Philosophy of Evolutionary Biology

The International Library of Essays on Evolutionary Thought
Series Editor: Neil Levy

Titles in the Series:

Philosophy of Evolutionary Biology: Volume I
Stefan Linquist

Evolutionary Psychology: Volume II
Stefan Linquist and Neil Levy

Evolutionary Ethics: Volume III
Neil Levy

The Evolution of Culture: Volume IV
Stefan Linquist

Intelligent Design and Religion as a Natural Phenomenon: Volume V
John S. Wilkins

Philosophy of Evolutionary Biology
Volume I

Edited by

Stefan Linquist

University of Guelph, Canada

ASHGATE

Published by
Ashgate Publishing Limited
Wey Court East
Union Road
Farnham
Surrey GU9 7PT
England

Ashgate Publishing Company
Suite 420
101 Cherry Street
Burlington, VT 05401-4405
USA

Ashgate website: http://www.ashgate.com

British Library Cataloguing in Publication Data
Philosophy of evolutionary biology.
Volume 1. – (International library of essays on evolutionary thought)
1. Evolution (Biology)--Philosophy.
I. Series II. Linquist, Stefan Paul.
576.8'01-dc22

Library of Congress Control Number: 2009938296

ISBN: 978 0 7546 2753 1

Mixed Sources
Product group from well-managed forests and other controlled sources
www.fsc.org Cert no. SGS-COC-2482
© 1996 Forest Stewardship Council

Printed and bound in Great Britain by
TJ International Ltd, Padstow, Cornwall

Contents

Acknowledgements

The editor and publishers wish to thank the following for permission to use copyright material.

Elsevier for the essays: David Sloan Wilson and Elliott Sober (1989), 'Reviving the Superorganism', *Journal of Theoretical Biology*, **136**, pp. 337–56. Copyright © 1989 Elsevier; Scott F. Gilbert, John M. Opitz and Rudolf A. Raff (1996), 'Resynthesizing Evolutionary and Developmental Biology', *Developmental Biology*, **173**, pp. 357–72. Copyright © 1996 Elsevier.

John Wiley and Sons for the essays: Peter Godfrey-Smith (1994), 'A Modern History Theory of Functions', *Noûs*, **28**, pp. 344–62. Copyright © 1994 John Wiley and Sons; Michael J. Wade and Susan Kalisz (1990), 'The Causes of Natural Selection', *Evolution*, **44**, pp. 1947–55. Copyright © 1990 John Wiley and Sons.

The Journal of Philosophy, P. E. Griffiths and R. Gray for the essay: Paul E. Griffiths and Russell Gray (1994), 'Developmental Systems and Evolutionary Explanation', *Journal of Philosophy*, **91**, pp. 277–304. Copyright © 1994 The Journal of Philosophy.

Oxford University Press for the essay: Frédéric Bouchard and Alex Rosenberg (2004), 'Fitness, Probability and the Principles of Natural Selection', *British Journal for the Philosophy of Science*, **55**, pp. 693–712. Copyright © 2004 Oxford University Press.

Springer for the essays: Ron Amundson and George V. Lauder (1994), 'Function without Purpose: The Uses of Causal Role Function in Evolutionary Biology', *Biology and Philosophy*, **9**, pp. 443–69. Copyright © 1994 Springer; Peter Godfrey-Smith (1999), 'Adaptationism and the Power of Selection', *Biology and Philosophy*, **14**, pp. 181–94. Copyright © 1994 Springer; Roberta L. Millstein (2002), 'Are Random Drift and Natural Selection Conceptually Distinct?', *Biology and Philosophy,* **17**, pp. 33–53. Copyright © 2002 Springer; Samir Okasha (2004), 'The "Averaging Fallacy" and the Levels of Selection', *Biology and Philosophy*, **19**, pp. 167–84. Copyright © 2004 Springer; Alex Rosenberg (1997), 'Reductionism Redux: Computing the Embryo', *Biology and Philosophy*, **12**, pp. 445–70. Copyright © 1997 Springer; Evelyn Fox Keller (1999), 'Understanding Development', *Biology and Philosophy*, **14**, pp**.** 321–30. Copyright © 1999 Springer; Manfred D. Laubichler and Günter P. Wagner (2001), 'How Molecular is Molecular Developmental Biology? A Reply to Alex Rosenberg's Reductionism Redux: Computing the Embryo', *Biology and Philosophy*, **16**, pp. 53–68. Copyright © 2001 Springer; Kim Sterelny, Kelly C. Smith and Michael Dickison (1996), 'The Extended Replicator', *Biology and Philosophy*, **11**, pp. 377–403. Copyright © 1996 Springer.

The University of Chicago Press for the essays: Karen Neander (1991), 'Functions as Selected Effects: The Conceptual Analyst's Defense', *Philosophy of Science*, **58**, pp. 168–84. Copyright

Series Preface

The theory of evolution is one of science's great achievements. Though to those outside science, it may seem that the theory is controversial, within science there is no controversy at all about its basic form. Moreover, the theory of evolution plays a pivotal role in guiding new research. 'Nothing in biology makes sense except in the light of evolution', Theodosius Dobzhansky famously wrote; the theory of evolution unifies disparate subfields of biology and generates testable predictions for each. The success of the theory and its explanatory fecundity for biology cannot be doubted. But might the theory also be capable of illuminating phenomena outside the direct purview of biology?

The volumes in this series are dedicated to exploring this question. They bring together some of the best writings of the past two decades which explore the relevance of evolution and evolutionarily-inspired thought to arenas of human life beyond the merely biological. Volumes focus on whether it is productive and illuminating to attempt to understand our most distinctive achievements and our most intimate features as evolved phenomena. Is the content of moral systems explained by evolution? To what extent are the processes of selection and reproduction that explain changes in gene frequencies also at work in explaining the reproduction of ideas? Can evolution shed light on why we think as we do, perceive as we do, even feel as we do? Might even our idea of God – and perhaps with it the perennial temptation to reject evolution in the name of religion – be explained by evolutionary thought?

Answering these questions requires not only a detailed grasp of the phenomena we aim to explain – the contours of religious thought, the features of morality, and so on – but also an understanding of the theory we aim to apply to the field. Though the theory of evolution is not itself controversial within science, there are lively controversies about its details. One volume of this theory is devoted to writings which illuminate these controversies and deepen our understanding of the mechanisms of evolution. It is only if we have an appreciation of how evolution works that we can begin to assess attempts to extend its reach to culture, to the mind, to morality and to religion.

The volumes are edited by experts in the philosophy of biology and include sensitive and thoughtful discussions of the material they contain. Naturally, in selecting the papers for inclusion, and given the large amount of high quality thought on the philosophy of biology, and on each of the topics covered by these volumes, it was necessary to make some hard choices. Each editor has chosen to focus on particular controversies within the field covered by their volume; on each topic, a range of views is canvassed (including the views of those who deny that evolution can contribute much to the understanding of non-biological features of human beings).

Evolution is our story; in coming to understand it, we come to understand ourselves. Readers of these volumes should be left with a deepened appreciation for the power and ambition of evolutionary thought, and with a greater understanding of what it means to be an evolved being.

NEIL LEVY
Florey Neuroscience Institutes, Australia and University of Oxford, UK

Introduction

The philosophy of evolutionary biology is a subdiscipline of the philosophy of science that focuses on a set of issues that emerged in the wake of Charles Darwin's *The Origin of Species* (1859). Most readers are undoubtedly familiar with Darwinian theory in broad outline if not in detail. At the very least, everyone will be acquainted with the two most basic tenets of evolutionary theory. The first principle states that all organisms are the products of descent with modification. Darwin argued that this principle provides the best explanation for what he called 'the unity of type', or the presence of shared anatomical and behavioural features among diverse organisms occupying different habitats. Darwin's second principle identifies one of the major causes of evolutionary change. The principle of natural selection states that populations will adapt to their environments whenever there are heritable phenotypic differences among organisms that differentially impact the reproductive fitness of those organisms. Darwin noted that an analogous process, artificial selection, can produce rapid evolutionary change within domestic plants and animals. He argued that the same kind of process operating in nature over extended periods explains why organisms are often so well 'designed' for survival in their environments.

Darwin's two principles have become so widely received that the non-specialist might be excused for thinking that evolutionary theory is no longer a controversial subject. Evolutionary biologists have gathered mountains of evidence in favour of both common descent and natural selection. Evolutionary biology has matured into an established science with its own textbooks, conferences and university courses. One might legitimately ask what sort of *philosophy* is there left to do in evolutionary biology?

In answering this question it is tempting to follow the lead of philosopher Alfred North Whitehead, who once described his discipline as 'a series of footnotes' to one of its greatest masters.

> The safest general characterization of the European philosophical tradition is that it consists of a series of footnotes to Plato. I do not mean the systematic scheme of thought which scholars have doubtfully extracted from his writings. I allude to the wealth of general ideas scattered through them. (1929, p. 39)

Just as Plato set the agenda for European philosophy over the last two millennia, so has Darwin laid the foundation for the philosophy of evolutionary biology for the foreseeable future. Although evolutionary theory has made significant progress since Darwin, many of the issues that he raised remain unresolved. The essays in this volume are organized around five such issues that have been hotly debated since Darwin first identified them.

The subject of Part I concerns the status of 'design talk' in biology. It is quite natural to interpret the parts of organisms as if they were the products of design. For example, a peacock's tail is said to have the function of attracting mates; the purpose of a zebra's stripes is to make it difficult for visual predators to pick one out of the herd. Such ascriptions of purpose made sense in the tradition of Natural Theology that predates Darwinism, where organisms

were assumed to be products of an intelligent designer. However, in viewing organisms as the outcome of 'blind' causal/mechanical processes Darwin did away with the literal idea of natural design. Yet Darwin and those following him continued to employ the language of purpose in describing organisms and their parts. This practice stands in contrast to other disciplines, like chemistry and physics, where the discovery of causal/mechanical explanations led to the abandonment of purpose-talk. Why then do biologists carry on with this language? Does the ontology of biological science differ from that of the physical sciences in containing mind-independent functions? If so, how might those functions be grounded within a naturalistic framework? The essays in Part I consider several attempts to provide a naturalistic account of biological function.

Part II concerns the relative strength of natural selection compared to other evolutionary processes. Darwin recognized that not all traits are best explained as the products of recent selection. Some traits are better explained in terms of the principle of descent. For example, many of the emotional expressions found in humans are also present in other primates. Darwin (1872) reasoned that the best explanation for why humans possess these expressions is because they were inherited from a common ancestor, not because they were selectively advantageous to our species. Critics of the 'adaptationist programme' in evolutionary biology argue that there is a systematic bias favouring adaptationist hypotheses over such alternative proposals. The problem is that adaptationist hypotheses are extremely easy to generate and at the same time difficult to test. It is often difficult to determine when a trait first evolved, what function (if any) it served in ancestral environments and whether it has undergone recent selection. The essays in Part II address these and other related issues surrounding the identification of adaptations and the relative strength of selection.

A third set of issues concerns the logical structure of evolutionary theory. Darwin's critics have long argued that a central theoretical concept in his theory, the concept of fitness, is devoid of empirical content. This objection stems from the observation that Darwinians often equate the fitness of an organism with number of offspring produced. As innocent looking as this definition may seem, it potentially turns the principle of natural selection into an empty tautology (see below). Philosophers have attempted to provide a careful formulation of the principle of natural selection that avoids the tautology problem. This exercise has generated a set of closely related questions. For instance, is natural selection a statistical theory like thermodynamics or a dynamical force theory like Newtonian mechanics? And if evolutionary theory is statistical in nature, then how can the processes of selection and genetic drift be distinguished from one another? Such issues relating to the nature of fitness and the logical structure of evolutionary theory are addressed in Part III.

Part IV addresses the 'levels of selection' debate. In *Origin of Species* Darwin recognized cooperative behaviour, in particular the 'self-sacrificial' behaviour of colonial insects, as 'the one special difficulty, which at first appeared to me insuperable, and actually fatal to my theory' (1859, p. 231). The problem as Darwin saw it was that selection acting on individual organisms couldn't possibly give rise to a trait like sterility which appears to reduce the fitness of an organism to zero. In his later writings Darwin (1871, p. 166) appealed to selection acting at the level of the group as a potential explanation for altruistic behaviour in humans. However, the idea of group selection has a controversial history in evolutionary biology. The essays appearing in Part IV document the recent trend in evolutionary biology to reconsider group selection after a long period of neglect.

Part V addresses the relationship between evolutionary theory and developmental biology. The field of developmental biology has only recently moved beyond the study of embryology as it existed in Darwin's day. As a result, a new set of 'developmentalist challenges' to the traditional Darwinian framework are currently attracting attention. Part V identifies three such challenges to received views about adaptationism, reductionism and the challenge to the idea of a genetic programme.

The remainder of this introduction is divided into five sections, each one a critical introduction to the corresponding part of the volume. But before turning to those discussions a few preliminary qualifications are in order.

First, it should be noted that the five topics addressed here do not exhaust the subject matter of the philosophy of evolutionary biology. In particular, no direct mention is made of the problem of species classification, which is as much an issue in ecology as it is in evolution. Interested readers will find several valuable collected volumes on this subject (Ereshefsky, 1992; Claridge *et al.*, 1997; Wilson, 1999). Nor does this volume address such interesting questions as whether there is direction or progress in evolution, or whether 'macro-evolutionary' patterns can be explained exclusively in terms of 'micro-evolutionary' theory. Other closely related subjects, including cultural evolution, evolutionary psychology, the moral implications of Darwinian theory and the relationship between evolution and religion, are represented in the companion volumes in this series. For a more general introduction to these and other issues in the philosophy of biology, Sober (1999), Sterelny and Griffiths (1999), and Rosenberg and McShea (2006) offer accessible introductions.

Finally, for any one of the topics addressed in this volume there are simply too many good essays available to choose from. An effort was made to select publications appearing in peer-reviewed journals (no collected volumes or book chapters) that have not been reprinted elsewhere. While this criterion helped narrow down the list somewhat, the selection process was still extremely difficult and many noteworthy publications had to be set aside for reprint elsewhere.

Function and Purpose in Biology

To illustrate how naturally 'function-talk' arises in biology, imagine working as a taxonomist for the London Natural History Museum in 1798. This was the year that the first platypus arrived – dried and stuffed – from Australia. Here was a creature with the most unlikely assemblage of traits: dense fur, webbed feet, horny spurs on its hind legs (the first specimen was male) and a rubbery bill resembling that of a duck. After realizing that it hadn't been stitched together, your first reaction would have surely been to ponder the functions of these traits. Why such dense fur in a subtropical species? Why the spurs? For what purpose might a mammal require a beak? A revealing fact about our approach to the biological world is that we assume such questions have answers. Upon discovering that platypuses use their bills to forage for insects in muddy river bottoms we infer that the bill has this application as its function. Notice that we are generally not willing to assign functions to inanimate objects in nature. Darwin was quite interested in how coral atolls are formed. He proposed that this process is driven by underwater volcanic activity. But at no point did Darwin assume that volcanoes have atoll formation as their function. Yet, according to Darwinian theory both platypus bills and coral atolls are the products of 'blind' causal/mechanical processes. So

what, if anything, justifies this discrepancy in the ways that we regard biological versus other physical objects?

Philosophers have developed four general accounts of biological function (but see Wouters, 2005 for a more extensive list). The first equates functions with an organism's selective history. More precisely, to say that 'the function of trait T is F' means, on this view, that T was naturally selected to perform F. This *selected effects* account of function has its origin in the work of Larry Wright (1973) and has been developed and defended by Peter Godfrey-Smith (1993, and Chapter 2, this volume), Ruth Millikan (1989), Karen Neander (2002, and Chapter 1, this volume) and Alex Rosenberg (1985), among others. The second tradition is the dispositional account of function. On this view, 'the function of T is F' means that T is disposed to F, and F-ing would promote the survival and reproductive potential of T's bearer. This view is associated with John Bigelow and Robert Pargetter (1987) and a more nuanced version has been developed by Arno Wouters (1995, 2003). According to the third *causal role* tradition, 'the function of T is F' means (roughly) that by ascribing F to T scientists gain explanatory insights into the mechanical workings of organisms that possess T. This view originated with Robert Cummins (1975) and has been developed by Ron Amundson and George Lauder (Chapter 3, this volume). The fourth view identifies functions that are implicit in the ways that biologists classify organisms and their parts by common descent. This view has been explored recently by Paul Griffiths (Chapter 4, this volume) among others. The remainder of this section considers each of these views in relation to the essays reprinted in this volume.

Biological Functions as Selected Effects (SE Theory)

In Chapter 1, 'Functions as Selected Effects: A Conceptual Analyst's Defense', Karen Neander defends the selected effects theory of function against several objections. One objection notes that prior to the discovery of natural selection biologists routinely ascribed functions to traits. Furthermore, non-biologists who are unfamiliar with Darwinian theory can easily classify traits as functional. So, there appears to be a sense of biological function that does not make reference to selective history. Therefore, it is argued, SE theory does not provide an adequate definition of this term.

Neander responds to the first objection by arguing that within biology the meaning of 'function' has changed over time. Prior to Darwin, functions were thought to be the products of divine creation; after Darwin, this notion was displaced by the SE concept. As for the second objection, Neander allows that non-experts might possess a concept of function distinct from that employed by professional biologists. However, she is primarily interested in what experts mean by this term. Her contention is that *within the field of biology* all legitimate functional ascriptions involve the identification of some selected effect.

Neander goes on to consider a more serious challenge to her position. A third objection to SE theory states that even contemporary biologists ascribe functions to traits without alluding to selective history. This objection is sometimes motivated by the use of some rather extravagant thought examples. Neander, in outlining this objection, imagines a case where a sort of 'winged lion' suddenly pops into existence through the spontaneous assemblage of molecules. As difficult as this scenario is to imagine, it is possible to anticipate how biologists would investigate such a creature. One of the things that biologists would most likely do is

classify the large, folded protrusions on its back as wings. However, notice that even this judgement involves the application of a functional concept. To classify something as a wing, it is argued, implies that its purpose is to generate flight. Such a classification would make sense to biologists, the critic of SE theory maintains, even though they are aware that this organism's wings were not the product of natural selection. Therefore, biological practice apparently involves a concept of function that is not captured by SE theory.

Neander responds to this objection by flatly denying that the relevant biological experts would be capable of assigning functions to this organism's traits:

> I contend that we could not reliably place them [lion 'wings'] in any category until we knew or could infer the lions' history. And if we were to somehow discover that the lions had no history, and were the result of an accidental and freak collision of atoms, they would definitely not belong in any of our familiar functional categories … Without history the usual biological/functional norms do not apply. (p. 15)

This is a controversial claim. In a moment we shall consider some actual cases where biologists appear to be making functional ascriptions that explicitly do not refer to selective history. If such practices are widespread in biology, then Neander is simply mistaken that biologists rely exclusively on an SE concept of function.

It is also important to note that by 'history' Neander means *selective* history. In the previous passage Neander fails to distinguish between two different historical modes of categorization. In addition to selective history biologists also categorize traits by their homology or descent (Griffiths, Chapter 4, this volume). Consider the front appendages of a typical penguin. Biologists classify these structures as wings, not paddles, even though penguins use them for swimming instead of flight. In classifying these structures as wings biologists are alluding to their history, but not their recent selective history. It is the fact that these structures share a common ancestry with the front appendages of other birds that makes them wings and not paddles. Categorization by ancestry, or homology, is a form of historical categorization that (arguably) does not presuppose anything about the selected function of the trait under investigation. This objection will be discussed in more detail momentarily.

In other cases, biologists appear to ascribe functions to traits while making no reference to history whatsoever. As Peter Godfrey-Smith explains in Chapter 2, the pioneering animal behaviour researcher Nico Tinbergen famously identified four different questions that one might ask of a trait: (1) What are the physiological mechanisms that produce the trait? (2) How does the trait develop? (3) What is the evolutionary history of the trait? And (4) what is the trait's current function? It is the difference between the third and the fourth questions that is of potential relevance to Neander's position. According to Tinbergen (1963), identifying the current function of a trait involves examining how it contributes to an organism's current survival and reproduction. It is entirely possible that a trait is currently used for purposes that differ from the functions for which it was selected (see Part II on adaptationism). So, contrary to the claim that all biological functions are selected effects, Tinbergen's notion of current function appears to be logically independent of an organism's evolutionary history.

There are three general ways of accommodating Tinbergen's notion of current function. First, one might adopt the pluralist position that there are multiple senses of function at play in biology. The question then is what are those distinct senses and how do they relate to

one another (see below). Alternatively, one might challenge the accuracy of Tinbergen's distinction. Perhaps he is wrong in assuming that there is a difference between current functions and evolutionary functions. This is the position that Neander adopts (2002; Neander and Rosenberg, forthcoming). Alternatively, some argue that there is an even more basic sense of function that Tinbergen's different categories can be reduced to. Let us now consider a version of the latter approach, which attempts to define functions in terms of an organism's biological goals.

Dispositional Functions

According to the dispositional account of functions, the function of a trait is determined by its *capacity* to promote the survival and reproduction of its bearer (regardless of whether it was actually selected to do so). One advantage of the dispositional view is that it provides a basis for ascribing functions to traits that have no selective history or to traits whose history is unknown. Even the wings of an 'instant lion' could legitimately be ascribed a function on this view, provided that they promote its goals of survival and reproduction. (Note that the organism need not consciously entertain these goals. Even a mindless bacterium has survival and reproduction as its 'goal' in this sense.) The dispositional account can potentially unify Tinbergen's different senses of function. On this view, the evolutionary function of a trait is the disposition it previously had to promote survival and reproduction in ancestral environments. By contrast, the current function of a trait is its tendency to promote survival and reproduction in the organism's current and foreseeable environments.

As plausible as it might seem, the dispositional theory faces serious challenges. As Godfrey-Smith explains an organism's disposition to survive and reproduce is always relative to some environment or other. There is no such thing as a trait that enhances survival and reproduction s*impliciter*. Thus, a dispositional theorist must explain which environments are the relevant ones for determining the function of a trait. John Bigelow and Robert Pargetter (1987), who defend this view, claim that the relevant context for identifying functions is an organism's 'natural habitat'. But what counts as natural here? If by 'natural' one means the environments in which the organism has been selected to survive and reproduce, the dispositional account collapses into the SE view.

Perhaps Tinbergen's notion of current function can be brought to the rescue. Tinbergen equates the current function of a trait with its utility in the environment currently occupied by the organism and with environments that are 'not too dissimilar' from the current one. However, Godfrey-Smith objects that even this qualification leaves the dispositional view inadequately constrained. For example, imagine an organism with a heart that is so malformed it can only survive on remote islands with no predators and plentiful resources. Do we want to say that its heart is functioning normally just because the organism *could* survive under these exceptional circumstances? How many possible cases like this are there? One advantage of SE theory is that it placed fairly stringent restrictions on the range of environments under consideration – just those environments inhabited by the organism's ancestors count. Godfrey-Smith offers a further amendment to the SE view, suggesting that the set of relevant environments should be restricted to those which the organism has encountered in its *recent* selective history. This is a welcome qualification in that it provides a basis for saying that vestigial traits, like the appendix, do not have a function because they have not been recently selected for.

Causal Role (CR) Functions

The notion of a causal role function is associated with the philosopher Robert Cummins, who was interested in a particular explanatory strategy that is employed in the life sciences. To *functionally analyse* some capacity C of a given system S involves describing how the components of S and their individual capacities k contribute to the production of C. For example, the platypus bill has prey detection as one of its higher level capacities (C). In performing this capacity it employs several types of cell involved in mechanical and electro-reception.

One type is a mechanoreceptor called a push rod because it is composed of a rod-like pillar of epithelium that crosses the whole epithelial thickness. It has an array of sensory neurons at its base that would function to signal displacements of the rod produced by impulses delivered to the free end (Pettigrew, 1999).

Here, a type of sensory neuron is identified as having the function (k) of signalling displacement in another functionally individuated component, the push rod. These lower level functions are the ones that Cummins is interested in. Functional explanations of this type seem to answer a different sort of question than SE functions. Instead of explaining why a certain structure exists, CR functions explain how it works.

In Chapter 3, 'Function Without Purpose: The Uses of Causal Role Function in Evolutionary Biology', Ron Amundson and George Lauder argue that CR functions are identified in a wide range of biological disciplines. For example, in the field of functional morphology it is standard practice to analyse traits like the vertebrate jaw into its structural and functional components. Importantly, Amundson and Lauder maintain that these functions are distinct in kind from SE functions, because their analysis does not hinge on facts about an organism's selective history:

> While the decision to analyze the jaw may have been motivated by a knowledge of its biological role … that knowledge plays no part in the analysis itself. The biological role of the jaw system does not influence the function which the component muscle is analyzed to have … Even more remote from functional analysis are hypotheses regarding selective pressures, or any other explanations of why the jaw has its present capacities. (p. 49)

Let's grant Amundson and Lauder the claim that functional anatomists do not make reference, either implicit or explicit, to selected effects when functionally analysing traits. The question remains, why view these capacities as a type of *function* as opposed to mere capacities? We have already seen that the ascription of a function to a trait implies some kind of goal or purpose. SE theorists cash out the relevant sense of purpose in terms of the organism's selective history. Dispositional theorists equate the purpose of a trait with the 'forward looking' goal of promoting survival and reproduction. However, the causal role theorist appears to be equating the purpose of a trait with the goals of an investigator. It is the functional anatomist's curiosity about 'how it works' that fixes the lower level functions that she identifies. This view opens the door to some rather odd sounding functions. Suppose that an investigator applies Cummins' analysis to coral atoll formation. Suddenly, underwater volcanoes have atoll formation as one of their functions. Likewise, to the geologist, tectonic plate movements can have the destruction of buildings as one of their functions. Such functional ascriptions sound odd, and (arguably) are not legitimate objects of scientific investigation.

There are at least three ways for the CR theorist to address this objection. The first strategy places limits on the sorts of things that lend themselves to functional analysis. Cummins states that a functional analysis of some capacity C is 'interesting and informative' only if (1) the analysing capacities k are simpler and different from C, and if (2) the system S possesses a complex degree of organization. Amundson and Lauder appeal to this qualification in an attempt to avoid the sorts of counterexamples mentioned earlier: 'the geological structures which result in earthquakes might be complex, but the "organization" of these structures *vis a vis* their explanation of the capacity of the earth to quake is not' (p. 50). Unfortunately, however, neither Cummins nor Amundson and Lauder elaborate on what 'complex organization' amounts to. This is a bit of a worry, since both of these terms are notoriously vague. Amundson and Lauder further argue that in practice CR functions will be restricted to biological systems, because natural selection is the only process capable of producing systems in nature that are sufficiently complex and organized to call for a functional analysis. But even this claim is questionable. Coral atolls display a degree of physical organization that is arguably quite complex. So, the CR account seems unable to explain why biological systems are the only ones in nature that possess functions.

A second strategy for controlling the spread of CR functions in nature is to leave it up to scientists to decide which functions are legitimate and which ones are not. On this view there might be no single criterion that applies across different scientific contexts for deciding whether a given capacity is a function or not. But this does not mean that scientists are arbitrarily identifying CR functions. Instead, it might turn out that different groups of scientists employ distinct criteria for identifying CR functions, all of which are strict. In this case the concept of a CR function is not necessarily too permissive.

A third sort of reply involves biting the bullet: some inanimate systems really do have functions, we just aren't used to thinking of them this way. This is the approach that Godfrey-Smith favours: 'once Cummins' functions have been recognized and the explanatory mode which utilizes them has been understood, they should be allowed to roam freely, even into the farthest periphery' (1993, p. 200). So volcanoes might in fact have atoll formation as one of their functions and earthquakes potentially have the purpose of knocking down buildings. According to Godfrey-Smith, the question is not whether such ascriptions sound bizarre. The important question, he claims, is whether such functional ascriptions generate scientific insights. If ascribing functions to volcanoes leads to discoveries about how they work then, intuitions aside, volcanoes have functions.

Classification by Homology and 'Boundary Disputes'

It would appear that the two primary concepts of function identified so far (CR and SE) play distinct but complementary explanatory roles in biology. One might describe their respective roles as follows. SE functions explain why a certain trait exists, while CR functions explain how a given trait works. However, we have already touched on an example that raises problems with the first half of this statement. Recall that penguin appendages are classified as wings even though they have paddling, not flight, as their recent adaptive function. Earlier I argued that classifying these structures as wings involves an appeal to their ancestry. It is because penguin wings are ancestrally related to the front appendages of flying birds that they qualify as wings and not paddles. Categorization in this case appears to be by common descent, not

selective history. Nor does this mode of categorization appear to be identifying CR functions, because classification by homology need not occur in the context of Cummins-style functional analysis. Hence, Paul Griffiths argues in Chapter 4 that classification by homology can involve a distinct kind of functional ascription, not reducible to either SE or CR functions.

Neander (2002) and Neander and Rosenberg (forthcoming) object to this analysis. They argue that classification by homology ultimately presupposes classification by selective history. This point is well illustrated by the penguin example. Although penguins do not currently use their wings for flight, these structures descend from a trait (the wing of an ancestral bird) that had flight as its selected function. Had penguin wings descended from a trait with a different function, they would not be classified as wings. Thus, Neander and Rosenberg argue that classification by homology always presupposes classification by selected effect.

In response, one might argue that there are many examples of classification by homology that do not trace back to some prior selective function. Paul Griffiths illustrates why this is the case. He explains that there are two distinct ways in which different traits can be identified as homologous. *Taxic homology* is a relationship among traits that share the same gross morphology because of their common descent. For example, the penguin wing has the same gross morphological characteristics (for example bone structure) as the wings of all birds. *Serial homology* is a relationship among the repeated, often differentiated structures of a single organism, defined by their developmental or genetic origin. For example, the vertebrae in the human spine are serially homologous, because they involve the repetition of a single developmental programme. The important point here is that serial homology provides a basis for classifying traits into types that is based on developmental history, not selective history. The reason that biologists classify neck vertebrae into the same category as vertebrae in the lower back is because they are serially homologous. If, instead, neck bones developed from a protrusion at the base of the skull that later fused with the spinal column, then anatomists would classify these as two distinct traits. Griffiths notes that classification by serial homology is logically independent of classification by selective history. Two serially homologous traits might fuse together, as in the case just imagined, to form a single functional unit. Alternatively, a single serially homologous structure might be repeated in different places within the same organism to produce structures with distinct adaptive functions. Therefore, he claims, not all biological classification is ultimately based on selected effects. Griffiths goes on to note that some traits have no selective history to speak of; they are 'spandrels' (see below) or developmental by-products of other traits. However, all traits have a homology and can be categorized accordingly. Therefore, Griffiths concludes that classification by homology is more basic or fundamental to biology than classification by selected effects.

Notice that Griffiths' final move turns on a question about the relative strength of natural selection compared to other processes like 'phylogenetic inertia' and developmental constraint. If natural selection is a relatively weak force, then there will be many traits that are classified by homology but which are lacking a selective function. In this case, the argument that classification by homology is more fundamental than classification by homology has some bite. However, if selection is a fairly pervasive force in nature then it becomes difficult, at least in practice, to tease apart selective functions from CR or homology functions. This brings us to the next topic in this volume: how 'important' is natural selection compared to other biological processes, and how can the effects of selection on organisms be identified?

Adaptationism, Optimality and Adaptive Co-variation

For over forty years biologists and philosophers have debated the sorts of evidence required to justify adaptationist hypotheses.[1] In 1966 the evolutionary biologist George Williams placed the concept of adaptation under scrutiny, stating that 'evolutionary adaptation is a special and onerous concept that should not be used unnecessarily, and an effect should not be called a function unless it is clearly produced by design and not by chance' (1966, vii). Williams appreciated that adaptationist hypotheses are easy to generate but difficult to test. They are easy to generate because it is typically possible to come up with indefinitely many scenarios under which a given trait might have been selectively favoured. The platypus bill, for example, might have been helpful in flipping over rivals during combat, it might have been attractive to mates, or maybe the bill offered hydrodynamic advantages in evading predators. Such hypotheses are difficult to test partly because they make claims about the distant past. Rarely does one know exactly which selective pressures were present in ancestral environments. Nor is it always possible to identify the evolutionary precursors for a trait or the amount of available variation in ancestral populations. Without adequate supporting evidence, Williams cautioned, adaptationist hypotheses should be regarded with scepticism.

This conservative stance was vigorously defended in 1979 by Steven Jay Gould and Richard Lewontin in an unusually titled essay, 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme'. Gould and Lewontin accused many of their fellow evolutionary biologists of flouting Williams' cautionary principle: 'faced with an organism, [they] tend to break it into parts and tell adaptive stories as if trade-offs among competing, well-designed parts were the only constraint upon perfection for each trait'(1979, p. 78). The central problem with this strategy of 'atomizing' an organism into independently optimized parts is that it ignores developmental and phylogenetic constraints on biological form. Gould and Lewontin illustrated this point with an architectural analogy. As they explain, spandrels are the triangular structures formed at the intersection between two adjacent arches and have no direct architectural function. However, in cathedrals like the Basilica of San Marco, they are typically adorned with religious iconography that is smoothly incorporated into the dome itself. One might naturally assume, therefore, that spandrels serve a decorative function. But this interpretation gets things the wrong way around. Spandrels came first and the iconography was an afterthought. This analogy is meant to show that just as one requires information about architectural constraints to identify the function of spandrels, so must one understand development and phylogeny to determine whether a trait is an adaptation for some purpose.

Gould and Lewontin accused their fellow evolutionists of subscribing to an 'adaptationist programme' that systematically ignores developmental constraints, genetic drift, mutation,

[1] I make a standard terminological distinction between 'adaptive' and 'adaptation'. To say that a trait is adaptive means that it can be viewed as having a function in some environment, but not necessarily that it was selected for that function. For example, the nose is adaptive for holding up spectacles. To identify a trait as an adaptation is to make a historical claim: that it was 'shaped' by natural selection for a particular function. For example, the vertebrate retina is an adaptation for visual perception. Generally speaking it is misleading to ask whether some trait is an adaptation without specifying what it is potentially an adaptation for. This is misleading because a given structure can have different functions in different organisms at different times.

genetic linkage and other non-selectionist influences on biological form. One of the more serious charges against the adaptationist programme was that it is unfalsifiable in practice: 'If one adaptive argument fails, try another … assume that another must exist … [or] attribute failure to imperfect understanding of where the organism lives and what it does' (Gould and Lewontin, 1979, p. 586). Gould and Lewontin caricatured this methodology as a 'Panglossian paradigm'. Like Voltaire's fictional Dr Pangloss who remained convinced that events always turn out for the best, the adaptationist clings to her cherished assumption that each trait is an optimal solution to some adaptive problem or other. Gould and Lewontin identify several alternative explanations for why a trait might appear adaptive without being an adaptation. For example, some phenotypically plastic traits are the result of an organism's accommodation to its environment over its lifetime. Other adaptive traits are the product of social transmission or individual learning. Still others merely appear adaptive, but are in fact the developmental by-products of other traits that have been selectively favoured. Until each of these alternative explanations have been ruled out, Gould and Lewontin argue, one cannot legitimately infer adaptation from apparent adaptiveness.

This critique of adaptationism has generated considerable controversy, beginning with Ernst Mayr (1983) who accused Gould and Lewontin of constructing a straw man. Practising evolutionists do not, as these authors claim, 'use *consistency* with natural selection as the sole criterion and consider their work done when they concoct a plausible story' (Gould and Lewontin, 1979, p. 588). In fact, Mayr argued, adaptationists apply rigorous experimental techniques to test –and sometimes falsify – their hypotheses. Mayr notes that the methodology ridiculed by Gould and Lewontin, of trying another hypothesis when the first one fails, is standard practice in disciplines like chemistry, physics, geology and archaeology. Mayr also notes that historical sciences face an 'epistemological dilemma': how to prove that a particular hypothesis is superior to the alternatives given limited access to the past. Few adaptationist hypotheses can be established with certainty, even the most elaborate trait might have arisen by chance. So adaptationist hypotheses are always defeasible in light of further information. But if adaptationism is to be condemned on these grounds, Mayr argues, one must adopt the same stance toward all historical sciences.

In reply, Gould and Lewontin might argue that Mayr is missing the point of the unfalsifiability objection. The falsification of particular adaptationist hypotheses does not threaten the over-arching assumption that a trait is some kind of adaptation or other. If adaptationist hypotheses can be churned out ad infinitum, then other non-adaptationist alternatives will be systematically ignored. As Gould and Lewontin put it, 'the range of adaptive stories is as wide as our minds are fertile' (1979, p. 587), and there appears to be no direct analogue to the adaptationist perspective in other historical sciences.

Arguably, this reply to Mayr is easily overstated. For one thing, Gould and Lewontin fail to distinguish between plausible and implausible adaptationist scenarios. Often, after three or four speculations about the selected function of a trait an adaptationist will run out of *plausible* alternatives. At that point any sane adaptationist will begin considering non-adaptationist alternatives. It is true that there is no well-defined point at which one should stop entertaining adaptationist hypotheses and begin considering alternatives: there is no precise 'stopping rule' for adaptationism. However, it doesn't follow that adaptationists are destined to a sort of tunnel vision that restricts non-adaptationist considerations from view.

Mayr offers an additional argument in favour of adaptationism. He claims that the best way to approach biological systems is by generating and testing adaptationist hypotheses *as if* all traits are the products of selection. The justification for adopting the 'design stance' (as Dennett, 1995 calls it) is that this approach is unparalleled in its utility as a scientific hypothesis generator. Interestingly, this argument views the fecundity of adaptationist thinking as virtue, not a vice. If adaptationist thinking is an unusually fruitful source of new hypotheses, then why not exploit the adaptationist programme for this purpose?

The debate over how to carry out the adaptationist programme remains one of the most controversial subjects in the philosophy of evolutionary biology. Recently, Peter Godfrey-Smith (Chapter 5) has distinguished three, logically distinct forms of adaptationism. This distinction has been particularly helpful in clarifying various threads of the adaptationist debate.

One form of adaptationism is what Godfrey-Smith calls explanatory adaptationism. On this view the phenomenon of 'design-fit' (the adaptedness of an organism to its environment) is the most interesting and important problem in biology. In the words of Richard Dawkins,

> Of course, large quantities of evolutionary change may be non-adaptive, in which case these alternative theories may well be important in parts of evolution, but only in the boring parts of evolution, not the parts concerned with what is special about life as opposed to non-life. (1986, p. 303)

Godfrey-Smith objects to this thesis, arguing that explanatory adaptationism comes down to a matter of individual or partisan interest. For example, questions of design-fit are particularly important to philosophers and theologians interested in the design argument for a supernatural creator. Some regard the principle of natural selection with special reverence because it provides a way of replacing teleological explanations with causal-mechanical ones. However, as Godfrey-Smith notes, these are personal reasons for privileging selectionist explanations, they do not extend to biology at large.

The second form of adaptationism identified by Godfrey-Smith is methodological adaptationism. On this view the concept of adaptation is a useful 'organizing principle' for conducting biological research. Mayr's heuristic argument (above) appears to be in line with methodological adaptationism, so is the idea that adaptationist thinking is a useful 'hypothesis generator'. Provided that methodological adaptationism is not pursued at the neglect of alternative, non-adaptationist research strategies there appears to be no reason to reject this thesis. Much more controversial is the third form of adaptationism. Godfrey-Smith describes empirical adaptationism as the view that,

> 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. (2001, p. 336)

Although this is an empirical claim about the importance of selection compared to other factors, opinions vary about how to test this thesis.

One strategy for testing empirical adaptationism involves the use of optimality analysis to identify highly adaptive traits. Optimality analysis is a modelling technique for estimating which of several competing traits offers the most efficient or 'optimal' solution to an adaptive

problem. The construction of an optimality model begins with an estimate of the costs and benefits associated with the focal trait. For example, suppose that the platypus bill requires C units of energy per centimetre to develop, and that each centimetre of bill length increases prey intake by B energy units up to some threshold. Based on these estimates there will always be some point at which B-C is maximized. Perhaps 14 cm is the optimal bill length after which point the return on energetic investment diminishes. A key assumption underlying this approach is that if selection is the dominant force acting on a trait, then that trait will eventually reach its optimal value. So, if a representative sample of platypus bills shows little variation around the optimal 14 cm mark, it is a safe bet, optimality theorists argue, that selection has been the primary force responsible for shaping this trait.

In Chapter 6, 'Optimality Models and the Test of Adaptationism', Steven Orzack and Elliott Sober outline the logical structure of this approach in detail. They begin by identifying three adaptationist theses of increasing strength:

(U) Natural selection played some role in the evolution of some trait T.
(I) Natural selection was an important cause of the evolution of T.
(O) Natural selection is a sufficient explanation of the evolution of T, and T is locally optimal.

Propositions U and I are fairly moderate and uncontroversial theses. Orzak and Sober define adaptationism as the belief that O is true for most traits (cf. Godfrey-Smith's description of empirical adaptationism). They suggest that O is true if there is a 'censored model' that predicts, to a degree of statistical accuracy, the distribution of T among members of a population. They explain that a censored model is one that ignores all factors besides selection. Optimality models fall into this category. The reason that a censored model can provide a sufficient explanation for T, they argue, is because in some cases other factors do not enhance the model's predictive power. So although non-selective factors might contribute to the evolution of T, if those factors offer no additional predictive value they can be ignored for explanatory purposes. Orzack and Sober propose that their version of adaptationism can be tested by undertaking an optimality analysis for a wide range of traits. If most of those traits can be explained by a censored model then, they argue, adaptationism is true.

One of the more controversial aspects of this proposal concerns the sort of evidence Orzack and Sober require for the confirmation of an optimality model. Not only must a population exhibit a statistically significant degree of fit to the model's predictions, but the population must exhibit no phenotypic variation in the optimal trait (unless predicted by the model). In Chapter 7 Robert Brandon and Mark Rausher argue that the second condition is too strong. Even traits that have been optimized by selection will exhibit some variation due to mutational and environmental influences on development. The invariance requirement also rules out the possibility of several optimal phenotypes in a single population. Furthermore, even if a trait exhibits more variation than is predicted by an optimality model, it doesn't follow that the trait is not an adaptation. The last point touches on a more general criticism raised by Brandon and Rausher. They argue that the concepts of optimality and adaptation are logically and biologically independent of one another. First, a trait's being optimal does not imply that it was shaped predominantly by selection. In some cases a trait becomes advantageous only after it has reached a certain frequency in the population, for example by drift. So there

are cases where a trait conforms with the predictions of a censored model, but where I (not O) is the appropriate explanation. Second, Brandon and Rausher argue that a trait's being predominantly shaped by selection does not entail that it is optimal. 'There are numerous situations in which natural selection can be the sole evolutionary process in operation but in which the locally optimal phenotype will not evolve to fixation' (p. 136). On these grounds Brandon and Rausher argue that Orzack and Sober's formulation of the optimality approach is inadequate (see Orzack and Sober, 1996 for a reply to their criticisms).

A second strategy for identifying adaptations takes some forms of individual variation as evidence in favour of adaptation. In Chapter 8 Michael Wade and Susan Kalisz propose to test adaptationist hypotheses by looking for adaptive co-variation among phenotypic traits of ecological factors that impact fitness. This strategy assumes that the intensity of selection on a given trait varies within a given population. For example, a population might experience varying degrees of predation pressure over its home range. If that population exhibits concomitant variation in some relevant trait, like running speed or skittishness, this is prima facie evidence that the trait is under selection pressure. Wade and Kalisz propose to test such hypotheses using manipulation experiments. For example, one might transplant individuals from a high to a low predator regime, and vice versa, while recording their survival rates over a period of time. If survival and reproduction rates change as predicted, this provides strong evidence in favour of a selective hypothesis. Some of the most convincing demonstrations of natural selection in wild populations have employed this approach (Endler, 1986; Breden and Wade, 1989). In cases where it is difficult or impossible to artificially manipulate the relevant variables, nature sometimes conducts the manipulation for us. Such 'natural experiments' are another important source of evidence for adaptation. However, the 'co-variation strategy' is of limited use when it comes to traits that show little inter-individual variation, traits that are no longer under selection pressure or traits that impact fitness in different ways now than they did in the past. Many purported human psychological adaptations are thought to have evolved under social and ecological conditions distinct from those of today (Barkow, Cosmides and Tooby, 1992). Without a careful reconstruction of ancestral environments the co-variation method cannot be applied to such cases.

Fitness, Drift and the Forces of Evolution

When Darwin first presented the principle of natural selection he drew on observations that would have been familiar to his readers. He noted that in all natural populations there is a 'struggle for existence' which is evidenced by the fact that only a fraction of the organisms born each generation survive to reproductive age. He also observed that within all populations there is variation in phenotypic traits that contribute differentially to those organisms' chances of survival and reproduction. Darwin further noted that parents tend to pass on phenotypic traits to their offspring. Putting these ideas together, Darwin reasoned that over successive generations of selection for the fittest individuals, a population becomes increasingly adapted to its environment. It is clear that Darwin took this thesis to apply generally to all organic life forms. Natural selection is a process that occurs throughout the biological world, not just within a few select species. The generality of Darwin's principle has suggested to many that natural selection is a law of nature that applies not only to biological populations on earth but potentially to any entities found anywhere in the universe. However, if the principle of

natural selection is a law of nature, it should be possible to state this principle in general terms. Such a general presentation would ideally reveal what all of the different instances of natural selection have in common. By analogy, Newton's three laws of motion revealed that the same processes governing the orbits of heavenly bodies also dictate the movement of objects here on earth. Likewise, a general statement of the principle of natural selection would potentially unify our understanding of adaptation across such diverse systems as plants, bacteria, animals and even human cultures.

However, the attempt to provide a general statement of the principle of natural selection has been fraught with difficulties. To get a sense of these issues it is helpful to begin with one of the most popular formulations of natural selection, which was presented by the evolutionary biologist Richard Lewontin (1970). Lewontin proposed that the principle of natural selection has three ingredients:

1. *Phenotypic variation*. Different individuals in a population have different morphologies, physiologies and behaviours.
2. *Differential fitness*. Different phenotypes have different rates of survival and reproduction in different environments.
3. *Fitness is heritable*. Offspring are more similar to their parents in phenotype than they are to the population mean.

Notice that the principle, thus stated, makes no ontological commitments regarding the kinds of entities that are capable of undergoing natural selection. Provided that a given population of entities satisfies these three conditions, those entities are capable of evolving and adapting.

However, Lewontin's three conditions are not jointly sufficient for evolution and adaptation. It is not *necessarily* the case that a population of entities exhibiting heritable phenotypic differences in fitness will therefore become increasingly adapted to its environment. For example, suppose that within a population of platypuses the longer an individual's bill, the more adept it is at capturing prey. Provided that there is heritable variation in this trait, one would predict that mean bill length will increase in this population. However, imagine that due to a freak occurrence all of the individuals with the 'long bill gene' are indiscriminately removed from the population before reproducing (perhaps because these individuals found their way into the traps of Natural History Museum collectors). In this scenario the conditions for natural selection are satisfied but the population does not become increasingly adapted: since the fittest organisms are removed from the population, mean bill length does not increase as expected. Biologists tend to conceptualize such events at the population level as a kind of drift, where the mean value of the population's phenotype (in this case average bill length) deviates or 'drifts' away from the predicted value. Any statement of the principle of natural selection must allow for the occurrence of drift, because such 'indiscriminate sampling events' occur reliably in nature. One way of accommodating drift into the principle of natural selection is to state the principle in probabilistic terms. Hence, a revised version of the principle of natural selection states that if a population of entities exhibits (1) phenotypic variation, (2) fitness differences and (3) heritability, then that population will *probably* become more adapted to its environment over successive generations.

We are now in a better position to appreciate some of the philosophical issues surrounding general formulations of Darwin's principle. The principle as it has just been stated contains two

particularly controversial terms. First, as was noted earlier, the concept of fitness is notoriously problematic. Although it is sometimes possible to identify what makes a particular organism more or less fit in a given environment (for example perhaps longer billed platypuses are relatively fit in most Australian streams), it is extremely difficult to say what property all relatively fit organisms have in common. If one looks at the way that biologists actually measure fitness in natural populations, it usually involves counting the average number of viable offspring produced by each phenotypic variant. It is therefore tempting to define fitness accordingly as the number of offspring or 'replicates' produced by each kind of variant in a population. On this view the phrase, 'x is fitter than y' means that x produces more offspring or replicates than y. For reasons that have already been noted, however, this definition is inadequate. Defining fitness in terms of the actual number of replicates produced by each phenotype renders empty tautologies instead of meaningful explanations. Suppose one asks, why in a population of peacocks did individuals with the longest tails produce the most offspring? The answer, presumably, is that individuals with the longest tails are the most fit. But if fitness is defined in terms of the number of self-replicates a phenotype produced, then this 'explanation' states that peacocks with the longest tails had the most offspring because they produced the greatest number of replicates. Notice that this merely restates the original observation.

The following subsection will examine some of the alternative candidate definitions of fitness in more detail. But first, consider the second problematic concept in our revised formulation of the principle of natural selection. Recall that in order to account for the occurrence of drift, it is necessary to state this principle in probabilistic terms: if Lewontin's three conditions are satisfied by a population, then that population will *probably* become increasingly adapted over time. Philosophers of biology have taken a keen interest in the nature of these probabilities. Some argue that the probability term is just a reflection of our ignorance (Rosenberg, 1988, 1994, 2001b; Horan 1994). Take the previous example where all of the longest billed platypuses were indiscriminately removed from the population before reproducing. One might argue that as a matter of fact those individuals turned out not to be the fittest members of the population, because in fact they produced no offspring. (Notice, however, that this objection seems to equate fitness with the actual number of replicates an organism produces, which has already been revealed as problematic). The important point here is that it might have been possible, at least in principle, to predict that these organisms would not survive to reproductive age. Perhaps the same gene that codes for a long bill also makes these animals more inquisitive, and an inquisitive platypus is less fit than a non-inquisitive one in an environment laden with platypus traps. On this view, a statistical term needs to be included in any statement of the principle of natural selection only because humans have a limited understanding of actual fitness values. In a moment we shall consider some objections to this view.

A Closer Look at Fitness

Recall that the tautology problem arises when fitness is defined in terms of the actual number of offspring produced by an organism over its lifetime. Such definitions that equate fitness with actual reproductive output become problematic in the context of certain evolutionary explanations. The problem arises when one attempts to explain an episode of selection by

appealing to fitness differences as the cause. If fitness is defined, say, as 'the average per capita lifetime contribution of individuals of [a] genotype to the population' (Futuyma, 1998, p. 366), then one cannot invoke fitness in a causal explanation for why individuals of one genotype make a greater average per capita contribution than individuals of some other genotype.

It is perhaps worth noting that biologists employ alternative definitions and measurements of fitness. In some cases, one must look further ahead to future generations before the reproductive benefits of a phenotype become apparent. So sometimes fitness is measured not in terms of the individual's net offspring but instead in terms of the total number of its offspring's descendants. Biologists also recognize that fitness is not an intrinsic property like mass or age. Rather, a given type of organism can be more or less fit depending on which environment it inhabits. It is also important to keep in mind that in order for evolution by natural selection to occur, there must be fitness *differences* among members of the population. Thus, biologists are typically interested in an individuals' relative fitness, the number of offspring it produces compared to other organisms in the same environment. However, provided that relative fitness is cashed out in terms of the *actual* differences in reproductive output among two or more organisms in an environment, the tautology problem hasn't gone anywhere.

Some philosophers and biologists have attempted to avoid this problem by equating fitness not with its effects on an organism's reproductive output, but instead with the 'design properties' of an organism that cause it to produce a given number of offspring. In principle, this strategy would avoid the tautology problem. For example, describing how relative bill length in platypuses contributes differentially to their prey capture rates would explain, for these organisms at least, why some phenotypes are more reproductively successful than others. Here fitness is equated with the optimal design of the bill for capturing prey. However, the problem with this strategy is that there appears to be no single principle of good design that applies across all organisms in all environments. But perhaps this isn't such a big problem. Perhaps fitness is best viewed as a higher-level property that supervenes on an enormous set of principles for good design. Notice, however, that this set of design principles is infinitely large, since environments are continually changing and organisms are constantly producing novel mutations. So, it is not just difficult in practice to specify all of the possible ways that organisms might be more or less well designed for their environments, it is impossible in principle to provide a complete list of all the possible ways that fitness might be realized (Rosenberg, 2001). Some philosophers have argued that, like it or not, this is just the sort of predicament that biologists have to work with. For example, Mary Williams (1973) Alex Rosenberg (1983) and Rosenberg and Williams (1986) have argued that 'fitness' should be regarded as a primitive term, one that cannot be analysed further but which plays a key role in evolutionary theory (see Sober, 1984, 1987 for criticisms of this proposal).

Other philosophers have recommended an alternative resolution to the tautology problem. Robert Brandon (1978) and Susan Mills and John Beatty, in Chapter 9, argue that fitness can be non-vacuously defined as a dispositional property. On this view, 'x is fitter than y in environment E' means that x has the disposition to produce more offspring than y in E. Perhaps the best way to understand this proposal is by drawing an analogy to other dispositional properties, like fragility or solubility. To say that a particular object is fragile means that under certain conditions it is disposed to shatter. Just as there are an infinite number of different physical arrangements of molecules that could realize this property, so too are there infinitely many different organism–environment relations that could realize fitness. Viewing fitness as a

dispositional property avoids the tautology problem because dispositions are a kind of cause. Suppose one asks, why did the vase shatter when it hit the floor? The answer is that it shattered because it was fragile. Here a dispositional property – fragility – is identified as having made a causal contribution to the vase's shattering. Likewise, suppose one asks, why did long-tailed peacocks produce more offspring than short-tailed ones? The answer, according to Mills and Beatty, is that long-tailed peacocks have the disposition to produce more offspring than short-tailed ones. Once again, this disposition is cited as a causal factor in explaining differential reproductive rates among peacocks. Notice that on this view it makes perfect sense to say that x was fitter than y, but x did not have more offspring than y. This makes sense because dispositions are not always realized: not all fragile vases shatter, and not all relatively fit organisms have the most offspring every generation. In drawing this comparison between fitness and fragility it is important to note one important difference between these properties. Fragility is an 'all or nothing' property in the sense that an object that possesses this disposition either realizes it or it does not. Fitness, however, can be realized to varying degrees. Perhaps the best way to conceptualize fitness is as a probability distribution. A particular platypus with a bill of a certain length has a high probability of producing x offspring, a slightly lower chance of producing x − 1 offspring, an extremely low chance of producing x + 1,000 offspring and so on. This point will become crucial when discussing the differences between fitness and drift below. Admittedly, evolutionary explanations that appeal to dispositional fitness are fairly 'thin' in the sense that they provide the minimum amount of causal detail about why some organisms are more reproductively successful than others. However, the advantage of this minimalist definition is that it is unifying. The dispositional account reveals what all of the many cases of fitness differences among organisms have in common.

A Closer Look at Drift and the Implications for Evolutionary Theory

In Chapter 10, 'Chance and Natural Selection', John Beatty explores some of the implications of defining the principle of natural selection in probabilistic terms. He invites us to imagine a version of Kettlewell's (1955) famous example of selection in moths due to industrial melanism. In Beatty's scenario there are two types of moth, dark and light, inhabiting a forest containing predominantly dark trees. In this environment dark moths have higher fitness than light ones. That is, dark moths are disposed to survive and reproduce at a greater frequency than light moths because they are less noticeable to predators against a background of predominantly dark trees. However, things don't go as expected. At generation G1 more dark moths just happen to land on light trees than on dark trees and their frequency decreases at G2. Beatty's question is as follows: was this change in frequency due to selection or drift? The answer, paradoxically, appears to be 'both'. Recall that fitness is being conceptualized here as a distribution of possible outcomes. Some reproductive outcomes are more likely than others, but all are within the range of possibilities caused by selection. So even if the reproductive output of dark moths at G2 is lower than one might expect, this does not rule out selection as the cause. On the other hand, according to Beatty's description of the scenario it appears that drift caused the dark moths to decrease in frequency. After all, it was due to their repeated *chance* encounters with light-coloured trees that fewer dark moths survived and reproduced. Beatty's only solution to this paradox is to concede that drift and selection are conceptually indistinguishable in certain cases.

One potential way of avoiding Beatty's paradox is to insist that the probabilities identified by the principle of natural selection are merely epistemic in nature. As noted earlier, it is conceivable that a complete understanding of an organism's traits plus a thorough understanding of its selective environment would allow one to predict, at least in principle, exactly how many offspring it will have in a given generation (Rosenberg, 1994, 2001). A complete understanding of the processes impacting survival and reproduction would therefore eliminate any need to treat fitness and drift as probabilistic concepts. Indeed, on this view drift appears to drop out of the equation entirely. Rosenberg (2001) does not think that such an understanding of evolutionary processes is attainable, however, because the set of organism–environment relations on which fitness supervenes is infinitely large. So evolutionary theory is unavoidably a statistical theory, but the probabilities it identifies have their origin in our epistemic limitations, not in any fundamental indeterminacy in evolutionary processes (see also Horan, 1994).

In Chapter 11, 'The Indeterministic Character of Evolutionary Theory: No "Hidden Variables Proof" but No Room for Determinism Either', Robert Brandon and Scott Carson defend the probabilistic formulation of the principle of natural selection by arguing that evolutionary processes are fundamentally indeterministic. First, they offer a thought experiment designed to show that quantum mechanical events could in all likelihood influence evolutionary outcomes. In this scenario quantum-level events cause a point mutation in an individual organism which, with a little help from genetic drift, causes the population to transition from one population frequency to another. If quantum-level phenomena can 'percolate up' to the macro-biological level, Brandon and Carson argue, the concepts of fitness and drift are irreducibly probabilistic. They offer a second argument for this conclusion. Brandon and Carson define drift as any trans-generational change in gene or genotype frequency due to sampling error. They note that the likelihood of sampling error increases as a population decreases in size. In cases where the decrease in size is sufficiently dramatic, drift can be inevitable. Although one can predict that allele frequencies will deviate from the parental frequency under such conditions, it is not possible to predict the precise direction of that change. Therefore, they conclude, drift is a fundamentally indeterministic process. Brandon and Carson further argue that that natural selection, or the discriminate sampling of genes or genotypes, also contains an element of indeterminacy because it occurs in finite populations where there is always some deviation among trials from the expected outcome. Finally, Brandon and Carson claim that the assumption of a deterministic universe (above the level of quantum events) is entirely a matter of ungrounded faith that is actually undermined by the available evidence. In highly controlled experiments where all of the independent variables are held fixed over repeated trials, some variation in the observed outcomes is always observed. For instance, Brandon and Carson cite an artificial selection experiment where cloned plants reared under identical conditions nonetheless grow to different heights. The determinist would interpret such variation as being due to some 'hidden variable', an uncontrolled factor that was overlooked in the experimental set-up. However, 'the positing of these hidden differences is purely gratuitous; they are posited for no reason other than to save the deterministic character of the theory' (p. 231).

In Chapter 12, an important reply essay, Leslie Graves, Barbara Horan and Alex Rosenberg rebut all four of these arguments. First, while it might be true that quantum events can occasionally 'percolate up' to the macro-biological level, this cannot be the explanation for why evolutionary theory treats fitness and drift probabilistically. For one thing, the

circumstances that allow for quantum percolation are exceedingly rare. So only a small subset of unpredictable events impacting survival and reproduction (perhaps a vanishingly small subset) can be explained by quantum percolation. Furthermore, they note that physical systems (the ones that Newtonian physics describes) are no more immune to quantum percolation than biological ones. Yet, above the level of quantum mechanics, physical theories do a fine job of explaining physical systems by assuming that they are deterministic. If one does not assume that drift and selection are non-deterministic processes, then Brandon and Carson's second and third arguments simply miss their mark. That is, it is possible at least in principle to explain any evolutionary outcome with a set of deterministic laws, so the probabilities identified in biology are epistemic. Finally, Graves, Horan and Rosenberg rightfully take issue with the no hidden variables interpretation of the experimental set-up that Brandon and Carson describe. There are all sorts of factors that this experiment fails to control for: 'identical' clones can vary genetically as a result of mutation; two plants growing side by side in a greenhouse will never receive the *exact* same amount of water and light; there are millions of species of bacteria living in a given handful of soil and it is extremely difficult to ensure that two soil samples are identical in this respect. Knowing these facts, any biologist faced with the sort of variation in experimental outcomes that Brandon and Carson describe would naturally appeal to hidden variables as the most likely explanation.

Recall that the debate over whether drift is a fundamentally indeterministic process stemmed from an attempt to avoid Beatty's paradox. If drift is posited merely as a placeholder for the set of events that humans are incapable of predicting due to their limited knowledge of deterministic processes, then Beatty's paradox is 'merely' a reflection of our epistemic limitations. However, the philosopher Roberta Millstein argues that Beatty's paradox is not so easily swept aside. She notes that regardless of whether evolutionary probabilities are ultimately epistemic or ontological in nature, they are something that working biologists must contend with (Millstein, 1996).

In Chapter 13, 'Are Random Drift and Natural Selection Conceptually Distinct?', Millstein offers a solution to Beatty's paradox. She notes that Beatty and others have failed to distinguish between two basic conceptions of drift. In some cases, Millstein notes, drift is conceptualized as a process (or cause). This tends to be the conception that people have in mind when describing drift as a form of indiscriminate sampling. For example, earlier it was imagined that all of the long-billed platypus in a particular population were removed indiscriminately with respect to their bill sizes. In other cases, drift is conceptualized as an outcome. For example, if the ratio of long- to short-billed platypuses decreases in a particular generation, this outcome might be attributed to drift. Millstein argues that this distinction between the two conceptions of drift – as a process and as an outcome – is the key to avoiding Beatty's paradox. In the peppered moth example that Beatty used to present his argument, we were told that a significant portion of black moths happen to land on light-coloured trees, thereby causing them to fall prey to visual predators. Millstein notes that in this portion of the description there is no mention of indiscriminate sampling: dark moths are being selected against in their environment on the basis of a heritable trait. As far as this process is concerned, it is one of selection and not drift. Had Beatty simply described the example in terms of an evolutionary outcome (for example dark moths decreased in frequency) then indeed it would be impossible to determine without further information whether this outcome was caused by selection or drift. But it doesn't

follow that drift and selection are *conceptually* indistinguishable, provided that one considers the process that generated this outcome.

Millstein acknowledges that it might sound odd to say that selection caused a decrease in the proportion of the fittest individuals (dark moths). However, the first thing to note is that this outcome is perfectly acceptable so long as fitness is conceptualized dispositionally as a probability distribution (see above). To say that dark moths are fitter than light ones means, on this view, that dark moths are *more likely* to survive and reproduce than light ones; it doesn't guarantee that this outcome will be realized each and every generation. Millstein further notes that Beatty's example takes place over just one generation. By contrast, the impact of selection on a population is usually calculated over many successive generations. Thus, if the proportion of fittest individuals had decreased over an extended number of generations, then this would indeed be an odd (though not impossible) outcome. To summarize, Millstein's analysis of drift leads us to take a second look at the nature of selection and drift. Upon inspection, it becomes clear that Beatty is conflating processes and outcomes in his discussion, and that if one is careful to separate selection processes from selection outcomes and drift processes from drift outcomes, then Beatty's 'paradox' disappears.

The Statistical Interpretation and the Structure of Evolutionary Theory

In recent years there has been a growing interest in whether evolutionary theory is properly conceptualized as a dynamical theory (a theory of forces) or as a statistical theory (a theory that explains evolution in terms of probabilities). Until recently the received view has been that evolutionary theory is both of these things. Drift, selection and mutation have been regarded as independent forces (especially within the subfield of population genetics) whose respective influences on a population can be summed in an analogous fashion to Newtonian forces. At the same time, processes like drift and selection have been conceptualized as inherently probabilistic processes that are akin to a form of sampling error. However, in two recent and influential essays Mohan Matthen and André Ariew (2002) and Denis Walsh, Tim Lewens and Ariew (Chapter 14) argue that these two perspectives are incompatible with one another. Moreover, they claim that the statistical conception is superior to the dynamical view. Matthen and Ariew (2002) argue that the dynamical conception is internally incoherent because the purported forces of evolution cannot act independently of one another, and because there is no 'common currency' in which the magnitudes of these forces can be measured. There is no way to determine, for example, how much of an evolutionary change was due to selection and how much to drift. Matthen and Ariew's arguments have been addressed in a number of recent publications including Frédéric Bouchard and Alex Rosenberg (Chapter 15) and Stephens (2004). Due to limited space the details of this debate cannot be considered here. Instead, what follows is a selective summary of Walsh *et al*.'s central arguments for the statistical approach along with a few critical remarks.

In Chapter 14 Walsh, Lewens and Ariew first outline in general terms the differences between dynamical and statistical explanations. One important difference lies in how the two approaches explain the occurrence of error, or the failure of an outcome to satisfy the prediction of a model or theory. On a dynamical theory error is explained by positing what Brandon and Carson would call a 'hidden variable': some causally significant factor that the model or theory overlooked is generating the unexpected outcome. According to Walsh *et*

al., in a statistical theory error is 'a built-in feature of the laws of probability' (p. 280). More specifically, error is regarded on this view as a brute fact emerging out of the nature of the selection process in question and the laws of probability governing selection processes of that type. This is a controversial thesis. Walsh *et al*. do not take statistical error to emerge out of quantum-level indeterminacy, nor does it emerge out of our ignorance of hidden variables. Error is an 'irreducible' property of ensembles. It follows from this view, they argue, that a statistical explanation of some series of events is autonomous from a lower-level causal account of the same process.

To put this point in more familiar terms, think back to Rosenberg's suggestion that a complete account of organisms and their environments would allow one to predict, at least in principle, the exact evolutionary outcome at any given generation. Walsh *et al*. deny this claim. They assume that even with a complete explanation of the deterministic (or indeterministic) events that led to the birth, life and death of each and every platypus in a population, it would still be necessary to include a statistical term when describing the process at the population level. Advocates of this view often draw a comparison between evolutionary theory and the theory of statistical thermodynamics. In the latter science, a higher-level property like heat is defined as the mean molecular motion of the ensemble. By the same token, fitness is being regarded here as an ensemble-level property (a probability distribution) that supervenes on all of the lower-level births, lives and deaths of the organisms in a population.

At this point one might ask to see some evidence of these strong claims. Why think that fitness is irreducibly statistical in nature? What, if anything, are the explanatory goods that only a statistical interpretation of natural selection can deliver? Walsh *et al*. cite three questions which, they claim, a statistical explanation answers but a dynamical explanation does not (in presenting these points it is helpful to think of evolutionary outcomes as analogous to series of coin tosses or 'trials'):

1. Why, given a certain number of trials (for example coin flips), is the expected outcome to be expected?
2. Why did a series of trials result in the particular sequence of outcomes that it did (for example six heads, four tails)?
3. Why does the likelihood of approaching the expected distribution of outcomes (for example 50 per cent heads) increase as the number of trials increases?

Now, it is not difficult to see why Walsh *et al*. think that points 1 and 2 cannot be explained on a dynamical theory, nor is it clear that a statistical theory can answer them. Why do we expect a fair coin to come up heads 50 per cent of the times it is tossed? This expectation does not arise out of the laws of probability; it emerges out of the physical properties of the coin, the laws of physics and the range of initial conditions specifying how the coin might be flipped. Second, it is simply false that probability theory can explain why a series of three coins flips came up, say, H H T rather than T T H. Again, the explanation lies in the properties of the coin, the laws of physics and (importantly) the initial conditions realized by each token flip. Perhaps the only explanatory advantage that a statistical theory has over a dynamical account is that the former can explain why the likelihood of getting a representative sample increases with the size of the sample (or vice versa). The statistical theory explains this relation by the so-called law of large numbers. Walsh *et al*. take this to be a fairly significant point, because

they regard drift as fundamentally an indiscriminate sampling process over small populations. So, if a statistical theory alone can explain the relation between the size of a sample and its representativeness, and if drift is essentially the non-representative sampling of alleles from a population due to its size, then only a statistical theory can explain drift.

However, in Chapter 15 Bouchard and Rosenberg do not afford the statistical approach even this much autonomy. What explains the relation between the size and representativeness of a sample? According to them, '[t]he causal explanation of the divergence of a finite sample from the long-run relative frequency of coin tosses is to be sought in the fact that *the initial conditions of the actual coin tosses were not representative of the set of initial conditions that give rise to the long run*' (p. 311; my italics). While this might be true, it pushes the problem back to a question about initial conditions. We now require an explanation of why the initial conditions that generate each outcome are expected to become more representative as the number of trials increases. It remains unclear whether this propensity is best explained by the law of large numbers, a statistical truth, or if it somehow drops out of a detailed causal explanation of successive token trials (pp. 312–13).

One potential reason for resisting the statistical approach is if fitness differences *among individual organisms* have a distinctive impact on evolutionary outcomes. If so, then a statistical theory that averages over those individual differences loses some explanatory traction. Walsh *et al*. resist this objection. They present a pair of simple models in which heritable variation occurs either among organisms or among their traits. One of the important insights that this model conveys is that these two modes of variation can be independent of one another. Walsh *et al*. go on to argue that variation among traits, not individuals, is always the more important factor when measuring evolutionary change. Their argument for this conclusion is highly questionable. However, let us grant for the sake of argument that there are at least some types of evolutionary change that are best explained in terms of differences in average trait fitness. Does it follow that these explanations are best construed statistically and not in terms of competing forces, as a dynamical theory would have it? The implication certainly isn't obvious. One could regard the average fitness values specified in 1B and 2B of Walsh *et al*.'s model as selective forces that predict a certain change in trait frequencies. In model 1B the forces are zero; this is why no change occurs. In model 2B the differences among selective forces explain the change in trait frequency. Walsh *et al*. suggest that matters will not be so simple in real populations where both selection and drift operate simultaneously. The problem, they claim, is that selection and drift are indistinguishable as forces. In echoing Beatty's paradox, they claim that 'One reason to suppose that we cannot decompose drift and selection into component forces is that the same processes that cause natural selection cause drift' (p. 289). Hence, 'The answer [to this paradox] ought to be clear. In so far as the trait fitnesses fail to predict or explain the changes in trait frequencies, there is error. We are even told where the error occurs; dark moths land more frequently on light trees than would be predicted by their relative frequencies … this is a case of drift' (p. 290).

Notice that these authors appear to be making the exact same error that Millstein identified in Beatty's original presentation. Namely, they are failing to distinguish between the two alternative conceptions of drift and selection. Recall that both of these phenomena can be regarded either as a cause of evolutionary change (a process) or as a certain type of evolutionary outcome. While drift and selection differ in the first respect, as causes, they do not necessarily differ in their effects on a population. Beatty's paradox arose because of a

failure to distinguish processes from outcomes. Likewise, in the present context, it is possible to view the *processes* of drift and selection as being governed by forces or what Elliott Sober (1984) calls 'source laws'. At the same time, it is possible to acknowledge that when talking about drift and selection as *outcomes* it is sometimes empirically impossible to distinguish them. In short, the point is that it doesn't follow from the fact that drift and selection are often described probabilistically when viewed as outcomes that they cannot be viewed as deterministic forces when viewed as processes or causes (see Stephens, 2004).

To summarize: the principle of natural selection is deceptively simple in its standard formulation as heritable variation in phenotypes with fitness differences. In order to avoid the tautology problem one must not define fitness in terms of the actual reproductive output of each variant in a population. The best current definition of fitness equates this property with the disposition to have a certain number of offspring in a given environment. However, this means that the principle of natural selection is inherently probabilistic. When it comes to interpreting probabilistic terms in biology, it is important to distinguish between processes and outcomes. When drift and selection are defined as outcomes it is difficult to distinguish drift from selection. Viewing drift and selection as outcomes also makes them seem less like law-like forces and more like summary statistics that cannot be identified except at the population level. Recognizing that drift and selection can be viewed as processes, however, provides a principled way of distinguishing drift and selection and better accords with the view that they are causal forces.

Selfish Genes, Altruistic Organisms and the Levels of the Selection Debate

Prior to the 1960s it was considered acceptable to think of social groups, species and even whole ecosystems as highly adapted, functionally integrated units. Individuals were seen as playing the role of organs in these higher-level 'super-organisms' and individual traits were interpreted as having evolved for the good of the group. The basic model of group selection employed by Darwin (1871) and others (Wynne-Edwards, 1962) was just the familiar principle of natural selection 'bumped up' to the level of groups. If there is heritable variation among groups in the traits impacting their survival and reproduction, then more successful groups will reproduce more rapidly than less successful ones. Group selection was thought to proceed via a budding process with parent populations increasing in size and splitting into daughter groups. The kinds of traits likely to contribute to this process, for example altruism and cooperation, are often difficult to explain as individual-level adaptations. Since these traits seemed widespread in nature, group selection was considered the predominant mechanism by which traits evolve. For reasons that will now become clear this view is generally known as *naive* group selection theory.

Group selection came under severe attack in 1966 with the publication of G.C. Williams' *Adaptation and Natural Selection*, a book regarded by many as second in importance only to Darwin's *The Origin of Species*. Williams argued, first, that the fitness of a group is typically reducible to the individual fitness of its members. Take a group-level property like the average running speed of a population of deer. If faster deer are individually fitter than slower ones, then over time the average speed of the group will increase. Clearly selection is not acting on the group in such cases: the increase in group fitness is an epiphenomenon of selection acting on individuals. Cases like this one give rise to what we might call the '*reductionist challenge*'

to group selection. The challenge is to identify a case where an evolutionary change at the level of the group *cannot* be explained in terms of selection acting directly on individual group members or some lower-level entity, like genes.

This is not the only challenge Williams raised for group selection. He further noted that even a modest amount of migration among groups, a normal occurrence in nature, will suppress intergroup variation. If selection at any level requires variation among the entities at that level, and if variation among groups disappears under normal levels of migration, then group selection cannot occur except under rare circumstances. Call this the *migration problem*. Williams' third objection states that groups composed of altruistic individuals are vulnerable to 'subversion from within'. An altruistic trait is defined as one that decreases the fitness of its bearer while increasing the fitness of other, unrelated individuals. If a group of altruists is invaded by a selfish individual (one who lacks the altruistic trait), then selfish types will by definition produce more offspring than altruists and eventually dominate the population. It is important to recognize that the *subversion problem* (as it is called) arises out of the differential rates at which groups and individuals reproduce. In the time it takes for a group to bud-off into daughters there will presumably have been several turns of the generational wheel. So the process favouring group-beneficial traits is relatively weak, Williams argued, compared to the process favouring individual adaptations.

These objections were so influential in discrediting group selection that before long it had joined ranks with Lamarckian inheritance theory – the very mention of a trait's being 'good for the group' was enough to invite charges of woolly-headed thinking. At around the same time an alternative perspective began to take hold. The gene's eye view, also known as *gene centrism*, identifies the gene and not the organism or group as the fundamental unit in which reproductive costs and benefits are to be calculated (Hamilton, 1964; Williams, 1966; Trivers, 1971; Dawkins, 1982). The extent to which organisms should behave 'altruistically' towards one another is, on this view, directly proportional to the amount of genetic material they share. One of gene centrism's big success stories was its explanation of eusociality (the existence of sterile casts) in insects, a problem that haunted Darwin because there seems to be no way that individual-level selection could favour such extreme forms of self-sacrifice. As W.D. Hamilton explained in a series of elegant models, the reason that eusociality evolved in insects has to do with their unique haplodiploid genetic system. An individual ant or wasp gets more of her genes into the next generation by producing sister clones than by diluting her genetic material to create daughters via sexual recombination. So, what appears as an extreme form of altruism at the individual level turns out to be extreme selfishness at the level of the gene. This was seen as an enormous victory for gene centrism, and another nail the in the coffin for group selection.

To this day gene centrism remains the dominant perspective in the biological and social sciences. However, in recent years a growing number of group selection advocates have re-emerged under the banner of a more nuanced view called *multi-level selection (MLS) theory* (Wilson, 1980, Sober, 1984, Brandon, 1990; Sober and Wilson, 1994, 1998; Okasha, 2006). One of the most important advances of multi-level selection theory over naive group selection involves a revised conception of the group. On the naive view groups were regarded much like large individuals with fairly well-defined spatial boundaries, more or less continuous membership, and uninterrupted continuity over generations. David Wilson (1980) replaced this traditional notion with the idea of a trait group. Trait groups are individuated not by their

spatio-temporal boundaries, but rather by fitness-relevant interactions among their members. For example, a particular tree might form a more closely knit trait group with its insect pollinators than it does with the other plants living in the same patch of forest. In Wilson's terms, the tree and its pollinators share a 'common fate' with respect to the trait of pollination. On the trait-group model selection is capable of acting both on the trait group as a whole (for example the tree and pollinator pair) as well as among members within the group – hence the name, *multi-level* selection theory.

Note that a single individual can belong to multiple trait groups. This is especially clear in cases where organisms form different trait groups for discrete stages of their life cycle. In Chapter 16 Wilson and Sober imagine a population of insects that spend a portion of their life cycle segregated into pairs in pools of water. A proportion of these larvae express a trait that detoxifies the water, enhancing both their own fitness as well as the fitness of their fellow pond-mates. Since both detoxifying and non-detoxifying larvae benefit equally from this trait, pond-mates belong to same trait group during this phase of their life cycle. However, when those larvae emerge as adults they might express different traits that pick out different trait groups, for example if they aggregate into mating pairs. So trait groups are more permissive than traditional groups in a variety of ways: they do not require discrete boundaries, they can exist for brief stage of an organism's life cycle, they can disband and reform over successive generations, and they can vary in strength (the cells in a tree form a more tightly knit trait group than the tree and its pollinators). The remainder of this section considers whether MLS theory, armed with this notion of a trait group, can avoid Williams' three objections.

MLS Theory and the Migration and Subversion Problems

Notice that the migration problem does not arise in Wilson and Sober's aquatic insect example. Individual larvae do not squirm out of polluted ponds to join their detoxifying brethren. One might therefore object that this example begs the question against the migration problem: it assumes that no migration occurs, but in fact this process is common in natural populations.

However, the point of Wilson's trait-group model is to show that individuals need not form permanently segregated groups in order for group selection to occur. Historically, both proponents of naive group selection and their critics operated with an unnecessarily robust concept of a group. Groups were regarded as 'bonded' collections of individuals whose boundaries remained more or less constant over the entire lives of their members. On the trait-group model, a group can form during a specific phase of the life cycle and group selection can act during this (potentially brief) period. For example, there is no expectation that pairs of aquatic insects will continue to travel together after growing wings and leaving their pool. This revised conception of groups mitigates the migration problem. If the phase during which a population is segregated into groups is relatively brief, then migration is less likely to suppress variation among those groups.

A second reason for thinking that the migration problem might be overstated is that it assumes a strong form of genetic determinism. Williams assumed that when a migrant joins a group it expresses whichever behaviour, selfish or altruistic, is encoded in its genes. Since migration reduces genetic variation among groups, he argued; it reduces phenotypic variation as well. However, there are good reasons for resisting this argument. Many organisms are conditional cooperators, they are capable of withholding cooperation in the midst of selfish

individuals and engaging in cooperating in the company of fellow cooperators. Dugatkin (1997) provides evidence of conditional cooperation in guppies, and Sober and Wilson (1998) suggest that this strategy is commonly found across the animal kingdom. If members of a population are capable of conditional cooperation then it is possible to maintain intergroup variation despite high levels of migration, provided that organisms switch strategies to resemble the group they migrate into.

MLS theory's response to the subversion problem also draws on its revised conception of the group. Recall that on the naive group selection model groups were thought to have fairly permanent boundaries with low rates of intergroup migration. Hence, if a selfish mutant appears in a group of altruists, it is inevitable that it would eventually come to dominate the population. By contrast, revised group selection models allow for high rates of shuffling and reassortment of group members. Provided that groups with a high proportion of altruists have a sufficiently higher fitness as groups, compared to groups containing selfish individuals, it is possible to obtain a steady equilibrium among altruists and selfish types. Hence, subversion from within does not eliminate altruism (see Sober and Wilson, 1998). However, this result has only been demonstrated mathematically and has not yet been shown to obtain in the wild.

Wilson and Sober's trait-group model marks an important advance in the levels of selection debate by broadening the range of conditions in which group selection can occur. As we have seen, the extent to which migration and subversion pose a threat to group selection depends on some fairly technical empirical questions that cannot be resolved a priori. There is also a set of purely philosophical questions that bear on this debate. One of these is the question of whether kin selection is a form of group selection or an alternative to it. Another important philosophical question, to which we now turn, is whether the reductionist challenge can be met and, if so, under what conditions.

Pluralism, Broad Individualism and the Averaging Fallacy

Wilson and Sober's aquatic insect example is aimed directly at the reductionist challenge. Recall that the individual-level selectionist calculates fitness in terms of the expected reproductive output of individual organisms, for instance individual deer. The point of the aquatic larvae example is to show that anyone who adopts this view must admit that groups can also be fitness bearers. The example parallels a case where two alleles are paired at a given locus in a population of individuals. Each pond (locus) contains one of three possible 'genotypes': two detoxifying larvae (AA), two non-detoxifying larvae (aa) or a 'heterozygous' combination of the two (Aa). This situation is structurally equivalent to a case where a dominant allele (R) contributes to running speed in deer, with selection favouring the homozygous dominant (RR) and the heterozygous (Rr) individuals over the (slower) recessive (rr). Since these two cases – the pond insects and genotypes – are structurally equivalent, if one allows that fitness is a property of genes then one must also allow (under the right conditions) that fitness is a property of groups as well.

However, gene-centrists have an available reply to this argument. They claim that such examples do not take the reduction to genes far enough. From the gene's eye perspective, fitness is not a property of genotypes; fitness is a property of *individual alleles*. In order to calculate the fitness coefficient of the R allele one must average the reproductive output

of RR and Rr genotypes weighted by their frequency in the population. By extension, to calculate the fitness of detoxifying (A) larvae one must average the number of adult larvae emerging from AA and Aa pools, weighted by their frequencies. This calculation suggests that the reductionist challenge has not been met: the fitness of larval groups (or genotypes) is reducible to the weighted fitness coefficients of their constituent members (or alleles). This view has come to be known as broad individualism (Sterelny and Kitcher, 1988).

MLS theorists offer three kinds of response to this argument. Some are convinced that MLS theory and broad individualism offer equivalent descriptions of the same process. This pluralist position has recently gained considerable popularity among philosophers and biologists (Dugatkin and Reeve, 1994, Kerr and Godfrey-Smith, 2002; Sterelny, Chapter 17, this volume). Others argue that the broad individualist strategy of calculating fitness as a weighted average of individual alleles is guilty of an 'averaging fallacy' (Sober and Wilson, 1998; Okasha, Chapter 18, this volume). A third response states that the two views do not make the same predictions, pluralism is false and the gene-centrist description simply gets things wrong (Brandon and Nijhout, Chapter 19, this volume). Let us now consider each of these options.

Gene-centrists tend to prefer broad individualism to the multi-level perspective because it is allegedly more parsimonious and has wider explanatory breadth (but see Sober, 1990 on parsimony). However, pluralists like Kim Sterelny argue that both perspectives have heuristic value. Some cases lend themselves to a broad-individualist description, others make more sense from a multi-level perspective; but ultimately, Sterelny argues, these two theoretical frameworks offer equivalent descriptions of the same underlying process (Dugatkin and Reeve, 1994; Kerr and Godfrey-Smith 2002).

There are reasons for resisting the pluralist approach. If two theories cite different causal processes in explaining the same outcome they cannot be equivalent. Selection operates at multiple levels, at just one level, or neither. So, one might ask the pluralist: which is it?

MLS theorists argue that broad individualism distorts the causal facts by committing the 'averaging fallacy'. The fallacy allegedly occurs when the broad individualist attempts to calculate fitness as a weighted average of alleles among groups. The averaging approach always results in selfish individuals having a higher fitness than altruists, because selfish individuals receive all the benefits of living among altruists while paying none of the costs. However, under certain conditions this calculation is mistaken: when between group fitness differences are sufficiently large and when groups engage in high rates of reassortment, altruism can persist indefinitely. So, the averaging approach is fallacious because it can generate false predictions, and the problem stems from its failure to calculate fitness differences among groups as a separate value from fitness differences within groups.

In Chapter 18, 'The "Averaging Fallacy" and the Levels of Selection', Samir Okasha offers an insightful diagnosis of this objection. As Okasha notes, the averaging fallacy is usually presented in the context of highly idealized examples. This makes it difficult to determine when exactly it is fallacious to calculate fitness as an average among groups. Clearly there are at least *some* cases where averaging is permitted. If one wants to calculate the relative fitness of fast versus slow deer there is no harm in averaging among groups. So what distinguishes these cases from ones where averaging is a problem?

Okasha argues that MLS theorists have invoked two different sets of criteria in their presentation of the averaging fallacy. One condition is what Okasha calls the monotonicity

condition, and it is supposedly violated in cases like sickle cell anaemia where the heterozygote (Aa) has higher fitness than either homozygote (AA or aa). Sober and Lewontin (1982) argue that in such cases it is necessary to calculate fitness at the level of the genotype, not at the level of the individual allele, or else one is unable to explain why the harmful (aa) allele persists. In other contexts, however, MLS theorists claim that the averaging fallacy is committed when two conditions hold. Namely, when (a) the fitness of an organism depends on the composition of its group and (b) groups vary in fitness. In these cases, it is argued, one must calculate fitness within and among groups separately.

To show that these are two distinct sorts of cases, Okasha first imagines a scenario where heterozygote superiority is 'bumped up' to the level of the organism. In the context of the aquatic larvae example, suppose that the 'heterozygous' combination of one detoxifying and one non-detoxifying larva has the highest relative fitness among pools. Now the monotonicity condition is violated. But the problem seems to stem from the violation of (a) and (b) – there is nothing special about the monotonicity condition per se. Note also that monotonicity was not violated in the original presentation of the aquatic insect example. Yet Wilson and Sober argued that averaging individual fitnesses among polls was an instance of the averaging fallacy. Therefore, if averaging is fallacious only when the monotonicity condition is violated, then it is not fallacious in standard evolution of altruism models.

Okasha offers an interesting solution to this incompatibility problem. He distinguishes between a selection *process* that occurs at the level of the individual allele and a gene-centrist *perspective* of selection processes that occur at other levels. Basically, conditions a and b identify when selection is acting above the level of the individual allele. However, the monotonicity condition identifies when it is possible to describe this process in terms of alleles and their selective environments. This proposal puts an interesting spin on the reductionist challenge. Initially, it was assumed that if fitness differences among groups (or individuals) could be reduced to fitness differences among individuals (or genes), then selection must be occurring at the lower level. Okasha's suggestion implies that gene-centrism can be a useful and illuminating mode of description even when it gets the causal facts wrong.

In Chapter 19, 'The Empirical Nonequivalence of Genic and Genotypic Models of Selection', Robert Brandon and Fred Nijhout offer a 'decisive' refutation of the pluralist and gene-centric approaches. Not only do the two perspectives make incompatible empirical predictions, they claim, gene-centrism tends to get it wrong. Their argument is somewhat mathematically detailed but can be stated informally quite easily. In a case of heterozygote superiority gene selectionists assume that once the two alleles reach equilibrium of 1:1 A:a, selection stops happening. But in any biologically realistic case, when selection is relaxed drift continues to occur. The larger the population, the stronger the influence of drift in the absence of selection. Therefore, gene-centric and genotypic models make different predictions about how a population will behave at equilibrium in the presence of drift. Gene-centric models predict that gene frequencies will drift away from equilibrium in even fairly sizable populations; whereas on genotypic models drift kicks in only when population sizes become quite low. Brandon and Nijhout argue that the predictions of gene-centric models are simply false. Any actual case of heterozygote superiority will conform more closely to the genotypic than to gene-centric models. In other words: 'genic selectionism is committed to describing equilibria that are actively maintained by selection as equilibria in which selection is not active. But the behaviors of such equilibria differ [thanks to differential responses to drift], and there is no

room for pluralism here' (p. 393). These authors demonstrate that this conclusion holds for cases other than heterozygote superiority, including heterozygote inferiority and when each phenotype is under the control of multiple loci. Although Brandon and Nijhout do not extend these models to the level of organisms within groups, it would be interesting to see whether their conclusions extend to higher levels of the biological hierarchy.

Three Challenges From Developmental Biology

For most of the twentieth century the disciplines of developmental and evolutionary biology had little to do with one another. Their separation was legitimized partly by the belief that these fields address different types of questions. Evolutionists were portrayed by Ernst Mayr (1961) as pursuing 'ultimate questions'. These are questions about processes like drift, selection and migration that occur over entire populations and span many generations. Developmentalists were thought to address 'proximate questions', about processes that occur within individual organisms over a single life cycle. Although the two sets of questions are clearly related – for example proximate mechanisms provide the phenotypic variation on which selection acts – many biologists and philosophers have considered it expedient to investigate the one set of questions independently of the other. Another reason for the disunity between these disciplines is that, historically, their scientific maturation was out of synch. While evolutionary biology was undergoing its modern synthesis with Mendelian genetics during the 1930s and 1940s, developmental biology languished in a state of relative obscurity. Although evolutionists expected that developmental biology would eventually explain how organisms transition from single cells to highly differentiated adults, most considered it unlikely that these details would require significant revisions to evolutionary theory. Since selection acts on genetic differences via phenotypic differences, it was assumed, any intervening processes could be safely bracketed. It therefore became customary for evolutionists to think of development as a 'black box' which takes genetic and environmental factors as inputs and delivers adult phenotypes out the other side.

In the twenty-first century, developmental biology has undergone some spectacular advances. Philosophers and biologists are now starting to rethink the traditional conception of how the two disciplines are related. The essays collected in Part V identify three 'developmentalist challenges' to traditional concepts in evolutionary biology. The first challenge stems from the idea that there are 'rules of development' that play a central role in determining biological form. A better appreciation of these rules is thought to challenge the 'adaptationism programme' (described earlier). The second challenge is directed at a form of reductionism which holds that all of developmental biology can be reduced to the level of molecular genetics. Proponents of the evolutionary developmental biology (Evo-Devo) movement argue that developmental generalizations constitute an autonomous level of biological explanation that cannot be reduced to the level of molecules interacting with genes. The third challenge stems from the growing realization that epigenetic (non-genetic) factors play a significant role in 'programming' an organism. Proponents of developmental systems theory challenge the traditional notion of the gene as the primary unit of evolutionary change.

The Developmentalist Challenge to Adaptationism

Anti-adaptationists like Gould and Lewontin (1979) criticized the strategy of atomizing an organism into traits and then regarding those traits as independent solutions to adaptive problems (see above). They object that this strategy ignores the importance of developmental constraints on selection. They also argue that developmental explanations often trump adaptationist 'just so' stories. One of Gould's favourite examples is the human chin. It is perhaps tempting to suppose that the chin is an adaptation for attracting mates, maybe hominids who are capable of developing pronounced chins are resistant to parasitic infection so that this trait signals the presence of 'good genes' in a potential mate (more outlandish adaptationist proposals have been put forward). However, Gould claims that from a developmental perspective the chin is the inevitable by-product of two convergent growth fields in the jaw. In this case a lack of developmental knowledge potentially leads to both an incorrect atomization of the chin and to the generation of (supposedly) misguided hypotheses about its function. Several authors have followed Gould in assuming that developmental explanations are typically at odds with adaptationist explanations (for example Lloyd, 2005). This is an obvious, yet surprisingly common error. All traits – adaptations and non-adaptations alike – have developmental histories. So if developmental explanations were opposed to adaptationist explanations as a general rule, no traits could be explained as adaptations. In order to avoid false dichotomies it is helpful to think back to Tinbergen's distinction between adaptive, evolutionary and developmental explanations of the same trait. These explanatory projects usually complement rather than conflict with one another. The only time that a developmental explanation definitively trumps an adaptationist one is when a trait turns out to be so constrained that heritable variation is impossible.

The debate, then, is not over the existence of developmental constraints but rather over their strength and frequency. There is a spectrum of views on these issues. At one extreme are process structuralists who think that development is governed by a fixed set of laws that are far more influential than selection in determining organismal form. Advocates of this view point to examples of traits that show little variation among members of a taxonomic group. For example, phyllotaxy – the radial arrangement of parts in a growing plant – is a highly conserved feature that is unlikely to have been maintained by selection (Amundson, Chapter 20, this volume). At the other end of the spectrum are empirical adaptationists, who hold that development plays only a minor role in constraining organismal form. Adaptationists claim that there are few known examples of fixed developmental constraints, but many examples of selection producing exaggerated phenotypic characters (Brakefield and Roskam, 2006). The multitude of exotic morphs found in butterfly wings or the exaggerated secondary sexual characters of birds of paradise testify to selection's ability to mould and shape phenotypes. Adaptationists also point to examples where selection, not developmental constraint, is responsible for maintaining conserved traits that show little taxonomic variation (Freeman, 2000). Such examples reveal that organisms are less developmentally constrained than process structuralists would suggest. However, without some way of quantifying the frequency and magnitude of constraints it is difficult to make informed generalizations about their relative importance.

In Chapter 20, 'Two Concepts of Constraint: Adaptationism and the Challenge from Developmental Biology', Ron Amundson argues that the disagreement between process

structuralists and adaptationists has been overstated. These two camps have been talking at cross purposes because they adopt different conceptions of what a constraint is. From the adaptationist perspective a constraint is something that prohibits selection from optimizing a trait. From the developmental perspective a constraint is something that prevents certain traits or trait combinations from being produced. These concepts are logically distinct. A trait can be highly developmentally constrained while being optimally adapted to its environment. Even after these concepts have been clarified, however, there remains an unresolved empirical disagreement between adaptationist and developmental viewpoints. As Amundson notes, these perspectives disagree about the extent to which organisms are developmentally constrained in their ability to produce variation in biological form. This difference is made vivid by imagining what would happen if a population were freed from selection and allowed to undergo indefinite amounts of mutation and drift. Under these conditions process structuralists predict that lineages would undergo little phenotypic divergence and remain clustered together in morphospace. Adaptationists assume that lineages would diverge significantly from one another so that a wide range of the available morphospace would become occupied. It is currently difficult to say which of these two default assumptions is closer to the truth. However, some recent advances in the field of evolutionary developmental biology offer a fresh theoretical perspective for evaluating the developmentalist challenge to adaptationism.

The field of evolutionary developmental biology (Evo-Devo) is sometimes described as opening the 'black box' of development. Scientists working in this discipline investigate the ways that developmental processes guide and constrain evolution and the ways that developmental systems themselves have evolved. One of the most exciting advances in this field is the discovery of modular developmental units (see Bolker, 2000; Sterelny, 2000; Gilbert, Optiz and Raff, Chapter 21, this volume). For example, the imaginal discs of insects are a well-defined group of cells that interact with one another in the production of a limb. This developmental subsystem is modular in the sense that discs can be moved to different locations in the developing embryo, like pieces in a Lego set, without disrupting their developmental trajectories. If modularity is a widespread feature of developmental systems then the potential for exploring diverse regions of morphospace is greatly enhanced. In the 1940s R.B. Goldschmidt created 'homeotic mutant' flies whose antennae had been replaced by leg parts and whose halteres (balancing organs) had been replaced by wings. Not only do such experiments support the view that organisms are highly modular, they also highlight the fact that phenotypic combinations that are rarely found in nature can easily be produced in the lab. Such findings support the adaptationist contention that selection, not developmental constraint, is the primary reason why organisms do not occupy an even wider expanse of morphospace. However, opponents of modularity point out that these sorts of experiments have been conducted on a relatively small set of model organisms. Perhaps mammals, birds, reptiles and plants are nowhere near as modular as insects. The verdict is still out on these issues.

The Developmentalist Challenge to Reductionism

The question of whether biological processes are reducible to molecular processes has been debated by philosophers for over forty years. Initially the debate focused on whether Mendelian genetics is reducible to molecular genetics. More recently, evolutionary developmental

biology has identified a set of structures and processes that are allegedly irreducible to the molecular/genetic level. Before exploring these more recent debates it is helpful to put the developmentalist challenge to reductionism into historical perspective.

Kenneth Schaffner was the first to offer a detailed proposal for the reduction of Mendelian genetics to molecular biology. He argued that Mendel's two 'laws' of inheritance – the *law of segregation* and the *law of independent assortment* – can be derived (though admittedly with a bit of fudging) from molecular principles (Schaffner, 1967, 1969). Moreover, in classical cases of theory reduction, like the reduction of Newtonian to relativistic physics, the reducing theory is able to explain why the reduced theory gets certain details wrong. This appeared to hold true for molecular biology's relationship to Mendelian genetics as well, because molecular biology can explain why Mendel's two laws obtain in some cases and why they are violated in others.

However, David Hull (1974) objected to Schaffner's proposal, arguing that terms like 'gene', 'dominance', 'allele' and so on have no unique correlates in molecular biology. If these Mendelian terms cannot be matched with corresponding molecular terms then a strict derivation of Mendelian genetics from molecular biology is impossible. A decade later, Philip Kitcher (1984) offered his 'gory details' argument against molecular reductionism. Kitcher argued that Mendelian genetics provides a set of unifying principles that are explanatorily autonomous from lower-level molecular descriptions. For example, the law of independent assortment describes a 'pairing and separation schema' that can be instantiated by an open-ended set of molecular mechanisms. To understand why the transmission of some alleles conform to this schema it is not helpful, Kitcher claimed, to focus on the gory details at the molecular level. However, Ken Waters (1994a, 1994b) objected that these details are in fact important. Such phenomena as crossing over in meiosis raise questions that can only be explained at the molecular level. Waters defended a model of 'informal reduction', one aspect of which is that molecular biology provides a greater understanding of Mendelian genetics. Waters also challenged Hull's claim that there are no bridge principles linking terms in the two theories. The term 'gene' can be defined, Waters claimed, as any relatively short stretch of DNA that functions as a biochemical unit. However, Waters admits that his definition fails to distinguish between several different molecular candidates that might occupy the gene role. There has since been a lively debate over what the term 'gene' refers to (Kitcher, 1982; Burian, 1985; Neumann-Held 1998; Sarkar, 1998; Griffiths and Neumann-Held, 1999; Beurton *et al*., 2000; Gilbert, 2000; Griesemer, 2000; Keller, 2000; Moss, 2003; Stotz *et al*., 2004; Griffiths and Stotz, 2007).

The reductionism debate is playing out again in the context of molecular developmental biology. Only this time the central question is whether developmental regularities identified at the level of the embryo can be reduced or, more specifically, *computed* at the level of gene sequences and proteins. Developmentalists like Gilbert *et al*. (Chapter 21, this volume) and Laubichler and Wagner (Chapter 24, this volume) can be interpreted as presenting a version of Kitcher's explanatory autonomy thesis. They identify embryological structures above the level of the individual cell that supposedly play a significant role in governing gene expression. Although these structures are themselves partly composed of genes, their behaviour allegedly cannot be explained solely in terms of molecular inputs and outputs to a genetic programme. One such structure is the morphogenic field, described by Scott Gilbert, John Optiz and Rudolf Raff in Chapter 21. As these authors explain, the concept of a morphogenetic field

has its origin in the embryological movement of the 1930s. Embryologists at the time were interested in identifying 'the laws of ordered form' that determine how cell lineages organize and differentiate into phenotypic structures. The morphogenic field was proposed as an organizing structure that arranges cells into interacting developmental subunits, analogous to the way that a magnetic field might organize iron filings into patterns. The developmental fate of a cell is determined on this view by its position within a particular field, not by factors internal to it. Gilbert *et al.* complain that the concept of a morphogenic field was abandoned without good reason by proponents of the New Synthesis and their emphasis on genes as the prime movers of development. But the morphogenic field is making a comeback. Brian Goodwin (1982, 1994) is cited by Gilbert *et al*. as being partially responsible for the revival of the morphogenic field concept, 'However, this is a field that is outside developmental genetics and is actively opposed to gene action as being important in field functions' (p. 439). Gilbert and colleagues do not discount the importance of genes in establishing morphogenic fields. In fact, one of the best candidates for a morphogenic field is the protein gradient set up by the expression of Hox genes. These protein gradients convey information to groups of cells that determine which structural genes they express and therefore what types of cell they differentiate into. But genes are not everything here, or so it is argued. Morphogenic fields are thought to supply developing cells with essential positional information that is not genetically encoded. So an explanation of development will have to invoke these higher order structures and the rules that govern their behaviour in addition to the genes they interact with. Notice that the rules governing the activity of morphogenic fields are being assigned a degree of explanatory autonomy reminiscent of the autonomy Kitcher attributed to the laws of Mendelian genetics.

Not everyone views these matters the same way. In Chapter 22, 'Reductionism Redux: Computing the Embryo', Alex Rosenberg argues that the discovery of Hox genes and other advances in molecular developmental biology *strengthen* the case for reductionism. Instead of adopting the Nagelian model of theory reduction like Schaffner, Rosenberg claims that developmental patterns can be 'computed' from a description of the genes and proteins contained in an egg cell. There is a trivial sense in which any law-governed process can be computed. The trajectory of a bouncing ball can be computed using the laws of Newtonian physics and its initial conditions. But Rosenberg has something stronger and more explanation-like in mind. For some process to be computable, he explains, three conditions must obtain. First, the process must be *decidable*: there must be some mathematical function that maps initial states of the process onto its outcome. Second, the process must be *tractable* or decomposable into simpler subprocesses (bouncing balls probably violate this condition). Finally, the system that executes the process must have something analogous to a *generative grammar* or syntax: with a finite set of rules it should be possible to generate an indefinitely large set of possible outcomes. Turning to Hox genes, Rosenberg is impressed by the fact that this relatively small cluster of genes is highly conserved across disparate taxa (appearing everywhere from insects to mammals) and has roughly the same function in such unrelated species. Wherever Hox genes are found, their role is to generate segmentation in body parts. Rosenberg argues that these genes and the proteins they interact with satisfy his three conditions for computability. The process that they initiate is decidable, he claims, because Hox genes operate in a programme-like fashion, responding to molecular inputs and producing protein gradients as outputs. Second, the functions that these genes execute are decomposable into simpler subfunctions.

For example, the complex function of causing cells to differentiate into segmented body parts is partly explained by the simpler function of creating a protein gradient. Most importantly perhaps, Rosenberg argues that Hox genes are generative. Slight genetic modifications can generate an open-ended range of phenotypic modifications. For example, a modification to the *eyeless* gene can cause eyes to sprout at various locations over a fly's body. Rosenberg does not endorse the naive view that genes operate in isolation from extra-genetic factors. But he regards these extra-genetic factors as providing the inputs to a genetic programme. If one knows the programme embodied in a genome plus the extra-genetic factors present in the fertilized egg cell then, he claims, one has sufficient information to compute the embryo.

What about the developmentalist claim that higher-level embryological structures, like the morphogenic fields, are irreducible features of developmental systems? Rosenberg offers two objections to this thesis. First, he claims that embryological generalizations are non-informative. These 'explanations' merely describe how certain cell lineages proceed down particular developmental pathways; they do not cite the causal mechanisms that drive this process. Like Molière's dormitive virtues, these descriptions provide the explanandum for molecular developmental biology, never the explanans. Second, Rosenberg interprets the work of molecular developmental biologists Lewis Wolpert and Walter Gehring to show that talk of morphogenic fields (and other higher-level embryological structures) can be dispensed with by a mechanistic explanation of development. On this view, the so-called morphogenic field is just a gradient of Hox gene transcription factors. Rosenberg argues that it is possible to represent this gradient in computational terms, say, as a linear sequence of instructions. Likewise, since it also seems possible to compute the responses of individual cells to those instructions, Rosenberg argues that there is nothing standing in the way of a computational account of Hox gene function.

Objections to these arguments take issue both with Rosenberg's interpretation of the science and with his metaphysics. In Chapter 23 Evelyn Fox Keller distinguishes between a function that is computable by a physical system and one that is computable by us. Just because a function satisfies the former condition does not guarantee that it satisfies the latter: many law-governed systems are computable in principle but intractable to us. Since Rosenberg plans to substitute computability for the traditional Nagelian notion of reduction as derivation, he requires an account of computability that is sensitive to the limits of human representation. Arguably, Rosenberg's two other requirements of tractability and generativity supply the constraints that he needs. But then the issue becomes whether Hox gene expression is in fact a tractable, generative process. Keller points out that Hox genes are not the prime movers of development. Even they require inputs both from epigenetic factors and from other genes that can be distributed widely throughout the genome. In many cases these interactions involve complex feedback loops, not following the linear, algorithmic process that Rosenberg envisions. So it is misleading, she claims, to isolate this small component of gene regulation and treat it as a tractable, much less a complete explanation of segmentation. According to Keller, a complete account of this or any other aspect of development would reveal developmental systems to be more holistic and less decomposable than Rosenberg portrays them to be. In his defence, Rosenberg might reply that the verdict is still out on this issue. The fact that the Hox gene modules interact with other genetic and epigenetic factors should not cause the pendulum to swing over to what some might consider unbridled holism. Perhaps a complete

and tractable description of gene expression is waiting in the wings and our understanding of Hox genes marks the first piece in this reductionist puzzle.

Manfred Laubichler and Günter Wagner, in Chapter 24, take further issue with Rosenberg's portrayal of molecular developmental biology. One of their central objections states that it is impossible to assign determinate functions to genes without considering contextual factors: 'A molecular or genetic explanation of a developmental phenomenon … only makes sense if there is a strong causal relationship between particular molecular events and the phenomenon to be explained. Such a strong relation is necessary in order to assign a specific causal role to a gene or a molecule' (p. 489). They claim that the phenomenon of genetic redundancy – compensation for the loss of one gene by another that adopts its functional role – suggests that relations between genes and developmental functions is many-to-one. Therefore, they conclude, since some genes do not have specific developmental functions, developmental functions cannot always be reduced to the genetic level. This argument is unlikely to convince reductionists like Rosenberg. The reductionist does not require a one-to-one relation between molecular structures and developmental functions, many-to-one relations will do. If a given developmental process is executed by a range of different genes and molecular processes, then it should be possible to reduce that function to a disjunction (even a long one) of genes and molecular inputs and outputs. Unless Laubichler and Wagner are suggesting that redundancy occurs as a result of some kind of top-down causal process (which would require further argument) this objection does not undermine reductionism. Laubichler and Wagner state that, 'the question we are concerned with here is whether it is meaningful or even possible to assign a particular function to a gene … without referring to the larger molecular, cellular, and organismal context within which genes are expressed. We do not yet know the definitive answer to this question, but the current body of evidence is not encouraging' (p. 489). They go on to cite examples of recent molecular developmental explanations that appeal to positional information and other allegedly irreducible features of developmental systems. However, Rosenberg is unsympathetic to the inference that if molecular biologists currently find it necessary or informative to appeal to higher level structures, those structures are therefore irreducible features of developing systems. As Laubichler and Wagner concede, molecular biology is an incomplete science so a thoroughgoing reduction of developmental to molecular biology might one day become possible. However, these authors reply that Rosenberg's faith in reductionism is nothing more than that: 'In many instances molecules are the relevant entities … but to *a priori* assume that they always have to be the relevant level of description is to adopt a metaphysical position contrary to proper scientific conduct' (p. 489). Of course, Rosenberg offers a simple argument for why the molecular level is the relevant level of explanation for developmental phenomena: because causation occurs at the molecular level and good scientific explanation tracks causation. It is unclear whether Laubichler and Wagner are taking issue with the first claim or the second. If they think that causation sometimes occurs above the molecular level, then they need to show this with a more detailed argument. Simply citing developmental biologists who appeal to higher-level entities won't do. If they think that explanation should sometimes be divorced from causation, then they are committed to a version of developmental biology which, in their own words, 'no longer … represents the ontology of nature but rather one that represents the particular interest of a group of scientists' (p. 484).

Developmental Systems Theory (DST) and the Challenge to Genes as replicators

Regardless of whether one thinks that development is computable at the level of genes and molecules, both reductionists and anti-reductionists agree that organisms must inherit more than just DNA to develop. DNA is an inert molecule that must be embedded in a package of cellular machinery for its potential to be realized. Along with DNA large quantities of RNA and gene transcription proteins are transmitted epigenetically (outside the genome) from mother to offspring via the egg cell. These maternally inherited factors initiate cellular differentiation in the zygote and continue to exert an influence on the developing organism. Moving beyond the egg cell, many species pass on important behavioural traits to their offspring through social learning, for example the mating calls of many oscine birds. Mutualistic organisms like the gut bacteria required for digestion are also transmitted outside the genome. Many of these epigenetically inherited factors are transmitted with the same fidelity as genes and make similar contributions to phenotypic variation. Therefore it is possible to regard these factors collectively as an epigenetic inheritance system, often working in coordination with genes, on which selection also acts.

Despite widespread recognition that DNA is not the only source of inheritable phenotypic variation, genes continue to be assigned priority in developmental and evolutionary explanations. DNA is thought to encode a developmental 'programme' while epigenetic factors provide the 'hardware' for reading and implementing genetic information. Genes are commonly described as the unit of inheritance while epigenetic factors are considered a part of the developmental environment. Recently, these popular metaphors have been severely criticized by proponents of developmental systems theory (DST). DST offers an alternative framework for conceptualizing the role of genes in evolution and development. On this view there is no single locus of control or centralized programme that directs developmental processes. Instead, development is guided by a range of genetic, epigenetic and ecological factors that are distributed throughout the organism and its environment. Since genes are just one of the many essential ingredients for development, it is argued that they have no unique information-bearing role. Proponents of this view also resist the idea that traits are 'transmitted' or 'replicated' each generation as if they were passed on intact from parent to offspring. These metaphors gloss over the complicated processes involved in the construction of an organism, fostering the illusion that no developmental explanation is required for traits that are 'replicated' or 'genetically transmitted'.

In Chapter 25, 'Developmental Systems and Evolutionary Explanation', Paul Griffiths and Russell Gray articulate and defend many of DST's core themes. One of their most influential arguments is a defence of the 'parity thesis' which states that genes do not play a privileged role in development and evolution but are on par with epigenetic factors. Griffiths and Gray note that the traditional view that genes are the central players in evolution and development is defended on the grounds that genes, and genes alone, carry information. However, Griffiths and Gray argue that this claim does not hold up under any consistent application of available information theoretic frameworks. One such framework quantifies the information contained in a signal in terms of the correlation between the signal and its source when the relevant 'channel conditions' are held constant. For example, the amount of information contained in a radio transmission is measured by correlating changes in the signal with changes in its source when the relevant channel conditions (for example atmospheric conditions) are

fixed. Granted, it is possible to apply this framework in a way that assigns an information-bearing role to genes. The phenotype is regarded as a signal that carries information about the presence or absence of certain genes (the source) when the environment (channel) is held constant. But information theory is a flexible tool. As Griffiths and Gray point out, source and channel conditions can be swapped. So the phenotype (signal) can also be viewed as carrying information about the state of the environment (source) against a constant genetic background (channel). Information theory provides no principled basis for preferring one interpretation over the other. Therefore this framework fails to support the distinction between genetic and non-genetic influences on development.

Setting aside the parity argument and turning to the question of replication, Griffiths and Gray argue that a broad view of inheritance systems, one that recognizes the importance of epigenetic factors as well as genes, is inconsistent with the idea that replicators are the fundamental units of evolution. The replicator-centred framework carries at least two implications that DST advocates find objectionable. First, this framework attempts to draw a principled distinction between replicators and their environments which, for reasons just considered, appears to run afoul of the parity argument. Second, the replicator-centred approach regards genes as the beneficiaries of adaptation: the evolutionary costs and benefits of a trait are calculated in terms of how well (or poorly) that trait facilitates genetic replication. But since the phenotypic effects of a gene are contingent on its epigenetic and environmental context, DST theorists argue, the reproductive consequences of individual genes cannot be considered in isolation from these contextual factors.

Instead of viewing replicators as the units of evolution Griffths and Gray focus on the life cycle. A life cycle can be defined roughly as a sequence of developmental events that has evolved to perpetuate itself over successive generations. It is perhaps helpful to think of examples of life cycles that are composed of several discrete but interdependent stages, such the transition from egg to caterpillar to winged adult in butterflies. The central question from a DST perspective is how such processes are repeated each generation and what are the sources of inheritable variation that impact upon them. A subset of the factors that sustain a life cycle are produced each generation by the organism itself. These factors range from basic building blocks, like proteins and nucleic acids, to modifications an organism makes to its niche, like beaver dams and bird nests. In addition to these self-generated factors are ones that can be reliably found in the organism's environment, for example the presence of sunlight or nitrogen. Such 'passively inherited' factors are no less important than genetic and epigenetically inherited resources for maintaining stable, re-occurring life cycles. Therefore Griffiths and Gray regard both actively and passively inherited factors as components of an organism's 'evolutionary developmental system'. A third category of developmental influences are neither produced by the organism nor reliably found in its environment. For example, an encounter with a toxic plant can stunt an organism's growth; gaining access to a new prey item can lead a predator to modify its foraging behaviour. These contingent factors might contribute to a token individual's development and therefore be included in its developmental system. But since these developmental inputs are not reliably encountered by the organism's lineage they are not considered a part of the evolutionary unit in question.

In Chapter 26, 'The Extended Replicator', Kim Sterelny, Kelly Smith and Michael Dickison accept the importance of epigenetic factors in evolution and development, but argue that these insights can be accommodated within a replicator-centred framework. They petition

for an expanded account of replicators that includes both genetic and epigenetically inherited sources of phenotypic variation. However, they draw the line at passively inherited factors, which they regard as a part of the replicator's environment. On this view, only traits that play an active role in self-replication qualify as units of evolution. Maternally inherited cellular machinery, bird nests and beaver dams are included in this list; but sunlight, nitrogen and hermit crab shells are not. Sterelny *et al*. resist the parity argument by appealing to a semantic account of information, or what they call 'representational' information. On this view, a gene or epigenetic factor represents a particular developmental outcome if producing that outcome is its selected function. Many of the non-genetic influences on development that are identified by DST theorists (arguably) do not have these effects as their selected functions.

Sterelny *et al*. are also opposed to viewing life cycles as the fundamental units of evolution. One of their central objections claims that life cycles suffer from a boundary problem: Elvis Presley is a part of Sterelny's developmental system, but surely Elvis and Sterelny should not be regarded as components of the same unit of evolution. So where does one draw the boundaries of a developmental system? In anticipation of such objections Griffiths and Gray (Chapter 25, this volume) draw a distinction between a token developmental system and a developmental unit of evolution (see also Griffiths and Gray, 1997). A token developmental system is the set of resources that contribute to an individual's development. However, a developmental unit of evolution is a process that is sustained by only those developmental resources that are re-occurring or present over successive generations (recall the distinction between self-generated, passively inherited and *contingent* influences on development). Therefore Elvis does not count as unit of evolution. However, Sterelny *et al*. claim, DST provides no principled basis for ruling it out as such.

Sterelny *et al*. argue that the phenomenon of phenotypic plasticity raises a further set of boundary problems. Phenotypically plastic traits are ones that an organism predictably develops in response to certain environmental inputs, like the formation of a scar in response to an injury or the growth of calluses as a result of manual labour. These traits are importantly distinct, they argue, from such idiosyncratic traits as Sterelny's affinity for Elvis or the individual song of a lyrebird. Allegedly DST provides no principled grounds for making this distinction:

> The lyrebird's song is unique to each bird, for they are famous mimics, and pick up all manner of extraneous sounds, including those of humans, their animals and machines … Yet this does not seem to be an 'individual' trait in the same sense that a scarred hand is. Moreover, there is a sense in which the scar has an evolutionary explanation: scarring events are 'historically associated with' the human lineage. There is an evolutionary component of any individual scar construction. So we have our doubts about the robustness of their distinction. (p. 534)

But Griffths and Gray respond that, 'the Lyrebirds' song is an individual trait in *exactly* the same sense as the scarred hand' (1997, p. 476). In both cases, they claim, there is an evolutionary explanation for why humans and lyrebirds are disposed to develop these traits. These dispositions are capacities of their respective lineages regarded as evolutionary developmental systems. By contrast, particular scarred hands or the songs of individual lyrebirds are features of token developmental systems and therefore not relevant from an evolutionary perspective.

Sterelny *et al*. argue that just as life cycles lack discrete spatial boundaries, so too do they lack precise temporal boundaries. The problem is that different components of a developmental system cycle at different rates. Some developmental sequences are repeated several times each generation, such as the construction of bird nests and beaver dams. Other sequences span several generations, like termite mounds and marmot burrows. Without a distinctive 'bottleneck' that distinguishes one generation from another, Sterelny *et al*. argue, there is no way to distinguish evolutionary change in a lineage from developmental change in an individual. For example, what prevents someone from viewing each breeding season as giving rise to a new 'generation' of bird nests, or each generation of termites as developmental stages in the growth of a termite mound? Griffiths and Gray might respond that it comes down to a matter of perspective how one carves up a lineage into temporally defined individuals. Some perspectives have more explanatory utility than others. If most developmental patterns tend to coincide with organismal generations, then it is perhaps useful to individuate token evolutionary developmental systems accordingly. However, Sterelny *et al*. might reply that this strategy implicitly acknowledges the centrality of replicators, especially if it turns out that most developmental sequences coincide with generations of replicators. Griffiths and Gray (1997) argue that the extended replicator and DST perspectives are formally equivalent to one another. Although Sterelny *et al*. prefer the former for heuristic purposes, Griffiths and Gray argue that there are some contexts in which DST provides theoretical insights that are less obvious from a replicator-centred perspective.

Conclusion

We began this introduction by noting that contemporary debates in the philosophy of evolutionary biology are an outgrowth of the ideas that Darwin presented in 1859. It is clear that advances in evolutionary and developmental biology call for a refinement of some of Darwin's original ideas. There appear to be several notions of biological function employed by evolutionary biologists, not all of them reducible to selective history. Evolutionary biologists have developed several methods for identifying whether a trait is an adaptation and what it might be an adaptation for. However, the relative strength of adaptation compared to alternative processes like drift remains controversial. The precise formulation of the principle of natural selection is another topic of ongoing debate. However, an important advance on this front involves the distinction between drift and selection viewed as processes and as outcomes. The rise of multi-level selection theory involves a return to Darwin's position that selection can operate on groups as well as at lower levels of biological organization, like individual organisms and alleles. However, it remains to be seen whether selection also acts above the level of the group, for instance on entire species or biological communities. Finally, it remains controversial whether recent findings in developmental biology undermine the idea that genes are 'central players' in evolution and development. This is currently one of the most active areas of research within the philosophy of evolutionary biology.

Bibliography

Barkow, Jerome, Cosmides, Leda and Tooby, John (1992), The Adapted Mind: Evolutionary Psychology and the Generation of Culture, New York: Oxford University Press.

Beurton, Peter, Falk, Raphael and Rheinberger, Hans-Jörg (2000), The Concept of the Gene in Development and Evolution: Historical and Epistemological Perspectives, Cambridge: Cambridge University Press.

Bigelow, John and Pargetter, Robert (1987), 'Functions', Journal of Philosophy, 84, pp. 181–96.

Bolker, Jessica A. (2000), 'Modularity in Development and Why it Matters to Evo-Devo', American Zoologist, 40, pp. 770–776.

Brakefield, P.M. and Roskam, J.C. (2006), 'Exploring Evolutionary Constraints is a Task for an Integrative Evolutionary Biology', American Naturalist, 68(6), pp. S4–S13.

Brandon, Robert N. (1978), 'Adaptation and Evolutionary Theory', Studies in the History and Philosophy of Science, 9, pp. 181–206.

Brandon, Robert N. (1990), Adaptation and Environment, Princeton, NJ: Princeton University Press.

Breden, Felix and Wade, Michael J. (1989), 'Selection Within and Between Kin Groups of the Imported Willow Leaf Beetle', American Naturalist, 134, pp. 35–50.

Burian, Richard M. (1985), 'On Conceptual Change in Biology: The Case of the Gene', in David J. Depew and Bruce H. Weber (eds), Evolution at a Crossroads: The New Biology and the New Philosophy of Science, Cambridge, MA: MIT Press, pp. 21–42.

Claridge, Michael F., Dawah, Hassan A. and Wilson, Michael R. (1997), Species: The Units of Biodiversity, London: Chapman and Hall.

Cummins, Robert E. (1975), 'Functional Analysis', Journal of Philosophy, 72: 741–764.

Darwin, Charles R. (1859), On the Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life, London: John Murray.

Darwin, Charles R. (1871), The Descent of Man, and Selection in Relation to Sex, London: John Murray.

Darwin, Charles R. (1872), The Expression of the Emotions in Man and Animals, London: John Murray.

Dawkins, Richard (1982), The Extended Phenotype: The Long Reach of the Gene, New York: Oxford University Press.

Dawkins, Richard (1986), The Blind Watchmaker, New York: Norton.

Dennett, Daniel C. (1995), Darwin's Dangerous Idea: Evolution and the Meanings of Life, New York: Simon and Schuster.

Dugatkin, Lee (1997), Cooperation among Animals: An Evolutionary Perspective, Oxford: Oxford University Press.

Dugatkin, Lee and Reeve, H.K. (1994), 'Behavioral Ecology and Levels of Selection: Dissolving the Group Selection Controversy', Advances in the Study of Behavior, 23, pp. 101–33.

Endler, John (1986), Natural Selection in the Wild, Princeton, NJ: Princeton University Press.

Ereshefsky, Marc (1992), The Units of Evolution: Essays on the Nature of Species, Cambridge, MA: MIT Press.

Freeman, Patricia W. (2000), 'Macroevolution in Microchiroptera: Recoupling Morphology and Ecology with Phylogeny', Evolutionary Ecology Research, 2, pp. 317–35.

Futuyma, Douglas J. (1998), Evolutionary Biology (3rd edn), Sunderland, MA: Sinauer Associates.

Gilbert, Scott (2000), 'Genes Classical and Genes Developmental: The Different Use of Genes in Evolutionary Syntheses', in Peter Beurton, Raphael Falk and Hans-Jörg Rheinberger (eds), The Concept of the Gene in Development and Evolution: Historical and Epistemological Perspectives, Cambridge: Cambridge University Press, pp. 178–92.

Godfrey-Smith, P. (1993), 'Functions: Consensus without Unity', Pacific Philosophical Quarterly, 74, pp. 196–208.
Godfrey-Smith, P. (2001), 'Three Kinds of Adaptationism', in Steven H. Orzack and Elliott Sober (eds), Adaptationism and Optimality, Cambridge: Cambridge University Press, pp. 335–57.
Goldschmidt, Robert B. (1945), 'The structure of podoptera, a homeostic mutant in Drosophila naelanogaster', Journal of Morphology 77, pp. 71–101.
Goodwin, Brian (1982), 'Development and Evolution', Journal of Theoretical Biology, 97, pp. 43–55.
Goodwin, Brian (1994), How the Leopard Changed its Spots: The Evolution of Complexity, New York: Scribners.
Gould, Stephen J. and Lewontin, Richard C. (1979), 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique Of The Adaptationist Programme', Proceedings Of The Royal Society of London, Series B, 205: 581-598.
Griesemer, James (2000), 'Reproduction and the Reduction of Genetics', in Peter Beurton, Raphael Falk and Hans-Jörg Rheinberger (eds), The Concept of the Gene in Development and Evolution: Historical and Epistemological Perspectives, Cambridge: Cambridge University Press, pp. 240–85.
Griffiths, Paul. E. and Gray, Russell D. (1997), 'Replicator II: Judgement Day', Biology and Philosophy, 12(4), pp. 471–92.
Griffiths, Paul E. and Neumann-Held, Eva (1999), 'The Many Faces of the Gene', BioScience, 49(8), pp. 656–62.
Griffiths, Paul E. and Stotz, Karola (2007), 'Gene', in Michael Ruse and David Hull (eds), Cambridge Companion to Philosophy of Biology, Cambridge: Cambridge University Press, pp. 85–102.
Hamilton, W.D. (1964), 'The Genetical Evolution of Social Behavior, 1 and 2', Journal of Theoretical Biology, 7, pp. 1–52.
Horan, Barbara L. (1994), 'The Statistical Character of Evolutionary Theory', Philosophy of Science, 61, pp. 76–95.
Hull, David (1974), Philosophy of Biological Science, Englewood Cliffs, NJ: Prentice-Hall.
Keller, Evelyn Fox (2000), The Century of the Gene, Cambridge, MA: Harvard University Press.
Kerr, Ben and Godfrey-Smith, Peter (2002), 'Individualist and Multi-level Perspectives on Selection in Structured Populations', Biology and Philosophy, 17: 477–517.
Kettlewell, Bernard (1955), 'Selection Experiments on Industrial Melanism in the Lepidoptera', Heredity, 9, pp. 323–342.
Kitcher, Philip (1982), 'Genes', British Journal for the Philosophy of Science, 33, pp. 337–59.
Kitcher, Philip (1984), '1953 and All That: A Tale of Two Sciences', Philosophical Review, 93, pp. 335–73.
Lewontin, Richard C. (1970), 'The Units of Selection', Annual Review of Ecology and Systematics, 1, pp. 1–18.
Lloyd, Elisabeth (2005), The Case of the Female Orgasm: Bias in the Science of Evolution, Cambridge, MA: Harvard University Press.
Matthen, Mohan and Ariew, André (2002), 'Two Ways of Thinking about Fitness and Natural Selection', Journal of Philosophy, 99, pp. 55–83.
Mayr, Ernst (1961), 'Cause and Effect in Biology: Kinds of causes, predictability, and teleology are viewed by a practicing biologist', Science, 134: 1501–1506.
Mayr, Ernst (1983), 'How to Carry Out the Adaptationist Program?' The American Naturalist, 121: 324–334.
Millikan, R.G. (1989), 'In Defense of Proper Functions', Philosophy of Science, 56, pp. 288–302.
Millstein, Roberta L. (1996), 'Random Drift and the Omniscient Viewpoint', Philosophy of Science, 63(3) (Supplement), pp. S10–S18.
Moss, Lenny (2003), What Genes Can't Do, Cambridge, MA: MIT Press.

Neander, Karen (2002), 'Types of Traits: The Importance of Functional Homologues', in André Ariew, Robert Cummins and Mark Perlman (eds) Functions: New Readings in the Philosophy of Psychology and Biology, Oxford: Oxford University Press, pp. 390–416.

Neumann-Held, E.M. (1998), 'The Gene is Dead – Long Live the Gene: Conceptualising the Gene the Constructionist Way', in P. Koslowski (ed.), Sociobiology and Bioeconomics: The Theory of Evolution in Biological and Economic Theory, Berlin: Springer-Verlag, pp. 105–37.

Orzack, Steven H. and Sober, Elliott (1996), 'How to Formulate and Test Adaptationism', The American Naturalist, 148(1), pp. 202–10.

Okasha, Samir (2006), Evolution and the Levels of Selection, Oxford: Oxford University Press.

Pettigrew, John D. (1999), 'Electroreception in Monotremes', Journal of Experimental Biology, 202: 1447–1454.

Rosenberg, Alex (1985), The Structure of Biological Science, New York: Cambridge University Press.

Rosenberg, Alex (1988), 'Is the Theory of Natural Selection a Statistical Theory?', Canadian Journal of Philosophy (Supplement), 14, pp. 187–207.

Rosenberg, Alex (1994), Instrumental Biology or the Disunity of Science, Chicago: University of Chicago Press.

Rosenberg, Alex (2001a) 'Reductionism in a Historical Science', Philosophy of Science, 68, pp. 135–64.

Rosenberg, Alex (2001b) 'Indeterminism, Probability, and Randomness in Evolutionary Theory', Philosophy of Science, 68, pp. 536–44.

Rosenberg, Alex and McShae, Daniel (2006), Philosophy of Biology: A Contemporary Introduction, Abingdon: Routledge

Rosenberg, Alex and Neander, Karen (2009), 'Are Homologies (Selected Effect or Causal Role) Function Free?' Philosophy of Science, 76: 307–334.

Rosenberg, Alexander (1983), 'Fitness', Journal of Philosophy, 80: 457–474.

Rosenberg, Alexander (2001), 'How is Biological Explanation Possible', British Journal for the Philosopy of Science, 52: 735–760.

Rosenberg, Alexander and Williams, Mary B. (1986), 'Fitness as Primitive and Propensity', Philosophy of Science, 53: 412–418.

Sarkar, Sahotra (1998), Genetics and Reductionism, Cambridge: Cambridge University Press.

Schaffner, Kenneth (1967), 'Approaches to Reduction', Philosophy of Science, 34, pp. 137–47.

Schaffner, Kenneth (1969), 'The Watson-Crick Model and Reductionism', British Journal for the Philosophy of Science, 20, pp. 325–48.

Sober, Elliott (1984), 'Fact, Fiction, and Fitness: A Reply to Rosenberg', Journal of Philosophy, 81(7), pp. 372–83.

Sober, Elliott (1987), 'Does "Fitness" Fit the Facts?: A Reply to Williams and Rosenberg's Rejoinder', Journal of Philosophy, 84, pp. 220–223.

Sober, Elliott (1990), 'Let's Razor Ockham's Razor', in Dudley Knowles (ed.), Explanation and Its Limits, Cambridge: Cambridge University Press, pp. 73–94.

Sober, Elliott (1999), Philosophy of Biology, Boulder, CO: Westview Press.

Sober, Elliott and Lewontin, Richard C. (1982), 'Artifact, Cause and Genic Selection', Philosophy of Science, 49: 157–180.

Sober, Elliott and Wilson, David Sloan (1998), Unto Others: The Evolution and Psychology of Unselfish Behavior, Cambridge, MA: Harvard University Press.

Stephens, Christopher (2004), 'Selection, Drift, and the "Forces" of Evolution', Philosophy of Science, 71(4), pp. 550–70.

Sterelny, Kim (2000), 'Development, Evolution, and Adaptation', Philosophy of Science, 67, pp. S369–S387.

Sterelny, Kim and Kitcher, Philip (1988),'The Return of the Gene', Journal of Philosophy, 85: 339–36.

Sterelny, Kim and Griffiths, Paul E. (1999), Sex and Death: An Introduction to the Philosophy of Biology, Chicago: University of Chicago Press.

Stotz, Karola, Griffiths, Paul E. and Knight, Rob D. (2004), 'How Scientists Conceptualise Genes: An Empirical Study', Studies in History & Philosophy of Biological and Biomedical Sciences, 35(4), pp. 647–73.

Trivers, Robert (1971), 'The Evolution of Reciprocal Altruism', Quarterly Review of Biology, 46, pp. 35–57.

Waters, Kenneth (1994a), 'Why the Antireductionist Consensus Won't Survive the Case of Classical Mendelian Genetics', in E. Sober (ed.), Conceptual Issues in Evolutionary Biology, Cambridge, MA: MIT Press, pp. 401–17.

Waters, K. (1994b), 'Genes Made Molecular', Philosophy of Science, 61, pp. 163–85.

Whitehead, Alfred North (1929), Process and Reality, London, UK: Macmillan Publishing.

Williams, George C. (1966), Adaptation and Natural Selection, Princeton NJ: Princeton University Press.

Williams, Mary B. (1973), 'Falsifiable Predictions of Evolutionary Theory', Philosophy of Science, 40: 518–537.

Wilson, David Sloan (1980), The Natural Selection of Populations and Communities, Cummings, CA: Benjamin.

Wilson, David Sloan and Sober, Elliott (1994), 'Re-introducing Group Selection to the Human Behavioural Sciences', Behavioural and Brain Sciences, 17(4), pp. 585–654.

Wilson, Robert (1999), Species: New Interdisciplinary Essays, Cambridge, MA: MIT Press, pp. 49–90.

Wouters, Arno (2003), 'Four Notions of Biological Function', Studies in History and Philosophy of Biology and Biomedical Science, 34, pp. 633–68.

Wouters, Arno (2005), 'The Function Debate in Philosophy', Acta Biotheoretica, 53, pp. 123–51.

Wright, Larry (1973), 'Functions', Philosophical Review, 82, pp. 139–68.

Wynne-Edwards, V.C. (1962), Animal Dispersion in Relation to Social Behaviour, Edinburgh: Oliver & Boyd.

Name Index