change has been significant enough to render us often maladapted, it seems culture can evolve well beyond any constraints our adaptations might have been assumed to place on culture. And if culture is nothing more than the product of our minds, our minds can do a good deal more than what they have been adapted for. Understanding our cognitive adaptations *as adaptations* will shed very little light on the generation of modern cultures or the behaviour of modern humans. Research into human behaviour and culture would be preferable, methodologically, assuming very little about the mind, at least from an evolutionary point of view. Building into any theory the details of the evolutionary explanations for cognitive architecture will be superfluous to demands given that evolutionary understandings of cognitive architecture will do

no work for anyone interested in modern human behaviour or culture. If we accept that we live in environments mismatched to our adaptations, whatever constraints the evolved architecture of our mind imposes, on the Evolutionary Psychologist's account, they must be broad and uninformative.

4. Conclusion

Evolutionary Psychology has its foundations in neo-Darwinian evolutionary theory. However, Evolutionary Psychologists make stronger claims about evolution and development than are warranted given this foundation, and do so without justification. To establish the universality of our supposed evolved cognitive architecture, Tooby and Cosmides allow that this universality can be at the level of genetic programmes; however, universality is discussed only the level of cognitive architecture once the case for universality is made. Similarly, the commitment to the idea that “our modern skulls house a stone age mind” downplays the role of genetic switches and the potential for evolutionary changes in periods of time shorter than they will allow. Again, such switches have become uncontroversial in neo-Darwinian biology, and Tooby and Cosmides do not offer justification for the marginalisation of such phenomena.

Further, the adaptationist programme poses a range of problems specific to, or at least more acute for, Evolutionary Psychology. It is very difficult, if not often impossible, to determine the selection pressures our ancestors would have encountered in their social worlds. This means that Evolutionary Psychology must rely on modern humans as a major source of information about our evolutionary past. However, modern humans – whether living in industrialised cities or hunter-gatherer societies – often do not appear to behave in accordance with the kind of cognitive adaptations Evolutionary Psychologists propose. We have created and maintained very different cultures from those of our ancestors by Evolutionary Psychologists’ own admission. What this means is that Evolutionary Psychologists will find it difficult to use modern human behaviour – or inferences about our psychology based on this behaviour – as evidence for claims about our Pleistocene ancestors. Ultimately, the tension for Evolutionary Psychologists with regard to culture arises here because on the one hand, they want to argue that culture is generated and constrained by our evolved psychology, and on the other, that we are often poorly matched to modern

cultures. If Evolutionary Psychologists downgrade their commitment to the idea of mismatch, then they are forced to either explain why modern cultures really are like our EEA, or they are forced to move closer to DSS and abandon their massive modularity hypothesis. Alternatively, if they feel it is implausible that modern environments are really that close to our EEA (as they appear to do at least some of the time, when not discussing cultural universals), and want to keep the idea of mismatch, then they are forced to allow the SSSM to do most of the work in explaining modern human behaviour and cultures.

Although I have focused here on Evolutionary Psychology, these issues will pose a problem for any attempt to offer an evolutionary explanation for human psychology and behaviour that supposes we are adapted to ancestral environments and mismatched to the modern world. There is a requirement for any theory of this sort to explain why modern cultures are so different. And because modern human behaviour and culture is so different, it cannot be used to decide between different hypotheses about selection pressures in our ancestors' social environments. Even if we do possess cognitive adaptations for our EEA, it seems unclear that we behave in accordance with them. Thus, any approach to the evolution of human psychology and behaviour that suggests we have evolved to suit very different environments to the ones we currently inhabit will be faced with a large evidential gap.

In the following chapter I will explore a broad class of approaches to the evolution of culture known as gene-culture co-evolutionary accounts. I will suggest how elements of this sort of approach can be made compatible with a developmental systems perspective. Unlike Evolutionary Psychology, these approaches do not focus to the same extent on the details of psychology. However, especially in the case of the model of gene-culture co-evolution presented by Richerson and Boyd (2005), a model of human psychology not unlike that employed by Evolutionary Psychologists is assumed. In chapter seven I will examine a very different model of cognition. Together, chapters six and seven will present an alternative way to understand the evolution of cognition and culture which incorporates the developmentally informed view of evolution developed in part one of this thesis.

Chapter Six

Biology and Culture

1. Introduction

A different kind of evolutionary approach to culture has been developed under the heading of gene-culture co-evolution. I will briefly examine one version of this – memetics – and argue that this approach adds very little to our understanding of culture or cultural evolution. The second approach, developed by Richerson and Boyd (2005), is more promising. They develop models of cultural evolution that undermine the worry that vertical transmission and error-prone learning make cultural evolution unlikely. While I will argue that they have identified some plausible mechanisms that might underpin cultural evolution, I will reject their assumption of evolved psychological modules of the sort Evolutionary Psychology supposes, and their dichotomous genes/culture account.

Richerson and Boyd's work on models of selection provides a means of dealing with an objection raised against developmental systems theory and discussed in chapter four – that of vertical transmission. I will use their general approach to sketch an account of the evolution of developmental systems that can be partially constituted by culture. However, I will argue that evolutionary explanations will not provide a complete explanation for culture, and using the work of Hacking (1996) on the looping effects of human kinds, suggest that the study of culture cannot easily be made contiguous with the natural sciences.

2. Gene-Culture Co-Evolution

Evolutionary Psychology views culture as the product of evolved cognitive architecture, and maintains that culture can be explained primarily by understanding these cognitive mechanisms. DSS, on the other hand, is non-committal about the nature of our psychology, save that we have evolved the capacity to establish an adaptive way of living in novel environments and to learn these ways of living from others. It nonetheless views culture as underpinned by the drive of individuals to maximise their inclusive fitness. A third approach to evolution and culture takes culture as something to be

understood in its own right. It may be related to biology, but it is not tightly constrained by it, and it can be largely understood at the level of culture (rather than psychology or biology). Cultures or cultural variants can, on these accounts, evolve and form lineages:

... cultural diversity should be understood in many respects in the same way as biological diversity – that is, as the result of an evolutionary process. In particular, individual variation, and processes analogous to selection, should be seen as providing a historical basis for the present existence of various internally articulated and integrated cultural forms. (Dupré 2002: 140)

Possibly the most well-known version of this view is memetics, and I will briefly discuss this to begin with. Far more interesting however is an approach developed by Richerson and Boyd (2005) which applies population thinking to the issue of cultural evolution. I will discuss this approach, and although I will ultimately disagree with some basic assumptions underlying it, I will suggest that a number of the tools developed by Richerson and Boyd can be of use to DST. Finally, I will sketch a picture of culture and cultural evolution from a DST perspective.

2.1 Memes

Cultural evolution, it might be thought, shares many similarities with biological evolution. There can be variation within a population, there can be inheritance of culture (or aspects of culture) from one generation to the next (and from peer-to-peer – I will discuss this complication in more detail below), and some cultures/aspects of a culture will be more successful than others in terms of their longevity. Memetics, for instance, takes this approach. The idea here is that we can understand culture as a collection of memes – discrete, gene-like ideas that behave in much the same way as genes. They form a distinct inheritance channel, and are replicated in each generation. Some memes disappear quite quickly, while others seem to spread from person to person over longer periods of time (Dawkins 1989; Blackmore 1999).

However, Sperber (1996, 2000) argues that culture cannot be understood in this way; it is neither particulate, not replicated. Certainly some cultural variants may conform to something like the meme-like picture. Sperber uses the example of a chain letter, for instance. Like genes, successful memes are supposed to benefit themselves through promoting their own replication, and not necessarily benefit the vehicle for their replication. A chain letter which warns of dire repercussions for the person who does not forward the letter does not benefit the individual who passes on the letter, but ensures its propagation throughout the population. Something like this, Sperber grants, might be considered a meme. But, he argues, such examples are not representative of culture more generally. Culture cannot be understood as being composed of discrete, gene-like entities. Cultural variants are not inherited vertically on this picture, and beyond this, are not even inherited from one or two parents, but potentially from a large number of sources. Memes do not appear to form lineages:

In general, if you are serious in describing bits of culture – individual texts, pots, songs or individual abilities to produce them – as replications of earlier bits, then you should be willing to ask about any given token cultural item: of which previous token is it a direct replica? In most cases, however, you will be forced to conclude that each token is replica not of one parent token, nor (as in sexual reproduction) of two parent tokens, nor of any fixed number of parent tokens, but of an indefinite number of tokens some of which have played a much greater “parental” role than others. (Sperber 1996: 104)

If memes are analogous to genes, we should also expect them to be replicators. However, a problem often raised against the idea of memetics is that the “replication” of memes is highly error-prone. While genetic replication is quite faithful, people can misremember a recipe, or the lyrics to a song, and so on. Selective forces can be rendered ineffectual against a high enough mutation rate (Williams 1966; Sperber 2000; Richerson & Boyd 2005). Whatever selection acts upon must persist long enough for a trend to emerge. The lack of fidelity in the replication of memes (and the transmission of culture in other theories of

cultural evolution) is taken to count against the possibility of selective forces having any effect on cultural traits.

Dawkins (1999b) argues that this objection is wrong. Dawkins and Sperber discuss the same sort of issue, and I will use Sperber's example here. Sperber imagines two groups of people where one person from each group is shown a picture and asked to make a copy. They must then show their picture to the next person in the group and ask them to make a copy. The process continues until all the people in the group have drawn a picture. The first group is shown a familiar picture, which in Sperber's example is a star. The second group is shown a random pattern. Sperber's reasonable assumption is that the individuals in the group that began with a picture of the star will all produce a picture of a star, while the group that began with the squiggle will produce pictures that tend further and further away from the original picture. The first picture will be more easily reproduced than the second. Dawkins explains the difference in fidelity of the copying as a result of individuals in the first set following a particular set of instructions for drawing a star – the “genotype” for the star's “phenotype” (Sperber 2000) – while the individuals in the second set have no access to the instructions or “genotype” for the scribble, and can only try to reproduce the “phenotype.” In terms of the star, Dawkins argues, “the instructions are self-normalising. The code is error-correcting” (Dawkins 1999b; quoted in Sperber 2000). Errors are thus the result of individuals not having access to the instructions. Once we have access to the instructions, faithful replication can occur. Of course, the group that reproduced the picture of the star will not have produced perfect facsimiles of the star; there will be variations between each drawing. However, there will still be a recognisable star in each case. This, it is argued, is much like biological reproduction. The genotype is replicated, but the phenotype can vary as a result of environmental influences and the vagaries of development. Despite variations in the phenotype, replication is still occurring at the level of the genes. Similarly, while the stars may not be identical in each picture, the instructions followed will be the same. Thus, once we have access to the instructions we can replicate a meme for long periods of time and so the idea that memetic evolution lacks fidelity is thought to have been dealt with. However, as Sperber notes, this has only replaced one problem with another: “saying that the instructions are 'self-normalising' amounts to resolving a problem by

invoking a mystery” (2000). That is, how is it that the instructions are so much better replicated than the drawing? Dawkins has not adequately responded here to the worry that cultural transmission is too error-prone to allow for selection to operate.⁶⁹

Ultimately, the meme approach is beset by problems. Even if we accept, as Sperber and Richerson and Boyd do, that some things like chain letters exhibit meme-like (or gene-like) properties, it is far from clear that this is typical of all aspects of culture. Culture does not replicate itself in the way genes are supposed to (on a neo-Darwinian reading of genetic replication at least), and cultures are very difficult to partition up into neat particles for selection to act on. For example, Lewens notes:

Ideas stand in logical relations to each other. Whether an individual is able to acquire some belief, for example, depends on their related conceptual competencies. It is impossible to believe in the theory of relativity without understanding it, and one cannot understand it without holding many additional beliefs relating to physics. (Lewens 2007)

However, as Lewens, notes, genes cannot be properly understood in isolation from other genes either. At the very least, genes need to be understood in the context of the DNA sequence, and a DST perspective would suggest a far wider context again. Despite this, it might be argued by neo-Darwinians, we can still think about genes as being selected for and forming lineages. Similarly, while memes might only be understood fully when viewed in their broader context, we can still talk about and study memes.

Perhaps the most damning problem for memetics is that it tells us very little; it merely re-describes well-known phenomena in the language of memetics, but gives us no new insight into these phenomena (Lewens 2007). Some ideas or

⁶⁹ Sperber himself argues that the meme approach cannot handle this problem of replication, because replication is not what is occurring. The failure of memetics to recognise this is because it is relatively neutral about human psychology. Sperber prefers an account that is based on domain-specific competencies that allow us to infer the intentions of the artist of each star drawing, so that rather than copying instructions for the drawing, the next person is able to infer the intentions of the previous artist and draw the star themselves. Thus, the star is not replicated, but rather re-produced.

behaviours are very popular and persist for many generations, others do not. Memetics does not tell us why some memes succeed while others fail. Memetics appears to be descriptive (and descriptive of potentially only a very small aspect of culture) rather than explanatory. If we want to know *why* certain ideas succeed or fail, we need to look somewhere other than memetics.

2.2 Population Thinking

Richerson and Boyd (2005) offer a more promising account. They reject the idea of memes, agreeing with Sperber (and others) that cultural variants are rarely particulate and are not faithfully replicated. However, they disagree with Sperber that the problem of error rules out the possibility of cultural evolution.

Culture need only be seen as meme-like if our account of cultural evolution is to be analogous to neo-Darwinian evolution. However Richerson and Boyd suggest that cultural evolution will not be strictly analogous to (neo-Darwinian) biological evolution. Nor need it be. In the case of DST, what can be inherited can come from numerous sources, and these resources cannot be understood independently, but only in the context of the entire developmental system.

One reason for treating culture as a collection of discrete gene-like memes stems from an early criticism of Darwinian evolution theory by Jenkin (1867). The worry here was that a system of inheritance like the one Darwin (1905) proposed (pangenesis) would lead to the blending of traits and eventually the averaging out of any variation. The worry is that if cultural inheritance is not particulate blending will eliminate variation and leave nothing for selection to work on. However, Richerson and Boyd argue that this need not be a concern for cultural evolution. The issue here turns on rates of mutation. In genetic evolution, mutation rates are taken to be relatively low; however, with cultural evolution the tendency for error means that mutation rates will be high. This means that the variation in the population will not level out; errors will continue to introduce novel variants. Contrary to the worry, raised above, that high levels of error ruled out cultural evolution by natural selection, moderately high error may actually make possible cultural evolution.

Richerson and Boyd admit that very little is known about the kinds of cultural variants that will be subject to selection, or how exactly we should understand them, but argue that “if it were true that adaptive evolution depended

critically on the units of transmission, Darwin and all his followers would still be marking time, waiting for developmental work definitively showing how genes give rise to the properties of organisms” (2005: 81). Such considerations allow Richerson and Boyd to remain relatively agnostic about the exact nature of cultural variants that selection might work on, instead focussing on the mechanisms that might enable cultural evolution.

But although some error may in fact help a non-particulate theory of cultural evolution, if the error rate is too high no cultural variant will be visible for long enough to be subject to selection. Richerson and Boyd suggest that the problem of error-prone learning can be overcome by focussing on the population, rather than the individual, level. That is, errors can be made by individuals, but at the population level we still see cultural variants persist. Richerson and Boyd propose three mechanisms that will keep error in check: content bias that allows individuals to decide between variants based on their perceived value, conformist bias that means that individuals tend to adopt the most common cultural variant, and prestige bias that means that individuals tend to adopt the cultural variant adopted by successful members of the population. With these three mechanisms, they argue, errors made by individuals in learning can be minimised to the extent that selection is possible.

Content bias allows individuals to decide between variants in the population. One individual may have made some error in learning a particular recipe, for example.⁷⁰ Now there are, in this very simplistic example, two variants to choose between – the original recipe and the new one. Perhaps the new recipe is preferred because the food produced tastes nicer, or because it relies on more easily available ingredients. In such a case the new recipe may spread through the population. On the other hand, perhaps the new recipe is inferior. In this case, other individuals can choose to reject it and stick with the original recipe. The point here is that individuals may make errors, but this does not mean that errors will propagate throughout the population. If people

⁷⁰ “Error” here does not have to imply mistake on the part of the individual. Error refers only to some change in the original cultural variant. The individual in question may have consciously chosen a different recipe. Guided variation of this sort is another reason to reject any tight analogies with neo-Darwinian evolution.

generally prefer one way of doing things, for whatever reasons, even quite frequent errors do not have to undermine the persistence of cultural variants.

Conformity bias is a tendency for individuals to imitate the most common variant in the population. Errors that occur in any one individual's attempt to imitate the behaviour of others will fail to be imitated by others in a large enough population where individuals are exposed to a representative sample of that population. Individuals who have made errors will form a minority of any sample, thus a conformity bias will ensure that a new imitator will favour the common type.

Finally, a prestige bias is a tendency for individuals to imitate the successful. As an example, going to university may be a cultural variant adopted by teachers. Let's assume for this example students view teachers as successful and so want to imitate them by going to university too. At least some of those students will go on to become teachers and exert the same sort of influence over future generations of students. The cultural variant – going to university and becoming a teacher – persists. For prestige bias to take place, there needs to be a correlation between those things that signal success and those things that cause success. For example, some students may make an erroneous connection – perhaps they suppose the teacher's accent or dress sense, rather than education, is the key to success and instead copy these aspects of the teacher's behaviour. Such individuals will not go on to be teachers, and their cultural variant (the accent or dress sense) is not imitated by future students (though future students may also independently hit upon these variants). However, so long as there is a causal connection between the cultural variant in question and the success of the individual, this will have the effect of dampening error. Those who get it wrong will be less able to influence future generations than those who get it right. With these three mechanisms, approached from the level of the population rather than just the level of individual learning, Richerson and Boyd argue that cultural evolution can withstand a high degree of error-prone learning.

Of course, stability is important, but so too is variation. Richerson and Boyd allow two mechanisms for the generation of variation: cultural mutation and cultural drift (2005: 69). Cultural mutation is the idea of error already discussed: people misremembering or misunderstanding something, or purposefully doing something in a new way. Content bias may then allow

individuals to compare variants and decide on the best one. The already common variant may be preferred, but where an error leads to some perceived benefit, that may be adopted instead. Cultural drift is analogous to genetic drift and is the result of sampling errors. While Richerson and Boyd provide mechanisms that prevent error propagating throughout the population, these mechanisms are only thought to be most effective in larger groups, in smaller groups cultural drift will play a more prominent role. If only a small number of people in a group possess a skill (boat building in Richerson and Boyd's example), and should they all die young, the skill may die out with them. The larger the population, the smaller the role drift will play.

Richerson and Boyd believe these mechanisms ensure cultures have the requisite properties required for evolution through selection. Cultural variants can persist long enough to be visible to selection, there can be competition between variants, and mutations can arise. I will turn to some criticisms of this approach, but argue that developmental systems theory can take much from this approach.

2.3 Culture, Genes and Development

Cultural evolution, on this account, does not float completely free from genetic evolution in two ways. First, the biases discussed above (for conformity or prestige) are thought to be grounded in evolved cognitive mechanisms so that genetic evolution influences cultural evolution. Second, cultural evolution is thought to impact on the direction of genetic evolution. Richerson and Boyd are presenting here a gene-culture co-evolution theory similar to Lewontin's niche constructionism whereby when we alter our environments, we alter the selection pressures we are subject to.⁷¹ The development of dairy farming and the subsequent evolution of lactose tolerance in some human populations is the best known example of this. Richerson and Boyd also suggest that technologies invented by our ancestors affected the evolution of our morphology. Our hominid ancestors are thought to have been physically stronger and more robust than modern humans. Effective projectile weapons are suggested to have played a role in selection for less robust, but cheaper, physiques. Similarly, they argue that our

⁷¹ See chapter two for a more detailed discussion of this.

vocal tracts and auditory systems evolved in response to proto-languages, which allowed us to develop more nuanced, sophisticated languages, which in turn provided more selection pressure on our vocal tracts and auditory systems (2005: 193).

The evolution of cognitive mechanisms that leads to the kinds of biases Richerson and Boyd invoke in their theory are more problematic, in that these mechanisms are assumed to be of the same sort as are postulated by Evolutionary Psychology, and so are susceptible to many of the same objections as face Evolutionary Psychology and neo-Darwinian evolutionary biology.⁷² Key to Richerson and Boyd's approach is a dichotomy between biology and culture that we can also put pressure on. Developmental systems theory allows us to view recurrent environmental structures as part of the developmental resources of an organism. Culture, or aspects of a culture, can form structures that feature as developmental resources for humans. The developmental systems approach, unlike gene-culture co-evolutionary accounts, does not envisage a dichotomous inheritance system – culture forming one inheritance channel and genes the other. Richerson and Boyd's account conceives the biological and the cultural as distinct domains, though domains that may interact. This allows Richerson and Boyd to assume certain psychological mechanisms (modules for conformity bias, for prestige bias) are already given, with culture as inputs to, and output from, pre-existing structures. But such a picture is misleading. Ingold argues:

... behavioral dispositions are neither preconstituted genetically nor simply down-loaded onto the passively receptive individual from a superior source in society, but are rather formed in and through a process of ontogenetic development within a specific environmental context. (2001: 257)

Psychological mechanisms do not appear first, only then to be given inputs from the environment. The developmental picture assumed by Richerson and Boyd's account, whereby we have specific modules that tend us towards imitating

⁷² See chapter five for a discussion for why we should doubt the evolutionary reasons given for the appearance of such mechanisms, and chapter seven for reasons to doubt the picture of cognition implied by such an approach.

common behaviours, or behaviours of the successful, is as misleading as that in Evolutionary Psychology and neo-Darwinism more generally. Psychological architecture develops in a culture:

The notion that culture is transmissible from one generation to the next as a corpus of knowledge, independently of its application in the real world, is untenable for the simple reason that it rests on the impossible precondition of a ready-made cognitive architecture. In fact, I maintain, nothing is really transmitted at all. The growth of knowledge in the life history of a person is a result not of information transmission but of guided rediscovery, where what each generation contributes to the next are not rules and representations for the production of appropriate behavior but the specific conditions of development under which successors, growing up in a social world, can build up their own aptitudes and dispositions. (Ingold 2001: 272)

In a world where certain environmental structures are ubiquitous and available as developmental resources, members of the population may experience similar developmental trajectories. As a result, it is possible to accept that the sort of biasing phenomena Richerson and Boyd discuss do in fact occur, but to locate the causes of these phenomena elsewhere.

For example, conformity bias suggests that we imitate the most common forms of behaviour in the population. A young beaver will often be born into a group which has built, and maintains, a dam and lodge. This will structure the behaviour of the beaver in many ways. As they learn how to gather food, they will tend to use the artificial lake created by the dam to transport the food home. They will tend to use the lake to get closer to the food and thus avoid predation. They will behave like other beavers because they are born into environments in which behaving this way is an easily available course of action. Similarly, if pens and paper are readily at hand because people commonly use them to, say, solve complex mathematical problems, new generations may continue to use them to solve complex mathematical problems. Thus a behaviour – solving complex mathematical problems – is reproduced in the population. Once a particular behaviour that is constituted by aspects of the environment is already common,

the environmental structures in question will be readily available to new generations and thus increase the likelihood of the cognitive process or behaviour persisting.

Prestige bias may partially be explained in a similar way. That is, prestige bias may – in some cases at least – be a special case of conformity bias. Richerson and Boyd might explain a fashion for dressing like a particular celebrity as an example of people attempting to imitate already successful people. However, this need not be the case. For example, if a celebrity adopts a certain style of clothing, high street stores soon offer clothing of a similar type. Indeed, many high street stores only offer clothing of this type. To buy new clothes, we must buy clothing that resembles the clothing of the celebrity. The environment becomes structured in such a way that we may end up dressing in a similar manner to the celebrity, though without any specific intention to do so. The success of the celebrity explains the ubiquity of the clothing in shops, but it is not (necessarily) the case that a wish to imitate the celebrity explains why in fact people may end up doing so. Rather, it is a fact about the structure of the environment, and why it has been structured that way. Successful people may often have more opportunity to structure the environment in ways that suit them, and so, as with conformity bias, the ubiquity of the environmental structure may explain (at least some of the time) why there can be a trend to behave in ways similar to successful individuals. In this way, we can accept Richerson and Boyd's idea that cultural evolution is possible, without subscribing to the idea that we have evolved specific modules for conformist and prestige biases. The alternative I have sketched here focuses on how environmental structures can shape behaviours, but says very little about our psychology. This, however, will be the topic of the next chapter.

Concerns raised by Sterelny (2001) about the possibility of selection in circumstances where inheritance is not largely horizontal were aimed as much at theories of cultural evolution as DST. In chapter four I discussed some reasons why this did not have to be a problem for DST. Similarly, it need not be a problem here. The worry with respect to cultural evolution stems from considerations of the intergenerational stability of traits. It is generally assumed within neo-Darwinian evolutionary biology that for evolution by natural selection, offspring should more closely resemble their parents than strangers in

order that traits can be visible to selection and acted upon over numerous generations. With cultural evolution it is envisaged, not that children will necessarily resemble their parents, but that individuals in one generation will tend to resemble individuals in the next such that cultural practices persist over time. But because horizontal transmission is possible, new variants could arise and spread throughout a population in less time than it takes for one generation to replace another.

Richerson and Boyd's (2005) psychological biases, and the idea developed here concerning behaviours structured by ubiquitous environmental structures, provide mechanisms which may act to suppress some cultural variation and preserve stability. Although neither of these accounts assumes 100% vertical transmission, or anything close to this, so long as there are mechanisms which preserve some variants over others, intergenerational stability of variants can be preserved. More broadly, developmental resources may be inherited from a wide range of places, but this need not lead to the instability that undermines the possibility of selection so long as mechanisms exist that recreate or maintain the relevant environmental structures. Some individuals may make errors, but viewed from a population-level perspective such stability-preserving mechanisms may prevent such errors from propagating throughout the population. A conformity bias, created by the ability to use a particular aspect of the environment as a resource for development, combined with the ubiquity of that aspect of the environment, will tend individuals towards recreating behaviour or developmental trajectories already common in the population.

This suggests that evolutionary developmental systems can include aspects of culture among the inherited developmental resources for the developing system. Further, developmental systems will recreate many of those developmental resources, including the cultural resources, so that they are available for future generations (as well as peers).

3. The Limits of Theories of Cultural Evolution

The preceding arguments have been aimed at establishing the possibility of evolution by natural selection for developmental systems partially composed of resources in the cultural environment. This is not to say that all aspects of culture ought to, or even can, be understood in this way. Indeed we might still maintain

that, even with some mechanisms that promote stability, a good deal of culture is too transient to be visible to natural selection.

Dupré (2002) notes that cultural variants may be less persistent in modern societies in which there exists a good deal of interaction between different cultural traditions, and that any useful cultural taxonomy may break down in the modern world. This merging of cultures may lead to cultural change at a rate too rapid for selection processes to operate on. The study of evolution is, among other things, the study of the diversity of species. Theories of cultural evolution need to identify analogous taxonomic entities, and Dupré calls these entities “cultural species.” He argues that in the modern world such a taxonomy breaks down because people may “be divided by economic class, ethnic background, religious belief, geographic region, and no doubt many other factors, all of which more or less cross-classify the population” (2002: 145). Against this, however, are the case studies of Richerson and Boyd (2005) concerning Anabaptist communities in America and Canada, as well as farming communities descended from German-Catholic immigrants in the American Midwest. Such communities are, by and large, maintaining their distinctive cultures. This, it is argued by Richerson and Boyd, is due to the aspects of their cultural practices which act as mechanisms to preserve their identity. For instance, the Anabaptist practice in the United States of parochial schooling, as well as lack of exposure to television and a good deal of other modern technologies, means children inherit cultural practices only from other Anabaptists. Particularly striking is the tradition of allowing young people between the ages of sixteen and their early twenties (their “rumspringa”) to take the opportunity to live like other non-Anabaptists of the same age. During this time, the young people have a relatively free hand from their parents and community elders to behave as they want. However, if after this they submit to adult baptism, they will no longer be granted these freedoms. Serious social repercussions befall those who deviate from the community norms after baptism. This period of living out in the world, followed by a voluntary baptism and a stricter code of conduct to live by, has the effect of weeding out “outlaws” and ensuring those that return preserve and reinforce the Anabaptist cultural practices. So it is possible for distinctive cultural groups to maintain their

identity in the face of strong modernizing forces.⁷³ That said, such groups are undoubtedly atypical and Dupré's point that useful taxonomic distinctions will be difficult to make in the modern world seems right (2002: 144-147). This should not be taken to imply that cultural evolution, or the idea of a "cultural species," is irrelevant however: "historical significance is hardly to be demeaned when the underlying topic is evolution" (Dupré 2002: 147). Further:

... if contemporary culture is seen as historically resulting from the gradual hybridization of many earlier cultural species, it should be clear that this only emphasizes the importance of culturally transmitted properties. (Dupré 2002: 145)

A different sort of criticism of theories of cultural evolution has been developed by Sober (1991). He argues that models of cultural evolution tell us about the consequences of cultural inheritance systems and differential fitness, but not about the most interesting aspects of culture – the sources of cultural variants and what makes some successful while others fail. Some of the discussion above dealt with one explanation for the differential fitness of cultural variants in terms of conformity biases and ubiquitous environmental structures. However, this is hardly exhaustive (it does not, for example, explain how the environmental structure became ubiquitous in the first place). While some practices may be explained in the sorts of ways outlined above, many will not be. A theory of cultural evolution will not tell us everything about culture, or behaviour. Neither will theories of cultural evolution always be the most appropriate explanation for cultural phenomena, even in the domains in which such theories have some purchase. Approaches such as Richerson and Boyd's (2005) do not claim otherwise. They instead argue for a methodological pluralism; the richer narratives of historical analysis, say, can be "complementary, not competing" with the simpler, more abstract models of cultural evolution they propose:

⁷³ Note that in these examples cited by Richerson and Boyd, nothing hinges on evolved psychological modules guiding the preferences of individuals. Rather, it is the structures of the environments these individuals inhabit that preserves the cultural practices.

... such explanatory models are not laws but tools to be taken up or not as the situation warrants. Good models are like good tools, they are known to do a certain job reasonably well... (2005: 95)

We might add to this Oyama's advice that any analysis “should be conducted in the interests of the eventual synthesis of a complex, multi-levelled reality” (1998: 420). We ought not to forget that analysis at one level, using one set of tools, will at best give us an incomplete picture. Understanding how these different analyses interact and may (or may not) fit together is an important task, and any individual approach which tends to obscure or deny the need for multiple approaches to such complex phenomena will ultimately hinder useful research.

Accepting that aspects of our cultural lives cannot be given evolutionary explanations, does not entail that culture and biology are distinct. Much of our individual “biological” development is to be explained primarily with reference to the unique set of events that characterise an individual life – a scar, muscle mass, and so on. That aspects of cultures cannot be best understood in evolutionary terms does not mean culture is distinct from biology, any more than failing to have an evolutionary explanation for the particulars of individual development separates development from biology. Aspects of our cultural environments can constitute developmental systems even when they do not constitute evolutionary developmental systems.⁷⁴ Not everything about culture needs to be understood in evolutionary terms in order to view culture as part of the developmental system.

3.1 The Looping Effects of Human Kinds

Although I have spent some time dealing with how stability in phenotypes can come about, I want to now deal with a particular source of variation, what Hacking calls the looping effects of human kinds. This source of variation is important because it raises a fundamental difficulty for studying human evolution if such studies are conceived as being continuous with the study of the evolution of non-human animals.

⁷⁴ See chapter four for discussion of this distinction between developmental systems and evolutionary developmental systems.

Hacking (1996) suggests that studying humans – specifically, human kinds – is a very different activity to studying natural kinds. Hacking lists four essential criteria for a classification to be a human kind. The first is that this kind must be relevant to at least some of us. Second, the kind must be a way of sorting people, their behaviour and their actions. Third, we must want to gain knowledge about this kind. Finally, human behaviour and action is relevant only to the extent it is projected to form a type of person (child abuse is projected to form the idea of the child abuser, and so on); it is kinds of people that are important (1996: 354). Human kinds are different to natural kinds in at least one important way. This difference lies in the fact that human kinds result in looping effects, whereas natural kinds do not. Looping effects occur when a group of people are classified as a human kind, and their awareness of this classification alters their behaviour in some way. Because their behaviour has changed, they no longer behave in accordance with the criteria for membership of this kind. This results in the criteria for kind membership changing in order to track the group of individuals in question, and this can result in these individuals further altering their behaviour, and so on. Thus a feedback loop is created between the group of people being classified, and the classification criteria.

Human kinds (alcoholics, teenage mothers, child abusers) carry with them some judgment. For instance, the kind “alcoholic” might be thought to be a bad thing (morally, medically, etc.). Knowledge about these groups will allow us to intervene and stop behaviour we dislike, and also prevent people behaving in this way in the future. If the human kind is one we approve of, knowledge about it might help others become part of this group. In the case of the kind “alcoholic,” anyone finding themselves as a member of this kind may want to modify their behaviour in some way so as to avoid the negative judgement of others. Alternatively, or in addition, the individual classified as an alcoholic may not wish to be part of this kind. This may mean taking steps to hide the addiction, or it may mean seeking help to overcome it. Self-help groups such as Alcoholics Anonymous are one type of institution that helps define human kinds. These groups encourage an understanding of alcoholism as something that can be controlled with will power. The medical profession – another institution which helps define human kinds – has, at least in some quarters, taken to understanding alcoholism as a medical problem, and treated it likewise. The medical approach has claimed to have identified a group of people who were alcoholics but who went on to become moderate social drinkers – that is, they are *recovered* alcoholics. On the other hand, Alcoholics Anonymous claim that an alcoholic is never cured and so must avoid consuming alcohol at all times. At best they are always *recovering* alcoholics (Hacking 1996: 373). Different classifications have different effects on the behaviour of those classified and have created two different kinds of alcoholics as a result of looping effects. Those

individuals who followed the AA route may never again drink alcohol, and attend AA meetings for very many years. Recovered alcoholics, on the other hand, may indulge in occasional alcoholic drinks, and have no significant contact with the medical community in relation to alcohol consumption after they are “cured.”

Not all groups of individuals classified as a particular kind accept such classifications:

Classifications can change our evaluations of our personal worth, of the moral kind of person we are. Sometimes this means that people passively accept what experts say about them, and see themselves in that light. But feedback can direct itself in many ways... A classification imposed from above is rearranged by the people to whom it is supposed to apply. (Hacking 1999: 131)

Here Hacking notes the gay pride movement as an example of a group of people that actively took charge of their own categorisation. “Coming out” and more generally participating in the gay pride movement became a way to reject the moral and medical norms that had previously been associated with this kind. The gay liberation movement embraced their classification as “homosexual,” but in doing so, changed what it meant to be gay (Hacking 1996: 381). How people will respond to being classified is unpredictable, but once they are aware of being so classified, there will be a response of some sort. And it is this awareness and subsequent behaviour that changes those criteria for kind membership.

3.2 Obesity

Looping effects such as those discussed above demonstrate the role of classification by social institutions on behaviour. I want to focus here on the obese as a human kind. The looping effects generated by this human kind demonstrate how awareness of classification not only alters behaviour, but the entire developmental system.

The category “the obese” meets Hacking’s criteria for human kinds. First, as a kind, the obese and obesity is relevant to at least some of us. Discussion of how to put a halt to the perceived rise in rates of obesity dominates not only academic and medical research, but television, radio and newspapers all carry regular stories on new drugs and the latest public health policy, as well as Byzantine advice on what we should eat. Second, it is a way of sorting people. There are a variety of measures, but perhaps the most popular is the BMI scale, and these measures sort people in to one of four groups: obese, overweight, ideal, and underweight. Third, this is a group of people we want to know more about,

as evidenced by the large number of academic journals devoted to the issue (e.g. *International Journal of Obesity, Obesity, Journal of Diabetes, Obesity and Metabolism*) as well as more general medical journals such as the *British Medical Journal* and the *New England Journal of Medicine*, which regularly feature articles on the topic. Finally, “obese” is not merely a description of a phenotype, but is projected to form the idea of a person. Research suggests that obese people are often thought to be lazy, dirty, ugly, overly emotional, asexual, sloppy and weak (Young & Powell 1985: 234). Teachers may have lower opinions of their obese students (Dwyer et al. 1970: 276). This attitude is also apparent in the medical community. For instance, Young and Powell discuss a study in which doctors professed to not wanting to advise and/or treat obese patients because “the physicians viewed obesity as an indicator of several undesirable qualities, including lack of control” (1985: 235). In their own research Young and Powell assessed the attitude of mental health workers to obese individuals. They presented clinicians with case histories and one of three possible photos. In one photo, the person was in the “ideal” weight range, in the next, the same person’s image was distorted to make her appear overweight, and in the last, she is made to appear obese. They found that the clinicians they questioned were willing to treat individuals regardless of their weight. However, evaluations of the case histories that accompanied the picture of the obese women were far more negative than the same case history with the slim and overweight versions of the woman. Some of the symptoms more likely to be assigned to the obese woman included agitation, emotional behaviour, impaired judgement, inadequate hygiene, inappropriate behaviour, obsessive compulsive behaviour, and self-injurious behaviour (Young & Powell 1985: 238). The obese appear to meet Hacking’s criteria for a human kind.

The next step in Hacking’s argument requires that the group categorised is both aware of, and reacts to, this classification. I want to argue here that being classified as “obese” may be leading those individuals to develop bad health, thus the idea that to be obese is to be unhealthy may be a self-fulfilling prophecy as a result of these looping effects.

There is no straight-forward correlation between health and weight (Gard & Wright 2005; Ernsberger & Koletsky 1999). Thin or “ideal” does not equate to good health but neither does fat automatically equate to bad health. The risks

associated with various measurements of weight (BMI being a favourite) are not easily generalised beyond the cohort studied. Epidemiological studies which have attempted to discover a correlation between weight and mortality and morbidity have had varying successes. Some studies show a correlation, others show little, while some even show an inverse correlation. The accuracy of predictions based on BMI has been argued to be a function of age, sex, and ethnicity (Gard & Wright 2005: 93-94; Kaplan 2000: 139).

The issues discussed above, as well as aesthetic judgements made about fatness in mainstream Western media, will often be well known to the obese. It is unsurprising then that those classified as obese might wish to lose weight. Calorie controlled diets are one way individuals attempt to control or reduce their weight. However, calorie controlled diets are increasingly being recognised as ineffective for anything more than short term weight loss: “The desired permanent solution to the problem of long-term weight maintenance... still seems far off” (Jeffery et al. 2000: 14). What this means is that many individuals engage in weight cycling (sometimes referred to as yo-yo dieting). This occurs where an individual loses weight only to regain it again, and so is forced to try and lose weight once more, followed by the inevitable regain, and so on. Recent research has suggested that weight cycling may lead to an increased risk of health problems (e.g. Ernsberger & Koletsky 1999; Berg 1999; Kassirer & Angell 1998). Problems thought to arise from weight cycling include elevated blood pressure, a reduction in the level of high density lipoprotein cholesterol (the “good” cholesterol), a reduction in the body’s reserves of omega-3 fatty acids, and an increased risk for gall bladder disease, kidney cancer, breast cancer and cardiovascular disease (Gaesser 2003). It has been suggested that those studies which suggest some degree of correlation between obesity and mortality and morbidity do so because those in the sample group who are overweight and obese are those most likely to embark upon weight cycling behaviour (Ernsberger & Koletsky 1999). For instance, in a study of young American nurses, there was a definite risk associated with obesity. However, Ernsberger and Koletsky have suggested, plausibly, that this group is very likely to engage in weight loss practices (1999: 224). Further, those in this group who are in the overweight or obese range are more likely to adopt a diet than those who are in the ideal range. Given this then, and the risks thought to be associated with

weight cycling, it seems unsurprising that obesity should be linked with increased mortality and morbidity for this group. Obesity, in some of these cases at least, will not be the direct cause of the associated risk – though of course being classified as “overweight” or “obese” may be why this group diet. Ernsberger and Koletsky (1999) compared the findings of a large number of studies, including the young nurse study mentioned above, with the likelihood of weight cycling being a common phenomenon within the group under scrutiny. From their analysis, there appears to be an inverse relation between the degree of risk associated with obesity and the prevalence of weight cycling. Those studies which looked at people who are very unlikely to diet find either no risk associated with excess body fat, or find an inverse relationship between body fat and risk levels. Groups where dieting was more likely were subject to a much higher risk of ill health. This suggests that, at least some of the time, bad health and obesity may be related to attempts to become thinner rather than being overweight or obese *per se*. That is, bad health in the obese might sometimes be the result of being classified as “obese.”

But although classifying individuals as obese may lead to a greater risk of poor health, this does not mean the original classification will remain unchanged. The ways in which this poor health manifests itself may differ from the ways in which such poor health was supposed to manifest itself given the differences between the actual aetiology and the supposed aetiology of the health problems. Further, the health risks of obesity may be felt more acutely by some groups depending on whether, and to what extent, they engage in weight cycling. The fat acceptance movement, through US groups like the National Association to Advance Fat Acceptance, may also have some effect on the behaviour of this group similar to the gay liberation movement. These factors, and no doubt others, may all lead to changes in this group of individuals. The classification will need to change in order to track this group.

Looping effects will occur any time a subsection of the population is classified as a particular kind and this classification comes with numerous judgements. As a result, that group of people, aware of being scrutinised and judged by others, and perhaps judging themselves, becomes motivated to modify their behaviour, joining Alcoholics Anonymous, enlisting the help of the medical profession, or dieting. This new behaviour may result in reclassification to a better kind, from “alcoholic” to “recovered alcoholic” for instance. It may also result in behaviour that is hoped exemplifies membership of that kind, as exemplified by the gay liberation movement, if

belonging to such a kind is considered a source of pride. In the case of the obese, a presumption that obesity entails poor health may be leading this group, or at least some in this group, to poor health. But whether those classified act so as to distance themselves from the original classification, as in the case of recovered alcoholics and the obese, or whether they embrace the classification but seek to change the judgements associated with it, as in the case of the gay pride movement, this has the effect of altering the original classification. The kind “alcoholic” may now be a group that can be medically treated, while the kind “homosexual” may no longer be considered a psychological illness. When this happens a new body of knowledge must be found, a new human kind may be created (“recovered alcoholics”), and an old one is modified (the alcoholic can now be medically treated).

In the physical and natural sciences, we can make predictions based on our understanding of natural kinds because they tend to behave in predictable ways. Human kinds are unstable and unpredictable, as a result of the awareness of those so classified of our interest in them. The characteristics by which we define a human kind (types of behaviour, physical appearance, etc.) are not stable – the very attempt to study this group may lead to its instability. The result of this is that we cannot treat human kinds like natural kinds – we cannot make law-like predictions about human kinds in the way we can about natural kinds. Human kinds are not natural kinds – not even “messy” natural kinds (Hacking 1996: 362). For this reason, argues Hacking, the social and human sciences cannot be seen as part of the same project as the physical sciences.

3.3 Objections

Cooper (2004) has objected to the idea that there is a meaningful difference between human and natural kinds on the grounds given by Hacking. She claims that feedback is not restricted to human kinds, and that the classification of natural kinds also results in feedback. Cooper discusses the effects of the classification of marijuana as illegal. Because it is illegal to grow marijuana, it tends to be grown in dark places such as wardrobes and attics, and this has caused an alteration in the physical appearance of marijuana. This, claims Cooper, is an example of our classificatory practices altering a natural kind (2004: 78). While I do not wish to deny that this classification has altered the world in some way here, it is not because marijuana is classified as marijuana that this happened, but rather because it is classified as *illegal*. Classifying marijuana as illegal may have altered something about marijuana, but this differs from the phenomena discussed by Hacking in a number of ways. First, it is marijuana’s classification as “illegal” rather than as “marijuana” that has brought about this change. That is, it is a legal kind rather than a natural kind that has had this effect. Second, the changes that have occurred in the appearance of marijuana have not resulted in a change in the definition of either the natural or the legal kind. That is, there has been no feedback between the classification and the behaviour of the thing classified. It is these looping effects that are important – it is these that cause the instability of human kinds. Cooper has not demonstrated that natural kinds exhibit a similar instability.

Ereshefsky (2004) has raised a similar objection to Hacking's distinction between human kinds and natural kinds, specifically other biological kinds in this instance. He suggests that social

structures can affect non-human animals too: “if a monkey is classified by his troop as dominant, then he is permitted to engage in certain activities” (2004: 915). We can acknowledge Ereshefsky's point that social structures may affect the lives of non-human animals without undermining Hacking's argument. It may be that the classificatory practices of the troop do alter the behaviour of the monkey. What is not clear is whether the monkey's new behaviour alters the classificatory practices of the troop. If the monkey should fail to live up to his dominant status, it is questionable (and perhaps unanswerable) whether the rest of the troop would alter their concept of “dominant” in response. What is vital for Hacking's argument in this context is not the initial effects of classification, but how these effects feed back to alter the classification system. The point is not, ultimately, one about causal relationships in the world (though it does depend on them), but rather about classification and about what appears to be a genuine difference between the natural and the human sciences. Looping effects do not appear to occur when we study parts of the non-human world, but do seem to feature in our attempts to study humans.

4. Looping Effect, DST and Theories of Cultural Evolution

A theory of cultural evolution may provide an account of trends in cultural practices, and offer insights into why some cultural variants persist while others do not. The kinds of mathematical models that Richerson and Boyd have devised to test many of their theories often generate unexpected results, such as the idea that we do not need primarily vertical transmission for selection. Population thinking picks up on trends that can emerge that might not always be discernible at the level of individuals. Richerson and Boyd's approach to cultural evolution, just like population genetics, say very little about the traits being tracked, but unlike population genetics they provide some explanation for why some cultures might die out, while others flourish based on considerations of population structure. It does not offer any explanation for errors. Mutations in genetics can be assumed to be random, but cultural mutations will not always be. Some mutations will be intended; variation can be guided and perhaps explained. However, providing explanations for this will not be within the ability of theories of cultural evolution; cultural evolution is at best a partial explanation of cultural trends.

This sort of approach may seem more powerful if we assume the sharp distinction between biology and culture inherent in Richerson and Boyd's account. The gene-culture co-evolutionary account allows that culture can change the selective landscape for genes, but does not allow for any evolutionarily significant interaction between culture and biology in any

individual lifetime. On such an account, the details of development can be skipped over. Development is easy to black box when it only involves the realisation of coded instructions in the genotype and adds little of evolutionary importance to the final phenotypic outcome. Culture only interacts with fully-formed cognitive architecture in these sorts of accounts, acting as an input into cognitive mechanisms, but never acting as a resource for their development. And although culture may affect genetic evolution, genetic evolution is slow, and mostly what will be tracked is cultural change. Culture and biology are thus kept as distinct, occasionally interacting, domains.

A more robust version of development, however, undermines the idea that we can understand cultural evolution as floating free of biology in any individual's life time:

What this means, in general terms, is that the forms and capacities of humans and other organisms are attributable, in the final analysis, not to genetic inheritance but to the generative potentials of the developmental system... that is, the entire system of relations constituted by the presence of the organism in a particular environment. (Ingold 2001: 261)

A developmental systems perspective can offer a genuinely developmental account of evolution, and can account for some of the trends Richerson and Boyd picked up on with their account. And though Richerson and Boyd's account does not do justice to the complexity of development, they have developed a model for evolution by selection that allows for horizontal transmission that could be of great use to a developmental systems perspective.

The picture of culture that emerges from this perspective is quite different. Rather than understanding cultural inheritance as the transmission of information, as do Richerson and Boyd (2005), cultural inheritance involves each new generation inheriting the resources to recreate that culture:

The growth of knowledge in the life history of a person is a result not of information transmission but of guided rediscovery, where what each generation contributes to the next are not rules and representations for the production of appropriate behavior but the specific conditions of

development under which successors, growing up in a social world, can build up their own aptitudes and dispositions. (Ingold 2001: 272)

The looping effects created by the human kind “the obese” demonstrate the problems with taking biology as given, and culture as something that is added on top. Neither “biology” nor “culture” is given. Rather both develop anew in each generation as parts of a population of developmental systems.

What this means for those studying the evolution and development of humans, particularly where the focus is behaviour and culture, is that the kinds being tracked change. Of course, evolution is the study of change, but in traditional studies, the change is not considered to arise from the activity of studying the evolutionary trends. For human kinds, on the other hand, change may not just be the result of evolutionary forces distinct from the researchers' activities, but may be in virtue of being studied. An example of this may be found in Morton et al. (2006). They demonstrate the effects of research into supposed gender differences. Participants in their study were shown a contrived article detailing research into brain function. The researchers generated four different articles, but any individual participant saw only one. The articles reached conclusions that flattered either male or female cognitive abilities, and were presented as either a piece of stereotypical neuroscience research (complete with pictures of MRI scans) or as a piece of stereotypical social science research. The researchers found – perhaps unsurprisingly – that participants tended to support the continuation of the research funding where the results flattered the participant's gender. Further, participants were more likely to accept the conclusions of the research when it was presented as a stereotypical piece of neuroscientific research. This suggests that, for instance, when Evolutionary Psychologists discuss gender differences, they may not be mapping out species-wide gender differences produced by natural selection, rather they may be, in part, creating these differences as a result of the sort of looping effects discussed above.⁷⁵

If we understand culture and biology as combining together to create developmental systems, the changes to human kinds brought about via the kind of looping effects discussed here will not just affect behaviour and culture, but anatomy and physiology too, as illustrated in the obesity example. This poses unique difficulties for those who wish to study humans, and challenges the idea that the study of humans can be made continuous with the natural sciences. These sorts of phenomena also point to the problem with viewing biology and culture as separate spheres. Looping effects do not necessarily manifest themselves just at the level of behaviour or culture, but can be felt throughout the developmental system. Because DST eschews predefined boundaries (beneath/beyond the skin, biology/culture, gene/environment), it can provide a framework for exploring further the complex interactions and feedback loops, even those feedback loops created by DST-inspired research.

⁷⁵ Dupré (2008) makes a similar point.

In this chapter I have primarily discussed biology, behaviour and culture. Where Richerson and Boyd made use of cognitive mechanisms to explain a bias towards reproducing the behaviour of the majority or of successful individuals, I suggested that instead of cognitive mechanisms, at least some of these phenomena could be explained by cultural structures guiding behaviour. But not everything we do can be explained in this way, and some account of cognition is required. In the next chapter, I will discuss a different model of cognition – extended cognition – to that supposed by both Evolutionary Psychologists and Richerson and Boyd.

Chapter Seven

The Evolution of Cognition

1. Introduction

Evolutionary Psychology offers an internalist model of the mind. Cognitive processes are constituted only by structures and events within the brain. Against the internalist model of cognition is an externalist model known as, among other things, the extended mind hypothesis.⁷⁶ Here the idea is that cognitive processes are constituted not just by structures within the brain, but also by action and aspects of the environment. I will outline this alternative view here, and defend it against some recent criticism. Attempts have been made by Rowlands (1999) and Menary (2007) to give an evolutionary explanation for this extended cognition hypothesis. However, I will argue here that these attempts are flawed; Rowlands' argument does not go far enough, and both his and Menary's approaches are wrong-headed in their adoption of an extended phenotype approach to this issue. As argued in chapter two, the extended phenotype approach privileges the internal, whereas the extended mind hypothesis suggests a framework for understanding cognition where no such privileging takes place. Given this, proponents of the extended mind hypothesis undermine their position by seeking to give an evolutionary account of extended cognitive processes in terms of the extended phenotype hypothesis. Instead I will argue that a developmental systems perspective offers an evolutionary account far more sympathetic to the model of cognition entailed by the extended mind hypothesis. Further, developmental systems theory allows us to deal with some worries about the vulnerability of extended cognitive processes compared to internal processes.

2. Cognition

Hurley describes one key distinction that can be drawn between the externalist positions available in the debate about cognition: there are “what”-externalist positions and “how”-externalist positions (Hurley forthcoming). “What”-externalism is an account of mental content that maintains that mental content is

⁷⁶ The extended mind hypothesis has also been referred to as environmentalism (Rowlands 1999) and cognitive integration (Menary 2007).

individuated by the external environment. Although I will say a little more about this position later on, it is not the focus of this chapter. “How”-externalism (also known as enabling externalism and vehicle externalism), the focus of this chapter, is an account of cognitive processes that maintains that such processes are not always constituted exclusively by structures and mechanisms internal to the brain.⁷⁷ The “how”-externalist maintains that we cannot give a full account of the nature of (at least some of) our cognitive processes without including aspects of the environment in our account. While internalist accounts may include reference to the environment, this will not be in terms of the environment *constituting* cognitive processes. Rather, where the internalist makes reference to the environment it is in the sense of the environment providing specific “inputs” for cognitive processes to work on.

The what/how distinction is also referred to as the content/vehicle distinction. The what/how or content/vehicle distinction picks out a genuine difference, but this does not mean that particular externalist accounts deal only with either the content or the vehicle issue; sometimes both are run together. The position one takes on the question of how cognitive abilities are enabled may influence or inform the position one takes on the question of the nature of mental content, and vice versa, and so it may be natural to treat them together in this way. However it is important to keep in mind the distinction between these two positions. Criticism of one, for instance, does not usually translate into criticism of the other. My focus here will primarily be on the “how” variety of externalism.

There are a few things that externalism of this sort is not. It is not the claim that the mind is a “blank slate” or “silly putty” (Pinker 2002), or that the environment does all the work in cognition. Rather it is a claim about how much of the work in cognition can be attributed to the brain alone, and how much must be understood in terms of a brain, a particular environment, and the interaction with that environment through bodily actions. It is also not the claim that all

⁷⁷ The sense of externalism used here differs from that in chapter two. Externalism, as defined there following Godfrey-Smith (1996), is an explanatory strategy that cites states of affairs external to the organism to explain structures or processes of the organism. This approach does not deny that factors internal to the organism might also have a role to play, but the focus is squarely on the external. The “how”-externalism discussed here does not fit with this taxonomy in that it attempts to explain cognition by citing both internal and external factors.

cognitive processes or capacities are necessarily constituted by aspects of the environment and/or bodily actions. The externalist may grant that some cognitive processes may be constituted by internal structures and mechanisms. “How”-externalism maintains that at least some cognitive processes will be constituted by features of the environment (Hurley forthcoming). Externalism of this sort suggests that attempts to explicate the nature of cognitive processes only in terms of structures and mechanisms internal to the brain will, in at least some cases, be inadequate.

2.1 The Extended Mind

Clark and Chalmers (1998) argue that internalist approaches to cognition are mistaken and instead suggest that cognition must be understood as something that is constituted by more than just the brain, including features of the external environment. Their approach falls under the “how”-externalism approach, in that they are discussing how it is cognition is achieved, rather than what it is that cognition is about. They split cognition into three categories: cognitive processes, cognitive states and experiences. Cognitive processes “consist in operations whereby one semantically evaluable state, or group of such states, is transformed into another state or group of states” (Rowlands 2003: 129). That is, cognitive processes are thought to be operations on cognitive states. It is cognitive processes and cognitive states that Clark and Chalmers are interested in for their argument. They are happy to allow that experiences (such as pain) may be best understood internally. I will outline the two thought experiments offered by Clark and Chalmers (1998) and the lessons they seek to draw from them, before dealing with some criticisms of their argument.

2.1.1 Tetris

Clark and Chalmers (1998) begin with a thought-experiment which seeks to test our intuitions about what might constitute cognitive processes. They ask us to imagine three scenarios involving a game much like Tetris. The aim of this game is to fit various shapes into sockets appropriate for those shapes. In many instances, the shapes may need to be rotated to achieve the best fit. In the first scenario, the subject is asked to rotate shapes in order to work out where the shape should be placed relying solely on internal capacities. In the second

scenario, the subject may rotate the shape on the screen by pressing a button. In the final scenario, the subject has been fitted with a neural implant that can perform the rotation. In the first case, it is the brain alone that enables the player to solve the cognitive task of correctly orientating the shape. In the second case, it is the process of physically manipulating the shape via the button that, along with the processes internal to the subject's brain, enables the cognitive task to be completed. In the final case, it is the neural implant, and the processes internal to the subject's brain that enables the task to be completed.

Clark and Chalmers then ask “how much *cognition* is present in these cases?” (1998: 9, italics in original). They argue that, despite obvious differences, there are also some important commonalities between these scenarios. In particular, the different ways the subject has of playing the game are all cognitive and all equally so. If, they argue, we accept that the brain-implant pairing counts as a case of cognition, then we have no good grounds for excluding the case involving physical manipulation. The neural implant does much the same work as pressing a button and having a computer generate images of the shape at various degrees of rotation, in that, like the computer, it reduces the workload for the brain. If the brain-implant coupling can count as a cognitive system, then so too should the brain-body-computer coupling. The conclusion Clark and Chalmers draw from this is that cognition can be dependent on processes that extend beyond the brain to include, for instance, the activity of pressing the button. Pressing the button to rotate the shape is part of what enables cognition in virtue of the fact that it is a constituent of a cognitive process.

Clark and Chalmers use this thought experiment to motivate a parity of reasoning about cognitive processes:

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. (Clark & Chalmers 1998: 8)

2.1.2 Otto's Notebook

Clark and Chalmers push their argument further. Many cognitive processes are unavailable to us in conscious experience; we are simply unaware of many of the things we do. That such things might be extended into the world is one thing, but those aspects of cognition we are aware of do seem, perhaps intuitively, to be located beneath our skin. However, Clark and Chalmers also claim that aspects of cognition, such as memory and belief, are extended into the world.

Clark and Chalmers use another thought experiment to try to demonstrate that external features can partly constitute the processes that result in mental states such as beliefs. Inga and Otto are informed that there is an exhibition in the Museum of Modern Art, and both wish to attend. Inga recalls that the museum is located on 53rd street, and sets off in that direction. Otto however suffers from Alzheimer's disease and cannot recall the location of the museum in the manner of Inga. Instead, Otto has a notebook in which he records any information he thinks he might require as he comes across it, including the location of the Museum of Modern Art. Once Otto checks the notebook, he too believes that the museum is located on 53rd street and sets off in that direction. Clark and Chalmers argue that the role the notebook plays for Otto is analogous to the role played by Inga's biological memory. Otto's interaction with his notebook and Inga's interaction with her biological memory are such that they both come to have an occurrent belief about the location of the museum and subsequently move in the direction they believe the museum to lie.

Although Clark and Chalmers conclude that “a belief is simply not in the head” (Clark & Chalmers 1998: 14), their argument does not demonstrate that beliefs have locations, either internally or extended out into the environment. Rather, this is an argument concerning *how* it is Otto comes to form certain beliefs. Just as certain parts of Inga's brain may feature in an explanation of her ability to locate the museum, so Otto's notebook will feature in any explanation of his ability to locate the museum. And just as those parts of Inga's brain might be thought to somehow encode information that enables Inga to locate the museum, so Otto's notebook can be thought to encode information that enables him to locate the museum. The details of how it is that biological memory actually stores information is not thought to be important here, rather what is important are the functional similarities between Inga's biological memory and

Otto's notebook. Both Inga's and Otto's abilities require physical objects (internal structures and mechanisms alone in Inga's case, or a combination of internal structures and mechanisms, a body with which to interact with the notebook, and the notebook itself for Otto), and require information (encoded in biological memory or encoded in the notebook). Both the objects and the information play a very similar role in the cognitive systems of Inga and Otto, and given this, Clark and Chalmers argue that this allows us to think of Otto's notebook as functionally equivalent to Inga's biological memory.

Where matters get somewhat more complicated is when we consider the non-occurrent beliefs of Otto and Inga. When Inga is not actively thinking about the location of the museum, she is not undergoing an occurrent belief. She has the ability to recall the location of the museum when required to do so. Given this, argue Clark and Chalmers, it would seem odd to deny that Inga has no belief about the location of the museum when she isn't actively contemplating it. So, even when Inga isn't actively thinking about the location of the museum, it would seem she has a standing, or dispositional belief, about the location of the museum. What about Otto? Does he have a dispositional belief about the museum when he is not consulting his notebook about it? Clark and Chalmers think so: "the information in the notebook functions just like the information constituting an ordinary non-occurrent belief; it just happens that this information lies beyond the skin" (1998: 13). More precisely, the notebook contains information that partly constitutes non-occurrent beliefs. The notebook must be in the right sort of relationship with Otto for this to be the case. It must be a constant feature of Otto's life, and he ought to consult it in most situations where he needs to recall something – much the same as Inga's biological memory must be a constant to her in order for it to count as a functioning memory. It also needs to be easily accessed, without any difficulty – again, just as Inga's biological memory needs to be. Otto's notebook must be a trustworthy source of information for it to function as a memory for Otto – again, just as Inga's memory might be considered dysfunctional if she were to always doubt whatever she recalled. Clark and Chalmers argue that if the notebook is a constant feature of Otto's life, if it is easily accessible and trustworthy, then it is genuinely functioning in the same way (or at least in the same relevant ways) as Inga's biological memory. It is the relationship that Otto has with the notebook, as

much as the information it might contain, that allows Clark and Chalmers to conclude that the notebook partly constitutes Otto's dispositional beliefs, or rather, partly constitutes the enabling conditions for dispositional beliefs. Given this, the notebook does not contain non-occurrent beliefs, but rather, the system that consists of Otto and the notebook is one that has dispositional beliefs.

2.1.3 Constituting Cognition

Clark and Chalmers are not arguing that anything cognition depends upon ought to be considered a constituent of cognition. To be considered a constituent of cognition, certain conditions must be met. First of all, Clark and Chalmers utilise a distinction originally made by Kirsh and Maglio (1994) between epistemic and pragmatic action. Whereas pragmatic actions alter the world in order to bring one physically closer to a goal, epistemic actions alter the world in order to enable cognitive processes or make them more efficient. Building a wall might count as a purely pragmatic action, as the wall is the end in itself. Counting on one's fingers will constitute an epistemic action on this distinction. Whatever alterations to the physical world occur because of this action (change in posture, position, etc.), these alterations are not the point, but merely a means to an end – adding numbers in this instance. The same action can have both pragmatic and epistemic aspects. Rotating the shape in the Tetris game is an example of this. The shape must have the correct orientation in order to fit into the available socket. So rotating it helps to ensure the pragmatic goal of putting the shape in the correct place is fulfilled. But the shape is rotated many more times than is strictly necessary to achieve the correct orientation for the available socket, even by expert players of the game (Kirsh & Maglio 1994). Given that speed is of the essence in playing this game, this suggests that rotating the shape is not just a pragmatic action. That is, the players do not rotate the shapes merely so that they will fit into the available sockets. Rather, they rotate the shapes for both pragmatic and epistemic ends. Rotation not only ensures the block has the correct orientation for the available socket, rotation allows the player to establish what the correct orientation ought to be.

Kirsh and Maglio claim that epistemic actions have three effects: they reduce the memory load on the individual, they reduce the number of steps needed to be performed internally, and they reduce the chance of mistakes being

made (1994: 514). Pressing the button to rotate the shape in Tetris alters the world in such a way that it is usually much easier and faster to determine the socket into which the shape should be placed. Epistemic actions are a means to better perform the task at hand, or a means to perform the task at all.

Such epistemic actions, argue Clark and Chalmers, demand “spread of epistemic credit” (1998: 10). There appears to be no reason, on this account, to rule out any process as being part of a larger cognitive process based solely on the fact that it occurs outside of the head. If this sub-process were internal and considered, unproblematically, part of a larger cognitive process, then this sub-process ought to be considered part of a larger cognitive process regardless of its actual location. To claim otherwise is to assume that the internal/external divide is significant, but without independent justification, this assumption is groundless. The onus is then on the internalist to explain why internal processes ought to be given greater weight, or a distinct ontological status.

Clark and Chalmers make a further qualification of what ought to be considered part of any cognitive processes based on a distinction between what they term passive and active externalism. They identify the kind of externalism developed by Putnam (1975) as an example of passive externalism, which stands in contrast to their own active externalism. Passive externalism, they argue, only requires that features of the world play a role in cognition as part of some long causal chain, beginning with that feature of the world and eventually ending in the aspects of cognition of interest. They utilise Putnam's twin earth example to make their point. We are asked to imagine a twin earth which in all its details resembles earth, with the one exception of the fluid that fills its seas and rivers and so on. Rather than this stuff being composed of H₂O, it is composed of a different compound, XYZ. Inhabitants of this planet call it “water,” and it is indistinguishable from our water at the macro-level. My twin-earth counterpart and myself will both utter the phrase “I believe water is wet,” but Putnam argues that the same belief is not expressed in both instances. My belief that water is wet is the product of a long history of engagement with H₂O and so, when I utter this phrase on twin-earth, I am referring to H₂O. My twin-earth counterpart, on the other hand, is referring to XYZ when she utters this phrase. What explains the difference between our beliefs is the difference in our respective histories and environments. Thus, the argument goes, what it is that makes my belief different

from my twin's cannot be adequately accounted for with reference only to states of affairs internal to brains. Instead, explaining something like belief requires an externalist explanation, in this case, one involving a world with H₂O, and another with XYZ. This sort of externalism is “what”-externalism, or content externalism, as opposed to the “how”- or vehicle externalism I have discussed so far. However, this difference does not bear on the point Clark and Chalmers are trying to make here.

Clark and Chalmers argue that this sort of externalism is passive. This is because if we imagine, unbeknownst to me, I was transported to twin earth and was asked if water is wet by a twin earth inhabitant, I would continue to express the belief that water is wet and I would continue (for a while at least) to refer to H₂O, not XYZ when I expressed the belief that water is wet. The change in my environment would not have any immediate effect on my beliefs concerning water. The relationship between those aspects of the environment that must feature in any explanation of mental content are, according to Clark and Chalmers, “distal and historical, at the other end of a lengthy causal chain... Because of their distal nature, they play no role in driving the cognitive process in the here-and-now” (1998: 10). It is a historical relationship between H₂O and myself that explains why I refer to H₂O when I state that water is wet and a change in my current environment does not alter that historical relationship.

The active/passive distinction also need not align with the internal/external distinction. Internal and external states of affairs may play a role in the causal history of a given cognitive process or state, but only at the end of a long causal chain. It may be possible, for instance, to understand certain processes essential for normal cognitive development largely by reference to internal states of affairs, and if so, these processes will be causally related to other cognitive processes available in adulthood. Nonetheless, these early developmental processes may be passive features of cognition by adulthood. It may no longer matter whether they continue to operate, and indeed, they may have only occurred for a short period of time during development.

Active features, on the other hand, play a crucial role in the performance of any cognitive task such that, if they were removed, this would affect the performance of this task, often for the worse. For example, it may be easier to sum large numbers with a pen and paper than without. The extended cognition

hypothesis suggests that the activity of writing in this situation constitutes an epistemic action, and that the pen, paper, and writing are all constituents of the cognitive process. The ability to perform complex calculations arises out of the interaction between the mathematician, the pen and the paper. This cognitive ability can only be explained by crossing the internal/external boundary. If either the pen or paper was removed, or I was impeded in some way from writing, it would immediately impact my ability to perform the calculation. The action of writing, the pen, and the paper all play active roles: “the relevant parts of the world are in the loop, not dangling at the end of a long causal chain... The external features here are just as causally relevant as typical internal features of the brain” (Clark & Chalmers 1998: 11).

So, the two criteria Clark and Chalmers lay out for external constituents of cognition are that, first, any actions that are to count as constituents of cognition must be epistemic actions. Second, any aspects of the environment (and, presumably, any epistemic actions) that are to count as constituents of cognition must be causally involved in cognition in real time, such that their removal would have an immediate impact on cognition.

2.2 Objections

2.2.1 Clark and Chalmers' Objections

Clark and Chalmers anticipate some objections to their argument. Unlike brains, features of the environment are often transient. There are many instances in which I may have to perform some calculation without pen and paper; they will not always be available. On the other hand, the imagined objection goes, wherever we go we always have the full range of our internal resources available to us. Clark and Chalmers offer two sorts of response to this objection. The first is to argue that while pen and paper may count as contingent aspects of the environment, it is not difficult to imagine a technological innovation that enabled devices to become permanently fixed to our bodies. It is also the case, though Clark and Chalmers do not raise this point explicitly in the context of this objection, that at least some aspects of our environment that may feature as constituents of our cognitive processes and states are not contingent in this way. For instance, our own fingers seem quite robustly available to us if we should

want to count on them. A second line to take is to argue against the assumption that our internal cognitive resources are as reliably accessible to us as this objection implies. In the extreme case, brain injury can impact our ability to perform cognitive tasks. More run of the mill scenarios involve sleep and intoxication, both scenarios cutting off or impeding cognitive processes. This suggests that the criterion that a constituent of a cognitive process ought to be constantly available to us is too demanding. The occasional unavailability of an external feature of the world, or its susceptibility to damage, does not, on its own, rule out the possibility that the feature may act as a constituent of a cognitive process. Rather, a constituent of cognition ought, at best, to be reliably present. Clark and Chalmers conclude that this sort of objection to the extended mind hypothesis will not work.

A further point that can be made here is that while pen and paper may not always be available to us, neither is the ability to perform complex calculations. That is, with a sufficiently complex calculation, it will not be possible to complete it without the pen and paper. The ability to perform complex calculations may depend on these external items. Our brains may be more available to us than pen and paper, but the abilities we are trying to explain may not be. If the ability is not always present, then there seems to be no justification for a demand that the resources that underpin the ability to always be present. At the very least, the onus is on the internalist to provide some justification for the claim that what is to count as a cognitive resource ought to be more reliably present than the cognitive process the resource is supposed to constitute.

2.2.2 The Internalist Understanding of Tetris

Clark and Chalmers argue that, of the three scenarios outlined in the Tetris example, no one case is somehow less cognitive than any other. This stands in contrast to an internalist reading of the situation. Very broadly, the internalist maintains that these three different scenarios involve different amounts of cognitive effort. Mentally rotating the shape involves the greatest amount of cognitive work. Pressing the button to rotate the shape on screen only requires cognition to decide to press the button, and to make decisions based on the visual inputs received from the screen. If part of the task is completed by processes occurring beyond the skin (pressing a button to rotate a shape on a screen), the

cognitive workload has been reduced, thus there is less cognition involved. Work has been done beyond the skin in order that this problem is solved, but that work is not aptly titled “cognitive” according to the internalist.

Clark and Chalmers agree with the internalist that offloading onto the environment in this manner may reduce the brain’s workload, but deny that this makes the process any less cognitive. That is, the subject's internal processes may not have to deal with as large a workload. Instead of the burden falling entirely on internal processes, it is spread out between internal and external processes. The workload is distributed across the entire extended process. And contrary to the claims of the internalist, the extended processes that compensate for the reduced load on internal processes are to be considered just as cognitive as the internal processes they replace.

Not all extended cognitive processes will reduce the workload on internal processes. Where a cognitive task is faced that cannot be completed except by utilising external features, there is no equivalent of the first scenario in the Tetris example (mentally rotating the shape). The only options available are to fail to complete the task, or to complete it using an extended cognitive system. In such a situation, there is no internal workload to reduce by utilising the environment in this way.

Further, some instances of extended cognitive processes may not necessarily reduce the workload if by using extended cognitive processes the cognitive task is transformed. An entirely internal approach to dealing with the cognitive task may require quite different processes than the internal sub-processes of an extended process. Take the example of Otto's notebook again. The internal processes underpinning Inga's ability to recall the location of the museum may be quite different to the internal processes Otto calls upon to use his notebook. Although Inga must carry the heavier workload in terms of biological memory storage, it is difficult to adjudicate whether it is Inga or Otto who is required to carry the heavier workload in terms of the internal cognitive processes that underpin their respective abilities to recall the location of the museum; Inga must access biological memory, while Otto must process perceptual information, and so on. It may be reasonable to assume that allowing features of the environment to enable cognition will in fact reduce the workload

on internal resources, but this need not necessarily be the case. I will discuss this point in more depth below.

2.2.3 The Causal-Constitution Fallacy

No one, presumably, would deny that performing calculations with pen and paper plays some role in the ability to deal with cognitive tasks. However, an internalist might argue that while writing, the pen, and the paper all *causally* contribute to cognitive processes, they do not *constitute* them; this is the line taken by Adams and Aizawa (2001), for instance. Just because *A* depends on *B*, this does not entail that *A* is identical with *B*, nor does it entail that *B* is a part of *A*. To claim that because cognition causally depends on aspects of the environment entails cognition is (in part) constituted by those aspects of the environment is to make a mistake: “a process *P* may actively interact with its environment, but this does not mean that *P* extends into its environment” (Adams & Aizawa 2001: 56). Adams and Aizawa use the following example to make this case: The process of blood filtration occurs in the kidneys. This process is causally related to the process of pumping blood in the heart, the size of the blood vessels, and so on. The causal dependence of the process in the kidneys on the processes in the heart “does not even make a prima facie case for the view filtration occurs throughout the circulatory system, rather than the kidney alone” (Adams & Aizawa 2001: 56). In the same way, their argument runs, Clark and Chalmers have conflated the causes of cognition with those things that constitute cognition. The pen, paper and activity of writing are causally related to cognitive processes, but this does not mean that they are constituents of cognitive processes.

The issue here then is to decide what should count as merely a cause of cognition and what should count as a constituent of cognition. What is at stake is how we demarcate cognition. The internalist accuses the extended mind theorist of conflating causation with constitution, and so committing a casual-constitution fallacy (Adams & Aizawa 2001). There are two versions of this argument against the extended mind hypothesis. The weaker version just claims that the extended mind hypothesis conflates constitutive role with causal role, but offers no grounds on which to judge the difference between constitutive and causal role beyond the assumption that the causal/constitutive distinction aligns with the

external/internal distinction. Such a criticism merely begs the question; whether the external/internal distinction marks any significant boundary in terms of cognition is the very issue at stake.⁷⁸ The stronger version of this criticism, the one developed by Aizawa and Adams (2001), offers grounds from which to adjudicate on the causal/constitutive distinction based on their definition of cognition which is independent of the internal/external divide.

Adams and Aizawa believe that the sorts of cases highlighted by Clark and Chalmers as examples of extended cognitive processes are not in fact cognitive. This is not because they cross the internal/external boundary, but because they lack “the mark of the cognitive” (2001: 46). The mark of the cognitive, according to Adams and Aizawa, is non-derived content. Words on a page only acquire their meaning because, according to Adams and Aizawa, we imbue them with this meaning. The content these words contain is derived from us as readers. On the other hand, our mental states are thought to have meaning that is not derived in this way; meaning originates in these mental states:

... the cognitive states in normal cognitive agents do not derive their meanings from conventions or social practices... it is not by anyone's convention that a state in a human brain is part of a person's thought that the cat is on the mat. (Adams & Aiwaza 2001: 48)

Only where we find non-derived content do we find genuine cognition:

... this means that the skull does not constitute a theoretically significant boundary for cognitive science. More specifically, it means that being inside the brain cannot be the mark of the cognitive. This seems to be true and obvious. The bounds of cognition must be found by finding the mark of the cognitive, then seeing what sorts of processes in the world have that mark. Following this method, we see that, as a matter of contingent fact, the cognitive processes we find in the real world all

⁷⁸ A similar objection might be made regarding DST – that is conflates the dependence of the developmental system on features of the environment with constitution of the developmental system by these features of the environment. As with the extended mind hypothesis, to avoid begging the question, the onus is on the critic to explain why the skin boundary should be considered to align with the causal/constitution distinction.

happen to be brain bound. It appears to be just a contingent empirical fact that cognitive processes are not transcorporeal processes. (Adams & Aizawa 2001: 46)

There are a number of problems with this approach. The first is that this analysis appears to apply to cognitive states rather than cognitive processes. Elsewhere Adams and Aizawa state that their view is that “at least some components of cognitive states require some intrinsic content, where states in Otto's notebook, video games, and most mundane tools do not” (2008: 50). To return to Adams and Aizawa's example of blood filtration in the kidneys, the process of blood filtration does not have all the same properties as blood, in either its filtered or non-filtered state. It is an operation or transformation of blood from its non-filtered to its filtered state. Similarly, cognitive processes do not have all of the same properties as the cognitive states they operate on or produce. Cognitive states are semantically evaluable, but the processes that enable such states are not. Thus the derived/non-derived distinction is not applicable to cognitive processes.

Even if we allow that there is some derived/non-derived distinction that differentiates Otto's case from Inga's, we need not admit that this undermines the claim that cognition is occurring in Otto's case. Adams and Aizawa are not applying a fair comparison when they compare words on a page and mental states. They compare mental states with wholly underived content (Inga's memories) with a state whose content is wholly derived (Otto's notebook). But what the externalist is concerned with is not just the writing in the notebook, but the entire process involving Otto and his notebook. For Otto to have the belief about the museum's location, his own internal capacities are also put to use. It is the pairing of Otto and the notebook that constitute the cognitive process, and not the notebook alone. If non-derived content can only be generated by minds, then there is certainly a mind involved. The issue concerns what constitutes the mind – Otto's brain or the pairing of Otto and the notebook? However meaning is generated, we cannot refute the extended mind hypothesis by assuming it is generated internally as this is to beg the question. Perhaps it is the case that the notebook cannot contain non-derived meaning alone, but perhaps non-derived meaning can be realised in the pairing of Otto and his notebook. The cognitive

process that generated Otto's belief would not have existed without the notebook and Otto. And recall that Adams and Aizawa only require that “some components of cognitive states require some intrinsic content” (2008: 50). Even if we grant them that intrinsic or non-derived content can only be generated internally, this does not seem to offer a principled reason to doubt that the Otto-notebook system is a cognitive one. Some components of this system will be internal ones, capable of generating non-derived content, and this is all Adams and Aizawa ask for.

Further, as Hurley notes, there is no uncontroversial definition of cognition: “criteria of the mental or cognitive vary widely (if not wildly) across theorists; it isn't even clear what agreed work such criteria should do” (Hurley forthcoming). She suggests that instead of beginning with a definition of cognition and then using this to decide whether aspects of the external environment should feature as causal or constitutive factors in explanations of how cognitive tasks are performed, we ought to see if good explanations for the performance of cognitive tasks can include just the internal, or whether better explanations can be found by including the external:

The issues between internalism and externalism should be solved bottom up by good scientific practice, not by advance metaphysics: by seeing whether any good psychological explanations are externalist, not by deciding on a criterion of the mental and using it to sort explanations as constitutive or not. (Hurley forthcoming)

Hurley understands the internalist/externalist debate in the philosophy of mind to be one primarily concerned with explanation. The internalist argues that good explanations in psychology will be generated by assuming an internalist model of cognition. The externalist, on the other hand, argues that good explanations in psychology will, at least some of the time, assume an externalist model of the mind. Indeed, Clark and Chalmers suggest that not only is an externalist approach to the completion of a mathematical calculation using pen and paper more adequately explained with their active externalist approach, but that it is also a more simple explanation:

... one could always try to explain my actions in terms of internal processes and a long series of “inputs” and “actions”, but this explanation would be needlessly complex. If an isomorphic process were going on in my head, we would feel no urge to characterize it in this cumbersome way. (Clark & Chalmers 1998: 12)

If good explanation is what motivates this debate, then it perhaps ought to proceed by seeing what explanations can be generated on both the internalist and externalist approach and deciding which among them is best, be they internalist or externalist explanations.

Cognitive processes rely on structures and mechanisms that can be located in the brain, body and environment of the agent. It is only in the interaction of these structures and mechanisms that certain cognitive processes are possible (e.g. Otto and his notebook). In other situations, the interactions between internal features of the brain and external features of the world create alternative cognitive processes to entirely internal processes (e.g. the various Tetris scenarios). This much is uncontroversial. Where the externalist differs from the internalist is to claim that some cognitive processes are constituted by the interaction of these various internal and external components. Various criteria need to be met before any feature of the world should be considered a component of cognition. It ought to be relatively reliably available to the agent. If action is to be considered a constituent, it ought to be an epistemic rather than pragmatic action. Any epistemic action or external feature of the world ought to be implicated in the performance of the cognitive ability to such an extent that removing this feature of the environment will result in an immediate impairment or removal of the cognitive ability. If these criteria are met, then a parity of reasoning suggests we ought to consider these actions or external features components of cognition.

3. The Evolution of Extended Cognitive Processes

Rowlands (1999) and Menary (2007) both offer evolutionary justifications for the idea that cognition might be extended into the world. Rowlands' argument is in terms of a cost-benefit analysis, while Menary adopts Dawkins' extended phenotype hypothesis to provide an evolutionary justification for his

“manipulation hypothesis.” I will first outline Rowlands' account and examine some potential objections to this view. Rowlands' argument faces a number of problems, only some of which are surmountable, but others that are not. Rowlands' cost-benefit approach will be shown to only provide grounds for the evolution of extended cognitive processes in a very narrow set of circumstances, narrower than he anticipates. I will outline other circumstances in which we might expect natural selection to favour extended processes. Thus, I will provide a stronger evolutionary justification for extended cognitive processes.

I will then outline Menary's account and examine what work the extended phenotype hypothesis does for it. Both Rowlands' and Menary's accounts rely on the “manipulation hypothesis,” and this accords well with the extended phenotype hypothesis. However, I will argue that neither the extended phenotype hypothesis nor the manipulation thesis is an appropriate tool for understanding extended cognition.

3.1 The Cost-Benefit Analysis

Rowlands argues that extended cognitive processes are likely to be selectively favoured over highly internalised processes, based on a cost-benefit analysis. This stems from his more general point that any offloading an organism can do onto the environment will tend to increase that organism's relative fitness because it reduces the overall energy expenditure of the skin-bound organism.

Any trait can cost the organism energy resources in two ways. The first is in the energy it takes to develop a structure, mechanism or behaviour – the implementation cost.⁷⁹ The second is in the energy it takes to maintain and use a structure, mechanism or behaviour – the performance cost. These two energy sinks combine to give the total amount of energy expended by the organism in virtue of possessing a given traits. The amount of energy an organism can acquire is dependent on two factors; the first concerns the amount of energy available in the environment, while the second concerns the upper limit on the organism's rate of absorption of this energy. Thus there will always be a finite amount of energy available for the organism to use. What energy is acquired by

⁷⁹ Ascertaining the implementation cost of any trait will present a technical difficulty for Rowlands' cost-benefit analysis.

the organism must then be used to supply the different energy sinks – the various structures, mechanisms and behaviours – that comprise the organism's system. Given the upper limit placed on energy coming into the system, the various energy sinks must compete for limited resources. The less energy required by a given structure, mechanism or behaviour, the greater the amount of energy available for the rest of the organism's system.

Rowlands maintains that offloading onto the environment will tend to reduce the amount of energy required for a given extended process when compared to an internal process and this will result in an increase in the fitness of the organism relying on the extended process. An extended process will rely on features of the environment which will not incur any implementation or performance costs for the organism. An organism which relies on an extended process will only incur the costs of those internal structures or mechanisms which couple with external features of the world. We can imagine two organisms, the first of which performs a given task by relying on an extended process, while the second relies solely on an internal process. In all other respects these organisms are identical. They assimilate the same amount of energy from their environment, and the same amount of competition occurs internally for these energy resources. The organism which does not have to incur the implementation and performance costs of those external constituents of the process will tend to expend less of its overall energy, thus having more energy to apportion to the rest of its system. The organism which must incur the implementation and performance costs associated with all the constituents of the internal process will tend to have less energy to make available to the rest of its system. This means that the organism which utilises extended processes will have more energy resources than the organism that utilises solely internal processes. Thus, the organism which relies on extended processes for the performance of some task will tend to be fitter than an organism which relies only on internal processes. Such considerations lead Rowlands to his “barking dog principle”:

If it is necessary for an organism to be able to perform a given adaptive task T, then it is differentially selectively *disadvantageous* for that organism to develop internal mechanisms sufficient for the performance

of T when it is possible for the organism to perform T by way of a combination of internal mechanisms and manipulation of the external environment. (1999: 80, italics in original)⁸⁰

Given this, we should expect there will tend to be selection for extended processes over exclusively internal processes.

Rowlands anticipates several objections to his argument. Two in particular I will deal with here involve what Rowland terms “hidden costs” and “hidden benefits.” Rowlands responds to both arguments, but I will identify stronger versions of these objections that Rowlands has not considered. I will ultimately argue that although these stronger arguments pose problems for Rowlands' account, they do not undermine the attempt to talk about extended cognitive processes in evolutionary terms. Further, in considering these issues, we can identify a wider set of circumstances in which extended cognitive processes might arise than presumed in Rowlands' account.

3.1.1 The Hidden Costs Argument

Rowlands' imagined critic suggests that although it may look as though extended processes incur fewer costs in terms of energy use, there may be costs we have not considered. Rowlands offers the following example. Hooking my car up to a neighbour's car ensures a free ride; however, there will be implementation costs involved in this. I might have to persuade the neighbour to allow me to do this, or return the favour (or perform a bigger favour) for my neighbour at some later point. I may not come out of this deal any better off than if I had just driven my own car, and I might even come off worse. Extended processes may incur hidden costs and thus, the argument goes, we cannot presume extended processes are less energy-hungry than internal processes.

As Rowlands points out, this argument cuts both ways. It may be the case that there are hidden costs involved in external processes, but this may also be true of internal processes. If we can say anything about the energy costs involved

⁸⁰ Rowlands bases the name for this principle on an inversion of the adage *why keep a dog if you are going to bark yourself?*: “... if you do have a dog, then you do not have to bark yourself. And getting your dog to do your barking for you will save you considerable investment of resources (i.e. energy)” (1999: 79).

in internal processes, we should be able to do the same about extended processes. The task is to establish the costs involved in both internal and extended processes as best as possible: “the possibility of hidden costs cannot be allowed to function as a sort of general and nebulous methodological angst” (Rowlands 1999: 93).

3.1.2 The Hidden Benefits Argument

The hidden benefits argument suggests that while there might be some short term benefit in relying on an extended rather than an internal cognitive process (perhaps in terms of energy savings), in the long term the entirely internal approach may have some extra benefits which ultimately make it a better option. That is, the pay off associated with an internal process is at least enough to compensate for any extra implementation and performance costs. Given this, we should expect entirely internal approaches to have been selected over extended approaches. Rowlands offers an example of this line of thought. Suppose we were faced with the task of having to lift a weight. The extended process might involve asking someone else to lift it, while the internal process might involve developing strong enough muscles to be able to lift the weight without assistance. Let's assume that there are less implementation and performance costs involved in persuading someone else to lift the weight (though such an assumption is not unproblematic). Nonetheless, being able to lift the weight unaided might be more selectively advantageous because there will be a bigger pay off in the long run: you may be able to lift other things, you may be able to fight off predators as a result, and so on. That is, if more adaptive problems can be solved by the internal process, it may actually be more cost effective. If the energy gained compensates for the energy lost in implementation and performance, then the internal solution – developing stronger muscles – will be selectively advantageous.

Rowlands identifies two problems with this line of argument. First, just as with the hidden costs argument, the argument that there may be hidden benefits associated with an internal process is just as applicable to an extended process. Being able to persuade someone to lift a weight is a skill that could generalise to many other tasks, just as being able to lift heavy objects might. Second, Rowlands suggests that this sort of argument misunderstands the nature of natural selection:

Evolution crosses each bridge as it comes to it, and has no conception of the possibility of further bridges. Now, while it might be a good thing for an organism to develop internal structures to perform a given task since these structures might give it the capacity to perform other important tasks also, the fact that it is a good thing cannot possibly be recognized by evolution. (Rowlands 1999: 94)

It is certainly the case that natural selection cannot anticipate solutions to problems organisms might face in the future. In this sense, Rowlands is right to object to hidden benefits arguments where it is assumed the hidden benefits will arise out of situations that are not yet realised in the organism's life and have not typically occurred in that lineage. Natural selection, for instance, could not have equipped a lineage with an adaptation for coping with industrialised cities prior to the development of such cities. However, given the weight-lifting example outlined by Rowlands above, it is not clear this sort of scenario is all he is limiting his claim to. The “hidden benefits” associated with developing the muscle strength to lift the weight unaided do not seem of the same sort to the “hidden benefits” associated with being equipped to deal with crossing busy roads. That is, there are many obvious situations where increased muscle strength would be useful within that organism’s lifetime. It is not clear that the ability to deal with busy roads would be useful to an organism in the absence of busy roads. It is only future generations of this population that would benefit from such an adaptation. But while natural selection cannot equip organisms with adaptations for scenarios that have not yet occurred, it can certainly favour a trait that causes a short term disadvantage to the organism if it ultimately increases its chances of survival and reproduction over its life time. But both internal and extended processes can deliver a short-term disadvantage but a long-term advantage to the organism. That there might be hidden benefits associated with an internal process equally applies to extended processes.

The hidden benefits argument, like the hidden costs argument, seems to come down to an argument about methodology. It suggests that we might overlook certain benefits associated with an internal process because we are focussed only on the role of that process in a given task. Should that process also

play a role in other tasks, our cost-benefit analysis will be too narrowly focussed and so not reflect the actual benefits associated with this process. It is possible for natural selection to select for a trait with a higher implementation and performance cost but which ultimately offers a bigger pay off for the organism because that trait has a wider range of uses.⁸¹ *If* it should turn out that developing stronger muscles confers a greater long-term survival and reproductive advantage to a person than the ability to persuade someone to lift the weight then this is the sort of thing natural selection can act on. Equally, persuading someone to lift the weight may be selected for even if (as is more likely) it requires a greater implementation and maintenance cost if it confers a greater long-term survival and reproductive advantage.

Of course, as with the hidden costs arguments, we ought not to allow this hidden benefits argument to undermine the entire project. If we are to adopt a cost-benefit analysis, we must do our best to identify all costs and benefits over the lifespan of the organism. That we may miss some cost or benefit is always a possibility, but it applies to both internal and extended processes. There is nothing in the methodology of this approach that makes this inevitable, and there is no systematic bias towards either internal or extended processes that will mean our results are hopelessly skewed. Missing something in this way is merely the risk we take with any attempt to find out about the world. The possibility of getting it wrong is not enough to undermine such an endeavour.

3.1.3 The Strong Hidden Benefits Argument

While the weak hidden benefits argument seems to be concerned with an internal process allowing for greater generalisability, and thus conferring more benefits in virtue of allowing the organism to solve more adaptive problems, the strong hidden benefits argument makes a slightly different point. A less costly external process may be selectively advantageous in the short term, but if the environmental feature that acts as a constituent of this process is not reliably available, then selection may eventually favour an internal process in virtue of

⁸¹ This is possible, but there is an upper limit. If the implementation and performance costs, say, are so great at some early point in life that the organism is vulnerable and thus is likely to fail to reproduce, then the long-term benefits need to be very high for such a strategy to be selected for. At some point, the risks will outweigh the benefits.

the fact that it allows the organism to solve the adaptive problem with greater regularity. The selective advantage will be most strongly felt by the internal process, according to this objection, if the adaptive problem is particularly pressing, or if the environmental feature is likely to be absent for periods of time greater than the life span of the individual organism.

Certain conditions would still have to be met for this to happen. In particular, the implementation and performance costs associated with the internal processes would have to cause less of a disadvantage to the organism than the disadvantage felt if the task could not be completed. An internal process would also have to be a possibility (more on this below). But for the moment let's assume that the disadvantage felt by the organism in not being able to solve a particular problem is greater than the costs of developing and maintaining an internal process which allows the problem to be solved.

In such circumstances, and assuming an internal process is developmentally possible and the required variation exists in the population, we should expect the internal process to be selected for, and the extended process to be selected against and even eliminated if the problem to be solved is of enough urgency. So, if beavers had historically spent enough time in environments with no rivers to dam, or only sporadic access to rivers, we might expect selection for alternative ways to solve the problem of avoiding predators and transporting food. And some of these alternatives might rely on structures, and mechanisms that could be characterised as internal. Thus, the argument goes, where environments have fluctuated such that the relevant feature has been absent, we should expect selection to have favoured internal processes.

The strong hidden benefits argument need not be a problem for the idea of extended cognitive processes, or for building an evolutionary account of the emergence of such phenomena. I want to argue here not so much that Rowlands' account is wrong, but that his account does not go far enough. He has set out some of the conditions for the possible evolution of extended cognitive processes, but these are just a small set of the conditions that would allow such processes to evolve. Even if the strong hidden benefits argument might suggest that extended processes are more vulnerable in virtue of their reliance on contingent environmental features (and we can contest even this, see below), and as a result, such processes might not emerge as frequently as Rowlands assumes,

there are far more conditions in which such processes could emerge, and thus we need not assume that extended cognitive processes will be atypical.⁸² A more pressing concern for Rowlands' project of providing an evolutionary justification for the existence of extended cognitive processes will then be explored. I will argue that Lewontin's niche construction approach to evolution suggests that the cost-benefit analysis employed by Rowlands is inappropriate for considering organism-environment interactions and their evolutionary consequences.

3.1.4 Changing Environments

One worry that might be raised here is that extended processes will be more vulnerable than internal processes in virtue of their reliance on features of a transient environment. If internal processes have been even fractionally more reliably available then, the thought goes, we should expect selection to favour internal processes. In the earlier discussion of Clark and Chalmers (1998), I discussed a version of this concern. Neural circuitry might be thought to be reliably available, whereas aspects of the environment might be considered more transient, and this might somehow motivate a distinction between the internal (constituents of cognitive processes) and the external (inputs for cognitive processes). Clark and Chalmers' responded to this objection by demonstrating that internal processes are not necessarily as reliable as we might tend to think. Intoxication and sleep were both used as examples to demonstrate that even entirely internal cognitive processes are not always available to us. And a condition of external features of the world counting as constituents of such extended cognitive processes was that they were very reliably available to us.

Over evolutionary time, however, the worry might be that even very slight differences in the reliability of an internal process over an extended process might be sufficient for the selection of the internal process over the extended one. However, this need not be the case. Stabilising selection can allow a variety of traits to co-exist within a population. Let's imagine a population comprised of two types of organism. One relies on a fully internal but costly process, while the other relies on an extended and cheap process. When the aspect of the environment relevant to the extended process is present, the organism which

⁸² Though note that externalists only need establish that some aspects of cognition are extended to defeat strong internalist arguments.

relies on the extended process is favoured by selection. When this feature of the environment is absent, the organism which relies on the internal process is favoured. Assuming the adaptive problem these processes allow the organism to solve are not so pressing that the organism that cannot solve it is almost certain to die or fail to reproduce (and similarly assuming the costs of the internal process are not so great that the organism that relies on such a process is almost certain to die or fail to reproduce), selection will switch between favouring the extended and internal process such that organisms that rely on the extended process and organisms that rely on the internal process both exist in the population. That is, there would be two distinct adaptations in the population for solving the given adaptive task.

The sort of objection to extended processes that concerns the transience of environmental features assumes that only one external solution is available. There is no justification for such an assumption; at any time, there may be a variety of extended processes (and a variety of internal processes) on offer. Natural selection cannot favour an internal process over an extended one if an internal process is not available. It will also not opt for an internal over an extended process if there is another less costly external process on offer. So if an environmental feature that partly comprises an extended process is unreliably present, another more reliable aspect of the environment may end up partially constituting a different extended process. And if this new extended process is less costly than whatever internal processes exist in the population, then we should expect selection to favour this new extended process instead. The choice is never between internal and extended processes in general, but between actual instantiations of both types of processes where they should happen to appear.

Let's grant for the moment that external processes, because they depend on transient features of the environment, are less likely to be favoured by selection than internal processes. This alone need not undermine the idea that we should expect such extended processes to exist, or to have played a role in evolution. In chapter three I discussed West-Eberhard's account of the evolutionary importance of adaptive plasticity and genetic accommodation. In her account, highly internalised traits are parasitic on traits that depend strongly on the environment. The idea here was that traits that initially depend heavily on the environment can, after a sufficient amount of time and with the right variation in

the population, develop independently of the relevant aspect of the environment. A form of behaviour, for instance, may start out in a population by being learnt. Some general plastic response of the organism, itself perhaps the product of natural selection, will be responsible for this. Variation in the population may mean that some organisms can acquire this behaviour with less exposure to the stimuli than most of their conspecifics and these individuals will be favoured by selection (again, assuming whatever variation is involved does not drain resources such that its cost is greater than the benefit felt by acquiring the behaviour more quickly). Further phenotypic variants may arise whereby even less learning is required, and again, with the usual provisos, selection will favour them. This can continue until the trait no longer needs to be learnt, or only needs the minimum of exposure to the right kind of situation. The initially entirely learnt behaviour in the population can often create the right sort of selective environment for the subsequent variants to be favoured. Take this hypothetical scenario. Some birds learn to remove the tops from bottles of milk in order to drink the milk inside. Now let's imagine that such birds move from a diet that had primarily been based on insects to one which, through learning this behaviour, now becomes one based primarily on milk. The ability to catch insects no longer matters, and perhaps such skills atrophy. In such a situation, there is a far greater selection pressure for variants that can open bottle tops with greater ease than there was when insects formed the main component of the diet. The learnt behaviour changes the selective landscape.

The usual way plastic responses are described is in terms of an internal adaptation that receives the correct sort of input from the environment thus triggering an adaptive response. This is the sort of picture that chapter four on developmental systems theory, as well as the earlier part of this chapter on extended cognitive processes attempted to undermine. Rather than understanding plastic responses as internal adaptations that produce the correct output for a given input, we can understand at least some such plastic responses as extended processes that are constituted by those features of the environment we have been characterising as “inputs.”

The point of all this is as follows: Even if we could establish that internal processes were generally favoured by natural selection over extended ones (so, for instance, if we accepted as the example above presumes, that internal

solutions will be quicker and more effective), this would not entail that there would be no extended processes. Internal processes might only come about because of extended processes and, given this, extended processes would play an important role in driving cognitive evolution. Given how recent, in evolutionary terms, many of what appear to be uniquely human cognitive abilities are, there are good grounds for thinking that many of them may rely on extended processes.

3.1.5 Plasticity

We have assumed so far that the transience of environmental features is a difficulty to be overcome. However, this need not be the case. Where different adaptive problems are faced in different environments, reliance on cyclically fluctuating environmental features can be an advantage. For example, depending on when the meadow vole *Microtus pennsylvanicus* is born, its coat will either be thick or sparse. If the pup is born at the later stages of the summer, it will have a thicker coat, thus providing it with some protection during the winter. If the pup is born in the spring, it will not need such a thick coat to see it through the summer and so can devote resources to other aspects of development. The coat thickness of the vole pup at birth is influenced by the amount of melatonin the pup is exposed to, which is produced by the mother in response to day length (Gluckman et al. 2005). Day length is a variable feature of the environment, but far from this being a problem to be overcome, is in fact very useful. The vole pup does not need to waste energy developing a thicker coat when such a coat is not required. The high amounts of melatonin received from its mother in utero by a vole pup born later in the summer provides the right resources for a thick coat to develop. Extended processes of any kind may rely on features of the environment that fluctuate – in the vole’s case, melatonin produced by its mother – but in some situations, the extended processes may only be required when that feature is present in the environment. As an example, the ability to solve complex mathematical problems may be more pressing in literate, industrialised societies. In other societies, for instance those of our stone age ancestors, the nuances of differential calculus may not only be less pressing, but a distraction from more important concerns. Societies in which one might need to perform some complex calculation are also environments in which one is most likely to find pen and

paper – or the equivalents thereof – in order to enable the completion of such a task.⁸³ Not all extended processes may have this feature of only being required in the same environments as the requisite resources are to be found, but some may. In such circumstances, the changing environment does not pose a problem for extended processes. Rather, extended processes may be more adaptive in changeable environments than internal processes. The implementation and performance costs for a suitable internal process for any given environment may even be less than the implementation and performance costs for a plastic, extended process. However, the organism must provide a number of different internal processes to deal with the different states of the environment, even though some such states will not be instantiated in a given organism's lifetime; the organism will incur unnecessary implementation costs, as well as costs associated with maintenance. These will form another sort of hidden cost. Thus, the fact that certain features of the environment may not be reliably present, far from threatening the idea of extended processes, actually provides us with some support for the idea that extended processes may be selectively advantageous. Later I will look at arguments from developmental systems theory concerning extended inheritance that will also allow us to question whether we need to consider all aspects of the environment transient.

3.1.6 Development

Even in situations where some posited internal process might be more selectively advantageous this does not entail such a process will be selected. Evolution can only work with what actually arises. Should the requisite variation not occur in the population, there can be no selection for it. Phenotypic variation will not just require the right sort of genetic variation, or even the right sort of variation in developmental resources, but also that the developmental pathway of the organism is such that a new structure or mechanism can develop without disrupting other aspects of the developing system. Related to this, an extended

⁸³ Of course, it is perfectly possible – even likely – that the ability to solve complex mathematical problems is not the outcome of natural selection at all. This example is only designed to point out that the importance of cognitive skills (and skills more generally) are very much context-dependent, and this very context might provide the resources required to enable these skills. This is not a necessary relationship. Rather it only needs, as a contingent fact, to have been the case for selection to have favoured a skill based on such an extended process.

and internal process may also differ in the timing of their development. An extended and an internal process may demand the same amount of energy for implementation. However, it may be that such implementation energy is more easily spared at different points in development. If the internal processes require this energy at a point where the developing system cannot provide it, then the structures that underpin the process cannot develop. Of course, this point applies equally to the internal aspects of extended processes; the requisite resources may not be available for the development of these processes when they are required either. We have, of course, no reason to assume that there would be only one internal process that could handle the task, and so just because one process is ruled out by such developmental considerations does not entail all internal processes will be. However, this does raise a problem for Rowlands' strategy of cost-benefit analysis. The straight-forward comparison of implementation and performance costs for extended and internal processes does not take into account developmental considerations such as timing and so cannot provide enough information to decide which of two proposed processes will be favoured by natural selection.

There are further problems for Rowlands' cost-benefit analysis. Recall his barking dog principle:

If it is necessary for an organism to be able to perform a given adaptive task T, then it is differentially selectively *disadvantageous* for that organism to develop internal mechanisms sufficient for the performance of T when it is possible for the organism to perform T by way of a combination of internal mechanisms and manipulation of the external environment. (1999: 80, italics in original)

This principle suggests extended processes will always, or generally, be favoured by selection over internal processes. This is based on Rowlands' argument that extended processes will usually be less costly for the organism than internal processes. However, can we assume that this will in fact usually be the case? Where we can expect this assumption to be warranted is in the following case. Consider two organisms, one of which relies on an entirely internal process, the other on an extended process. There is overlap between the internal and extended

processes such that all the internal components of the extended process are also components of the internal process. The extended process involves the common internal sub-processes plus some external sub-processes, while the internal process involves the common internal sub-processes plus some other internal sub-processes. The organism that relies on the fully internal process then has to find implementation and performance energy for the common set of sub-processes, plus an extra set of sub-processes. The organism with the extended process only needs to find implementation and performance energy for the types of sub-processes the two organisms share in common. Thus, the organism that relies on the extended process incurs fewer energy costs. So, if internal and extended processes differed from one another in this way, we might be justified in assuming, along with Rowlands, that extended processes would be favoured by natural selection on the grounds of demanding fewer implementation and performance costs.

However, we have no reason to assume that an extended process will differ from an internal process in just this way. The internal components of the extended process may have little in common with any of the components of the purely internal process. Rowlands adds another principle – the principle of the non-obvious character of evolved internal mechanisms:

For the performance of a given task T, and for any internal mechanism M which has evolved in organism O and which, when combined with suitable environmental manipulation on the part of O, allows O to perform T, the nature of M is not always obvious on the basis of T.
(1999: 81)

When an organism completes a task by relying on external features of the world, those internal processes involved will only form sub-processes of that larger extended process. Given there is no requirement to assume that the sub-processes must share all or many of the same properties as the larger processes, the sub-processes may look quite different to the larger process. If we identify a task that an organism is required to complete, we may be able to identify the solution to the task when we see it. For example, if an organism has natural predators, a solution to this problem might involve sharp claws to fight off predators. There

are many specific strategies to avoid predation (and a combination of such strategies is both possible and common), but we can be fairly sure, according to Rowlands, that we can identify such solutions, or at least make educated guesses about these strategies, where they arise by merely looking at the organism if such strategies are entirely internal. We cannot say the same about the sub-process or components of a solution. They may look quite different to the actual solution. A cog from a tin-opener does not appear an obvious component of a solution to the problem of opening tinned goods. Rowlands highlights this point with the example of the beaver. The beaver's dam solves, or helps solve, the problem of transporting food and avoiding predators. The dam is difficult for predators to negotiate. The beavers' lodge built in its centre can only be accessed from beneath the water, and the structure is very strong. The beaver's main source of food is tree bark and substances surrounding it. Beavers forage for food, and transport it back to their lodge to consume, in part at least to avoid predation. Because they create lakes behind their dams, they create a long shoreline from which they can forage without spending too much time exposed to predators. Further, when they have found food, they can transport it on the deeper water created by the dam back to the lodge. The dam helps the beaver avoid predators and transport food. The dam can be considered a component of the beaver's extended processes of catching food and avoiding predators. Essential to the beaver's ability to build a dam is their long flat teeth. This serves as a nice example of the non-obvious character of evolved internal mechanisms. The long, flat shape of the beaver's teeth is not an obvious solution to the problem of avoiding predators. We should not expect the internal aspects of extended solutions to be the obvious solutions to the task at hand; we should not expect this because they are not solutions to the task, they are merely components of the solution, and there is no necessity for components to have properties similar to the properties of the larger system. Of course, once we take a step back and look at the organism in its environment, we may find a something that "looks" like a solution to a given adaptive problem.

Consider Clark and Chalmers' Tetris example again. The scenario in which the player pressed a button to rotate the shape on the screen involved internal sub-processes such as those involved in deciding to press the button, and those involved in moving the players' arm and fingers. In considering what a solution

to the problem of fitting shapes into sockets might look like, following Rowlands' principle of the non-obvious character of evolved internal mechanisms, these internal processes do not look like obvious candidates. At least, they do not look like obvious solutions when considered in isolation. When we look at the bigger picture, at the entire extended process, how such sub-processes allow the player to complete the task becomes clearer. The sorts of internal processes at work in this example of extended cognitive process seem, by and large, quite different to the sorts of processes that need to be employed in the case where the player rotates the shape mentally.

If extended processes rely on quite different sub-processes to those utilised by fully internal processes, we cannot make this assumption underlying Rowlands' barking dog principle; we are not comparing like with like. We can accept that if the same (or very similar) sort of process was to be extended, this will reduce energy costs for the organism, but we cannot do so if the internal components of the extended process are not of the same (or of a very similar) sort to some of the components of the entirely internal process. And it is not, I take it, unreasonable to assume that extended processes will often, if not usually, involve set of components that have little overlap with the components that comprise entirely internal processes (see below for a more detailed discussion on this point). Rowlands' barking dog principle appears to be at odds with his principle of the non-obvious character of evolved internal mechanisms. In cases where the entirely internal processes are markedly different than the internal components of extended processes, we cannot make any assumptions about which process will be the most energy hungry. We can establish this fact on a case by case basis, as each process is actually realised in the world, but we cannot *a priori* take extended processes to be cheaper solutions to internal processes.

In summary, Rowlands' cost-benefit analysis may not provide good grounds for expecting the evolution of extended processes, but that does not mean we cannot expect extended processes to have evolved. Indeed, there may be a wider range of circumstances in which extended processes have evolved than those Rowlands' analysis suggests. His account points to extended processes being favoured when they are cheaper, but this need not be the case. It may be that a given extended process required greater implementation energy, but that it required this energy at

a time the organism could spare it, while the internal process required its implementation energy at a point in development when the organism could not spare it. Concerns about the transience of the environment and thus the selective disadvantage to organisms relying on extended processes fail to undermine the case for the evolution of extended processes on three grounds. First, the choice need not be between one internal and one extended process. There may be multiple extended (or internal) processes instantiated in the population. Just because one extended process fails due to environmental changeability does not mean another extended process will not be favoured instead, especially if it is cheaper than any of the internal options. Second, even if an internal process is preferable in a given scenario, this process may only have become selectively advantageous because of the pre-existence of the extended process. Natural selection may only come to favour internal processes because extended processes altered the selective landscape in some way. This would suggest that, at least in some cases, extended processes are primary as drivers of cognitive evolution. Third, there is no reason to assume internal processes will always be favoured when the relevant aspects of the environment are transient. Stabilising selection could allow at least two adaptations, relying on either an entirely internal or extended process to exist in the population. The transience of certain aspects of the environment may even work to the advantage of extended processes, as the discussion of adaptive plasticity demonstrated. All of which is to say that, in considering the circumstances in which extended processes could have come about through evolution, there are more issues to bear in mind than just those that pertain to a straightforward cost-benefit analysis. So even if we cannot always assume that extended processes will be cheaper in the manner of Rowlands, we can consider a far wider range of circumstances in which extended processes might be selected for, and thus we have good grounds to expect some extended processes to arise as a result of natural selection.

3.2 The Manipulation Thesis and the Extended Phenotype Hypothesis

Menary (2007) develops a version of “how”-externalism called “cognitive integration” that, he argues, avoids the pitfalls of Clark and Chalmers' extended mind hypothesis as well as providing more detail of the ways in which cognitive processes can be understood as being constituted by aspects of the environment. In particular, he describes and extends what he calls the manipulation thesis, also evident in Rowlands (1999) and Clark and Chalmers' (1998) work. This

manipulation thesis, argues Menary, can be given an evolutionary justification by incorporating Dawkins' extended phenotype hypothesis. I will argue here that in adopting the manipulation thesis, Menary (and, by extension, Rowlands) have fallen foul of the very problem Menary identifies with Clark and Chalmers' account of extended cognition. Further, this demonstrates the unsuitability of attempting to understand the evolution of extended cognitive processes in terms of the extended phenotype hypothesis.

3.2.1 The Parity Principle

Menary's criticisms of Clark and Chalmers' account of extended cognition centres around what has come to be known as their “parity principle”:

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. (Clark & Chalmers 1998: 8)

Menary (2007: 55-59) notes that critics have taken this to imply that extended processes are of the same or very similar type as internal processes. Menary quotes Adams and Aizawa as stating that Clark and Chalmers “contend that the active causal processes that extend into the environment are just like the ones found in intracranial cognition” (Menary 2007: 56; Adams and Aizawa 2001: 56). Menary suggests that such a reading of the parity principle is ultimately damaging to the idea of extended cognition⁸⁴:

⁸⁴ It is not always clear what Menary's position on the parity principle is. On the one hand, he suggests Adams and Aizawa (2001) have misinterpreted the principle. Nonetheless, a good deal of his argument relies on rejecting the parity principle and replacing it with the manipulation thesis. Indeed, he states: “a major difference between extended mind style arguments and cognitive integration is that the latter does not depend upon the parity principle” (2007: 57). But this would only appear to be correct if by parity principle we mean Adams and Aizawa's misinterpreted version of the parity principle. Menary's cognitive integration seems entirely to accord with, even presume, the parity principle. Menary's suggestion that cognitive integration does not depend on the parity principle appears to implicitly endorse Adams & Aizawa's misreading of the principle.

... this version of the parity principle is fatally flawed because it assumes the very position it is meant to displace. The “extended mind” and the parity principle encourage us to think of an internal cognitive system that is extended outwards into the world. Hence it implicitly endorses a picture of a discrete cognitive agent, some of whose cognitive processes get extended out into the world. (2007: 56)

Menary worries, along with Adams and Aizawa, that the parity principle implies that extended cognitive processes count as cognitive processes only in virtue of their resemblance (in terms of structures and mechanisms) to internal processes. Such an approach to extended processes would ultimately be self-defeating for two reasons. First, there are unlikely to be very many extended processes which do share this degree of resemblance to internal processes. Second, even if such extended processes were found to exist, they would be considered cognitive only in virtue of their resemblance to internal cognitive processes. Thus we would be prioritising internal cognitive processes by allowing internal processes to define, in Adams and Aizawa's terms, the “mark of the cognitive.” The idea that extended cognitive processes can only be understood *after* we have come to understand internal cognitive processes seems to fundamentally undermine the various forms of externalism about cognition which seek to undermine accounts which, *a priori*, privilege the internal. Given this, Menary argues that the parity principle ought to be abandoned.

As discussed above, it is unlikely that there are many extended processes that share much in common with internal processes that achieve the same goal (solve the appropriate cognitive task). Rather, extended processes are likely to differ in a great many respects. If Clark and Chalmers' argument hung on this point, then their case for extended cognition would be a very weak one. However, it is not the case that their argument relies on this claim. Taking just the two examples Clark and Chalmers use to highlight their case, Tetris and Otto's notebook, it is very clear that they are suggesting that the extended and internal processes are quite different in their detail. The player that rotates the shapes in the Tetris game mentally is doing something quite different than the player that rotates the shape using the button. These processes differ considerably in the structures and mechanisms that comprise them. Similarly, Otto in

consulting his notebook is relying on structures, mechanisms and actions quite differ to Inga when she uses her biological memory. There can be no suspicion that Clark and Chalmers were suggesting internal and external processes were of the same sort in terms of the structures and mechanisms that comprised them:

...our claim is not that the processes in Otto and Inga are identical, or even similar, in terms of their detailed implementation. It is simply that, in respect of the role that the long-term encodings play in guiding current response, both modes of storage can be seen to be supporting dispositional beliefs. (Clark 2008)

Cognitive processes are, on this account, functionally specified. Where an internal and an external process can underpin an ability to solve a cognitive task, though they do this in quite different ways, we can consider them to be of the same (functionally-specified) sort. This is a much weaker commitment than the sort implied by Adams and Aizawa (2001) which suggested a much closer form of similarity between internal and external processes. This criticism, then, does no harm to the extended mind hypothesis.

The “parity principle” might better be named the “parity of reasoning about cognitive processes principle.” Clark and Chalmers are not asking us to think of internal and extended processes as the same sort of thing (in terms of their components, though they can be considered to be of the same sort when functionally defined). What Clark and Chalmers are asking us to do is to apply the same *reasoning* about cognition to a wider variety of cases than perhaps we are used to. There are similarities here with the parity of reasoning called for by developmental systems theorists. This does not amount to the claim that all developmental resources play exactly the same role in development. Rather, it is the claim that we should not, in general, privilege any resource for development over others (though, of course, any particular resource may play a greater or lesser role in a particular developmental process and outcome). We should not assume, *in advance*, that one set of resources are more important, or fundamental, than any other. Clark and Chalmers are making a similar point. If the only reason to discount a process as cognitive is that some of the components for the processes were located beyond the skin, then we would be *a priori*

privileging a certain sort of cognitive resource (internal processes) over others (extended processes). If the task were completed by means of some internal process and, without knowing the details of this process, we were nonetheless happy to term that process cognitive, there seems to be little grounds, in general, to imagine that extended processes that allow us to complete the same task could not be termed cognitive too. I say “in general” here, because specific conditions will have to be met concerning reliable availability and so on, as detailed earlier. But there appear to be no *a priori* grounds to exclude extended processes as cognitive in virtue of the fact that they are extended.

Menary also criticises the sense of functional similarity employed by Clark and Chalmers:

Otto and his notebook do not really function in the same kind of way that Inga does when she has immediate recall from biological memory. There are genuine and important differences in the way that memories are stored internally and externally and these differences matter to how the memories are processed. (2007: 59)

This is undoubtedly correct.⁸⁵ But it misses Clark and Chalmers' objective in stating this parity principle. Functional similarity at the level Clark and Chalmers are working on may be less useful in explaining cognitive processes, but functional similarity at this level is important for establishing their point that there are no *a priori* grounds to rule out a process as cognitive just because some of the components of the process happen to be external to the organism. Again, the parity they are calling for is a parity of reasoning, not a claim that we must treat all internal and external processes as of the same sort in their particular details.

In fact, Adams and Aizawa's (2001) and Menary's (2007) reading of the parity principle is exactly what Clark wanted to avoid:

... far from requiring any deep similarity between inner and outer processes, the parity claim was specifically meant to undermine any

⁸⁵ Clark and Chalmers (1998) appear to accept this point too given they discuss the differences in the details of Otto's and Inga's cognitive performances.

tendency to think that the shape of the present-day, human inner processes sets some bar... on what should count as part of a genuinely cognitive process. (Clark 2008: 114)

It may be the case that others have (mis)understood the parity principle in the way Adams and Aizawa (2001) appear to. Rowlands' cost-benefit analysis, for instance, worked best only when we assumed that the extended processes were structurally very similar to internal processes. However, to ditch the parity principle on the grounds that it has been misinterpreted would be to throw the baby out with the bath water. Ultimately, the parity principle makes a fairly small claim (albeit one with potentially large consequences): do not allow the location of the components of a process to factor into decisions about the cognitive status of that process. This claim is one that Menary ultimately endorses in his own work.

3.2.2 Privileging the Internal

The extended mind hypothesis involves a commitment to a parity of reasoning about what might constitute a resource for a cognitive process. This position rejects attempts to privilege the brain over bodily actions and environmental structures in cognition. However, I want to argue here that there are two related ways in which an implicit privileging of the internal can creep into extended mind accounts. The first is in terms of the manipulation thesis, and the second is in terms of the adoption of Dawkins' (1999a) extended phenotype hypothesis. Clark and Chalmers (1998), Rowlands (1999), and Menary (2007) all make use of the notion of "manipulation," but it is only Menary that explores the extended phenotype hypothesis. Given this, it is Menary's account I will focus on here.

Menary takes Rowlands' articulation of the manipulation thesis as a starting point for his own analysis:

... cognitive processes are not located exclusively inside the skin of the cognizing organisms because such processes are, in part, made up of physical or bodily *manipulation* of structures in the environments of such organisms. (Rowlands 1999: 23, quoted in Menary 2007: 83, italics in original)

Cognitive processes can include manipulations of the environment. I might press a button to rotate a shape in Tetris, or manipulate a pen to solve a maths problem. Menary notes that manipulation of an environmental structure is not enough to identify cognitive processes (2007: 84). In chapters two and four, for example, I discussed ways in which organisms were reciprocally coupled to their environments, but none of these cases are appropriately described as cognitive. Such cases of non-cognitive reciprocal couplings are termed “biological couplings” by Menary. At a minimum, to count as a *cognitive* reciprocal coupling, manipulations would have to meet the criteria for epistemic actions. That is, they would have to alter the world as a means to the end of solving a cognitive task.⁸⁶

Note that my pressing a button or manipulating a pen will not fully describe the external aspects of the cognitive processes here. In the Tetris game, the shape must move in response to my pressing the button, and I must be able to see this. And of course, internal sub-processes will also be involved throughout. That is, my pressing the button is only one part of a cognitive process, and beyond that, only one part of the extended aspect of this particular cognitive process. Many acts of manipulation will require feedback from that which is being manipulated in order to guide future action. However, while there may be a reciprocal causal coupling involved in any manipulation, there is something that manipulates (controls, guides) something else. The manipulation thesis does not suggest that both environmental structure and agent manipulate one another. It is the agent that manipulates the environmental structure, not vice versa. Although manipulation may require a causal symmetry, there is nonetheless an asymmetry in terms of control. Both internal processes and external processes causally affect one another, but there is an ontological difference between the internal and external. The internal manipulates, while the external is manipulated. This is reminiscent of the privileging of the gene in neo-Darwinian evolutionary

⁸⁶ As well as epistemic actions, Menary offers two other classes of reciprocal couplings that can be considered cognitive. The first is “self-correcting action” which involves “the use of language and exogenous props to direct and structure practical actions in completing tasks. The second is “cognitive practice” which involves “manipulations of external representational and notational systems regulated by cognitive norms” (Menary 2007: 84).

biology. Although all manner of developmental resources, located beneath and beyond the skin, are thought to causally affect phenotypes, non-genetic resources are thought to merely causally affect phenotypes, while genes also exert some degree of control over phenotypes in virtue of their informational role in development. The manipulation thesis too suggests that both internal and external structures can play a causal role in the cognitive outcome, but it is the internal that is privileged over the external in terms of control. Because manipulation focuses on an agent who does something to the world, it supports an important boundary between a skin-bound organism and the relevant aspects of the world. What this results in is a tendency to neglect or downgrade the role of external structures. But the point of the extended mind hypothesis is that these aspects of the world are part of cognitive processes and we risk generating inadequate explanations while we neglect external components of cognitive processes.

Privileging the internal makes it easier to neglect the contribution made by external structures. This neglect can be seen in Menary's adoption of Dawkins' (1999a) extended phenotype hypothesis. Menary suggests that the extended phenotype hypothesis provides an evolutionary justification for viewing cognition as extending beyond the brain. Recall from chapter two that the extended phenotype hypothesis is concerned with explaining the close relationship between organisms and aspects of the world. The idea is this: If the phenotype is what genes create and use to ensure their own replication, then a beaver's dam seems to fit this picture. A beaver that builds a better dam than its conspecifics may be favoured by natural selection. And genes may influence, on this account, the building of the dam, just as they might influence the development of the beaver's strong, flat front teeth. Just as two beavers can differ in the effectiveness of their teeth as a result of possessing different alleles, so two dams might differ due to the beavers possessing different alleles. If this is the case, goes this line of thought, why can't we also think of the dam itself, rather than just the dam building behaviour, as another effect of the genes?

... it is clear that the phenotypic effects of genes extend beyond the body of the organism housing the gene, incorporating aspects of the environmental niche into the organismic system as extended

phenotypes. (Menary 2007: 110)

The point raised in chapter two about the extended phenotype was that, even accepting for the moment the strong gene-centric approach it relies on, it only describes one half of a more complex causal interaction. That is, it describes how the environment, via natural selection, can mould the organism such that aspects of the environment become part of its extended phenotype. What it fails to describe is how, in interacting with, and altering, its environment, the organism alters the selection pressures it will be exposed to. That is, the extended phenotype account fails to adequately incorporate reciprocal coupling.

In part, this is because the extended phenotype hypothesis, as with other gene-centric accounts, places genes in control of development, behaviour and even the extended phenotype. Phenotypes, extended or otherwise, do not emerge from the interactions of the many resources involved in development; rather they are created by genes to enable their own replication. This view encourages a neglect of other causal relations. Just as development is conceived as a black box from which emerges the phenotype, the gene-centric approach may result in a similar picture for the extended phenotype. This time the black box incorporates everything up to the adaptation that emerges from it – the beaver's dam for example. This approach does not deny that some two-way interaction between aspects of the environment forming the extended phenotype and the traditional organism is possible. However given the presumption of genes as the source of control, there will tend to be a systematic neglect of other causal relations due in part to their presumed irrelevance, to evolution by natural selection. If the environment is viewed as autonomous when thinking in terms of natural selection, and when genes are considered the only unit of selection, two-way interactions between organisms and their environments will seem unimportant, at least from an evolutionary point of view (see chapter two for a more detailed discussion of these points).

Similarly, by focussing on manipulation in extended cognitive processes, rather than avoiding the underlying prioritising of the internal supposed by a misreading of Clark and Chalmers' parity principle, Menary falls foul of this very error. Menary states that “if we accept the picture of a cognitive agent as implementing a discrete cognitive system, before they ever encounter an external

vehicle, then we will have accepted the very picture of cognition we set out to reject” (2007: 63). Yet the manipulation thesis assumes a discrete cognitive system prior to an extended process, and thus runs counter to Menary’s larger point. Menary’s manipulation thesis, when viewed in isolation, works well with the extended phenotype hypothesis, but neither is appropriate for an account of extended cognitive processes that seeks to avoid either an explicit or implicit internalist bias.

3.3 Developmental Systems Theory

Developmental systems theory provides an evolutionary account of the close, reciprocal coupling an organism can have with aspects of its environment that is neglected by taking an extended phenotype view. Within developmental systems, outcomes emerge from the interactions of component parts. Control of these outcomes is located at the level of the system as a whole, rather than in any of its components. Cognitive processes, as understood from within the extended mind approach, appear to share similar properties. Indeed, if the developmental systems view is right, it should be no surprise that cognitive processes operate in a similar manner. If brains and cognition can be considered to be, at least in part, the products of evolution, then it should come as no surprise that some aspects of cognitive systems mirror other aspects of the biological system. And if other parts of the developing system seem to be composed of processes and resources both internal and external to the skin of the organism, this is even more likely to be true in the case of cognition. There is a wide consensus that humans display a high degree of plasticity and environmental responsiveness when it comes to cognition; evolutionary psychologists spend a good deal of time trying to explain this fact (evoked and epidemiological culture are supposed to help here – see chapter five for a full discussion of this). If we accept that other aspects of biological systems can be constituted by resources located in the environment, the case for extended cognition seems all the stronger.

Developmental systems theory can add more support to the extended mind hypothesis by allowing us to further ameliorate concerns about the transience of environmental features for such an account. Recall that a potential concern for the extended mind hypothesis arose because environmental features were thought to be less reliable than internal ones, and that this would result in selection

favouring those organisms that relied on entirely internal resources over those that also relied on external resources to solve a given problem. Developmental systems theory allows us to challenge the assumption that environmental structures will always be more transient than internal developmental resources.

The expanded notion of inheritance that forms a key commitment of developmental systems theory suggests that we can expect aspects of the environment to be robust and reliably present for each generation. Offspring do not inherit naked DNA, nor do they inherit just a cell. Rather, they can inherit a wide range of developmental resources from DNA, to cell membranes, gut bacteria, food preferences, nests, dams, linguistic communities, cultural practices and social structures and institutions. The environmental structures need not be viewed as entirely autonomous parts of the world that organisms can latch on to but not create or sustain. Those aspects of our environments that constitute cognitive processes may continue to do so not through sheer luck, but because one generation passes the necessary resources to the next generation, and this allows them to recreate or sustain the required environmental structures. These environmental structures may form part of our inheritance, just as genes do.

Some of our cognitive processes will be far too modern to have been selected for, particularly those aspects that rely on cultural artefacts such as computers (as in the Tetris example). These sorts of features of the environment that partly constitute cognitive processes cannot be considered part of the evolutionary developmental system (though they may still be considered part of the developmental system – see chapter four). Nonetheless, even if such structures are too recent for any significant selection to have occurred, this does not undermine the idea that they may be inherited. However, other aspects of our environment may have been around for long enough, and been involved in our ancestors' lives, to such an extent that we may consider them part of the evolutionary developmental system.

External representations may fall into this category. Donald (1991) charts the cognitive evolution of modern humans through the different types of representational system characteristic of *Homo sapiens* at different times. Indeed, on this account, each representation system makes possible the next step in cognitive evolution, and so plays an important role in driving cognitive evolution.

Donald argues that the first step in the evolution of modern human cognitive abilities came from moving away from the purely episodic mind that characterises the rest of the primates to mimetic culture.⁸⁷ Episodic culture is characterised by the highest form of memory representation available; that is, episodic memory. Episodic memory is “memory for specific episodes in life, that is, events with a specific time-space locus... The important feature of this type of memory is its concrete, perceptual nature and its retention of specific episodic details” (Donald 1991: 150). Donald contrasts this with procedural memory – a more archaic form of memory – where what is remembered is only the generalities. Both procedural and episodic memory are common in many different animals, and apes appear to have the most developed episodic memories in non-human animals. This allows them to discern discrete events, even when such events are quite complex. However, where modern and older human cultures really differ from other primates is in their ability to utilise representations. What Donald terms mimetic culture is supposed to signal the break from the purely procedural and episodic memory of the non-human primates to something beginning to resemble modern human culture. Mimetic culture is so-called because mimesis becomes “the dominant or governing mode of representation” (Donald 1991: 162). This is a culture that is non-linguistic yet heavily reliant on intentional representational acts. Mimesis does not refer to reflexive, automatic acts – the kinds of gestures that occur when modern humans speak – nor to acts of imitation or mimicry. It also does not include mimicry, which aims at an exact copy of the original behaviour, while imitation involves a close, but not exact, copy of the original behaviour. What makes mimesis different to both of these cases is that mimesis does not require very close resemblances between the act of mimesis and the original act or event and, further, mimesis is fundamentally concerned with representation. Mimicry and imitation are not trying to represent some action or event, they are attempts to actually replicate the act or event (or almost replicate it in the case of imitation):

⁸⁷ Donald (1991) discusses a total of three key transitions that he argues led to modern human cognitive abilities. After the appearance of mimetic culture came mythic culture. This corresponds to the appearance of spoken language. After this came what Donald terms external symbolic storage and the theoretic culture it enabled. This corresponds to the appearance of early forms of writing, and Donald's account on this third stage shares much in common with accounts of extended cognition that focus on cultural artefacts as constituents of cognitive processes.

Mimesis adds a representational dimension to imitation. It usually incorporates both mimicry and imitation to a higher end, that of re-enacting and re-presenting an event or relationship... mimesis is fundamentally different from imitation and mimicry in that it involves the invention of intentional representations. When there is an audience to interpret the action, mimesis also serves the purpose of social communication. However, mimesis may simply represent the event to oneself, for the purpose of rehearsing and refining a skill: the act itself may be analyzed, re-enacted and re-analyzed, that is, represented to oneself. (Donald 1991: 169)

Mimesis, argues Donald, played a vital role in the culture of early hominids. As well as allowing the rehearsing and finessing of a skill, it enabled social communication, reciprocal mimetic games, social conformity and coordination, group mimetic acts and pedagogy. Donald offers ritual dance as an example of reciprocal mimetic games. If one individual mimetically represents an event or act, others may decide to reproduce these acts too. This can lead to some kind of conformity. The ritual dance arises out of a sense of how the dance should be done. Such rituals could easily extend beyond dance and characterise a good many other aspects of such a culture. Hunting, as Donald notes, requires a good deal of coordination – from making tools, to catching and killing prey, and finally butchering the animal: “it is likely that mimesis was the basis, if not the formative element, behind this new cooperative, specialized social organization... it would have enabled purposive signalling” (Donald 1991: 175).

The mimetic representational system transformed early hominid culture. Mimetic representations, although they can be used purely for one's own benefit, if Donald is right, appear also to have played a role as an external feature of the environment for the audience in helping to structure and coordinate complex activities. If mimesis can lead to coordinated activities such as ritual dances and hunting, the vehicles of these external representations may, in some instances, count as constituents of cognitive processes. And if this is the case, this indicates that a coupling between our ancestors and the vehicles of such representations is a long one.

Being born into such a culture, one inherits the system of representations that characterise it. Mimesis continues to play an important role in modern societies; for example, it appears in opera, theatre, and in children's play. Humans have had a long history of engaging with vehicles of mimetic representations and, if cognitive processes can have such vehicles as components, mimetic representational systems might then feature as resources for evolutionary developmental systems.

Regardless of whether we can in fact say that any particular extended cognitive process is a product of natural selection, we can still acknowledge that parents actively pass on structures in the environment to their offspring, thus ensuring features external to the organism's skin are not as transient as might be supposed. We are born into highly structured environments that can provide resources for the initial development of cognitive abilities (Clark and Chalmers might call these features of the environment passive features of extended cognition), as well as constituents of cognitive processes. Further, offspring reconstruct or maintain structures both beneath and beyond their skin in each generation such that resources for development can persist for potentially evolutionarily significant periods of time.

Clark lays out the aims of the extended mind approach:

We do indeed seek to carve nature at the most causally relevant joints, a task not accomplished by elevating anatomic or metabolic boundaries into make-or-break cognitive ones... As philosophers and cognitive scientists we can (and should) practice the art of flipping between... different perspectives, treating each as a lens apt to draw attention to certain features, regularities, and contributions while making it harder to spot others, or to give them their problem-solving due. (2007: 191-192)

Developmental systems theory also builds in a parity of reasoning about what may be causally relevant. Further, it encourages “flipping” between perspectives when viewing a problem. We might first hold all but one constituent of a system still, but DST encourages us to look again at those parts of the system that formed the background conditions and assess their causal role too. Thus, the developmental systems perspective provides an account of evolution and

development that is consistent with both the methodology and the model of cognition of the extended mind approach.

Conclusion

In my introduction I stated that I wanted to explicate the role of development in evolutionary considerations. I have argued that development is misconstrued and neglected in neo-Darwinian evolutionary biology. The developmental processes of an individual organism tend to be viewed as the output of the evolutionary process rather than as something that may contribute to a lineage's evolutionary trajectory. This view results in a commitment to the gene as the sole unit of inheritance and the only entity capable of having a history of selection. Genes with a history of selection are understood to embody information about phenotypes in a way that marks them out as distinct from all other developmental resources. The informational gene then removes the need to investigate development in order to understand evolution; development is just the realisation of the instructions encoded in the genes. Where development is not outright neglected, its role is misconstrued. This occurs when facts about development are thought to account only for constraints on evolution by natural selection. This model, one shared with process structuralism, is built on a false dichotomy between development and natural selection. However, I argued that, because natural selection can act to conserve traits and developmental processes can help explain evolution by natural selection, that this model was inadequate.

Evolutionary developmental biology also views the gene as a fundamentally different sort of developmental resource, but it rejects the false dichotomy between development and natural selection. Evolutionary developmental biologists argue that it is developmental processes that enable evolutionary innovation and that it is these processes that must be understood if we are to provide adequate evolutionary explanations. Development is construed in this approach as involving more than the realisation of a genetic programme and is characterised by emergent phenomena at various levels of biological organisation. While the neo-Darwinist might grant that such emergent phenomena may be a feature of development, evolutionary developmental biologists go further by claiming these phenomena are evolutionarily significant and thus worthy of investigation. Evo-devo makes a strong case for the

importance of development in evolutionary considerations but faces a fundamental problem. Mainstream or synthetic evo-devo continues to treat the gene as the sole unit of hereditary transmission, despite its recognition that more than genes are inherited, and thus evo-devo privileges the gene in development. This, I have argued, is based on a circular argument. The gene is thought to play an instructional, not merely causal, role because it is considered to be the unit of hereditary transmission, and it is the unit of hereditary transmission because it is thought to play this directive, not merely causal, role in development.

DST also brings development to the fore in its approach to understanding evolution. Unlike evo-devo, DST rejects the idea that the gene is the unit of hereditary transmission and instead argues for a much broader understanding of what might bring about heritable similarity. Offspring can inherit entities found within the zygote such as DNA methylation patterns and cytoplasmic polarities, but also extra-cellular entities such as gut bacteria, nests, and social groups. This allows the role of the environment in development to be incorporated into an account of evolution. Further, it is not only discrete entities that are inherited but also the interactions between them. This amounts to a rejection of the notion that first there is hereditary transmission, and only then does the phenotype emerge as a result of development; hereditary “transmission” does not precede development. This expanded notion of inheritance offers no justification for viewing the gene as ontologically distinct from other developmental resources and, as a result, preformationism is eschewed. Development is instead characterised as the product of the interactions between developmental resources such that none of the individual resources controls development and instead control is distributed throughout the system.

The developmental systems perspective rejects moves to *a priori* privilege genes over other developmental resources or, in more general terms, to privilege the internal over the external. This is not the unworkable holism feared by critics; DST allows us to acknowledge the distinct contributions of different developmental resources and acknowledges that not all resources that constitute the developmental system will be particularly important in any given developmental outcome. The details of the role cell membranes play in development will differ in many ways from the details of the role a linguistic community may play in development. But by beginning with the assumption that

there is no ontological distinction between developmental resources, DST ensures the role played by developmental resources is a matter for investigation. Rather than being black boxed, development becomes an important part of the explanandum.

The developmental systems approach avoids two problems faced by the evolutionary developmental biologist. First, it does not fall prey to the circularity inherent in the evo-devo position that recognises more than genes are inherited, but which nonetheless maintains an ontologically distinct role for genes. Second, it is able to integrate the role of the environment in development into evolutionary considerations. Evo-devo has tended to focus on processes internal to the skin-bound organism and has not yet incorporated the role of the environment in development into broader evolutionary considerations, a state of affairs lamented by some evolutionary developmental biologists (Gilbert & Bolker 2003). DST provides a framework to evaluate the full role of the environment in development and the evolutionary consequences of this. The developmental systems approach rejects the idea that an organism's environment is largely autonomous – an idea implicit in much of the adaptationist programme – and acknowledges that the skin-bound organism constructs and maintains its relationships with structures in the environment. But DST can be improved by taking on board some of the work of evolutionary developmental biologists. For instance, the hierarchical model of selection can be extended to include selection working at the level of the different modules identified by evolutionary developmental biologists.

Taking a developmental systems perspective results in a very different account of the evolution of culture and cognition. I examined a number of other approaches to these issues which, to a greater or lesser degree, relied on the neo-Darwinian model of development and evolution, as well as an internalist model of cognition, and argued that these approaches faced serious problems. The combination of neo-Darwinian evolutionary biology and an internalist model of the mind leads Evolutionary Psychologists to argue that we possess “stone age minds.” However, the basic tension that arises from the alleged mismatch between the evolved cognitive capacities of Evolutionary Psychology and the modern world renders unfeasible attempts to explain culture as the product of cognitive capacities evolved to suit the EEA. The kind of gene-culture co-

evolution account offered by Richerson and Boyd (2005) does not suggest that culture is tightly constrained by evolved cognitive capacities and it may allow us to track evolutionary changes at the cultural level; however, it faces two difficulties. First, in adopting a gene-culture co-evolutionary model, they endorse an account of evolution which continues to exclude any significant role for development. Genetic evolution and cultural evolution represent two distinct realms, though they can influence one another, while all other developmental resources and the interactions between them are overlooked. Second, it relies on cognitive modules that resemble the kind of modules favoured by Evolutionary Psychologists. Against both of these internalist accounts, I argued that we can offer a more adequate account of the evolution of cognition by adopting the extended mind hypothesis. This hypothesis suggests that our cognitive capacities are best understood as being constituted by resources both beneath our skin and in the environment. The implication of this model of cognition is that, to the extent that we inhabit environments that differ from our Pleistocene ancestors, our cognitive processes will differ from those of our ancestors.

A fundamental problem remains, however, for the study of the evolution of cognition and culture. Because human kinds differ from natural kinds in that they can result in looping effects such that these kinds are unstable, studying the evolution of human cognition and culture will differ from studying the evolution of non-human organisms. Human kinds, as presented by Hacking, have tended to focus on the change brought about in the behaviour of the individuals so classed; I looked at the case of the obese to demonstrate how the developmental system more broadly can be affected in this way. When we study humans, what we are trying to explain may change as a result of our research, and this change will often be in ways we cannot predict. This places a limitation on how well we can come to understand modern culture and cognition in its specifics. However, the developmental systems approach combined with the extended mind hypothesis can at least help us understand why human kinds differ from natural kinds in this way.

Combining DST with the extended mind hypothesis changes our view of the relationship between biology, psychology and culture. Evolutionary Psychologists view culture primarily as the output of psychological modules evolved to suit our EEA, while gene-culture evolutionary accounts tend to view

biology and culture as distinct though interacting. In both instances, cognitive modules are taken as given; culture may then act as an input to these modules. The developmental systems perspective, on the other hand, rejects the claim that such modules are given prior to any interaction with culture and the environment more generally. Rather, features of the cultural and non-cultural environment may act as resources for development. Further, the extended mind hypothesis suggests that structures in the environment might not only partly constitute developmental processes, but may also partly constitute the cognitive processes themselves. Taken together, the developmental systems approach and the extended mind hypothesis change the explanandum for the study of the evolution of cognition, and in doing so, may provide a more productive and coherent framework in which to conduct research.

References

- Adams, F. & K. Aizawa. 2001. The Bounds of Cognition. *Philosophical Psychology* 14: 43-64.
- Adams, F. & K. Aizawa. 2008. *The Bounds of Cognition*. Malden MA: Blackwell.
- Amundson, R. 1994. Two Concepts of Constraint: Adaptationism and the Challenge from Developmental Biology. *Philosophy of Science* 61: 556-578.
- Ariew, A. 1999. Innateness is Canalization: In Defense of a Developmental Account of Innateness. In *Where Biology Meets Psychology: Philosophical Essays*, ed. V. G. Hardcastle. Cambridge, Mass.: MIT Press.
- Arthur, W. 1987. *Theories of Life: Darwin, Mendel and Beyond*. Middlesex: England.
- Barkow, J. H., L. Cosmides, & J. Tooby. 1992. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford: Oxford University Press.
- Berg, F. M. 1999. Health Risks Associated with Weight Loss and Obesity Treatment Programs. *Journal of Social Issues* 55(2): 277-297.
- Blackmore, S. J. 1999. *The Meme Machine*. Oxford: Oxford University Press.
- Bolker, J. 2000. Modularity in Development and Why it Matters to Evo-Devo. *American Zoologist* 40: 770-776.
- Bourke, R. M. 2009. The History of Agriculture in Papua New Guinea. In *Food and Agriculture in Papua New Guinea*, eds. R. M. Bourke & T. Harwood. Canberra: ANU Press.
- Brandon, R. N. 1996. *Concepts and Methods in Evolutionary Biology*. Cambridge: Cambridge University Press.
- Brandon, R. 1990. *Adaptation and Environment*. Princeton: Princeton University Press.
- Brunnander, B. 2007. What is Natural Selection? *Biology and Philosophy* 22: 231-246.

- Buller, D. J. 2005. *Adapting Minds: Evolutionary Psychology and the Persistent Quest for Human Nature*. Cambridge, Mass.: MIT Press.
- Buller, D. J. & V. G. Hardcastle. 2000. Evolutionary Psychology Meet Developmental Neurobiology: Against Promiscuous Modularity. *Brain and Mind* 1: 307-325.
- Burian, R. M. 2005. *The Epistemology of Development, Evolution, and Genetics*. New York: Cambridge University Press.
- Buss, D. M. 1992. On the Use and Misuse of Darwinism in the Study of Behaviour. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, eds. J. Barkow, L. Cosmides & J. Tooby. Oxford: Oxford University Press.
- Buss, D. M. 1999. *Evolutionary Psychology: The New Science of the Mind*. Boston, Mass.: Allyn & Bacon.
- Buss, D. M., M. G. Haselton, T. K. Shakelford, A. L. Bleske, and J. C. Wakefield. 1998. Adaptations, Exaptations, and Spandrels. *American Psychologist* 53: 533-548.
- Buss, D. M., T. K. Shakelford, A. L. Bleske, & J. C. Wakefield. 1999. Interactionism, Flexibility, and Inferences about the Past. *American Psychologist* 54: 443-445
- Clark, A. 1997. *Being There: Putting Brain, Body, and World Together Again*. Cambridge, Mass.; London: MIT Press.
- Clark, A. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
- Clark, A. & D. Chalmers. 1998. The Extended Mind. *Analysis* 58: 7-19.
- Cooper, R. 2004. Why Hacking is Wrong about Human Kinds. *The British Journal of the Philosophy of Science* 55: 73-85.
- Cosmides, L. & J. Tooby. 1997. Evolutionary Psychology: A Primer. Available at <http://www.psych.ucsb.edu/research/cep/primer.html> Accessed on 20/10/09.
- Coyne, J. A. 2009. Evolution's Challenge to Genetics. *Nature* 457: 382-383.
- Crick, F. 1970. Central Dogma of Molecular Biology. *Nature* 227: 561-563.
- Crook, J. H. & S. J. Crook. 1988. Tibetan Polyandry: Problems of Adaptations and Fitness. In *Human Reproductive Behaviour: A Darwinian*

- Perspective*, eds. L. Betzig, M. Borgerhoff Mulder & P. Turke. New York: Cambridge University Press.
- Daly, M. & M. Wilson. 1988. Evolutionary Social Psychology and Family Homicide. *Science* 242: 519-524.
- Darwin, C. 1859/1996. *On the Origin of Species*. Oxford: Oxford University Press.
- Darwin, C. 1905. *The Variation of Animals and Plants under Domestication*, ed. F. Darwin. London: John Murray.
- Dawkins, R. 1981. *In Defence of Selfish Genes*. *Philosophy* 56: 556-573.
- Dawkins, R. 1986. *The Blind Watchmaker*. London: Penguin.
- Dawkins, R. 1989. *The Selfish Gene*. 2nd edition. Oxford: Oxford University Press.
- Dawkins, R. 1998. Universal Darwinism. In *The Philosophy of Biology*, eds. D. L. Hull & M. Ruse. Oxford: Oxford University Press.
- Dawkins, R. 1999a. *The Extended Phenotype: The Long Reach of the Gene*. Oxford: Oxford University Press.
- Dawkins, R. 1999b. Introduction to *The Meme Machine*, S. Blakemore.
- Donald, M. 1991. *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Cambridge, Mass.: Harvard University Press.
- Dretske, F. 1981. *Knowledge and the Flow of Information*. Cambridge, Mass.: MIT Press.
- Dupré, J. 2001. *Human Nature and the Limits of Science*. Oxford: Oxford University Press.
- Dupré, J. 2002. *Humans and Other Animals*. Oxford: Clarendon Press.
- Dupré, J. 2005. Are Their Genes? In *Philosophy, Biology and Life* (Royal Institute of Philosophy Supplements), ed. A. O'Hear. Cambridge: Cambridge University Press.
- Dupré, J. 2008. Against Maladaptationism: or, What is Wrong with Evolutionary Psychology? In *Knowledge as Social Order: Rethinking the Sociology of Barry Barnes*, ed. M. Mazzetti. Aldershot: Ashgate.
- Dwyer, J. T., Feldman, J. J. & Mayer, J. 1970. The Social Psychology of Dieting. *Journal of Health and Social Behaviour*. 11(4): 269-287.
- Ereshefsky, M. 2004. Bridging the Gap between Human Kinds and Natural Kinds. *Philosophy of Science* 71: 912-921.

- Ernsberger, P. & R. J. Koletsky. 1999. Biomedical Rationale for a Wellness Approach to Obesity: An Alternative to a Focus on Weight Loss. *Journal of Social Issues* 55(2): 221-260.
- Gaesser, G. A. 2003. Is it Necessary to be Thin to be Healthy? *Harvard Health Policy Review* 4(2): 40-47.
- Gard, M. & J. Wright. 2005. *The Obesity Epidemic: Science, Morality and Ideology*. Oxon: Routledge.
- Gilbert, S. F. 2001. Ecological Developmental Biology: Developmental Biology Meets the Real World. *Developmental Biology* 233: 1-12.
- Gilbert, S. F. 2003. Evo-Devo, Devo-Evo, and Devgen-Popgen. *Biology and Philosophy* 18: 347-352.
- Gilbert, S. F. & J. A. Bolker. 2003. Ecological Developmental Biology. *Evolution and Development* 5: 3-8.
- Gluckman, P. D, M. A. Hanson, & H. G. Spencer. 2005. Predictive Adaptive Responses and Human Evolution. *Trends in Ecology and Evolution* 20: 527-533.
- Godfrey-Smith, P. 1989. Misinformation. *The Canadian Journal of Philosophy* 19(5): 533-550.
- Godfrey-Smith, P. 1996. *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.
- Godfrey-Smith, P. 2001a. Three Kinds of Adaptationism. In *Adaptationism and Optimality*, eds. S. H. Orzack & E. Sober. Cambridge: Cambridge University Press.
- Godfrey-Smith, P. 2001b. On the Status and Explanatory Structure of Developmental Systems. In *Cycles of Contingency: Developmental Systems and Evolution*, eds. S. Oyama, P. E. Griffiths & R. D. Gray. Cambridge, Mass.: MIT Press.
- Godfrey-Smith, P. 2007. Conditions for Evolution by Natural Selection. *Journal of Philosophy* 104: 489-516.
- Godfrey-Smith, P. 2008. Information in Biology. *The Cambridge Companion to the Philosophy of Biology*, eds. D. Hull & M. Ruse. Cambridge: Cambridge University Press.
- Godfrey-Smith, P. 2009. *Darwinian Populations and Natural Selection*. Oxford: Oxford University Press.

- Goodwin, B. C. 2001. *How the Leopard Changed its Spots: The Evolution of Complexity*. Princeton: Princeton University Press.
- Gould, S. J. & R. Lewontin. 1979. The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme. *Proceedings of the Royal Society of London B* 205: 581-598.
- Gould, S. J. & E. S. Vrba. 1998. Exaptation – A Missing Term in the Science of Form. In *Philosophy of Biology*, eds. D. L. Hull & M. Ruse. Oxford: Oxford University Press.
- Griffiths, P. E. 1996a. The Historical Turn in the Study of Adaptationism. *The British Journal for the Philosophy of Science* 47: 511-532.
- Griffiths, P. E. 1996b. Darwinism, Process Structuralism and Natural Kinds. *Philosophy of Science* (Proceedings) 63: S1-S9.
- Griffiths, P. E. 2001. Genetic Information: A Metaphor in Search of a Theory. *Philosophy of Science* 68(3): 394-412.
- Griffiths, P. E. 2007. The Phenomena of Homology. *Biology and Philosophy* 22: 643-658.
- Griffiths, P. E. & R. D. Gray. 1994. Developmental Systems and Evolutionary Explanations. *Journal of Philosophy* 91(6): 277-304.
- Griffiths, P. E. & R. D. Gray. 1997. Replicator II: Judgement Day. *Biology and Philosophy* 12: 471-492.
- Griffiths, P. E. & R. D. Gray. 2001. Darwinism and Developmental Systems. In *Cycles of Contingency: Developmental Systems and Evolution*, eds. S. Oyama, P. E. Griffiths & R. D. Gray. Cambridge, Mass.: MIT Press.
- Griffiths, P. E. & R. D. Gray. 2004. The Developmental Systems Perspective: Organism-Environment Systems as Units of Development and Evolution. In *Phenotypic Integration: Studying the Ecology and Evolution of Complex Phenotypes*, eds. M. Pigliucci & K. Preston. Oxford: Oxford University Press.
- Griffiths, P. E. & R. D. Gray. 2005. Three Ways to Misunderstand Developmental Systems Theory. *Biology and Philosophy* 20: 417-425.
- Griffiths, P. E. & R. D. Knight. 1998. What is the Developmentalist Challenge? *Philosophy of Science* 65: 253-258.
- Griffiths, P. E. & K. Stotz. 2000. How the Mind Grows: A Developmental Perspective on the Biology of Cognition. *Synthese* 122: 29-51.

- Hacking, I. 1996. The Looping Effects of Human Kinds. In *Causal Cognition*, eds. D. Sperber & A. Premack. Oxford: Clarendon Press.
- Hacking, I. 1999. *The Social Construction of What?* Cambridge, Mass.: Harvard University Press.
- Hall, B. K. 1992. *Evolutionary Developmental Biology*. 1st Edition. London: Chapman & Hall.
- Hall, B. K. 1998. Epigenetics: Regulation not Replication. *Journal of Evolutionary Biology* 11: 201-205.
- Hall, B. K. 1999. *Evolutionary Developmental Biology*. Second Edition. Dordrecht: Kluwer.
- Hall, B. K. 2000. Evo-devo or Devo-Evo – Does it Matter? *Evolution and Development* 2: 177-178.
- Hall, B. K. 2001. The Gene is not Dead, Merely Orphaned and Seeking a Home. *Evolution and Development* 3: 225-228.
- Hall, B. K. 2003. Unlocking the Black Box between Genotype and Phenotype: Cell Condensations as Morphogenetic (Modular) Units. *Biology and Philosophy* 18: 219-247.
- Ho, M.-W. 1988. On Not Holding Nature Still: Evolution by Process, not Consequence. In *Evolutionary Processes and Metaphors*, eds. M.-W. Ho & S. W. Fox. Chichester: Wiley.
- Hull, D. L. 1980. Individuality and Selection. *Annual Review of Ecology and Systematics* 11: 311-332.
- Hurley, S. Forthcoming. Varieties of Externalism. In *The Extended Mind*, ed. R. Menary. Ashgate.
- Ingold, T. 2001. From Complementarity to Obviation: On Dissolving the Boundaries between Social and Biological Anthropology, Archaeology, and Psychology. In *Cycles of Contingency: Developmental Systems and Evolution*, eds. S. Oyama, P. E. Griffiths & R. D. Gray. Cambridge, Mass.: MIT Press.
- Jablonka, E. & M. Lamb. 2005. *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge, Mass.: MIT Press.
- Jeffery, R. W. et al. 2000. Long-term Maintenance of Weight Loss: Current Status. *Health Psychology* 19(1): 5-16.

- Jenkin, F. 1867. (Review of) *The Origin of Species*. *The North British Review* 46: 277-318.
- Kaplan, J. M. 2000. *The Limits and Lies of Human Genetic Research: Dangers for Social Policy*. New York: Routledge.
- Kassirer, J. P. & M. Angell. 1998. Losing Weight – An Ill-fated New Year’s Resolution. *The New England Journal of Medicine* 338(1): 52-54.
- Keller, E. F. 2001. Beyond the Gene but Beneath the Skin. In *Cycles of Contingency: Developmental Systems and Evolution*, eds. S. Oyama, P. E. Griffiths & R. D. Gray. Cambridge, Mass.: MIT Press.
- Kimura, M. 1968. Evolutionary Rate at the Molecular Level. *Nature* 217: 624-626
- Kirsh, D. & P. Maglio. 1994. On Distinguishing Epistemic from Pragmatic Action. *Cognitive Science* 18: 513-549.
- Levy, N. 2004. Evolutionary Psychology, Human Universals, and the Standard Social Science Model. *Biology and Philosophy* 19: 459-472.
- Lewens, T. 2007. Cultural Evolution. *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta. Available at <http://plato.stanford.edu/entries/evolution-cultural/> Accessed on 20/10/09
- Lewontin, R. C. 1978. Adaptation. *Scientific American* 239(3): 156-169.
- Lewontin, R. C. 2000. *The Triple Helix: Gene, Organism and Environment*. Cambridge, Mass.: Harvard University Press.
- Lewontin, R. C. 2001. Gene, Organism and Environment. In *Cycles of Contingency: Developmental Systems and Evolution*, eds. S. Oyama, P. E. Griffiths & R. D. Gray. Cambridge, Mass.: MIT Press.
- Lloyd, E. A. 2001. Units and Levels of Selection: An Anatomy of the Units of Selection Debate. In *Thinking about Evolution: Historical, Philosophical, and Political Perspectives*, eds. R. S. Singh, C. B. Krimbas, D. B. Paul, & J. Beatty. Cambridge: Cambridge University Press.
- Lloyd, E. A. 2005. Units and Levels of Selection. *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta. Available at <http://plato.stanford.edu/entries/evolution-cultural/> Accessed on 20/10/09
- Lutnesky, M. M. F. 1994. Density-Dependent Protogynous Sex Change in Territorial-Haremic Fishes: Models and Evidence. *International Society for Behavioral Ecology* 5: 375-383.

- Maynard Smith, J. 2000. The Concept of Information in Biology. *Philosophy of Science* 67(2): 177-194.
- Maynard Smith, J. & E. Szathmáry. 1995. *The Major Transitions in Evolution*. Oxford: Oxford University Press.
- Mayr, E. 2001. *What Evolution Is*. London: Phoenix.
- Menary, R. 2007. *Cognitive Integration: Mind and Cognition Unbound*. Hampshire: Palgrave Macmillan.
- Morange, M. 2000. Gene Function. *Comptes Rendus de l'Academie des Sciences* 323: 1147-1153.
- Morton, T. A., S. A. Haslam, T. Postmes, & M. K. Ryan. 2006. *We Value What Values Us: The Appeal of Identity Affirming Science*. *Political Psychology* 27: 823-838.
- Moss, L. 2004. *What Genes Can't Do*. Cambridge, Mass.: MIT Press.
- Nijhout, H. F. 2001. The Ontogeny of Phenotypes. In *Cycles of Contingency: Developmental Systems and Evolution*, eds. S. Oyama, P. E. Griffiths & R. D. Gray. Cambridge, Mass.: MIT Press.
- Noss, A. J., & B. S. Hewlett. 2001. The Contexts of Female Hunting in Central Africa. *American Anthropologist* 103: 1024-1040.
- Odling-Smee, F. J., K. N. Laland, & M. W. Feldman. 2003. *Niche Construction: The Neglected Process in Evolution*. Princeton University Press.
- Oyama, S. 2000a. *The Ontogeny of Information*. Duke University Press.
- Oyama, S. 2000b. *Evolution's Eye: A Systems View of the Biology-Culture Divide*. Durham, NC: Duke University Press.
- Oyama, S. 2000c. Causal Democracy and Causal Contributions in Developmental Systems Theory. *Philosophy of Science* 67: S332-347.
- Oyama, S., P. E. Griffiths & R. D. Gray. 2001. Introduction: What is Developmental Systems Theory? In *Cycles of Contingency: Developmental Systems and Evolution*, eds. S. Oyama, P. E. Griffiths & R. D. Gray. Cambridge, Mass.: MIT Press.
- Pinker, S. 2002. *The Blank Slate: The Modern Denial of Human Nature*. London: Penguin.
- Pinker, S. 2004. *The Language Instinct: The New Science of Language and Mind*. London: Penguin.

- Pinker, S. & Bloom, P. 1992. Natural Language and Natural Selection. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, eds. J. Barkow, L. Cosmides & J. Tooby. Oxford: Oxford University Press.
- Pinto-Correia, C. 1997. *The Ovary of Eve: Egg and Sperm and Preformation*. Chicago: Chicago University Press.
- Putnam, H. 1975. The Meaning of "Meaning." *Minnesota Studies in the Philosophy of Science* 7: 131-193.
- Richerson, P. J. & R. Boyd. 2005. *Not by Genes Alone: How Culture Transformed Human Evolution*. Chicago: The University of Chicago Press.
- Richerson, R. C. 2007. *Evolutionary Psychology as Maladapted Psychology*. Cambridge, Mass.: MIT Press.
- Ridley, M. 1997. *The Origin of Virtue*. London: Penguin.
- Robert, J. S. 2002. How Developmental is Evolutionary Developmental Biology? *Biology and Philosophy* 17: 591-611.
- Robert, J. S. 2004. *Embryology, Epigenesis, and Evolution: Taking Development Seriously*. Cambridge: Cambridge University Press.
- Robert, J. S., B. K. Hall, & W. M. Olson. 2001. Bridging the Gap between Developmental Systems Theory and Evolutionary Developmental Biology. *BioEssays* 23: 954-962.
- Rowlands, M. 1999. *The Body in Mind: Understanding Cognitive Processes*. Cambridge: Cambridge University Press.
- Rowlands, M. 2003. *Externalism: Putting Mind and World Back Together Again*. Chesham: Acumen.
- Samuels, R. 2000. Massively Modular Minds: Evolutionary Psychology and Cognitive Architecture. In *Evolution and the Human Mind*, ed. P. Carruthers. Cambridge: Cambridge University Press.
- Sarkar, S. & J. S. Robert. 2003. Introduction. *Biology and Philosophy* 18: 209-217.
- Schlosser, G. 2002. Modularity and Units of Evolution. *Theory in Biosciences* 121: 1-80.
- Schwenk, K. & G. P. Wagner. 2004. The Relativism of Constraints on Phenotypic Evolution. In *Phenotypic Integration: Studying the Ecology*

- and Evolution of Complex Phenotypes*, eds. M. Pigliucci & K. Preston. Oxford: Oxford University Press.
- Silverman, I. & M. Eals. 1992. Sex Differences in Spatial Abilities: Evolutionary Theory and Data. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, eds. J. Barkow, L. Cosmides & J. Tooby. Oxford: Oxford University Press.
- Silvertown, J. & D. Charlesworth. 2001. *Introduction to Plant Population Biology*. 4th Edition. Oxford: Blackwell Scientific Publications.
- Sober, E. 1984. *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. London: MIT Press.
- Sober, E. 1991. Models of Cultural Evolution. In *Trees of Life: Essays in Philosophy of Biology*, ed. P. E. Griffiths. Dordrecht: Kluwer.
- Sober, E. & Wilson, D. S. 1998. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, Mass.: Harvard University Press.
- Sperber, D. 1996. *Explaining Culture: A Naturalistic Approach*. Cambridge, Mass.: Blackwell.
- Sperber, D. 2000. An Objection to the Memetic Approach to Culture. In *Darwinizing Culture: The Status of Memetics as a Science*. Oxford: Oxford University Press.
- Sterelny, K. 2000. Development, Evolution, and Adaptation. *Philosophy of Science (Proceedings)* 67: S369-S387.
- Sterelny, K. 2001. Niche Construction, Developmental Systems, and the Extended Replicator. In *Cycles of Contingency: Developmental Systems and Evolution*, eds. S. Oyama, P. E. Griffiths & R. D. Gray. Cambridge, Mass.: MIT Press.
- Sterelny, K. & P. E. Griffiths. 1999. *Sex and Death: An Introduction to the Philosophy of Biology*. Chicago, Ill.: University of Chicago Press.
- Sterelny, K., K. C. Smith & M. Dickison. 1996. The Extended Replicator. *Biology and Philosophy* 11: 377-403.
- Stotz, K., P. E. Griffiths, R. Knight. 2004. How Biologists Conceptualise Genes: An Empirical Study. *Studies in History and Philosophy of Biology* 35: 647-673.
- Symons, D. 1992. On the Use and Misuse of Darwinism in the Study of Human Behavior. In *The Adapted Mind: Evolutionary Psychology and the*

- Generation of Culture*, eds. J. Barkow, L. Cosmides & J. Tooby. Oxford: Oxford University Press.
- Száthmary, E. & J. Maynard Smith. 1995. The Major Evolutionary Transitions. *Nature* 374: 227-232.
- Tooby, J. & L. Cosmides. 1989. Evolutionary Psychology and the Generation of Culture, Part I: Theoretical Considerations. *Ethology and Sociobiology* 10: 20-49.
- Tooby, J. & L. Cosmides. 1990. On the Universality of Human Nature and the Uniqueness of the Individual. *Journal of Personality* 58: 17-67.
- Tooby, J. & L. Cosmides. 1992. The Psychological Foundations of Culture. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, eds. J. Barkow, L. Cosmides & J. Tooby. Oxford: Oxford University Press.
- Turner, J. S. 2004. Extended Phenotypes and Extended Organisms. *Biology and Philosophy* 19: 327-352.
- Waddington, C. H. 1942. Canalization of Development and the Inheritance of Acquired Characters. *Nature* 150: 563-565.
- Wagner, G. P. 2000. What is the Promise of Developmental Evolution? Part I: Why is Developmental Evolution Necessary to Explain Evolutionary Innovations? *Journal of Experimental Zoology (Mol Dev Evol)* 288: 95-98.
- Wagner, G. P., C.-H. Chiu, & M. Laubichler. 2000. Developmental Evolution as a Mechanistic Science: The Inference from Developmental Mechanisms to Evolutionary Processes. *American Zoologist* 40: 819-831.
- Webster, G. & B. C. Goodwin. 1982. The Origin of Species: A Structuralist Approach. *Journal of Social and Biological Structures* 5: 15-47.
- West-Eberhard, M. J. 1998. Evolution in the Light of Developmental and Cell Biology, and Vice Versa. *Proceedings of the National Academy of Sciences USA* 95: 8417-8419.
- West-Eberhard, M. J. 2003. *Developmental Plasticity and Evolution*. New York: Oxford University Press.
- West-Eberhard, M. J. 2005. Phenotypic Accommodation: Adaptive Innovation Due to Developmental Plasticity. *Journal of Experimental Zoology* 304B: 610-618.

- Williams, G. C. 1966. *Adaptation and Natural Selection*. Princeton: Princeton University Press.
- Wilson, M. & M. Daly. 1992. The Man Who Mistook His Wife for a Chattel. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, eds. J. Barkow, L. Cosmides & J. Tooby. Oxford: Oxford University Press.
- Wright, S. 1932. The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution. *Proceedings of the Sixth Annual Congress of Genetics* 1: 356-366. Reprinted in Provine, W. B. 1986. *Sewall Wright: Selected Papers*. Chicago: University of Chicago.
- Young, L. & B. Powell. 1985. The Effects of Obesity on the Clinical Judgements of Mental Health Professionals. *Journal of Health and Social Behaviour*. 26(3): 233-246.