Spring 1-1-2011

Biological Mechanisms and Evolution

Michael Zerella
University of Colorado at Boulder, mjzerella@gmail.com

BIOLOGICAL MECHANISMS AND EVOLUTION

by

MICHAEL ZERELLA

B.S., University of Arizona, 1993

M.S., California Polytechnic University, Pomona, 1995

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirement for the degree of
Doctor of Philosophy
Department of Philosophy
2011

This thesis entitled:
Biological Mechanisms and Evolution
written by Michael Zerella
has been approved for the Department of Philosophy

Professor Carol Cleland

Assistant Professor Robert Hanna

Associate Professor John Jackson

Associate Professor Bradley Monton

Professor Robert Rupert

Date

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

ABSTRACT

Zerella, Michael (Ph.D., Philosophy)
Biological Mechanisms and Evolution
Thesis directed by Professor Carol Cleland

Recent work in philosophy of biology has challenged the traditional assumption in philosophy of science that proper scientific explanations and predictions must be grounded in laws of nature that hold with necessity. Philosophical work on mechanisms in biology and cognitive science has shown that the traditional view, on its own, is too restrictive to account for scientific explanations and predictions in the complex, contingent realm of living systems. In the first part of my dissertation, I defend the view that a mechanistic approach to biology is preferable to the traditional approach, and I defend my own realist account of biological mechanisms. While most authors seem to consider the mechanistic approach in the life sciences to be a heuristic device rather than an objective description of nature, I argue that biological mechanisms are real in the same sense that planets, molecules, and life itself are real. Specifically, I argue that a biological mechanism is a structure or process that is part of and maintained by a living organism, and works to promote continuation of the ongoing living process of which the mechanism is a part. So, as long as organisms and life are real, biological mechanisms are also real.

In the second part of my dissertation, I apply the mechanistic approach defended in Part One to the three process that together result in biological evolution: hereditary reproduction, generation of new variations, and natural selection. I show that the mechanistic approach helps clarify our understanding of how each process works (and does not work) on its own, and how each contributes to biological evolution. I consider whether each process, along with the physical realizers that bring it about, meet the criteria for being a biological mechanism as laid

out in Part One. All three processes are often referred to as mechanisms, but I conclude that while there are biological mechanisms for achieving heredity and generation of new variations, the process of natural selection does not meet the criteria for being a biological mechanism, nor does the overall process of biological evolution.

ACKNOWLEDGEMENTS

I would like to thank the Department of Philosophy and the University of Colorado for allowing me the opportunity to pursue a doctorate in philosophy. With their training, guidance, and financial support, I have made great strides academically and as a teacher. In particular, I would like to thank my committee chair, Carol Cleland, for her many comments and editorial contributions, and for her enduring enthusiasm.

I would also like to thank the rest of my committee, Robert Rupert, Bradley Monton, Robert Hanna, and Eric Chwang, and John Jackson, for their support and generous comments on my work. I also am indebted to my fellow graduate students for helping me through my time in Boulder, both academically and socially. Many fine people come through this department, and I have made some very good friends. Special thanks go to the departmental staff, Maureen Detmer, Karen Sites, and Timothy Putnam for their consistently helpful assistance and for their friendship.

I also thank my parents, Diana and Joseph, for their unwavering love and support. I never would have been able to make it through this program (or anything else) without them.

CONTENTS

INTRODUCTION

This dissertation is a work in philosophy of science and philosophy of biology. Science is in the business of exploring and understanding natural phenomena, and with that understanding, scientists can explain how and why those phenomena occur. Furthermore, they are also well equipped to generate predictions about what will happen under current conditions as well as to offer counterfactual statements about what might happen if conditions were other than they actually are. Since it is clear that scientists, by and large, are successful at understanding nature, I believe that philosophy of science should pay close attention to the actual practices of scientists, including the conceptual framework and ontological commitments under which they achieve such success. That said, it is also true that investigations of nature remain incomplete, and scientific understanding is regularly revised and improved. With this in mind, one of my central goals is to describe the thinking of mainstream biologists, and then propose and defend a variety of updates and improvements. More precisely, my dual goals for this dissertation are 1) to better understand how biology works as a field of science, and 2) to better understand the processes at work in evolution by natural selection. In this introductory section I summarize my project and my approach to understanding of biological evolution.

My first two chapters establish that scientists in general, and biologists in particular, no longer see themselves as searching for laws of nature. Instead, they are searching for mechanisms of nature. This does not, by itself, imply that there are no laws of nature or that scientists are not committed to their existence. Rather, my position is that while there may be laws of nature underpinning the basic causal structure of the universe, discovery of a law is not a major goal of science as it is practiced today. In order to understand nature, scientists are

searching for mechanisms. In Chapter One, I describe the nature of biological research in order to argue that scientific explanations and predictions in biology are not based on laws of nature, but rather on biological mechanisms. I call this the mechanistic approach, and in Chapter Two, I defend my own account of biological mechanisms at the heart of the mechanistic approach. A biological mechanism, I argue, is any process that 1) achieves a biological function, where a biological function promotes the continuation of the life of which the process is a part, and 2) is composed of entities that achieve their function non-aggregatively and are part of the organism whose life is being supported by the mechanism. The remaining chapters assume the mechanistic approach in understanding and describing biological systems.

Chapters Three, Four, and Five are organized around my belief that adaptive evolution is the result of three separate processes working in tandem: hereditary reproduction, generation of new variations, and natural selection. Evolution in general is the accumulation of changes over time, and in biological evolution the most important changes that accumulate are, by and large, adaptive changes. In order for ongoing accumulation of adaptive changes, all three processes must be present and active (Lewontin, 1970; Lloyd, 2005).

Chapter Three covers hereditary reproduction. This is a necessary element of adaptive evolution because in order for changes to accumulate, any adaptive traits that help an individual succeed in its environment must be reliably inherited by offspring. Chapter Four concerns generation of new variations, as there must be a source of new traits in order for an evolutionary process to continue indefinitely. Without new traits, a population will eventually settle on the most adaptive set of already existing traits and no further adaptive evolution will occur.[1] Finally, Chapter Five concerns natural selection, which occurs whenever an individual is struggling to

[1] Plus, without a way to generate new traits, the population may be unlikely to survive any major shift in environmental conditions.

survive in its environment. That struggle necessarily involves complex, ongoing interactions between the individual and its environment, and if the interactions don't destroy the individual, then it has been naturally selected for continued survival and may go on to reproduce.[2]

A central argument of Chapters Three, Four, and Five is that while all three processes are needed for long-term, adaptive evolution, each of the three can occur on its own, without either of the other two. However, in life as we know it, all three processes usually are intimately intertwined, so I will take some time to describe what I take to be their individual contributions to adaptive evolution and what can (and cannot) happen when only one or two of the three processes are in place. This will set the stage for the more detailed discussions in Chapters Three, Four, and Five.

If only one of the three elements of adaptive evolution is in place, then while some interesting things might occur, long-term adaptive evolution in very unlikely to be one of them. Suppose we have only diversity among individuals in a population, but no hereditary reproduction and no environmental interactions affecting survival. In order to consider what would happen in such a situation, let us set aside how it could be achieved or whether that is even possible for any actual organisms.[3] If there is no reproduction of any kind, then there cannot be any new individuals (barring immigration) and the population cannot change by growing. If there is reproduction, but it is not hereditary reproduction, then while the population could grow, any traits possessed by new individuals would be unrelated to any of their parents' traits, and so there could be no reliable accumulation of changes and, therefore, no evolution.

---

[2] Many authors include reproduction in their accounts of natural selection, but I maintain that reproduction is a separate process from natural selection. This is defended in Chapter Five.

[3] If we want to flesh out the possibility more thoroughly, science fiction may provide some inspiration. Consider some kind of fusion powered life forms existing in interstellar space and with an on-board fuel supply. They face no specific dangers or lack of resources, and do not have much in the way of environmental interactions, so they are not subjected to natural selection.

There only would be ever-increasing amounts of variation. Further, if there is diversity but no interaction affecting survival (no natural selection), then the concept of adaptedness does not even apply because adaptedness necessarily involves an ability to survive in an otherwise hostile environment. Actual biological organisms are described as well adapted only because they are able to survive despite environmental challenges such as oxidation, freezing, over-heating, dilution, desiccation, and attack from other organisms. If there are no environmental interactions like these that affect continued survival, then there is no sense in which individuals can be adapted to their environments because there is nothing to adapt to. Therefore, adaptive evolution would be impossible. For all of these reasons, generation of new variations by itself is insufficient for adaptive evolution.

Next, suppose we have only hereditary reproduction without any diversity or environmental interaction affecting survival. As before, if there is no environmental interaction affecting survival, then there is no sense in which individuals can be adapted (or maladapted), and so there can be no adaptive evolution. Further, if there is no diversity, then even if individuals are successfully reproducing and the population increases, all individuals have the same traits and reproductive capacities, so there is no change in the population (other than its size), and so there is no adaptive evolution. Therefore, hereditary reproduction, by itself, is not sufficient for adaptive evolution.

Now suppose we have environmental interaction affecting survival (natural selection), but no diversity or hereditary reproduction. In this situation, there can be a sense in which individuals are adapted (or maladapted) to their environments because an individual's set of traits can have an affect on its likelihood of survival. However, if there is no diversity, then even if some individuals end up surviving while others die, the population will not change, except to

get smaller. There will be no new traits or changes in frequencies of existing traits, and so there will be no adaptive evolution.[4] Further, if there is no hereditary reproduction, then even if the existing individuals are very well adapted to their environments, they cannot create offspring with the same or similar traits. If they cannot reproduce at all, then the population will either remain static or decline to extinction. If they can reproduce, but not with heredity, then any adaptive traits of parents will not be preserved in offspring, and there can be no adaptive evolution. Therefore, natural selection by itself is not sufficient for ongoing adaptive evolution.

We also can consider what can happen when only two out of the three elements are in place. If hereditary reproduction and a source of variations are in place, then a population can increase its numbers, and it can even accumulate changes. However, the evolution of such a population would not be adaptive evolution in the relevant sense. As already discussed above, if there is no natural selection, then there are no positive or negative consequences of having one set of traits over another. Therefore, there is no sense in which one individual is any better adapted than any other, and so there can be no adaptive evolution. Rather, most evolutionary changes would be roughly random.[5] The only directionality to any evolution in this population would be toward greater reproductive potential. If reproductive potential—the speed of reproduction and number of offspring per generation—is one of the traits that varies between individuals, then there can be evolution of ever-increasing reproductive potential because the faster reproducers will outpace the slower reproducers.[6] Although I maintain that reproduction is separate from natural selection, this phenomenon is similar to natural selection, and is sometimes

[4] In addition, if there is no adaptive evolution in this population, then the origin of these already-adapted individuals is left unexplained.

[5] Of course, potential changes are always constrained by what reproduction and development are able to support.

[6] This is closely related to "fitness," which is a measure of reproductive success.

called “reproductive selection”. However, reproductive selection is not the same as “viability selection,” and only the latter is central to adaptive evolution. Reproductive selection does require a minimum amount of viability, but reproductive selection by itself does not promote evolution of environmental adaptations. In fact, reproductive potential can be increased by diverting energy or other resources toward reproduction at the expense of individual viability. Therefore, if there is no viability selection favoring individuals that are better adapted to surviving in their environments, evolution will not lead to increasing adaptedness even if it leads to increasing rates of reproduction.[7]

Now suppose we have hereditary reproduction and environmental interaction affecting survival or reproduction, but no diversity. In that case, even though individuals may be adapted (or mal-adapted) to their environments, there are no differences in adaptedness between individuals, so the population cannot experience adaptive evolution and may very likely go extinct, especially if the environment is changing. Some individuals may survive and reproduce while others do not, but the traits possessed by individuals in the population cannot change over time if there is no diversity and no source of any new variations. In such a scenario, extinction seems very likely, especially in a changing environment.

Finally, suppose we have variation and environmental interaction affecting survival, but no hereditary reproduction. In this scenario there actually could be some change in average adaptedness in the population, but there is not likely to be ongoing adaptive evolution. A source of new variations means that some individuals might be better adapted to the environment than others, and when less adapted individuals begin to die at a faster rate, the population will be left with a greater proportion of better-adapted individuals than it started with. However, as long as

[7] This means that viability selection and reproductive selection usually are opposing evolutionary tendencies, and long-term continuation of a lineage usually requires balance between the two, rather than maximization of either one by itself.

the individuals have limited life spans and a limited ability to generate new variations in themselves, the population will be unable to continue any further adaptive evolution. This amounts to what Dawkins (1980) calls "one-off selection" in which interactions between environment and individuals leads to elimination of some and preservation of others, and then that's it. However, if the individuals remaining after one round of selection are able to continue generating new variations in themselves without reproducing, it could be that some of the newly-generated traits are more adaptive than others, and the process of differential natural selection between individuals could continue, resulting in a remaining population that, on average, is even better adapted than it was before. Thus, there could be some adaptive evolution even without reproduction. Without a source of new individuals, however, even this process is destined to terminate when the last individual finally dies.

Only when all three processes occur together can a population sustain long-term adaptive evolution, and obviously this is the case in actual biological populations.[8] In most biological species, all three processes are intimately bound up with one another so that it can be difficult to distinguish any one of the processes from the other two. Nevertheless, I maintain that they all are separate and distinguishable and can be fully characterized individually. This is done in Chapters Three, Four, and Five.

In Chapter Three, I clarify contemporary views on biological heredity, including historical sources of the gene-centric concepts and assumptions that typically frame most of the work on heredity. I then argue that proper understanding of heredity, supported by the mechanistic approach, makes clear why the traditional, gene-centric approach is too restrictive. I then present a variety of non-genetic processes of heredity and assess the extent to which they

[8] Other possible explanations for adaptation, such as divine intervention and Lamarckian evolution, are discussed in Chapter Five.

are biological mechanisms on par with genetic mechanisms. I also use the mechanistic approach to diagnose some special challenges faced by Developmental Systems Theory, which has become a prominent alternative to gene centrism.

In Chapter Four, I clarify contemporary views on generation of new variations, including historical sources of current concepts and assumptions, and the reasons for rejection of Lamarckian evolution. I then argue that proper understanding of generation of new variations, supported by the mechanistic approach, makes clear why the traditional, gene-centric approach inappropriately ignores non-genetic mechanisms that generate new variations. I also argue in favor of a reintroduction of directed generation of variation and inheritance of acquired characteristics, which are ideas traditionally associated with Lamarckian evolution.

In Chapter Five I clarify contemporary views on natural selection, including historical sources of current concepts and assumptions. For the previous two processes, heredity and variation, debates usually have been concerned with how the processes are realized, not what the processes are. When it comes to natural selection, however, there continues to be a great deal of debate over process itself, how it works, and whether it makes sense to call it a mechanism. I argue that most of the disagreements over natural selection stem from a failure to consistently distinguish natural selection from reproduction or generation of new variations. My own account of natural selection maintains the proper distinctions, and I use that account to resolve some of the most prominent debates over natural selection.

CHAPTER ONE

THE MECHANISTIC APPROACH AND THE TRADITIONAL VIEW

1.1 Introduction

A dominant tradition in philosophy of science has maintained that scientists rely on laws of nature to support their explanations, predictions and counterfactuals. Indeed, scientific work rooted in Newtonian mechanics relied heavily on laws of nature to make sense of the world, but more recent work in all fields of science has moved away from explicit reference to laws of nature and instead is more likely to rely on so-called mechanisms. However, it is not entirely clear what scientists mean when they claim to be studying mechanisms, and scientists from different fields may not even mean the same thing by "mechanism." This first chapter is dedicated to sorting out these issues, and I defend a contemporary version of what I call "the mechanistic approach" for developing and justifying scientific explanations, predictions and counterfactual statements in the biological sciences. Although the mechanistic approach is often contrasted with a traditional view that relies on laws of nature and formal, logical arguments to generate and justify scientific explanations, predictions and counterfactuals, I argue that the mechanistic approach is compatible with and even complementary to the traditional approach.

In Section 1.2, I present the relevant history of the use of mechanism and laws of nature in science. I show that the mechanistic approach I advocate is not the same as what is often called the mechanical view that grew out of the work of Descartes, Boyle, and Newton. Unlike the mechanical view, the mechanistic approach is not necessarily linked to any particular theory from physics (such as Newtonian mechanics), is not synonymous with naturalism (though it may entail naturalism), and does not imply reduction of the special sciences to physics.

In Section 1.3, I describe more carefully the traditional view that relies on laws of nature used in formal, logical arguments to justify scientific explanations, predictions and counterfactuals. While I am not inclined to endorse that traditional view, I do not argue for its replacement. Rather, I argue that by itself, the traditional view provides an insufficient account of how scientists actually develop and justify their explanations, predictions, and counterfactuals, and that the mechanistic approach complements the traditional view by filling in gaps in the traditional view.

Section 1.2 History

In this section, I present a historical overview of the ways in which scientists and philosophers have formulated their understanding and explanations of nature. This is intended to provide context for the rest of the chapter rather than present any new historical work. However, within this overview I make the case that the mechanistic approach I defend is not necessarily linked to any particular theory from physics (such as Newtonian mechanics), is not synonymous with naturalism (though it typically comes with a commitment to naturalism), and does not imply reduction of special sciences to physics.

A naturalistic approach to characterizing and understanding nature has been a rival to supernatural views at least as far back as Ancient Greece. Atomists, such as Democritus and Epicurus believed that all natural phenomena are the result of particles of different sizes and shapes moving, combining, and recombining in different ways (see Boas, 1952 for an overview). This was a reductive view of nature in the sense that all macroscopic phenomena were thought to be explainable in terms of the action and interaction of tiny, indivisible particles.

An opposing, nonreductive view came from Aristotle, especially in the *Posterior Analytics*, who believed that in addition to any bits of material, our understanding and explanation of natural phenomenon must include irreducible formal causes (essences, substantial forms) and final causes (innate functions or purposes). In *De Anima*, for example, Aristotle explains that all living creatures are held together by their essences, or souls, and have the innate purpose of continuing to live and reproduce. The basic materials of which organisms are made is not enough, by itself, to cause organisms to exhibit their distinctive physical characteristics or their purposive behavior, and so, according to Aristotle, an explanation of life cannot be reduced to mere atomic interactions.

Another relevant contribution from Aristotle, found mainly in his *Posterior Analytics*, is the use of logical demonstration to explain natural phenomena. According to Aristotle's approach, scientific explanation of a natural phenomenon should take the form of a deductively sound argument, the conclusion of which is a description of the phenomenon. Although Aristotle's specific ontology of nature eventually was rejected, his general method of using logical arguments to generate scientific explanations has continued in modern philosophy of science (e.g. Hempel, 1962, 1965; Railton, 1978).

Aristotle's ontology eventually was rejected by the scientific community because, during the European Renaissance, scientific explanations in terms of Aristotelian substantial forms or teleological forces was having trouble explaining some the new knowledge being gathered by naturalistic philosophers. Consequently, the Aristotelian view was in need of radical updating or replacement by a better theoretical foundation. Alchemists tended to favor the former strategy, allowing for a greater variety of substantial forms, some of which could shift from material to material, thus accounting for some of the radical changes observed in chemical reactions. For

example, the alchemists' view famously implied the possibility of transmutation since the substantial forms associated with gold could move into some other substance, thus converting that substance into gold.

Those who wanted to reject and replace Aristotelian science, particularly Gassendi, Galileo, Descartes and Boyle, wanted to replace it with something very much like the reductive atomist view from ancient Greece in which macroscopic phenomena are explained in terms of the characteristics and interactions of microscopic particles of various shapes and sizes (Boas, 1952). That view was able to provide superior explanations and predictions for technologically important phenomena such as wind, air pressure, phase changes and the action of heat. Those examples are technologically important because they were (and still are) used to support mechanical innovation and industrialization. Thus, the resurgent atomic theory had a huge advantage over its rivals because of its ability to yield practical results for mechanical engineers, and so it slowly started to dominate scientific theorizing. Perhaps since it coincided with and supported the rise of mechanical technology, the reductive, atomist view became known as a mechanical view[9] of nature in which all natural phenomena could be understood and explained in terms of constant pushing, pulling, bumping, sticking, separating, combining, and recombining of microscopic particles of various shapes, sizes, and velocities. As the mechanical view became the dominant view of nature in the Western world, Aristotelian substantial forms and final causes, or anything similar, were expunged from science. Any theorizing not consistent with the new view was (and, to a large extent, continues to be) labeled religious, occult, magical, or otherwise non-scientific thinking, and therefore not likely to reflect the real world (Muenzinger, 1935; Mayr, 1982).

[9] I use the term '*mechanical* view' here (rather than 'mechanistic') because, at the time, the view is closely connected to Mechanics as a theory in physics, and I want to distinguish this mechanical view from the mechanistic view that I ultimately wish to defend.

Let us take a closer look at why the mechanical approach coming out of the 17th century was so successful. A big part of its power came from its reduction of complex processes to the collective action of simple entities that are describable in precise terms and whose relationships can be represented mathematically. Individual elements of a system could be understood in isolation, and once each element is understood individually, one can come to understand the whole, complex system. Then, since the relationships among the parts could be represented mathematically, the mechanical approach was able to generate very precise and accurate predictions and counterfactual statements that led to great success in manipulating nature and creating useful technology. Some of the relationships described in scientists's equations were so reliable that they were believed to be laws of nature in the sense that they were universal and exceptionless relations or regularities. As described more fully in Section 1.3, laws of nature can be used to support deductive arguments (what Aristotle called "demonstrations") that explain and predict natural phenomena. Although Aristotle's ontology had been rejected, it is difficult to deny that his deductive argumentation, if successful, can provide very strong justification for scientific explanations, predictions and counterfactuals, and since the proposed laws of nature fit neatly into that scheme, everything seemed to be working out very well for the mechanical view.

The mechanical view hit an apex with the success of Newtonian Mechanics. Newton's ability to describe and explain a large number and variety of complex phenomena led to a common belief that every natural phenomenon could be understood in terms of basic physics. In other words, all other fields of science could ultimately be reduced to just physics, once all of the details had finally been discovered. Further, since Newtonian Mechanics was overwhelmingly accepted as the correct theory of physics up until the early 20th century, the mechanical approach—and even naturalism in general—seemed to overlap completely with Newtonian

Mechanics in the minds of scientists (Muenzinger, 1935; Turner, 1940; Boas, 1952; Allen, 2005). This complete overlap between naturalism, the mechanical approach, and Newtonian mechanics became problematic when Newtonian mechanics finally was rejected in the early 20th century. General relativity, quantum mechanics, and molecular chemistry do not resemble the basic collisions and push-pull interactions of solid particles envisioned by the classical Newtonian view. But if Newtonian Mechanics got its ontology wrong, what should we conclude about naturalism and the mechanical view which were so closely tied to that rejected ontology?

Some argued that all three stood or fell together, and since Newtonian mechanics was wrong, the mechanical approach to biology and even naturalism in general also were wrong (Haldane, J.S., 1884, 1930; Oparin, 1964). Such was the opinion of vitalists, who resisted any attempts to reduce biological phenomena to physics. Biology, they argued, includes complex organic molecules all working together to achieve survival and reproduction. Individual, non-living bits of matter that are completely describable in terms of basic physics, are incapable of working together to produce the purposive behavior of living organisms, and so there must be some other constituent, a vital substance, force, or fluid, present in living creatures that gives them their cohesion and purposiveness. Thus, vitalists were harkening back to Aristotelian substantial forms and teleological forces as described above, as well as to Kant, who argued in his *Critique of Judgment* (1790) that our conception of life and living organisms necessarily includes a sense of purposiveness and design, and thus cannot be reduced to Newtonian, non-teleological terms. That approach put vitalists at odds with most naturalists, even those that accepted the failure of Newtonian Mechanics, since most naturalists assumed that teleological forces do not exist as part of the fundamental constituents of nature. As a result, vitalism was considered by many to be a departure from the naturalistic view and was eventually dismissed as

non-scientific, or even mystical (Turner, 1940; Mayr, 1982). The vitalists' cause was not helped by the fact that any explanations, predictions, or counterfactuals from their program were far less precise or useful when compared to alternatives based on the relatively simple ontology and methodology of physics. After all, vitalists explicitly rejected the ability of mathematics to describe living processes adequately, and when it comes to acceptance of theories in science, precision, accuracy and usefulness are key, and mathematization of a theory gives it a huge advantage in all those respects. Consequently, vitalism ended up fading from mainstream science and the reductionist viewpoint prevailed.

As mentioned above, contemporary physics includes non-classical properties and interactions that are very unlike those in classical Newtonian theory. However, general relativity and quantum mechanics are still expressible in mathematical terms, and so many of the important benefits of the Newtonian program have been maintained, especially the use of precise, mathematical formulations that can generate very clear and accurate descriptions of experimentally controlled phenomena. When used in deductive logic, those descriptions can generate clear and accurate explanations, predictions and counterfactuals, very much like classical physics. As a result, reductive[10] views of nature now are based on quantum mechanics. This, combined with the fact that the term "mechanics" remains attached to the contemporary theory, has led some authors to continue referring to a generally reductive view of nature that relies on fundamental physics as a "mechanical" or "mechanistic" view of nature (Muenzinger, 1935; Turner, 1940). This is especially true of non-reductionists in philosophy of biology or philosophy of mind (e.g. Haldane, 1930; Hanna, personal communication).

[10] To be clear, I am talking about the kind of reductionism that assumes all natural phenomena can be understood and explained using only basic physics.

Since it does not include Newtonian ontology, however, it is not clear to what extent contemporary physics and chemistry still represent the mechanical view. Reductionism by itself entails neither acceptance nor rejection of the mechanical view, and reductionist philosophers of science over the last century have tended not to speak much about mechanisms, either positively or negatively, so while there has not been acceptance of mechanisms, they have not been explicitly rejected either. Mitchell (1997, 2000) suggests that part of the reason mechanisms have not been included in the conversation is because most philosophy of science has relied heavily on deductive logic, and unlike laws of nature, it is not clear that mechanisms can be described adequately using only sentences or mathematical formulae, as standard systems of deductive logic require.[11] In any case, one option for naturalists after the rise of quantum mechanics has been to embrace the new physics and reject the mechanical view entirely because it is too closely connected to Newtonian Mechanics. The mechanical view, this reasoning goes, was a specific philosophical position developed by people committed to the view that all natural phenomena are the result of classical pushing, pulling, and colliding of solid bits of matter. Modern physics has shown us that the real constituents of matter do not conform to the classical conceptions, and so beliefs based on the classical conception of matter are wrong. The mechanical approach is based on the classical conception of matter, and so the mechanical approach is wrong (for example, Haldane, 1930; Oparin, 1964). According to this view, to claim that the mechanical view survives the collapse of Newtonian theory is akin to claiming that rigid three-dimensional space survives the shift to general relativity.

Critics of this line of reasoning point out that it overstates the link between the mechanical view and a particular theory of Mechanics (Muenzinger, 1935; Turner, 1940; Allen, 2005). For example, Turner (1940) argues that we must maintain the distinction between a

[11] This issue will be explored in much more detail in Section 1.3

particular theory of Mechanics on the one hand, and the mechanical (what I will eventually be calling the "mechanistic") approach on the other. Although they often have overlapped significantly, the mechanistic approach itself is not committed to any particular theory of Mechanics, whether it is Newtonian Mechanics or Quantum Mechanics. I agree with Turner (1940) and believe that while the original formulators of the mechanical view got their ontology wrong, it still is true that living organisms share many important similarities to human-built machines and it is important to acknowledge this fact and use it to aid our understanding of nature, regardless of which particular theory of physics turns out to be correct (see also Machamer, Darden, and Craver, 2000; Bechtel and Abrahamsen, 2005). In Section 1.3, I discuss the merits of the mechanistic approach as compared to the traditional approach that formulates explanations, predictions and counterfactuals rooted in laws of nature rather than natural mechanisms.

One more position worth mentioning came from so-called holists of the early 20th century. While some holists were perhaps best described as vitalists, Allen (2005) argues that holism can also be formulated in naturalistic terms. A holistic view is opposed to reductionism because, like vitalists, holists argue that focusing only on individual activities of the individual bits of matter that compose an organism will provide an insufficient understanding of the whole, living organism. In addition to the individual bits of matter, holists argue that one must also understand how those bits work together to accomplish complex tasks and how they are situated within and affected by the rest of the organism and its surrounding environment. Although nothing seemingly occult or supernatural is necessarily involved, a problem with the holistic view was that its properties, relations, and interactions are so complex, contingent, and shifting that they were not precisely describable, especially in mathematical terms. As a result, any

explanations, predictions, or counterfactuals from a holistic program were far less precise or useful, making it very difficult for holist to contribute anything substantial to biologists' research programs. Thus, by and large, reductionism remained dominant while holism faded away from mainstream biology and was considered little different from vitalism, even if some holists (such as J.S. Haldane) insisted that they were committed to naturalism (Allen, 2005). It should be noted that critics of holism did not necessarily ignore the complexity to which holists refer. As discussed in Section 1.3, that complexity is often expressed as *ceteris paribus* clauses that are attached to relatively simple statements of some scientific generalization or relation.

To summarize the most important points from this section so far: the most popular attitude in philosophy of science has been to maintain commitments to naturalism, reductionism, laws of nature, and the logical construction of scientific explanations, predictions, and counterfactuals. I call this "the traditional view" in philosophy of science. The traditional view has not included much talk of mechanisms partly because mechanisms are not readily representable as sentences in deductive logic, and partly because the mechanical view was too closely associated with the rejected theory of classical mechanics.

The mechanistic[12] approach has experienced a resurgence in philosophy of science in recent years, carving out a role for itself that is independent of the traditional view. The contemporary mechanistic approach is not committed to the old Newtonian ontology, and focuses on biological systems (and sub-systems) as complex, coordinated processes that work to achieve particular functions, such as ion regulation or ventilation. Only by understanding how all the parts of a biological system work together in context can we understand how the system is able to do what it does. In this way, the contemporary mechanistic approach I am advocating

[12] I use "mechanistic" rather than "mechanical" here because I am referring to the contemporary view rather than the classical view.

bears a resemblance to holism (Allen, 2005), except that we now are much better positioned than was early 20th century holism to use this approach to support successful and productive research programs.

The mechanistic approach in the biological sciences starts by recognizing that every organism is a complex system composed of many interconnected sub-systems that contribute to the organism's continued survival. Most of those life-sustaining subsystems are similar to human built machines in that their component materials are arranged and coordinated in very particular ways that allow them to perform their functions.[13] So similar to human built machines are these biological subsystems, that they are best described as biological mechanisms. As with human built mechanisms, in order to understand a biological mechanism—and so to explain how and why it works and to make predictions about its future activity—we must gain an understanding of the causal interactions occurring within that mechanism and how its activity interacts with the larger system of which it is a part. Such understanding was beyond the reach of biologists before the rise of molecular biology and biochemistry in the mid-20th century, and so the early holists described above were unable to flesh out their views. Today, however, it is commonplace for researchers to adopt the mechanistic approach. Indeed, failure to do so is often regarded as bad—or at least incomplete—science (Sarkar, 2005; Casadevall and Fang, 2009)

In Section 1.3, I discuss the use of laws of nature in philosophy of science in more detail, along with recent criticisms. I also show how the mechanistic approach is better positioned than the traditional view to addresses those criticisms successfully (MDC, 2000; Bechtel and Abrahamsen, 2005). While some authors argue that mechanisms should replace

[13] Living systems do not employ much of the classical push-pull mechanisms that were central to the mechanical view, so in that sense they may be disanalagous. However, the specific kind of causal interaction (push-pull vs. electrochemical vs. molecular vs. whatever else) is less relevant to the mechanistic view than is the spatiotemporal coordination of parts to achieve a function.

laws of nature as the justification for scientific explanations, predictions and counterfactuals (Bechtel and Abrahamsen, 2005), I argue, following Woodward (2003), that if one wishes to maintain a commitment to laws of nature, the mechanistic approach should not be seen as a rival view, but rather as complementary.

Section 1.3 The Traditional View and the Mechanistic Approach

In this section, I argue that the traditional view, by itself, provides an insufficient account of how scientists actually develop and justify their explanations, predictions and counterfactuals. However, I do not argue for its outright rejection or replacement. Rather, I argue for inclusion of the mechanistic approach as complementary to the traditional view.

The traditional view in philosophy of science maintains that in order to produce scientific explanations, predictions, or counterfactual statements, scientists must find laws of nature.[14] Without laws of nature, scientists would be restricted just to *describing* what exists and what happened, and they would be unable to justify any explanations for why things are the way they are, any predictions about what will happen, or any counterfactuals about what would have happened had things been different.

The kinds of laws of interest here are universal and exceptionless generalizations that support counterfactuals. The laws of logic have these features, and so not only do they support the validity of deductive arguments but they also help us explain *why* one particular conclusion rather than some other conclusion follows from the given premises. Without reference to laws of logic, all we could say is that the conclusion just always does follow those premises and we could not explain why it follows. Laws of logic are able to do this explanatory work because

[14] Of course, not all prominent and influential philosophers of science rely on laws of nature in this way, but the law-based approach is a dominant and enduring trend.

they are logically necessary. That is, we can explain that the conclusion follows from the premises because it is logically necessary that they follow. There may not be any further explanation for why the laws are the way they are, except to say that it is just plain impossible for things to be otherwise, and so "Because it's a law of logic" may be considered an end to explanation.

Laws of nature are thought to do for science what laws of logic do for logical deduction. Scientists should be able to tell us why, given certain initial conditions, a particular natural phenomenon is observed rather than some other phenomenon. Without laws of nature, the thinking goes, all we could say is that the relevant phenomenon just always does occur after the preceding conditions obtain, and we could say nothing about *why* it occurs. Paralleling the situation in deductive logic, "Because it's a law of nature" may be considered an end to any scientific explanation because we then know that it's just plain impossible for things to be otherwise, at least in a universe like this one. In order to capture their similarity to deductive logic, laws of nature are typically expressed as universally quantified conditionals such as $(x)(Ax \rightarrow Bx)$, or just $(A \rightarrow B)$. "A" represents some initial conditions and "B" represents the conditions that result, so the generic form of a law of nature says that whenever A happens, B happens. In this way, a law of nature can be used as a premise in a deductive argument that, when combined with other premises stating that the relevant initial conditions obtain, can generate the phenomenon as its conclusion.

According to this view—the traditional view—the above procedure can provide a scientific explanation, prediction, or counterfactual, depending on our needs. If the phenomenon has already been observed and we are interested in an explanation, then the conclusion (the phenomenon) is the explanandum while the premises (initial conditions plus any relevant law(s)

of nature) constitute the explanans. If, on the other hand, we believe we know the initial conditions and relevant laws of nature, then the conclusion would be a prediction. Counterfactuals are much like predictions except that one or more of the premises is hypothetical. So, according to the traditional view, a single scientific argument can be construed as either an explanation, a prediction, or a counterfactual, depending on our epistemic position relative to the premises and conclusion of the argument (Carnap, 1966; Hempel, 1962, 1965). This aspect of the traditional view is known as the thesis of structural identity, or the symmetry thesis (Hempel, 1965).

The traditional view has its roots in Aristotle's *Posterior Analytics* (see Section 1.2), and has been supported in one form or another by more recent philosophers of science such as Carnap (1966), Hempel (1962, 1965), Lewis (1973), Dretske (1977), Tooley (1977), Armstrong (1978), and Schiffer (1991). Although these authors disagree on many things, a common thread in their work is a commitment to formal systems and laws of nature to construct and to support scientific explanations, predictions and counterfactuals.

One initial problem faced by the traditional view is that laws of nature do not carry the power of logical necessity that laws of logic enjoy. Instead, the general characteristics of the universe, including any laws of nature, are thought to be logically contingent. To fix this problem, laws of nature are said to hold with *natural* (or nomological) necessity rather than with logical necessity. Natural necessity applies in our actual universe, but what's naturally necessary in our universe may not be naturally necessary in a universe with different characteristics, so rather than holding with logical necessity, a law of nature only holds necessarily in a universe just like ours. Scientific explanations, predictions, and counterfactuals typically carry the

assumption that we are talking about the actual universe (unless they explicitly say otherwise), so any laws they invoke should carry the natural necessity we require.

The best candidates for universal and exceptionless laws of nature come from fundamental physics, such as the law of conservation of energy or the Pauli exclusion principle. If such laws truly are universal and exceptionless[15] and hold with natural necessity, then they are capable of justifying explanations, predictions and counterfactuals from physics in the way the traditional view requires. However, when we get to the so-called special sciences such as biology and psychology, things become problematic. Regularities and generalizations discovered in the special sciences are riddled with exceptions and caveats, yet somehow they still must be able to justify special science explanations, predictions and counterfactuals. A pessimistic view might take the seemingly contingent nature of the special sciences to mean that their predictions, counterfactuals, and explanations are not properly justified (e.g. Rosenberg, 1994). However, it is difficult to deny that the special sciences have achieved a great deal of success, especially over the last several decades, and so there must be a way for their predictions, counterfactuals, and explanations to be properly justified despite their apparent contingency.

Laws in the special sciences are not universal in the same way as fundamental laws of physics because special science laws only apply to complex systems that may exist in only one very small part of the universe. They also appear contingent because a special science law may fail even when the appropriate complex system exists and is manifesting conditions to which the special science law typically applies. A very common approach to accommodating these features of special science laws and making them work like the laws of physics is to include *ceteris paribus* clauses. A *ceteris paribus* clause includes all of the background conditions

[15] As discussed below, there may be good reasons to doubt that even these are truly exceptionless and universal, as the traditional account appears to require.

needed for the generalization expressed by the law to hold. Following Sober (1997) we can represent *ceteris paribus* laws as C→(A→B). The conditional (A→B) is the proposed generalization—the law itself—and C represents the background conditions as a *ceteris paribus* clause. So, the generic form of a *ceteris paribus* (c.p.) law says that as long as the special conditions obtain, then if A happens, B happens. This formulation says that if the background conditions obtain, then the generalization is a law and can be used in arguments justifying the relevant special science explanations, predictions, and counterfactuals.

A critic of the traditional view might complain at this point by saying that even with c.p. clauses attached, special science laws still do not hold with natural necessity and so cannot be used to support scientific explanations. The idea is that the c.p. conditions obtain only contingently, and so the generalization expressed by the law obtains only contingently. A contingent law is not a naturally necessary law, and so special science laws that require c.p. clauses are not laws of nature. Rupert (2008) presents a somewhat modified version of the traditional account of c.p. laws that fixes that potential problem with c.p. laws by showing that the laws themselves actually *are* universal and exceptionless and hold with natural necessity. Rupert shifts the supposed contingency of c.p. laws into what he calls "combinatorial laws" that describe how the properties represented in the antecedent of a c.p. generalization come together to generate the consequent properties. This all requires some unpacking, so I'll begin with the simplest case and then build to c.p. laws as Rupert describes them. Instantiations of basic physical properties generate causal influences according to basic laws of nature and do so in a non-combinatorial way. For example, consider a system consisting of just two electrons, each instantiating a $^{-}1$ charge. Interaction between these two electrons is governed by a single law of

nature[16], and so a prediction about what will happen involves simple application of this single law. In a slightly more complex system with multiple basic physical properties being instantiated, there may be many causal influences all at work at the same time. Each individual causal influence works according to a law of nature, and the individual effects combine according to a "combinatorial law" of nature. Therefore, in order to predict the behavior of a complex system, we need to combine the effects of the different laws according to the appropriate combinatorial law. The best example comes from Newtonian mechanics in which the effects of multiple forces acting on an object are combined to determine the net force and the resulting acceleration. So far this is nothing particularly new or different from previous views. The difference comes when we move to c.p. laws.

Rupert (2008) argues that in many respects, c.p. laws in the special sciences operate in the same way as laws in fundamental physics. Paralleling the simplest case described above, a special science system that instantiates a single special science property will generate an instantiation of another special science property according to a law of nature, and that law will hold with natural necessity. Then, in a more complex system that instantiates multiple special science properties that are each governed by a different special science law, with the effects of each law combining according to a combinatorial law from the relevant special science to produce the final outcome. The contingency often associated with c.p. laws, according Rupert, lies in these combinatorial laws of the special sciences and their potential conflict with the combinatorial laws of physics. As I discuss below, Rupert's (2008) more thorough inclusion of the complex causal details included in the *ceteris paribus* clause of c.p. laws show how the mechanistic approach is complementary to commitment to laws of nature. The precise

[16] The relevant law comes from electromagnetism. The effect of other forces is assumed to be negligible in this case.

characterization of laws of nature and c.p. laws in particular is not settled, but whether one adopts Rupert's view or some other version of the traditional view, the mechanistic approach is still valuable.

Criticisms of the Traditional View and Benefits of the Mechanistic Approach

The traditional view has been challenged on many fronts. For example, some have claimed that the traditional view is insufficient because one can construct explanations that appear to meet the relevant criteria for explanation based on laws, but are intuitively irrelevant or incomplete as scientific explanations. Consider the following uncharitable example from Ruben (1990) in his argument against Hempel's (1962) Deductive-Nomological (D-N) model of scientific explanation:

1. All metals conduct electricity
2. Whatever conducts electricity is subject to gravitational attraction

∴3. All metals are subject to gravitational attraction (p. 182)

The argument is valid with true premises, and premise 2, as an exceptionless generalization, meets Hempel's (1962) conditions for laws of nature. Nevertheless, Ruben (1990) claims, the argument does not help to explain *why* metals are subject to gravitational attraction, and therefore the explanation is seriously incomplete. Ruben (1990) does not conclude that the traditional view is worthless, but rather that any explanations based on laws of nature must include statements of causation, rather than mere exceptionless regularities, in order to provide good scientific explanations. This view is shared by Railton (1978) who argues that an account of the complex causation—the mechanism—behind a phenomenon is needed to supplement any mere statements of regularity in order to explain satisfactorily how and why the phenomenon

occurs. This call for greater detail concerning the causal process by which a natural phenomenon is produced is a running theme in this section. A sufficient scientific explanation must include such causal details, and building on Railton (1978) and Ruben's (1990) suggestions, I argue below that the mechanistic approach is well suited to provide those causal details. First, however, I will discuss some other criticisms of the traditional view that will help to show why I believe the mechanistic approach is needed.

Cartwright (1983) also argues that the traditional view is insufficient, but unlike me, she argues for its complete rejection. Her reasoning is that if we rely on the traditional view to tell us what counts as a law of nature, then it turns out that there are no laws of nature. Cartwright points out that even the laws of fundamental physics do not appear to meet the standards established by the traditional view because in any real-life situation, the state of affairs we would predict using any proposed law of physics never obtains in exactly the way the laws say it should. Cartwright takes this to mean that the supposed laws of physics are just rough guides to nature rather than universal or exceptionless. A defender of the traditional view might reply to Cartwright by pointing out a prediction generated by a given law of physics is based on an idealized view of the world in which only that single law is operative. Under realistic circumstances, of course, a variety of laws of nature would be applicable, and so the pure prediction generated by any single law by itself is unlikely to be exactly correct. In response, Cartwright then argues that if we revise our view of laws in the way suggested—so that single laws all by themselves only apply in idealized circumstances—then they are no longer laws describing the real, contingency-riddled world, but only an imaginary, idealized world. Therefore, they are not really laws of nature as it actually is. Therefore, there are no universal, exceptionless laws of nature to provide the natural necessity that the traditional account requires.

It seems to me that whether one is inclined to accept the existence of laws of nature or not, Cartwright's views seem to be overly pessimistic. Scientists have invented a dazzling number and variety of clever ways to manipulate different aspects of whatever system they are studying in order to gain insight into the individual causal influences at work, and they have found great success in isolating, characterizing, and then re-combining various causal influences in order to produce very accurate and precise predictions, particularly in physics (Woodward, 1997, 2002a, 2003; Woodward and Hitchcock, 2003). It does seem true that, because of inevitable technological limitations, it is impossible in practice to generate exactly accurate predictions for even very simple systems, but that does not seem like enough reason to conclude that laws of nature do not exist. Our experimental apparatus and measuring devices will always have a finite level of precision, so any results obtained using that apparatus cannot be expected to match predictions *exactly*. Plus, some proposed laws of nature are inherently statistical while others, like the Heisenberg Uncertainty Principle, ensure that we cannot know everything about a given system. Given these factors, our lack of complete accuracy and precision is not sufficient justification for rejecting the existence of laws of nature, so Cartwright's (1983) arguments seem unconvincing (see also Rupert, 2008).

That is not to say that Cartwright's criticisms are completely off base. We should be concerned about the traditional requirement that laws be universal and exceptionless. Mitchell (1997, 2000, 2002), following Cartwright's general line of reasoning, points out that according to contemporary physics, even the most fundamental laws break down under certain extreme circumstances. For example, during the first tiny fractions of a second after the big bang, and even now within black holes, some proposed laws of nature do not seem to apply. The reason is that some values, like space-time curvature, become infinite, and this cause the relevant

equations to produce undefined results. Therefore, even the most fundamental laws of physics may fail to be universal or exceptionless, which means that even if there is such a thing as natural necessity, no known laws of physics have it.

Since there are good reasons to think that even the most basic laws of nature fail to be universal or exceptionless, Mitchell (1997) concludes that there is no such thing as natural necessity that can operate like logical necessity in justifying scientific explanations, predictions and counterfactuals. Unlike Cartwright, however, Mitchell does not wish to abandon laws of nature. Rather, she argues that all natural regularities actually lie on a continuum of contingency, and we should abandon the distinction between basic, exceptionless laws of nature and *ceteris paribus* laws. More specifically, Mitchell argues that the background conditions needed to support any supposed law of nature are contingent, with some conditions being more stable than others. For example, the background conditions required for laws of fundamental physics are very stable (though not perfectly stable) and so have a correspondingly low level of contingency. By contrast, the conditions needed for the laws of biology are much less stable because the conditions can come into and go out of existence. Plus, the same laws of biology can hold in a wide variety of different physical systems, and so it seems that laws from the special sciences come with a great deal more contingency. Therefore, a proper explanation for any actual phenomena must include more than just a description of the regularity itself. We also need to know something about the conditions needed for the regularity to occur (other than just that there are some conditions that need to occur, as expressed by a *ceteris paribus* clause) and how changes in these conditions will affect the final product. In physics, very little extra information beyond the regularity or generalization itself is usually needed to justify any explanations, predictions or counterfactuals. But when it comes to the special sciences, a great deal of extra

information about context, history, etc. may be needed in order to feel confident that the relevant laws are applicable. Woodward (1997, 2002a, 2003) makes a similar case when he argues that even if laws of nature and the traditional view are adequate for explanations in physics or chemistry, they are not sufficient for the special sciences because understanding the background conditions for a proposed law in the special sciences is at least as important as recognition of the law itself.

Mitchell's (1997, 2000, 2002) arguments, which are later mirrored in MDC (2000) and Bechtel and Abrahamsen (2005), stem from her dislike for the traditional use of logical systems to represent laws of nature and to generate explanations, predictions and counterfactuals. Systems of formal logic are poorly equipped to handle the contingency inherent in the special sciences, forcing it all into unexamined *ceteris paribus* clauses (or combinatorial laws). Mitchell (1997) argues that standard formulations of c.p. laws, such as C→(A→B) hide two kinds of contingency[17] that scientists find important. First, the set of conditions represented by "C" is itself messy, diverse, and historically contingent. It is historically contingent because, as discussed above, the appropriate conditions have not always existed in the past and may cease to exist in the future, and the universe could have unfolded such that the conditions never obtained at all (Beatty, 1994). The second sort of contingency hidden in c.p. laws, according to Mitchell (1997), lies in the material conditional within "C→" which represents a very complex web of causal interactions that is likely to be nonlinear and extremely sensitive to initial conditions (p.S472). That is, very small changes in C can result in a substantially different web of causal interactions and a different outcome. To be clear, the charge here is not that any *ceteris paribus*

[17] Mitchell (1997) uses the word "contingency," but it may be more appropriate to call this "complexity." To say that a c.p. law is contingent implies that even when the relevant conditions obtain, the corresponding law might not, but this is not quite what Mitchell has in mind. Rather, she is concerned with the complex details contained within the c.p. clause of a special science law that help us understand why the regularity occurs and why it sometimes does not happen when we might expect it to.

laws are wrong or incoherent. Rather, the charge is only that c.p. laws do little to aid our understanding of natural phenomena because they tell us only that a regularity occurs, not how or why the regularity occurs.

As Mitchell (1997) explains, "It is not *that* biological generalizations are contingent, but rather *how* they are contingent that is significant" (S472, emphases in original). Also, "To know when to rely on a generalization we need to know when it will apply, and this can be decided only from knowing under what specific conditions it has applied before" (S477). Then, as she puts it in Mitchell (2000), "The insights that scientists acquire about the causal structure of our world may be deformed by being squeezed into Boolean garb" (2000, 247), and "Because the [traditional] view has been wedded to representing epistemological relations (like explanation, prediction, confirmation) and causal relations in first order predicate logic, it has allowed a reification of the features of the representational apparatus to be imposed on the thing represented" (248). As described in more detail below, Mitchell's views are similar to the mechanistic approach, as Mitchell herself suggests in her example of "redundant mechanisms" in biological systems (2002, p. 338). As explained in Section 1.2, although a mechanistic view of nature has been around for centuries, the contemporary version of the mechanistic approach in biology sheds the anachronistic classical ontology and focuses on how parts of living organisms come together to achieve different functions, with emphases on why the mechanisms do what they do and how the outputs of mechanisms can be changed by outside factors. In other words, the mechanistic approach emphasizes the causal complexity with which Mitchell and other critics of the traditional view are so concerned.

Since I am considering the roles of the traditional view and the mechanistic approach in the sciences, it is worth noting how often scientists actually invoke laws in their explanations,

predictions and counterfactuals. After all, an explicit premise of the argument that philosophical investigation into laws of nature is important for philosophy of science is that scientists really are searching for laws to justify their scientific predictions, counterfactuals, and explanations. For example, in his discussion of c.p. laws, Fodor (1991) writes, "I am, in general, in favor of taking the sciences at face value: Psychology, meteorology, and geology (to say nothing of physics) keep announcing hedged laws" so we should seek to find a proper account of such laws (p. 22). Nearly all prominent authors on this subject, regardless of their position on the issues, seem to accept something like Fodor's premise that science concerns itself with the discovery of laws (Dretske, 1977; Schiffer, 1991; Tooley, 1977; Rupert, 2008). Since this premise is supposed to justify our attempts to characterize laws of nature, it is worth examining whether this premise is true. I believe that it is false.

I believe, along with Woodward (2003), that it is false that researchers in the special sciences, especially the biological sciences, think of themselves as searching for laws of nature at all, whether c.p. or non-c.p. laws. It is difficult, however, to evaluate the truth of this premise without a thorough polling of scientists' beliefs. In my own experience[18], I have found it rare for those working in the biological sciences to refer to laws of nature when discussing biology. Even when it does occur, any talk of laws of nature among biologists usually is restricted to supposed laws of physics and their relevance to certain biological processes, or about scaling laws related to, for example, diffusion rates, that ultimately come from physics.

To gain a rough idea, beyond my own personal experience, of whether scientists really are searching for laws, I conducted a variety of literature searches in different academic databases. Since I am most interested in science as it is practiced today, I restricted my search to

[18] This is, of course, of limited value as evidence, but I did live in the biologists' world for many years: taking classes, attending conferences, pursuing laboratory research, and ultimately earning a Master's degree in biology and teaching life science for several years.

books and articles published after 1980. In the biological sciences, around 1,500 peer-reviewed articles (less than 1%) contained the word "law," and a cursory review of the titles revealed that in the vast majority of those works, "law" was used to refer to legal codes or to a law of physics or a scaling law as described above. Only rarely did the title refer to a proposed law of nature restricted to life as we know it on earth. Instead of searching for laws, a number of recent authors claim that researchers in the special sciences are looking for mechanisms to justify their explanations, predictions and counterfactuals (e.g. Machamer, Darden and Craver, 2000; Glennan, 2002; Woodward, 2002b; Bechtel and Abrahamsen, 2005; Craver and Darden, 2005; Casadevall and Fang, 2009; Thaggard, 2010). Supporting this view is another literature search of titles in the biological sciences containing the word "mechanism" that yielded over 80,000 results (over 30% of total).

Put very briefly, as this is the focus of Chapter Two, a biological mechanism is a dynamic, physically instantiated process organized and maintained by an organism that performs a biological function, where a biological function is an output of the process that contributes to survival or reproduction of the organism that is maintaining the process. For example, a mammalian heart is constructed and maintained by an organism and the heart, in turn, contributes to the continued survival of the living organism of which the heart is a part. Once a mechanism has been identified and characterized, scientists can explain how and why the mechanism works and what we can expect of it in the future. This is the mechanistic approach that scientists rely on to justify their explanations, predictions, and counterfactuals. In fact, it is not uncommon to find criticisms of proposed scientific explanations that are based on a lack of any mechanism capable of producing the proposed phenomenon or regularity (Sarkar, 2005; Casadevall and Fang, 2009; Gould, 2002).

Although the mechanistic approach is sometimes portrayed as a rival to the traditional view (e.g. Bechtel and Abrahamsen, 2005), this is not necessarily so. I believe that a commitment to laws of nature and their role in scientific explanations, predictions and counterfactuals does not preclude acceptance of the mechanistic approach, and acceptance of the mechanistic approach does not preclude a commitment to laws of nature. Indeed, I believe that the mechanistic approach is complementary to the traditional view (see also MDC, 2000). For example, Rupert's (2008) component forces view of c.p. laws includes combinatorial laws governing the way the current state of a complex system produces a later state. The mechanistic approach can be complimentary to such a view by filling in the details concerning how component forces actually combine to produce their effects. Philosophers working on the mechanistic approach typically do not speak in terms of "component forces," but the mechanisms they describe could theoretically be analyzed in terms of the individual component forces produced by the microscopic bits that make up the mechanisms, though such a task would be extremely tedious and difficult.

As a real-world example of how the mechanistic approach can complement explanations in terms of laws, early geneticists cited the so-called Law of Independent Assortment and Law of Segregation in describing Mendelian inheritance patterns. However, while recognizing that a law of nature connects parental traits to offspring traits may be interesting and even quite useful to plant and animal breeders, it does not help us understand how or why inheritance works that way, nor does it tell us how or why it occasionally deviates from the standard patterns. Those questions can only be answered by investigating and understanding the mechanisms that produce the phenomena described by the laws.

Consistent with that example from genetics, Mitchell (1997, 2000) argues that understanding the complex, contingent circumstances associated with a given generalization is at least as important as discovering the generalization itself, and the mechanistic approach captures that sentiment very well. Rather than ending with the discovery of a regularity, the mechanistic approach *starts* with observation of a regularity or other interesting phenomenon in need of explanation, and then does the investigative work needed to discover the mechanism responsible for producing that regularity. We see a similar account in pro-mechanism author Lindley Darden (2002), who argues that researchers in the special sciences typically begin with a "mechanism schema," which Darden describes as "a truncated abstract description of a mechanism" that amounts to little more than a description of a regularity, and "can be filled with more specific descriptions of component entities and activities" (S358). The idea is that a mechanism schema may at first be highly speculative, based on a limited knowledge of the component entities or how such entities are able to produce the phenomenon. Scientists may, for example, "sketch hypothetical roles that components of the mechanism being sought are expected to carry out" (Darden, 2002, p.S360). Scientists then search for physical entities in the complex system under study that are capable of carrying out the hypothesized activities. If appropriate entities are found, they can be used to turn the mechanism schema into a more concrete model of the mechanism that can be used to generate better explanations, predictions, and counterfactuals. On the other hand (and just as importantly), if appropriate entities are not found, then scientists are forced to modify or abandon their original mechanism schema. In other words, failure to find entities capable of achieving a hypothesized activity forces scientists to acknowledge an error in their characterization of the phenomena that had attracted their interest in the first place.

Two examples can help illustrate Darden's (2002) account of how the mechanistic approach works in biology. First, consider the example of Lamarckian inheritance patterns. Lamarck (1809) hypothesized that interactions between an organism and its environment caused adaptive changes in the organism that were then inheritable by that organism's offspring. This is the famous inheritance of acquired characteristics that has come to be the most recognizable feature of Lamarckian evolution. In the 19th century, environmentally-induced adaptations and inheritance of acquired characteristics were both believed to occur in nature. In fact, even though Darwin's (1859) theory of evolution is often held up in opposition to Lamarck's, Darwin included both in his theory. However, once the mechanisms of cellular and molecular biology became better understood, biologists did not find any mechanisms that were able to support the proposed Lamarckian phenomena. That is, when more and more was discovered about how organisms and biological inheritance actually work, biologists sympathetic to Lamarckian evolution were left without any way to explain how organisms could manage to do what Lamarck required them to do. As discussed more thoroughly in Chapter Three, it was that lack of any plausible supporting mechanisms, as much as any direct evidence against Lamarck, that sank his theory of evolution (Jablonka and Lamb, 2005; Sarkar, 2005).

Contrast the failed case of Lamarckian inheritance with the successful genetic research program begun in the early 1900's with the rediscovery of Mendelian inheritance patterns. Scientists had found some intriguing regularities in the way certain traits were passed from parents to offspring, and while they recognized that there must be some kind of underlying physical process that supported the phenomena, technological hurdles limited them to merely describing the phenomenon. As described briefly above, before the underlying entities could be identified or characterized, the observed regularities were among the few ever to be called

biological laws: The Law of Segregation and the Law of Independent Assortment.[19] Unlike the case with Lamarck, improved understanding of cellular and molecular mechanisms revealed that they were well suited to producing the Mendelian phenomena, and the genetic research program became extremely successful and popular. In fact, as understanding of genetics progressed, any talk in terms of laws gave way to more nuanced, more accurate, and more broadly applicable explanations in terms of cellular and molecular mechanisms. Today, geneticists no longer speak in terms of laws of inheritance, and instead provide explanations, predictions and counterfactuals in terms of the underlying mechanisms, with special attention paid to all of the contingent factors that can cause the mechanisms to change or stop their activities.

A defender of the traditional view could argue that all of the geneticist's explanations, predictions and counterfactuals in terms of mechanisms can be rephrased using *ceteris paribus* laws of nature in order to ensure their proper justification. However, as I have argued above, such rephrasing does not constitute a replacement of the mechanistic approach. Rather, the mechanistic approach complements the traditional view by characterizing what, exactly, is contained in the *ceteris paribus* clause of a particular special science law and how it affects any explanations, predictions and counterfactuals.

Bechtel and Abrahamsen (2005) take up this point and spell out a number of advantages enjoyed by the mechanistic approach.[20] Like Mitchell (1997, 2000, 2002), Bechtel and Abrahamsen (2005) argue that using standard logic to express and understand biological

[19] Application of that label likely was aided by scientists' ability to mathematically model the phenomena and use their equations to generate remarkably accurate predictions, at least under very controlled circumstances.

[20] Bechtel and Abrahamsen (2005) take a harder line against the traditional view than I do. They argue that the mechanistic approach should replace the traditional view rather than compliment it, particularly in the biological sciences. Although I have argued that such a replacement may not be necessary, I believe that Bechtel and Abrahamsen's (2005) arguments are relevant because they help to describe the merits of the mechanistic approach very well.

regularities is too limiting (or at least too cumbersome). They argue that mechanistic models used by scientists include modes of representation and styles of reasoning that cannot be captured by traditional logical systems. For example, certain dynamic spatial or temporal relations that are represented relatively easily by diagrams (individually or in sequence) may not be expressible verbally or linearly, as required by standard logic. Beyond questions of representation, Bechtel and Abrahamsen (2005) claim that actual reasoning processes employed by scientists involve some exclusively visual information that is never converted into verbally expressible language. When thinking about biological processes, Bechtel and Abrahamsen contend that we construct non-verbal mental representations of mechanisms and then animate them forward and backward in time in order understand how they work and to make predictions about their behavior. In this sense, humans are mentally *simulating* the proposed mechanism rather than making formal inferences about it, and so the actual reasoning of scientists is better captured by the mechanistic approach and its use of diagrams and other non-verbal representations rather than by the traditional view with its verbal statements and logical inferences.

I believe that Bechtel and Abrahamsen (2005) may very well be correct, but even if they are wrong and it turns out that all non-linguistic representations used by scientists can be converted to verbal representations compatible with the traditional view, the result would be unusable (or at least unused) by scientists or others who are reasoning about these matters. The precise verbal descriptions of even a fairly simple diagram or graph are very long, complex and cumbersome, which is why nearly all science textbooks and research articles include nonverbal representations of the relevant material. Further, it seems that diagrams and mental simulations are in some sense more accurate or realistic representations of the complex systems under study

than are verbal descriptions. It may be difficult to spell out exactly how much more accurate or realistic diagrams and mental simulations are when compared to purely verbal descriptions, but the point is that we avoid linguistic descriptions of complex systems not merely because they are tedious and cumbersome, but also because they are more removed from the actual phenomena. It seems, then, that the mechanistic approach is able to match more closely the actual reasoning and successful explanations used by researchers in the special sciences.

Another advantage of the mechanistic approach, argue Bechtel and Abrahamsen (2005), lies in its inclusion of external context. The exact behavior of a mechanism may change when its context changes, often because a given mechanism tends to be embedded within a kind of hierarchy of mechanisms. That is, a given mechanism may be a component of some larger mechanism, and when the actions of that larger mechanism change, the actions of the component mechanism may also change (see also Craver, 2001; Craver and Bechtel, 2007) and so the mechanistic approach is able to incorporate context into its explanations more naturally than the traditional view. In Chapter Two I flesh out the mechanistic approach that I am advocating by describing the features a system must have to distinguish it as a biological mechanism. Then, in Chapters Three, Four, and Five, I use the mechanistic approach in addressing a variety of active debates in philosophy of biology.

CHAPTER TWO

BIOLOGICAL MECHANISMS

2.1 Introduction

It has become very common for scientists, especially biologists, to describe their work as a search for mechanisms (Craver and Darden, 2005; Casadevall and Fang, 2009). However, only very recently have philosophers begun addressing what biological mechanisms are and how to identify them. In distinguishing mechanisms from non-mechanisms, most recent authors try to identify special physical characteristics of mechanisms and special features of the actions of mechanisms versus non-mechanisms. In this chapter I build upon that recent work to more carefully describe biological mechanisms and the mechanistic approach. I argue that a biological mechanism is a dynamic system that is not a mere aggregate and serves a biological function. In order for a mechanism to serve a biological function, the mechanism must contribute to sustaining the life of an organism, and, in turn, the organism must maintain the existence and operation of the mechanism. As I explain in Section 2.4, my emphasis on functions introduces a form of natural teleology that previous authors have been reluctant to embrace.

2.2 Natural Mechanisms in Biology

Drawing from the historical overview in Section 1.1, there are at least two popular ways of thinking about natural mechanisms in contemporary science. One way is to think of *any* physical process or event as a mechanism, making the mechanistic approach little different from a general commitment to physicalism. According to this first view, not only are there mechanisms responsible for pumping blood and digesting food, but also mechanisms for splitting

radioactive nuclei, producing lightening, rusting iron, and keeping planets in orbit. Many scientists appear to hold something like this version of the mechanistic view, especially when they claim that proper scientific research should be uncovering mechanisms (for everything) rather than merely providing descriptions (e.g. Casadevall and Fang, 2009). In contrast to this extremely permissive notion of mechanisms is the version found in the biological sciences and in cognitive science that marks a distinction between physical objects or processes that are mechanisms and those that are not. According to this second view, living organisms make use of mechanisms of a different sort from non-living "mechanisms" included in the other, more permissive view.

The mechanistic approach used in the biological sciences takes seriously the analogy with human-built machines, like clocks and motors, which are composed of interacting parts coordinated so as to perform a function. Human-built mechanisms can be contrasted with static artifacts that are not mechanisms, like paperweights or doorstoppers. Despite its functionality, a paperweights not a mechanism (in the relevant sense) because it does not achieve its function through the dynamic, coordinated interaction of its parts. In other words, a paperweight is just a static lump rather than a mechanism. According to the mechanistic approach in biology, hearts are more similar to clocks and motors and so are properly considered natural mechanisms. By contrast, a gall stone, which is more like the paperweight, is not a natural mechanism because it is just a static lump of material.

Much of the contemporary efforts to clarify and develop the concept of a biological mechanism has been framed by Machamer, Darden and Craver (MDC) (2000), who build upon the work of Wimsatt (1974), Cummins (1975), and Salmon (1984, 1997), among others. MDC (2000) develop what they call a "dualist" account of natural mechanisms that incorporates the

physical entities that compose the mechanism and the overall activity that the mechanism achieves. Thus, MDC (2000) define a mechanism as "entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions" (2000, p.3). Their account is dualist because it resists attempts by previous authors to characterize mechanisms solely in terms of the individual physical entities, or solely in terms of the overall activity the mechanism performs. Both are necessary for a complete account of any mechanism because even if we think we've observed some activity, if there are no physical entities to perform the task, then there is no mechanism. And even if a variety of entities are present, if they do not perform the right kind of overall activity, then, again, there is no mechanism. I largely agree with this approach to characterizing mechanisms and will use it to build a more complete account.

Before unpacking MDC's (2000) account of natural mechanisms, however, I will mention a prominent, alternative account from Glennan (2002). Glennan, in describing what he calls the "complex systems" approach to natural mechanisms, believes that "a mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where interaction between parts can be characterized by direct, invariant, change-relating generalizations" (p. S344). In apparent contrast to MDC (2000), Glennan (2002) focuses on the parts of the mechanism (the entities) and their direct interactions and apparently ignores the overall activity of the mechanism. At least this is how Glennan's view is perceived (e.g. Tabery, 2003; Craver and Darden 2005; Barros, 2008).

Contrary to this received view, however, I believe that Glennan's (2002) account actually is as dualist as MDC's (2000). Although he does place a great deal of emphasis on the interaction of parts, the behaviors of Glennan's mechanisms appear to do the same work as

MDC's activities. For example, Glennan explains that "mechanisms are not mechanisms *simpliciter*, but mechanisms *for* behaviors" (2002, p. S344, emphases in original), which sounds very much like MDC's insistence that "it is artificial and impoverished to describe mechanisms solely in terms of entities, properties, interactions, inputs-outputs, and state changes over time. Mechanisms do things…and so ought to be described in terms of the activities of their entities" (MDC, 2000, p.5). Both accounts make it clear that including only the physical parts is insufficient for a full account of a mechanism. To drive the point home, we should note that immediately after introducing his definition of a mechanism, Glennan (2002) introduces what he calls a "mechanical model" which is a complete description of a mechanism. Tellingly, Glennan's mechanical model has two necessary components, "… (i) a description of the mechanism's behavior; and (ii) a description of the mechanism which accounts for that behavior" (2002, p. S347). While Glennan refers to "the mechanism" in the second part of his account of a mechanical model and keeps that separate from "the behavior," this should not automatically be taken to mean that a mechanism can be characterized independently of its behavior because, as seen in the previous quotation, "mechanisms are not mechanisms *simpliciter*, but mechanisms *for* behaviors." So it seems that when Glennan refers to "the mechanism that accounts for the behavior" he is talking about the actual, physical instantiation of the mechanism (the "entities" in MDC's terminology) and how that particular instantiation is able to produce that behavior. With this in mind, we can re-phrase Glennan's two-part description of a mechanical model as (i) a description of the mechanism's behavior; and (ii) a description how the mechanism is physically realized. This seems to capture Glennan's intentions while showing more clearly how his account is not really so different from MDC (2000). So, whether it is given the dualist label (MDC, 2000) or not (Glennan, 2002), a proper

account of natural mechanisms must include an activity (a behavior) and a set of entities that produce the activity.

Although I believe that MDC (2000) and Glennan (2002) have very similar accounts of natural mechanisms, in developing my own account of mechanisms in biology, I will use the terminology from MDC (2000) since they more explicitly acknowledge the insufficiency of an account formulated in terms of entities only. I claimed above that biologists mark a distinction between natural mechanisms and mere physical objects or processes, and so I must make clear the circumstances under which something counts as a mechanism rather than a non-mechanism. That is, I must characterize the conditions that the entities and activities of a system must meet in order for the system to qualify as a biological mechanism.

2.3 Entities

MDC (2000) do not have much to say about the entities except that they must exist and be capable of producing the proposed activity. They say only that "one identifies and individuates entities in terms of their properties and spatiotemporal location" (p. 5). This tells us that entities must be specifiable in physical terms but says nothing about whether there are any constraints on the kinds of entities or their arrangements. The only real guidance comes from their examples of entities that are parts of mechanisms, including ions, nucleotide bases, neurotransmitters, membrane ion channels, cells, and hearts. Perhaps MDC (2000) are thinking that their entities must meet some reasonably robust, commonsense notion of objecthood that is satisfied, at least, by individual molecules, proteins, protein complexes, and internal organs, but they do not say explicitly whether this is the case.

Glennan (2002) spends much more time characterizing the entities that compose mechanisms, but still ends up without clear guidelines. Glennan is trying to set himself apart from Railton (1978) and Salmon (1984) who seem willing to allow just about any physical, causal process to qualify as a mechanism. By contrast, Glennan (2002) says that the parts of a mechanism must be objects "in the most general sense…with a relatively high degree of robustness or stability" (p. S344) and cannot be composed of fleeting arrangements and interactions of parts. As an example of a non-mechanism, Glennan suggests a boy accidentally hitting a ball through a window. "While the sequence of events leading to the breaking of the window certainly involves some entities that are stable enough to be called objects," writes Glennan, "…the complex of these objects do not form a stable enough configuration to be called an object" and so the "Salmon/Railton mechanisms are [mere] sequences of interconnected events while complex systems mechanisms are things (or objects)" (p. S345). Here Glennan clearly states that a complex system itself must count as an object in order for it to be a mechanism, and by example he suggests some things with very stable configurations, such as watches, cells, and organisms.

On the other hand, Glennan later appears to soften his view by explaining that the examples on his list count as mechanisms because they are "systems consisting of stable arrangements of parts" and "the systems as a whole have stable dispositions," which seems a much more lenient requirement than full-blown objecthood. This interpretation is supported by Glenna's example of certain social groups in his list of mechanisms because social groups (p. S345) do not match our commonsense notion of an object at all. Glennan (2002) also spends some time explicating an even more tenuous and fleeting mechanism in his example of a "phone-calling chain" that disperses information via an organized set of telephone calls (p. S346).

Glennan argues that this is a mechanism that has people as some of its parts, but the problem is that while there may be a somewhat stable arrangement of the parts (insofar as the people involved know their roles and are ready to act when called upon), it is quite a stretch to think of the dispersed group of people as an object. There are some other examples of mechanisms that Glennan presents that are similarly problematic in that the proposed mechanism is not itself an object, but rather a collection of objects that operate together to achieve some behavior, such as the process that results in Mendel's law of segregation (p. S346) and neural components of information processing mechanisms in the brain which, Glennan admits, have parts that are spatially distributed and realized by different neurological components in different instances of the mechanism's operation. Taken together, then, it seems that Glennan intends for dispersed collections of parts to be able to interact as a mechanism, and that his claim that mechanisms should be objects is meant only to capture the idea that a mechanism should be a somewhat stable arrangement of parts that can consistently interact to produce the behavior.

Much clearer guidance on the composition of mechanisms comes from Craver (2001),[21] who builds on the work of Wimsatt (1986, 1997) in providing specific criteria that can be used to determine whether a group or system of physical items, the Xs, is a mere aggregation or is an actual mechanism. Recall the examples from the beginning of this section: Watches and hearts are mechanisms, while paperweights and gall stones are not. Paperweights and gall stones fail to be mechanisms, but not because they lack any activities. It is clear that paperweights hold down papers and gall stones block bile ducts. They also obviously are composed of entities since they both are material objects of some sort. The reason that paperweights and gall stones do not count as mechanisms is because of *how* entities are able to achieve their respective activities. Following Wimsatt (1986, 1997), Craver (2001) lays out four criteria under which a system is a

[21] Craver is the 'C' of MDC (2000)

mere aggregation of things rather than a mechanism. When considering a system with a certain output (activity), the system is a mere aggregate when the output of the system 1) will be produced even when the parts are rearranged and intersubstituted, 2) remains qualitatively similar with the addition or subtraction of parts, 3) will be produced even when the parts are disaggregated and reaggregated, and 4) is not produced as a result of cooperative or inhibitory interaction among the parts.

According to Craver's (and Wimsatt's) criteria, paperweights and gall stones are mere aggregations (not mechanisms) because their parts can be extensively rearranged and replaced and they will still do the same things they always did. In Wimsatt's terms, they achieve their function aggregatively. By contrast, a mammalian heart, in terms of its ability to pump blood, can suffer only very minor rearrangements of its parts without losing its ability to pump blood, and so its function is achieved non-aggregatively (not by mere aggregation of its parts) and is therefore a good candidate for being a mechanism. Of course, the actions of gall stones and paperweights are not completely impervious to rearrangements. Very radical rearrangement of their parts could ensure that a gall stone no longer blocks its bile duct or that a paperweight no longer holds down papers. This indicates that aggregativity is a vague property, but it can still be used to help distinguish mechanisms from non-mechanisms. The general idea is that the number of rearrangements that disrupt the function of a mechanism is much greater than the number of rearrangements that do not disrupt the function. For aggregative activities, on the other hand, only a relatively small percentage of possible rearrangements will cause a disruption. There may be a large gray area in which it is not clear whether a given system produces its activity aggregatively or non-aggregatively, but there are many clear cases on either side, and so the predicate is able to do substantial work in distinguishing mechanisms from non-mechanisms.

Notice that Craver's (2001) criteria are not meant to be criteria for objecthood, and so this account does not entail that mechanisms must be objects in any strict sense. It does, however, place some appropriate physical constraints on what can count as a mechanism rather than a mere aggregation. Craver's account applies very well to biological mechanisms since most biological mechanisms are composed and controlled largely by proteins, and proteins are very sensitive to rearrangement of their parts. Even small differences in amino acid sequence can alter the shape of a protein, and altering a protein's shape can easily change its activity. Even when macromolecules other than proteins are major constituents of a proposed biological mechanism, as in cell membranes (phospholipids) and chromosomes (DNA), Craver's criteria apply. Small changes in the parts of membranes or of chromosomes can change their activity in dramatic ways. Contrast this with features of an organism that are biological but not themselves mechanisms. For example, consider the layer of oil many land animals produce on their outer surface. While the production and excretion of the oil may be accomplished by a mechanism, the oil itself is not a mechanism. The parts can be rearranged in a wide variety of ways without significantly changing or eliminating its protective activity.

One objection to Craver's (2001) view may be that he is being selective about what he views as the parts (the entities). Take my example of oil on the surface of the skin. While it may be true that rearrangement of the individual oil molecules may not result in significant differences in the behavior or functionality of the oil layer, if we rearrange the atoms that compose the oil molecules, then we may disrupt the behavior of the material quite dramatically. Thus, a critic could claim that according to Craver's (2001) account, almost nothing would count as a mere aggregate because rearrangement of parts on the atomic or subatomic scale will always significantly disrupt the behavior of a system, whether it's a complex, dynamic system or just a

static, seemingly homogenous lump. Therefore, says the critic, Craver's non-aggregativity requirement fails to help us distinguish between a biological mechanism and a non-mechanism.

This objection to Craver (2001) is blunted by considering that every mechanism operates at a certain level of organization (Machamer, Darden, and Craver, 2000; Darden, 2002; Craver and Bechtel, 2007), and at each level there are entities that are understood to be basic. MDC (2000) call this "bottoming out" at a given level. A given level bottoms out when we reach entities whose properties are taken to be basic or unproblematic for explanation of items at that level. Oftentimes, the entities we find at the bottom are themselves mechanisms, but if so, they are mechanisms at a different (smaller) level of organization that contribute to the higher-level mechanism. In themselves, the activities of any lower-level mechanisms that compose a higher-level mechanism may be taken as basic when accounting for the high-level mechanism. This helps to blunt the objection to Craver (2001) because when we consider the level at which a proposed mechanism bottoms out, we need to ask whether rearrangement of those parts would disrupt the proposed mechanism's activity. If so, then the object or process potentially is a mechanism. If not, then it is a mere aggregate. In a heart, for example, the basic entities may include contractile proteins. The activity of the heart is very sensitive to rearrangement of those entities, and the actions of those proteins can be taken as basic in describing the heart. Perhaps those proteins also are mechanisms that are sensitive to the rearrangement of their amino acids[22], but amino acid sequence would not be relevant to a description of the heart as a mechanism and so they are below the bottom. By contrast, in descriptions of the layer of oil on our skin or a gall stone in a bile duct, where the molecular level is taken to be basic, rearrangement of component molecules will not significantly change the activity of either, and so their activities are

[22] Amino acid sequence is but one aspect of protein structure, as there can be a great deal of post-translational modification of proteins, including extensive cutting and re-arranging of parts, connection to other proteins, and addition of multiple saccharides.

aggregative. Therefore, such items do not have the structural characteristics required to be biological mechanisms. Mechanisms must have structure that, while perhaps falling short of full objecthood (whatever that may require), must still have a structure that does not allow for significant alterations without disruption of its activity.

2.4 Activities

Now I will consider the activities of mechanisms. According to MDC (2000), activities are simply what the entities accomplish. Of course, any given entity or process may end up doing a large variety of different things, so there needs to be a way to identify just one (or one small set) of those things as the activity, leaving the rest of the things that happen as mere side effects. I argue that the activity of a biological mechanism must be a biological function, where a biological function is an output that furthers the life of the organism of which the mechanism is a part. This general sentiment, if not its specific formulation, can be seen in Aristotle's *de Anima*, in which furtherance of life is an essential function of all creatures, and in Kant's *Critique of Judgment* in which he argues that we cannot understand something to be alive without it having the natural purpose of continued life. More recently, Dennet (1989, 1996) has argued that when examining living processes, we always adopt an "intentional stance" because we always assume that any living system acts with the intention of furthering its own life.

MDC (2000) hold something similar to my view in that they associate activities with functions, but their account of functions is more observer-dependent than I am comfortable with. They tell us that "[t]o see an activity as a function is to see it as a component in some mechanism, that is, to see it in a context that is taken to be important, vital, or otherwise significant" (MDC, 2000 p.6; see also Craver, 2001). MDC's account of functions is observer-

dependent in that a mechanism requires an activity that is the mechanism's function, and a function is determined by whatever the observers take "to be important, vital, or otherwise significant." Similarly, MDC (2000) also describe both the "set-up conditions" and the "termination conditions" for a mechanism in an observer dependent fashion. They explain that when scientists describe how a mechanism begins, they pick out a "static time slice" of an active, ongoing process as the starting point and explain how the conditions at that moment will be able to lead to the termination conditions (p. 11). Then, the termination conditions describe a "privileged endpoint" that is something that "we set out to understand or create" with the mechanism (p. 12). Thus, as long as a case can be made that a certain product or endpoint can be characterized as a function, then the responsible entity or process may be thought of as a mechanism.

Craver (2001) elaborates on MDC's view by explaining that since a given system (S) is likely to have a wide variety of outputs, mechanistic explanations are "ineliminably perspectival" and "rely upon shared background assumptions that S can ψ or that the ψ-ing of S is important, significant, or relevant. The idea that the heart is for making glub-blub noises only seems absurd until one is able to conjure (often with some contortion) a suitable mechanistic context for those glub-blub noises" (p.71). Glennan (2002) makes a very similar point when he explains that "A complex system may exhibit several different behaviors, and the decomposition of the system will depend on which behavior is under consideration… the heart qua pump may admit of a different decomposition than the heart qua noisemaker." Finally, Darden (2002) also reinforces this view in her discussion of how mechanisms are discovered. She explains that one way the search for a proposed mechanism may begin is when "some phenomenon to be produced, some start or end stage, and some entities and/or activities (or roles for them) are specified" or when

researchers "sketch hypothetical roles that components of the mechanism being sought are expected to carry out" (p. S360). Again, the work of identifying the function of a proposed mechanism seems entirely up to the creativity of the researcher.

Although they clearly acknowledge the importance of functions, I believe that functions are more central to a proper account of mechanisms than the above authors imply, and more careful focus on biological functions can help make the distinction between natural mechanisms and non-mechanisms less observer dependent. Before doing so, however, we should be clear about the relevance of this issue. If the mechanistic approach is merely a methodological tool that provides a useful heuristic for scientific investigation, then perhaps it is appropriate for functions to be completely observer-dependent. If it is merely a researcher's tool, then the researcher should feel free to use that tool in any way she thinks is useful. However, I am defending the view that the mechanistic approach more than just a useful heuristic. The mechanistic approach is useful because it captures something important about living organisms themselves, not just something about our peculiar way of understanding organisms.

There is an extensive literature concerning biological functions and teleology (e.g. Ayala, 1970; Wright, 1973, 1976; Cummins, 1975; Nagel, 1977; Falk, 1981; Nissen, 1993; Millikan, 1998). As mentioned in the historical overview in Section 1.1, modern, naturalistic science assumes that there are no vital forces, and aside from sentient animals and their intentional mental states, there are no fundamental teleological forces that drive simple physical matter toward a particular goal like life or adaptation. Therefore, when biologists (and philosophers of biology) speak of natural functions, they do not mean to invoke any Aristotelian final causes or vital forces. Instead, they prefer to invoke only standard causation as understood in physics and chemistry to account for the functionality of biological systems and subsystems. However, as

Kant suggested, it is not clear that causation as understood in physics or chemistry (what Aristotle would have called "efficient" causation) can be teleological by itself. Basic causal forces just do what they do, without any goals or functions.

To get a sense of the issue, let us start with a common (and incorrect) first attempt to characterize biological functionality. One might say that an item (X) in a living system (S) has function (F) if S does F, X contributes to the performance of F in S, and S would not be able to do F without X. For example, consider a specific living mammal. There is a heart that contributes to the pumping of blood in that mammal, and the mammal would not be able to pump blood without its heart. Therefore, that heart has the function of pumping blood in that animal. The problem with this relatively simple formulation, as pointed out by Wright (1973), is that it allows side effects to count as functions. For example, under normal environmental conditions, the heart produces a thumping sound when it is working inside the body. Production of the thumping sound meets all of the criteria above, but we are not at all inclined to think that the heart functions to make a thumping noise. This example shows that we need to add something else to our account of natural functions that distinguishes functions from side-effects.

Perhaps the most common strategy is to use an etiological (historical) account of biological functions based on evolution by natural selection (see, for example, Allen (2004), Ayala (1970), Cummins (1975), Falk (1981), Woodfield (1998), and Sober (2003)). Put very briefly, as this is the focus of Chapter Five, here is how evolution by natural selection works. Individuals with characteristics that happen to make them better suited to living in their environments will tend to survive and reproduce better than individuals without those favorable characteristics. As a result, the favorable characteristics will tend to become more prevalent in the population. For example, many mammals have fur that provides very good camouflage.

From among the diversity of fur colors in the ancestral population, individuals with fur that made them harder to detect by predators survived to produce offspring with a similar fur color. According to the etiological account of natural teleology, since the fur color was naturally selected because of camouflaging properties, the purpose, or function of the fur is to provide camouflage.

While the etiological account based on natural selection captures evolutionary reasoning fairly well, Lowell Nissen (1993) presents a serious problem. He writes,

> Suppose that, as a glacial period slowly approaches, the fur of rabbits gradually changes from brown to white. Regarding a certain generation, e.g. the fiftieth after the first snowfall, saying that a few rabbits had slightly lighter fur than the rest and also survived and reproduced at a slightly higher rate than the rest does not mean that they had lighter fur in order to live longer and reproduce better or that the lighter fur was supposed to protect from predators, neither does it justify or support such claims. If a historical account of one generation provides no grounds for such claims, neither will fifty repetitions of that account, nor a large number of such repetitions describing such events over a long span of time (p. 33).

As Nissen sees it, since teleology implies a goal toward which the system is working, teleology is necessarily forward-looking, while etiological accounts are exclusively backward-looking. By and large, new variations arise spontaneously and blindly (see Chapter Three), and when a new variation causes an organism to be better adapted than its conspecifics, that increase in adaptedness manifests *after* the new variation has appeared. This leads to a problem because when we claim that a new variation has an adaptive purpose we imply that the trait exists because the trait has that function. But if the function only manifested after the trait appeared,

then the function cannot be a reason for the trait's existence because that would require reverse causation from a future state to a present one. Aside from strange effects in quantum mechanics, reverse causation is considered impossible. Hence, argues Nissen (1993), the etiological account of natural teleology based on evolution by natural selection fails.

One route to salvaging the etiological account of natural teleology in the face of the above critique is represented by Karen Neander (1991). In order to avoid problems with reverse causation, Neander focuses on the *type* of trait in question as opposed to the individual token of that type. She writes, "It is the/a proper function of an item (X) of an organism (O) to do that which items of X's type did to contribute to the inclusive fitness of O's ancestors, and which caused the genotype, of which X is the phenotypic expression, to be selected by natural selection" (p.180). Consider again our example of hearts as blood pumps. Neander explains that it is clear that hearts in general have contributed to the fitness of animals in the past, and the reproductive success of a given animal's ancestors is a causal antecedent of the existence of the animal now. Therefore, the presently living animal's heart has the purpose of pumping blood because its heart belongs to the type of thing that contributed to its ancestor's survival and reproductive success by pumping their blood in the past.

Neander's account seems to capture something of the spirit of evolutionary explanations, but, as explained again by Nissen (1993), there remains a serious causal problem. In Neander's theory, the teleological relation between X and O (the item and the organism) is a relation between the types to which X and O belong. But types are abstractions from their instances, and as such do not have any causal power of their own. Therefore, an explanation of the existence of an adaptive trait cannot be based on the type to which the trait belongs because a *type* of thing cannot cause an *actual* thing to come about. So, rather than suffering from reverse-causation,

Neander's account fails to include any causation at all, and so she also fails to show how an etiological account based on natural selection can establish biological functionality.

In addition to the problems discussed above, an etiological account also runs the risk of excluding the first living organisms. The first organisms (or first living systems, if "organisms" is an inappropriate label) which emerged some 3.5 billion years ago, may not have had the proper history of natural selection required by etiological accounts, yet it seems inappropriate to claim that such organisms lacked any biological functions. A similar problem arises from a well-known thought experiment. Suppose that the lifeless muck of a stagnant swamp just happens to come together in such a way as to produce a living animal, sometimes called "swampman" (from Davidson (1987)). Swampman appears to be every bit alive as any other organism, so if other organisms possess some kind of teleology, then swampman must too, even though swampman fails to have a history comparable to other organisms.

Because of the problems with the etiological account, my own account of natural teleology does not rely on etiology at all, but rather focuses only presently existing biological systems. I believe that biological functionality comes from the fact that living systems are inherently directed toward continued living, where being alive is a special sort of property that exists in the world independent of any human (or human-like) observers. The precise differences between living and non-living systems are notoriously difficult to pin down (Schrödinger, 1945; Mayr, 1998; Kauffman, 2000; Korzeniewski, 2001; Cleland and Chyba, 2002). However, the existence of biology as a field of research is based on the uncontroversial assumption that there is a very real difference between life and nonlife. To be alive is a real property[23] of certain complex systems, and systems with that property are able to exhibit natural functionality because

[23] There may be a variety of ways to be alive, so perhaps "alive" is not a single property. However, all known life on earth appears to share in the same fundamental style of living so at least all terrestrial organisms share the same property of being alive.

their processes have a natural goal: continuance of that life. Continuance of a life is most straightforwardly achieved by the continued survival of the living organism, and so most processes carried out by an organism that aid survival have organismic survival as their function. However, reproductive processes also have biological functionality even though they do not contribute to the survival of the reproducing organism. In fact, reproduction often directly harms survival of the reproducing organism. That is why I chose to call the goal "continuance of that life" rather than "survival or reproduction." We tend to think of life in terms of individuated organisms, but the process of life is really an unbroken, ever-branching set of lineages dating back to life's origins over 3 billion years ago. Therefore, it is perfectly legitimate to say that both survival and reproduction by an individual organism contribute to furtherance of the life—the ongoing, living process—that is happening in that organism's lineage. So, most processes carried out by organisms serve the general function of promoting the continuance of the life of which the process is a part.

My account does not run into causal problems because there is an ongoing, reciprocal causal relationship between a mechanism with a biological function and the life to which the mechanism contributes. The living organism is responsible for the continuing existence of the process, and the process contributes to the continuing existence of the organism. More formally, some biological item X performs a biological function for organism O when X contributes to the continued life of O and the life of O contributes to the continued existence of X.

My account of biological functions requires that the entities that carry out the relevant process must be constructed and controlled by an organism. Many things external to an organism, such as rainfall and sunshine, may also contribute the organism's survival, but it is not appropriate to say that rainfall and sunshine have the *function* of aiding some particular plant. It

is appropriate to say that the leaves of that plant have the function of absorbing sunlight. The difference is that a plant's leaves are constructed, maintained, and regulated by the plant, while sunlight, rainfall, and other external processes are not. Interestingly, a process need not be part of an organism to have a biological function. For example, certain kinds of niche construction, such as bird nests and ant hills, are constructed and maintained by organisms in order to further the life of those organisms, and so serve that function even though they are not part of the organism itself.[24]

Although life and natural functionality are not observer-dependent properties of living systems, when it comes to investigating natural mechanisms, there still is an element of observer dependence at work. A certain mechanism may have more than one biological function, or the entities that compose a mechanism may also be involved in composing another mechanism, and so one researcher's characterization of a mechanism may depend on which function is under consideration. Thus there still is an element of observer dependence in the mechanistic approach because the activity to be studied is up to the researchers, as long as the activity can be shown to have a biological function.

A final issue to address in this section is the reliability of mechanisms. Recall the account from MDC (2000) who require that mechanisms be "productive of regular changes from start or set-up to finish or termination conditions" (2000, p.3). For his part, Glennan (2002) says that mechanisms produce their results by way of "direct, invariant, change-relating generalizations" between the parts of a mechanism (p. S344) and must be distinguished from "genuinely singular events" (p. S348). Calling the action of mechanisms "regular" or "invariant" suggests a very high level of reliability, but a strictly deterministic result where we get exactly

[24] I am not claiming that bird nests are mechanisms. Biological functionality is just one criterion for being a biological mechanism.

the same result every time a mechanism operates is far too strict of a criterion for at least two reasons. First, it may be that because of micro-physical variabilities, no actual mechanisms are ever able to achieve such an extreme level of reliability. Second, there are many examples of mechanisms that are *supposed* to produce variable results, such as lottery number selectors and sexual reproduction. Therefore, some amount of variation in the outputs of a mechanism must be allowed. Then, as long as the criteria I laid out above are met, then the concerns about unique or highly fleeting processes is alleviated. If the process is maintained and regulated by an organism and is directed toward continuance of the life of which the organism is a part, then the process will not be unique, fleeting, or random occurrence, but will instead produce regular results in the way desired.

In summary, a biological mechanism is composed of entities with an activity. The activity must be a biological function and the entities must achieve their function non-aggregatively. A mechanism has a biological function when the mechanism is part of an organism and positively contributes to continuing the organism's life and, in turn, the organism actively maintains the existence of the mechanism. The activity is achieved non-aggregatively when the entities must be properly coordinated in space and time in order to function. The entities of a biological mechanism cannot be rearranged in any significant way without disrupting the mechanism's ability to achieve its biological function.

CHAPTER THREE

HEREDITARY REPRODUCTION

3.1 Introduction

Evolution involves accumulation of changes over time, and since all organisms eventually die, organisms must be able to pass their traits on to their offspring in order for evolution to continue over the long term. Therefore, heredity is essential for adaptive evolution. There never has been significant controversy over whether organisms can reproduce with heredity. Instead, questions have concerned what kinds of traits are inheritable, and what physical entities are involved in heredity and how they work. Answers to these questions have had significant impacts on evolutionary theory, closing off some possibilities (such as Lamarckian evolution) while supporting others (such as neo-Darwinian evolution). Section 3.2, is a historical overview of research into biological heredity and how it has shaped the way scientists think about biological evolution.

Most of the really successful research into heredity has come from genetics, so genetics has dominated the field of heredity during most of the history of modern biology. The success of genetics led to a gene-centric view of biology, especially evolutionary biology, that overshadows the roles of non-genetic processes in biological heredity and evolution. In Sections 3.3 and 3.4 I use the mechanistic approach detailed in Chapter Two to dislodge the gene-centric approach and elevate non-genetic process to their rightful place in biological theory.

Before going any further, however, it should be noted that sometimes it can be difficult to distinguish a mechanism of heredity from a mechanism that produces new variations. There seem to be at least two reasons for this. First, many of the biological mechanisms that generate new variations work in conjunction with mechanisms of heredity, sometimes even using the very

same machinery (or very same "entities" in the language of MDC (2000)). For example, Meiosis in preparation for sexual reproduction includes steps that systematically introduce new combinations of traits into the resulting offspring. Thus, Meiosis is a good candidate for being a mechanism of heredity (see Section 3.4) and for being a mechanism for generating new variations (see Section 4.4). Perhaps it can be difficult to conceive of two separate kinds of mechanisms at work at the same time and realized by some of the same cellular machinery, but this is the view that I defend.

The second reason for the difficulty in distinguishing a mechanism of heredity from a mechanism for producing new variations is that, within evolutionary biology, a mechanism of heredity is only considered evolutionarily relevant if it can carry new variations to the next generation,[25] and so heredity and variation typically are discussed in tandem. For example, the gene-centric view, which still dominates most biological sciences in one form or another, assumes that genetic changes (those associated with nucleic acids) are the only changes that can be biologically inherited, and so genetic mechanisms of heredity (those that replicate and transfer nucleic acids) have come to be thought of as the only mechanisms of heredity. Just because a process cannot preserve changes over multiple generations does not mean that it is not a proper mechanism of heredity, and so I will not restrict my discussion only to processes that are thought to be evolutionarily relevant. Then, once other, non-genetic processes of heredity are accepted, it will become clear that at least some of them actually can preserve changes over many generations and so really are evolutionarily relevant after all.

[25] The reverse is also true: production of new variations is evolutionarily relevant only if the changes are inheritable.

3.2 Historical Development of Concepts and Terms

Although it had always been obvious that heredity is operative in biological reproduction, technological limitations severely hampered investigations into the physical process(es) by which heredity is achieved, at least until the rise of molecular biology in the mid-20th century. As a result, proper understanding of the actual mechanisms of heredity remained somewhat of a mystery. Nevertheless, biologists were able to observe and measure the phenomenon of heredity, and so were able to refine their account of the activity even when they only had rudimentary knowledge of the entities. Modern biology really started to take form near the turn of the 19th century, so I begin my account there.

At the end of the 18th century, most scientists assumed that species represented unchanging types of organisms. While some superficial variability could be observed in any species, if species are types, then there is a standard form or essence for every species, and members of the species should never stray too far from the standard form (Mayr, 1982; Gould, 2002).[26] According to this view, an account of the overall activity of any hereditary mechanism is relatively simple because all it needs to do is produce more of the same. While popular within the scientific community at the time, such typological thinking about species was tempered by plant and animal breeders who recognized that novel traits could emerge in a population and be passed on to offspring, and they knew that selective breeding could change a species quite radically. However, despite the plant and animal breeders' track record of changing species in dramatic ways, typologists were able to dismiss their results by pointing out that domesticated animals, when returned to the wild, tended to revert to their wild forms and to lose any special traits that had been promoted through artificial breeding programs (Mayr, 1982, 2001a). Thus,

[26] Occasional extreme deviations from the standard form were considered "monstrosities" and did not persist in the population.

the effects of artificial breeding programs, though sometimes quite dramatic, seemed superficial and readily reversible, and there was little reason to think that the very large, permanent changes needed for general evolution of one species into a new species was at all possible.

By the early 19th century, enough evidence from the fossil record had been gathered to indicate very strongly that some sort of long-term change in species on the planet had occurred, and so typologists needed to explain how this could occur. One option for typologists such as the influential anatomist and paleontologist George Cuvier, who wanted to avoid any kind of change in species, was to say that old species regularly went extinct and were somehow replaced by new species, perhaps via immigration from other areas (Mayr, 1982; Hull, 1984). Not all typologists were opposed to evolution of species, however. But those inclined toward evolutionism faced a problem because typology typically assumed that an organism's set of traits form a coherent whole, and that changing just one or a few traits would disrupt the essence of the organism to the point that it would not be viable (Mayr, 1982). Thus, gradual, step-by-step evolution was considered impossible because small changes would result in non-viable offspring. A prominent way of reconciling evolution with typology was to adopt some version of "saltationism" (Mayr, 1982; Dawkins, 1983) in which very rare but very radical changes (saltations) could appear suddenly in a lineage—even over a single generation—thus converting the species into a different viable type. Although no plausible mechanisms were suggested that could perform the activity proposed by saltationsts, it seemed consistent with the fossil record, which was relatively discontinuous at the time, and so saltationism (or something like it) remained entrenched among evolutionists in the mid-19th century (Mayr, 1982; Hull, 1984).

An obvious deviation from early 19th century typological thinking was J.B. Lamarck (1809) who presented one of the first full theories of biological evolution. Lamarck's theory

famously included a process of heredity in which certain characteristics acquired by a parent could be passed down to its offspring, thus leading to evolution of the lineage that, over many generations, resulted in very dramatic changes. Lamarck was unable to suggest any entities that could carry out his proposed hereditary activity, but that problem was common to all theorists, and so was not what caused the biggest problems for Lamarck during his lifetime. Rather, it was his commitment to gradual, continuous evolution of species from very simple to very complex, which required transitional forms, that led to his marginalization (Mayr, 1982; Hull, 1984). In particular, George Cuvier and Charles Lyell used their knowledge of the fossil record to show that although life on earth may have changed over time, it occurred in discontinuous jumps rather than through gradual evolution of lineages (Mayr, 1982; Hull, 1984). Further, Lamarck theorized that all lineages go through similar progressions from simple to complex forms, but Cuvier showed that there is no overall trend or plan in the history of life on earth. So, even as long-term and large-scale evolution of species was gaining serious attention, typological thinking still held sway.

The next prominent theory of biological evolution to challenge typological thinking was Darwin's in 1859. Like Lamarck, Darwin proposed gradual evolution of species rather than species stasis or discontinuous jumps, but unlike Lamarck, Darwin avoided postulating an overall trend toward complexity or any other kind of goal other than mere survival. In that way, Darwin was careful to avoid the penetrating criticisms of Lamarck made by Cuvier and Lyell. However, Darwin's gradual evolution of species still ran afoul of typological thinking, which partly contributed to the slow rate of acceptance of Darwinian's ideas.

Regardless of the differing views on species as types or on evolution, it was clear that hereditary mechanisms must cause offspring to resemble their parents. As mentioned above,

technological limitations severely constrained what could be known about the entities that carried out the activity of heredity. The most common hypotheses from before the 19th century usually involved some sort of "preformationism" in which offspring were provided with some preformed material, perhaps very small versions of their adult tissues and structures, that would grow into the adult form. By the late 18th century, advances in microscopy had allowed biologists to observe cell division and the early stages of embryological development. While they failed to observe anything like the smaller versions of their adult forms, they did find *some* cellular structures. Thus they were able to rule out the homunculus version[27] of preformationism, but not more sophisticated versions of preformationism such as Darwin's gemmule theory (Darwin, 1868; Mayr, 1982; Jablonka and Lamb, 2005) in which offspring inherited preformed gemmules that contribute to forming the offspring's tissues and organs. Thus, preformationism endured as a viable option in the 19th century, even though nobody had yet been able to verify the existence or specific characteristics of entities that could carry it out.

As an alternative to preformationism, some proposed instead that the action of some special, undifferentiated substance, often called 'germ plasm,' was responsible for hereditary activity (Mayr, 1982). It was not clear what the germ plasm was or how it worked, so when it was introduced it was little more than a placeholder for whatever entities actually composed the mechanism. However, that does not mean that biologists couldn't make some informed guesses about what hereditary activities could be performed by the germ plasm and which could not.

For example, in the late 1880's, August Weismann took note of observations showing that very early in embryological development, cells that eventually become those used in reproduction, what Weismann called "germ cells," are sequestered away from all other cells that

[27] The homunculus version of preformationism posited that offspring were provided with a small version of its adult form, and that development consisted of little more than growing of the homunculus.

compose the rest of the organism's developing body, dubbed "somatic cells." Weismann reasoned that since germs cells are isolated away from somatic cells, any changes to the somatic cells cannot be transferred to the germ cells. Then, since any changes acquired by a parent organism during its lifetime will only affect its somatic cells, its offspring cannot inherit any new characteristics the parent had acquired. This helped convince Weismann that the only inheritable changes were those generated in the germ cells or during the production of offspring. Weismann's views were given a boost when change-inducing processes were observed during Meiosis, which only happens to germ cells, but not in Mitosis, which is the method by which somatic cells divide. For example, before homologous chromosomes[28] are distributed into gametes, small segments of each member of the homologous pair physically overlap (cross over) and are exchanged, thus slightly shuffling the chromosomal material that ends up in each gamete. Then, when a randomly selected male gamete fuses with a female gamete, new homologous pairs of chromosomes are matched together to form an individual with a unique hereditary endowment not seen in either parent. If chromosomes carry hereditary information (a plausible but not fully established proposal at the time[29]), then all that shuffling and recombining of chromosomes is likely to be an important source of variation. And since all that mixing of hereditary material happens in parental germ cells rather than to somatic cells, the processes do not end up altering any parental traits. Therefore, any new traits that show up in the offspring were never displayed (were never acquired) by parents. Thus, Weismann argued, there was no inheritance of acquired characteristics.

---

[28] Homologous chromosomes are pairs of chromosomes that contain similar hereditary information. An individual's parents each contribute one member of every homologous pair.

[29] The details of how chromosomes contribute to heredity was not known until the mid-20th century, but observations of cellular activity had already made it clear that chromosomes were key to hereditary processes.

In addition to the germ cell / somatic cell distinction endorsed by Weismann, there is a closely related distinction, established a few decades later and attributed largely to Wilhelm Johansson, between an organism's *genotype* and its *phenotype*. The phenotype is the set of observable characteristics an organism possesses, while its genotype is the hereditary information (contained in the germ plasm) that the organism received from its parents and that directs development of the organism's phenotype. The term "genotype" comes from the notion that this stuff, whatever it is, is able to *gen*erate the proper traits in the offspring. The term "phenotype" refers to the observable *phen*omena—anatomy, physiology, behavior, etc.—that an organism displays (Downes, 2010). With this terminology in place, we can say that according to neo-Darwinism, changes to an organism's genotype can change its phenotype, but changes in an organism's phenotype that occur during the organism's lifetime cannot change its genotype, and therefore such acquired changes cannot be passed on to offspring.

Since Weismann could support his views with available knowledge of plausible hereditary mechanisms, his arguments eventually proved to be very effective against remaining supporters of inheritance of acquired characteristics. When Weismann's views were combined with Darwin's theory of evolution (minus Darwin's own views on heredity) the result was what came to be known as neo-Darwinism. Neo-Darwinian evolution was Darwinian because it involved common descent and evolution by natural selection, but it did not allow for inheritance of acquired characteristics (Mayr, 1982).

Weismann's version of heredity eventually won out, as discussed below, but in the late 19th century, neo-Darwinism was by no means the consensus view. Typologists opposed to any long term evolution still held sway, and even among evolutionists there was much doubt concerning the power of natural selection to drive the dramatic, long-term changes implied by

common descent. This is relevant to heredity because criticism of the effectiveness of natural selection was based in large part on the presumed inability of the hereditary mechanisms (which were still mostly mysterious) to preserve any new variations, even if they provided a selective advantage. That is, even if significant new variations could appear in some members of a species, and even if some of those new variations were advantageous, it was believed that those new variations would not be preserved over the generations and so would not cause the species overall to evolve. This view was supported by the common belief that heredity involved some amount of blending between parents. If heredity really did involve blending, then since most individuals would not possess the new, advantageous trait, the new trait would be diluted by successive rounds of blending reproduction until there was no trace of it left (Mayr, 1982).

Things began to change in the early 20th century, with the rediscovery of Gregor Mendel's mid-19th discovery that certain traits are inherited as discrete units and that the units can be mixed and matched in different ways during the process of sexual reproduction. Although this new work described the activity of heredity without directly studying the underlying physical processes, the proposed activity fit very neatly with what little was known of the entities (germ cells and chromosomes) that were thought to compose the hereditary mechanisms. In particular, observations of chromosome behavior inside of germ cells during Meiosis were compatible with Mendel's Laws of Segregation and Independent Assortment[30], and so it appeared that the Mendelian program was on the right track. This emerging view of heredity also was consistent with Weismann's views on hard inheritance, but it was not clear that

[30] The Law of Segregation states that the two alleles associated with a particular trait are separated (segregated) during gamete formation so that each gamete has one allele for each trait. The Law of Independent Assortment states that allele pairs for different traits are segregated independently from one another, so that segregation of one allele pair does not affect the way other allele pairs are segregated.

Mendelian inheritance could support evolution by natural selection, so neo-Darwinism remained a minority view until the 1930's.

It is worth taking a moment to assess this historical episode using the mechanistic approach. Mendelians had some important advantages over their rivals. Experimental manipulation and quantitative analysis of Mendelian inheritance patterns was relatively easy, and their research generated clear, reproducible results with interesting practical implications. Drawing upon Chapter One, we can say that the Mendelians were able to provide very good explanations, predictions and counterfactuals, and that this drove their success. Mendelians even claimed to make use of laws of nature that govern the way gene variants (alleles) are distributed to offspring.[31] We know now that Mendel's Laws no longer deserve such status, and it seems that this episode instead is more consistent with Darden's (2002) account of mechanism discovery. When searching for mechanisms, observation of activities guides the search for entities that can perform those activities. At the same time, even rudimentary knowledge of the entities involved in a mechanism can guide further understanding of the activities (MDC, 2000; Darden, 2002; Woodward, 2002b, 2003). As discussed above, Weismann used basic observations of the entities, namely germ cells and chromosomes, to argue that the activity of heredity was did not allow for inheritance of acquired characteristics (Mayr, 1980, 1982). The Mendelians' views on inheritance were also consistent with what was known of chromosome actions, and since they also were providing precise characterizations of the activity of heredity, their project took the lead. Contrast this with the relative failure of those who still defended some version of Lamarckian evolutionary processes, especially inheritance of acquired characteristics (IAC). Defenders of IAC were unable to point to satisfactory activities or entities to support their views, and so their research programs faded away in favor of Mendelian

[31] The proposed laws were the Law of Segregation and the Law of Independent Assortment.

inheritance (Burkhardt, 1984; Hull, 1984; Sarkar, 2005). In addition, since the mechanisms of inheritance did not appear compatible with natural selection, Darwinian evolution also remained a minority view.

With its successful use of the mechanistic approach, the science of heredity at the beginning of the 20th century came to be dominated by the Mendelian research program. Mendelians soon adopted the term 'genetics' as a label for their field, and they claimed to be studying "genes." As with 'genotype,' the term 'gene' (attributed to Wilhelm Johansson) comes from early ideas about the *gene*sis or *gene*ration of life. However, since geneticists at the time still knew very little about the entities that realized the mechanisms they were studying, they conceived of genes as abstract bearers of hereditary information associated with specific phenotypic traits (Mayr, 1982; Rheinberger and Müller-Wille, 2010).

According to early Mendelian genetics, a single trait, such as eye color, could be attributed to one gene, and researchers could track the inheritance of that gene through the generations by tracking the trait it produced. Further, geneticists eventually recognized that new, inheritable traits could spontaneously appear in a lineage and be preserved over multiple generations. Despite Hugo de Vries' previous use of the term 'mutation' to refer to large, whole-organism changes needed for saltationism, Thomas Hunt Morgan dubbed those new gene variants 'mutations' as the label stuck. In retrospect, we can see that the discovery of mutations helped pave the way for acceptance of neo-Darwinism, but since genetics was the more successful and respected research program, that acceptance would need to wait until natural selection was formally reconciled with genetics.

That process of reconciliation began when simple Mendelian genetics gave rise to population genetics. Rather than simply tracking inheritance patterns between parents and

offspring, population genetics tracks the frequency of different gene variants (alleles) in whole populations and tracks how allele frequencies change over time. Since they were studying population changes over time, their work provided a way for genetics to be connected to long-term evolution. The connection was not easy to establish, however because the process of heredity, as understood by Mendelians of the first decades of the 20th century, still seemed unable to preserve incremental changes needed for evolution by natural selection. By the 1930's and 1940's, however, a more mature field of population genetics used newly developed statistical techniques for modeling populations to show that genetics is indeed compatible with neo-Darwinian evolution by natural selection. The resulting combination of population genetics and neo-Darwinian evolution came to be known as the Modern Synthesis and is the basis of contemporary evolutionary theory.[32]

Soon after development of the Modern Synthesis, new methods for studying biological molecules finally began to gain access to the entities that carry out heredity, focusing mainly on chromosomes as the most likely entities for carrying our genetic inheritance. By the 1950's it was clear to geneticists that the activity of genes was accomplished by double-stranded DNA. Consequently, the term 'gene' became attached to specific stretches of DNA, in particular, stretches of DNA that are transcribable. In the process of transcription, DNA is used by cells to produce specific proteins that go on to play major roles in producing the phenotypic characteristics. So, according to the new understanding of the Modern Synthesis, transcribable sections of DNA—genes—carry the genotype that is responsible for producing the phenotype. Further, the term 'mutation,' which previously was used to refer to inheritable phenotypic

[32] The contemporary theory of evolution has modified and added to the Modern Synthesis so extensively that it is not clear whether today's version is the same theory in modified form, or whether we now have a completely different theory from the Modern Synthesis. This issue does not affect my arguments, so I take no stand on the matter.

changes, came to refer to changes in nucleotide sequence of genes because such sequence changes are the cause of those inheritable changes in phenotype (Rheinberger, 2000; Rheinberger and Müller-Wille, 2010). Consequently, the field of genetics, while still tracking changing allele frequencies with the techniques of population genetics, shifted much of its focus onto the molecular study of how protein-building information is stored in sequences of nucleic acids and how such sequences are passed on to and expressed in offspring.

The amazing success of this continuing line of research combining molecular genetics and population genetics led some in the biological community to assume that passing on of genes from parent to offspring is *the* process (or mechanism) of heredity and that all other cellular processes active in reproduction and development are just the background conditions needed for genes to do their work (Williams, 1966; Dawkins, 1976, 1980; for critiques see Gould, 2002; Robert, 2004; Jablonka and Lamb, 2005). Even more dramatically, this gene-centric approach tended to appropriate not just the study of inheritance, but all of evolutionary biology. For example, Theodosius Dobzhansky, a founder of the Modern Synthesis, helped popularize the notion that biological evolution is nothing more than changes in gene (allele) frequencies in a population over time. This positioning of population genetics at the center of evolutionary biology has had large consequences for biological theorizing, as will be discussed in Section 3.3 and in the following chapters.

As molecular biology progressed through the latter part of the 20th century, the gene-centric view remained dominant, but the concept of a gene and its role in heredity continued to change. Molecular geneticists discovered ever more complex gene interactions, making it very difficult for geneticists to associate specific genes with specific phenotypic traits, as was done in basic Mendelian and population genetics. In addition, the molecular concept of a gene simply as

a transcribable section of DNA was shaken by the discovery of diverse gene regulation processes. For example, there are extensive segments of DNA that control which genes are transcribed and when. Such segments of DNA do not carry information concerning any particular proteins but have nevertheless come to be called 'regulatory genes' because they are inheritable sections of DNA that actively contribute to the development of an organism's phenotype. Then, by the end of the 20th century, biologists had become aware of so-called 'epigenetic' processes, such as DNA methylation, that involve chemical modifications to existing DNA that can activate or inhibit specific genes, producing dramatic impacts on phenotype (Robert, 2004; Jablonka and Lamb, 1989, 2005). Discovery of all these gene regulation mechanisms, along with growing recognition of the importance of environmental conditions on development, led to rejection of the strongest forms of gene-centrism in which transcribable genes all but determined an organism's phenotype. Nevertheless, a softened form of gene-centrism continues to dominate biology, with other cellular processes and environmental conditions acting as mere support for gene actions. That is, while the roles of regulatory mechanisms and environmental conditions are acknowledged, explanations and predictions, especially in heredity and evolutionary biology continue to be couched in terms of gene actions (Robert, 2004).

Recent authors have argued that even a softer version of gene-centrism is unjustified (Landman, 1991; Robert, 2004; Jablonka and Lamb, 1989, 2005. They argue that other mechanisms of heredity exist besides genetic heredity, and some of those mechanisms are evolutionarily relevant. In the following sections I continue those efforts. I do this by first clarifying, in section 3.3, what activity a mechanism of heredity should be performing, regardless of what entities compose the mechanism. Then, in Section 3.4 I use that account of hereditary

activity to show that non-genetic heredity is as vitally important to successful reproduction as is genetic heredity.

3.3 Activity of Heredity

Recall that every biological mechanism must have an activity and entities that carry out that activity. Research programs like those discussed in Section 3.2 typically study both activity and entities, but considering the activity by itself can be very instructive. When we leave out the entities (e.g. genes, germ cells, etc.), then an account of the activity of heredity is fairly uncontroversial. As a first approximation, Lewontin (1970) described heredity merely as correlation between parent and offspring, a view echoed in Downes (2010). While that seems in the right direction, Lewontin's account allows for correlation due to chance while biological heredity is not something that can happen by chance. That is, if one individual produced another, and the second resembled the first only by chance, then the process was not of heredity, but mere luck. Instead, heredity is something that is caused to happen by the parent and is regulated by that parent. More precisely, heredity occurs when a process that is initiated and coordinated by the parent causes characteristics of the parent (or other recent ancestors) to appear in the offspring. The resemblance need not be exact, but in order for a process to have heredity as its activity, the process must make it likely that at least some characteristics of the parent appear again in the offspring in roughly the same form (Gould, 2002).

Before proceeding, a terminological clarification is needed for both 'replication' and 'reproduction' since both are used extensively in the literature on heredity. Roughly speaking, replication occurs when an individual (typically a single kind of molecule or polymer) serves as a template for making a high-fidelity copy of that individual. An uncontroversial example is so-

called semi-conservative replication of DNA in which the two complimentary strands of a double-stranded DNA polymer are separated, and both are used as templates for constructing new complimentary strands. The result is two double-stranded DNA molecules with nucleotide sequences roughly identical to the original. On the other hand, standard examples of reproduction all involve a more complex process in which an individual (typically a complex entity composed of many different kinds of molecules) is built to resemble another, pre-existing individual. Although the terms typically are used without much controversy, it seems that there is no clear, consistent distinction between reproduction and replication available in the philosophy of biology literature.

Evidence of confusion over the proper distinction between replication and reproduction can be seen in the highly influential work of David Hull (1980), who uses replication as a sort of catch-all for any copying or reproduction. Consequently, Hull does not distinguish between replication and reproduction per se, but between direct replication (e.g. the kind that involves DNA) and indirect replication (e.g. production of offspring). Hull (1980) starts out, roughly following Williams (1966) and Dawkins (1976), by distinguishing between a replicator, which is "an entity that passes on its structure largely intact, directly in replication,"(p. 318) and an interactor,[33] which is an organism that houses and supports the replicators. Hull then explains that while replicators "replicate themselves directly, but interact with increasingly inclusive environments only indirectly," interactors "interact with their effective environments directly, but usually replicate themselves only indirectly" (319). Hull then goes on to explain that while unicellular organisms would, under this view, be considered replicators, multicellular, sexually reproducing organisms would not because "replication at the level of sexual organisms is

[33] Dawkins (1976) referred to organisms as 'vehicles' for the genes in order to emphasize their subordinate role. Hull (1980) calls them 'interactors' in order to emphasize the notions that organisms interact with the environment while genes do not.

indirect" (p. 321). So, according to Hull, some entities can replicate without being replicators, as long as their replication is indirect. If something can replicate without being a replicator, then it seems that Hull is using 'replication' to refer to any kind of biological copying process, including what we normally call reproduction, as evidenced by his willingness to concede that reproduction of unicellular organisms probably counts as replication according to his view.

A more promising suggestion comes from Szathmáry and Maynard Smith (1997) and Wimsatt and Griesemer (2007) who describe reproduction as a composite process of "progeneration," where progeneration involves material overlap between the original and the copy. A potential problem I see with using this account to mark the distinction between reproduction and replication is that the proposed account of reproduction applies to DNA replication. Semiconservative DNA replication results in two new double-stranded DNA molecules, each with one old strand and one new strand. Thus, DNA replication involves substantial material overlap—progeneration—resulting in numerically distinct individuals, and therefore would count as reproduction rather than replication.

In the face of this inconsistency, I will use the terms in the following way. Formation of a whole, living organism from another organism I consider to be reproduction. Replication, on the other hand, is copying of a single kind of molecule that is either very simple or is a polymer composed of highly similar repeating units, where new units are added to a growing molecule in a way that matches the pattern of units present in a previously-existing molecule. A further difference is that replication is a single process performed on a single kind of molecule, while reproduction involves many different kinds of processes (including replication in most cases) that are coordinated so as to build something (a whole organism) composed of many different kinds molecules and structures.

With the above account of replication and reproduction in place, we can see that both may count as hereditary activities. Reproduction obviously is a process of heredity because it results in a new organism with traits similar to the organism(s) that produced it. Replication can also count as a hereditary activity when it is contributing to reproduction, as it does during Meiosis or in preparation for bacterial cell division. Recall from Section 2.3 that the activity of any biological mechanism must be a biological function, where a biological function is any output of a process that promotes the continued life of which the process is a part. Heredity counts as a biological function because it helps offspring to continue living by providing them with characteristics that helped the parents survive, and in every case of biological reproduction, at least some materials and processes that were part of the parent become part of the offspring, as captured by Szathmáry and Maynard Smith's (1997) notion of progeneration described above. Therefore, even though we distinguish parent and offspring as separate living individuals, offspring participate in the very same overall, multigenerational living process that has been sustained by all of their direct ancestors. Therefore, a process of biological heredity serves a biological function (as detailed in Section 2.4) because it promotes continuation of the life of which the process is a part. The only difference between heredity and most other biological functions (such as thermoregulation or nutrient acquisition that directly promote the continued living of a single organism) is that heredity promotes continuation of life by producing a new organism that is numerically distinct from its ancestors yet is still an offshoot of the very same life process in which its ancestors participated.

3.4 Mechanisms of Heredity

In this section I show how the mechanistic approach is able to characterize a mechanism of inheritance, and I show that when it is consistently applied, the mechanistic approach indicates that biological inheritance involves much more than just the action of genes. I also use the mechanistic approach to explain why developmental systems theory (DST), while capturing important facts about biological development, faces especially high hurdles in establishing a successful research program. Put briefly, the problem with DST is that it readily includes processes that are not biological mechanisms and such processes are less open to experimental manipulation than are mechanisms (Woodward, 1997, 2002b; Machamer, Darden and Craver, 2000; Darden, 2002). Consequently, DST is less able to supply satisfactory explanations, predictions and counterfactuals in support of a full research program.

Genetic Mechanisms of Inheritance

Although my primary goal for this chapter is to use the mechanistic approach to support arguments against gene-centrism, I first will use the mechanistic approach to explain how genetic heredity easily qualifies as a biological mechanism. While this insight into genetics may be somewhat unsurprising, it illustrates the effectiveness of the mechanistic approach in preparation for its application to more contentious cases. As discussed in Section 3.2, genetics, especially molecular genetics, has been extremely successful and has dominated the study of heredity for the last several decades, so the most obvious candidate for a biological mechanism of heredity comes from genetics and the action of nucleic acids. An important way for parents to cause their offspring to display parental traits is for parents to copy and then transmit some or all of their

genes (that is, the nucleic acids that constitute their genome) to their offspring. Therefore, any entities that perform those activities are good candidates for being mechanisms of heredity.

Replication of genetic material (typically DNA) is performed by a host of enzymes and cofactors that are precisely organized and coordinated with the polymer being replicated. The spatial positions of the enzymes and cofactors relative to each other cannot be changed significantly without sacrificing function, so the entities clearly meet the non-aggregativity criterion. Further, replication occurs intracellularly, so any entities directly involved in replication are parts of the organism whose genetic material is being replicated. Therefore, the entities involved in replication meet the mechanism criteria. Then, as made clear in Section 3.3, replication in preparation for reproduction contributes to the continuation of the ongoing life process in which the organism and its eventual offspring all participate. Therefore, replication in preparation for reproduction has an activity and entities that show it to be a biological mechanism of heredity. As a biological mechanism, subtle manipulation (Woodward, 1997, 2002b, 2003) of the enzymes and cofactors involved in replication have yielded deep insight into the process, and genetic mechanisms are regularly used by scientists to support useful explanations, predictions and counterfactuals. The same can be said about transmission of genes to offspring, whether that involves nothing more than coordinated cell division or it involves complex delivery and fusion of male gametes with female gametes.

The more contentious cases come when we consider the possibility of hereditary mechanisms aside from replication and transmission of genes. Genetics has been such a successful research program, any work on heredity have tended to treat the rest of the cell as mere background conditions for gene action (see Section 3.2). Gene-centrism has become so prevalent that, at least when speaking informally, the term ‘genetic’ has almost come to be

synonymous with biological heredity and is commonly used to distinguish biological heredity from environmental influences. Although defenders of gene-centrism acknowledge the ever-increasing complexity of the connection between genotype and phenotype and the importance of environmental factors, I argue, following Robert (2004) and Jablonka and Lamb (1989, 2005) that any kind of gene-centrism is unjustified, and that expanding our view of heredity to include non-genetic mechanisms will improve our biological explanations, predictions and counterfactuals.

An initial problem with any kind of gene-centrism is that, by themselves, genes can do very little. Simply transmitting them to offspring does not cause offspring to have traits similar to parental traits. In order for heredity to work, genes must be included within at least one fully functioning cell.[34] A fully functional cell is a complex network of interacting parts—ribosomes, mitochondria, plasma membranes, cell walls, cytoplasmic contents, as well as genes—that all work together to support the life of which they are a part. It certainly is true that a cell depends on proper functioning of its genes, but it also is true that genes only exist because of the proper functioning of the cells of which they are a part. Since the various different parts of a living cell are so interdependent, picking out any one component as causally prior is arbitrary. This means that the standard distinction between replicators and interactors, so popular in the literature on evolution by natural selection (e.g. Williams, 1966; Dawkins, 1976, 1980; Hull, 1980; Sterelny, 2001) is seriously mistaken. While it is important to recognize the unique role played by genes, and singling them out as the most important part of an organism may occasionally be a good research heuristic, it is important to remember that, in reality, no one part of a living organism is causally prior over all other parts. Dawkins' idea of the "selfish gene" (1976, 1980) famously

[34] This point is made clear by viruses that have a full genome but, since they do not have any metabolism of their own, are unable to make reproduce themselves. Only when a virus makes its way into a proper host cell and hijacks the host's internal machinery are viral genes able to contribute to viral reproduction.

gave rise to the suggestion that organisms are just genes' way of making more genes. While that may be a catchy aphorism, it is no more true than the belief that genes are nothing more than an organism's way of making another organism. Neither viewpoint is accurate because in a complex, dynamic, interconnected system like a living cell all parts are working for the mutual benefit of all other parts.

In order to establish the case against gene-centrism more fully, it is necessary to show how the arguments of the most influential advocates, Williams (1966) and Dawkins (1978, 1980, 1984), fail. It is important to do so because they, along with Hull (1980), are largely responsible for establishing the problematic replicator–interactor distinction into the literature. In arguing against gene-centrism I am following the work of Lloyd (1986), Mitchell (1987), Gould (2002), Jablonka and Lamb (2005), and Robert, (2004). Richard Dawkins (1978, 1980, 1984) begins his argument in favor of gene-centrism with the uncontroversial assumption that biological adaptations displayed by extant organisms have evolved by natural selection over many generations. Dawkins then argues that in order for evolution to occur over many generations, there must be something that physically persists over those generations. Evolution is happening to *something* over all those generations, and it can't be happening to individual organisms because they regularly die off while the process of evolution continues. The only things that physically persist through the process of reproduction, Dawkins believes, are the genes, and so it turns out that biological evolution is really evolution of genes. By contrast, he explains, all phenotypic characteristics are absent at one time or another during the process of reproduction, and most are passed on imprecisely, as when an offspring inherits a certain limb structure from its parents, but the exact size and shape of the offspring's limbs is somewhat different from either parent's. The genes, however, are inherited much more precisely and reliably. Sexual

reproduction may combine parental chromosomes in unique ways, but the individual genes that an offspring actually inherits are exact copies of genes found in its parents. Therefore, Dawkins concludes, genes persist through the generations while other aspects of organisms do not, and so adaptive biological evolution is evolution of genes.

In order for his description of evolution to be applicable to life in general rather than only to known, earthly life, Dawkins (1978, 1980, 1984) actually prefers to talk about replicators as the subjects of evolution, though he believes that replicators in earthly organisms happen to be genes made of nucleic acids. A replicator, in Dawkins's terminology, is any entity of which exact copies are made, and an *active* replicator is able to have some causal influence on its own likelihood of being replicated. For example, active replicators inside a biological organism influence their own likelihood of being replicated by influencing the organism's phenotype in a way that impacts survival and successful reproduction, where reproduction includes passing replicators on to offspring.

In contrast to active replicators, passive replicators have no effect on their own ability to be replicated. That is, a passive replicator has no effect on phenotype and merely gets copied and passed on to offsping along with any active replicators.[35] Second, Dawkins distinguishes between germ-line replicators and somatic-line replicators.[36] Germ-line replicators are those that are contained in the cells used for reproduction and are available to be passed on to the next generation. Somatic-line replicators, what Dawkins calls "dead end" replicators, are those contained in all of the somatic cells. In other words, dead end replicators in somatic cells are

[35] In actual organisms, passive replicators are sometimes called 'junk DNA' because they appear to have no functionality. Since cells do not seem able to identify and eliminate junk DNA, it all gets copied and transmitted along with active genes.

[36] This is essentially the same distinction between germ cells and somatic cells made by August Weismann and discussed in Section 3.2, but now expressed with Dawkins's terminology.

responsible for producing phenotypic characteristics, and hence are responsible for survival and reproductive success of the organism of which they are a part. Meanwhile, germ line replicators are passed on to offspring, thus ensuring that the replicators persist into the next generation. To drive his distinctions home, Dawkins refers to organisms, with all of their phenotypic characteristics, as mere "vehicles" in which replicators reside and whose phenotypic charactersitcs are largely determined by its active replicators. This centrality of replicators (genes) in Dawkins's understanding of biology and evolution is the source of his now famous "selfish gene" view in which any evolutionary adaptation must be something that benefits the genes since gene are what do the actual evolving.

This brings us to one of the central tensions within Dawkins's account of evolution as the story of active germ-line replicators. Recall that to be active, a replicator must have some causal influence on its own chances of being copied. The problem is that in multicellular, sexually reproducing animals, the germ-line replicators are not in somatic cell, and so are not the actual genes responsible for producing phenotypic traits and therefore cannot have any causal influence on the chances of their own replication and transmission to offspring. Far from being active, germ-line replicators just sit there in the gonads waiting to be used in sexual reproduction while replicators in somatic cells (dead end replicators, as Dawkins calls them) actually work to produce the phenotype. Further, any DNA molecules passed on to offspring are very unlikely to be the very same DNA molecules that were present in any of the offspring's ancestors. Rather, any DNA that actually ends up in a germ cell is a copy of a copy (of a copy of a copy…) of the DNA that was in the cells of that organism's ancestors. Therefore, the replicators used by earthly organisms (nucleic acids) do not phsycially persist over multiple generations in the way Dawkins suggests. This is a problem because Dawkins believes that whatever actually evolves

over the generations must physically persists over those generations, and since genes (nucleic acids) do not physically persist over multiple generations, genes cannot undergo adaptive evolution after all.

Dawkins recognizes this problem, and his solution is to appeal to gene types, rather than gene tokens, as the subject of evolution. A gene type is a segment of DNA with a specific nucleotide sequence, and since germ cells and somatic cells contain DNA with identical (or very similar) nucleotide sequences, we can say that somatic cells and germ cells have the same genes types. Thus, a single gene type can simultaneously be an active replicator and a germ line replicator that both influences its own chances of being passed on to the next generation and actually gets passed on. Therefore, gene types (rather than gene tokens) persist over multiple generations and evolve by natural selection.

The main problem with Dawkins's shift to gene types is that an active replicator is supposed to have a *causal* influence over its own chances of replication, and types cannot cause anything; only tokens can. Types are abstract universals, and as such cannot have causal influence in the world over and above the causal influence of the various tokens of that type. Therefore, Dawkins's appeal to types continues to run afoul of his need to have active replicators be part of the causal process that results in their own replication.

Another serious problem for Dawkins's account is that the move to gene types that he makes for replicators can also be made for his so-called vehicles (Mitchell, 1987). Dawkins originally concluded that genes are central to biology and evolution because they are the only physically continuous part of any lineage. But then, since no actual DNA molecules are physically preserved across generations, Dawkins appealed to gene types as what gets preserved. However, once this has been done to preserve the continuity of gene lineages, the same move can

be made to shore up the continuity of the rest of the organism, which Dawkins had dismissed as mere vehicle for the genes. That is, we can say that even though vehicle tokens are not continuous across generations, vehicle types certainly are.[37] This is especially true for unicellular organisms that reproduce via simple cell division with little or no development, an area of the biological world that Dawkins ignores almost entirely. Thus, there is nothing unique about genes in terms of persistence across generations, and so Dawkins-style gene-centrism is not justified. There is no question that gene-centrism has been an extremely useful research hueristic, but it does not capture the whole story of how heredity is achieved. That is, while genetic mechanisms of inheritance have supported very fruitful research programs, there may very well be other mechanisms of inheritance that, once recognized, can also be the subjects of fruitful research and can support more complete explanations of biological inheritance and and evolution.

Epigenetic Inheritance

Expression of genes (where expression of a gene amounts to its transcription into RNA that is translated into protein) can be controlled in a number of different ways. One common way that cells regulate expression of their genes involves sections of DNA that are not themselves transcribed into RNA, but instead influence the rate at which other, nearby sections of DNA are transcribed. Such "regulatory genes" can influence whether nearby genes are transcribed and at what rate, and so they have a dramatic effect on phenotype. Regulatory genes

[37] The point here is not to claim that organisms belong to any strict notion of types, as that may lead to a commitment to old typological thinking that viewed species as unchangeable (see Section 3.2). Rather, the point is that Dawkins' notion of types, whatever it is, can be applied to whole organisms as well as to genes, and so his arguments do not support his preferred conclusion.

and their position in the genome are inherited via genetic replication and tranferral mechanisms described above, so regulatory genes are inherited via familiar genetic inheritance mechanisms.

In addition to control by regulatory genes, however, gene expression can be influenced by so-called epigenetic control systems. The term 'epigenetic' once refered to a theory of biological development that was set in opposition to preformationism. Recall from Section 3.2 that preformationists, still common until the late 18th century, believed that development involved mere enlargement of tissues already present in a newly formed embryo. The opposing view at the time, dubbed 'epigenesis' (Mayr, 1982 p.106), was that an embryo did not contain pre-made tissues, but instead an offspring's own tissues and organs somehow developed out of an amorphous egg. As described in Section 3.2, preformationism faced a serious setback when microscopy revealed that the preformed tissues resembling the adult form were not present in a newly formed embryo. However, epigenesis was not in a better position to suggest any specific mechanisms by which heredity is assured, so epigenesis was not universally accepted either. In Mayr's (1982) description of the history, both sides were partially vindicated. Preformationism was correct insofar as offspring inherit a preformed genome (what Mayr calls a "genetic program"), but epigenesis was also correct in that embryos, relying on their inherited genome, must construct their tissues and organs anew. However, as knowledge of cell structures progressed, and especially after establishment of the Modern Synthesis in the early 20th century, 'preformationism' and 'epigenesis' both dropped out of the lexicon of biology in favor of the language of genetics.

In the latter part of the 20th century, however, 'epigenesis' was revived and used to refer to a new variety of genetic control systems. An epigenetic control system influences gene expression without changing the nucleotide sequence or genomic location of any somatic or

regulatory genes. One of the best known epigenetic control mechanisms is DNA methylation. When a stretch of DNA on a chromosome is methylated, it becomes less available to transcription mechanisms and so any genes in that location have a dramatically decreased chance of being transcribed. Biologists became aware of DNA methylation from its importance in embryonic development. When somatic cells differentiate into specialized tissues, they use DNA methylation to effectively disable genes that are are not relevant to their specialzed functions. For example, liver cells methylate stretches of DNA that contain genes relevevant to other kinds of specialized cells such as skin cells or neurons.

Proper methylation of DNA is important for development of phenotypic traits and appears to be accomplished by biolgical mechanisms, and so an important question is whether DNA methylation also acts as a mechanism of inheritance. That is, do parents use DNA methylation as a way to make their offspring look like them? There is some evidence that this occurs. For example, DNA methylation is responsible for allowing cells to differentiate between paternal and maternal chromosomes (Jablonka and Lamb, 2005), in some cases allowing only one to be used. In addition, at least some inheritable changes in plant lineages are attributed to DNA methylation. Jablonka and Lamb (1989, 2005) call inheritance of DNA methylation patterns a "non-genetic" mechanism of inheritance, but I think their designation is somewhat misleading because modern epigenetics is still concerned with genes and their regulation.

Cell Walls and Membranes

A better example of truly non-genetic inheritance comes from the inheritance of cell walls (Landman, 1991), which are the relatively rigid structures surrounding the cells of most non-animal organisms. Production of new cell walls in plants, algae, fungi, or bacteria is vital

for viability of offspring and is an important part of causing offspring to resemble their parents. Therefore, when new germ cells are built for use in reproduction, or when unicellular organisms with cell walls divide, contruction of new cell walls for the offspring is a hereditary activity. The entities involved in building new cell walls (mostly cytoskeletal components such as microtubules along with polymerized sugar molecules) must be precisely coordinated in time and space in order to build a new cell wall that matches the old cell wall (Salisbury and Ross, 1992; Campbell, 2004). This means that the entities cannot be rearranged or jumbled without loss of function, and so they achieve their function non-aggregatively. Since the entire process is accomplished within the cell itself, the entities certainly are part of the organism, and so the cell wall-building entities and their activity constitute a biological mechanism. And when the building of new cell walls is done as part of producing offspring, it is a mechanism of heredity.

At this point, a defender of gene-centrism might object that cell walls end up getting built only as a consequence of gene actions. That is, gene products such as cytoskeletal proteins and enzymes are responsible for building cell walls and so it seems justified to say that cell walls are built by genes. Therefore, genetic inheritance is still the mechanism of inheritance that really matters because all other feature follow from inheritance of genes. This gene-centric criticism fails because it is well known that genes are only one factor in determining the characteristics of cellular features such as cell walls. The way in which the wall elements assemble is not determined by genes, but rather by existing wall elements. Similar to crystal formation, new pieces of a cell wall are arranged according to the pre-existing pattern of cell wall pieces rather than to a pattern somehow encoded in the genes. In other words, while genes help produce some of the building blocks, they do not determine how the building blocks will be arranged in the cell wall. This means that non-lethal changes to the pre-existing cell wall pattern can cause any new

cell wall to be built to match the new pattern even though no genetic changes have occurred (Landman, 1991).

Finally, since building of new cell walls in preparation for reproduction is a mechanism of inheritance, there is the potential to preserve cell wall changes over multiple generations, which makes this mechanism evolutionarily important. If adaptively advantageous changes to cell walls can be inherited, then a lineage's cell wall features can evolve even when the lineages genome remains unchanged (Landman, 1991; Jablonka and Lamb, 1989, 2004).

A very similar acount can be given for inheritance of plasma membranes, which are the fatty layers that separate internal cell contents from the environment. Every offspring must be supplied with cell membranes resembling parental membranes in order for offspring to display parental traits. Indeed, any new membranes are always built from existing membranes in existing cells. New phospholipids[38] line up with those already in the membrane, and new glycoproteins and other membrane structures are delivered and inserted into the growing membrane, which increases the size of the cell to which this is happening. When the cell reaches a certain size, a finely choreographed sequence of protein actions constrict the cell near the center and split it into two cells in a process known as cytokinesis. The processes involved in membrane building and cytokinesis are relevantly similar to those described above for cell wall production, and also are biological mechanisms. Therefore, when new plasma membranes are made in construction of offsrping, any mechanisms involved are acting as mechanisms of heredity.

[38] Cell membranes are composed primarily of phospholipids, with a wide variety of channels, pumps, receptors, and other glycoprotein structures embedded throughout.

Cytoplasmic Inheritance

Another interesting example of non-genetic inheritance invovles transfer of cytoplasmic contents from parent to offspring. In their discussion of very early life forms, Segre, et. al. (1998), Segre and Lancet (2000), and Segre, Ben-Ali and Lancet (2000), call this transfer of cell contents from parent to offspring "compositional inheritance." As they explain, metabolic cycles, including any necessary ions, dissolved gasses, reaction substrates and catalysts are contained within a cell's cytoplasm. Cytoplasmic contents are maintained within strict concentrations limits by the collective action of various molecualr channels and pumps embedded in cell membranes. While the amino acid sequences of the proteins that make up channels and pumps are coded for by genes, the relative concetrations of cytoplasmic solutes are not encoded in any genes. When it comes time for a cell to reproduce cytokinesis will split the cell roughly in half, thus providing each daughter cell with a full compliment of ions and molecules needed to continue all vital metabolic processes. Metabolism has long been considered a necessary feature of any living system (e.g. Sagan, 1970; Luisi, 1998; Boden, 1999; Cleland and Chyba, 2002), so inheritance of metabolic cycles occuring in the cytoplasm is absolutely necessary for successful reproduction. Therefore, every offspring must receive at least one cell with cytoplasm very similar to parental cells in order for the offspring to resemble the parent (see also Landman, 1991; Jablonka and Lamb, 2005).

Segre, et. al. (1998), Segre and Lancet (2000), and Segre, Ben-Ali and Lancet (2000) refer to compositional inheritance as a mechanism, and consider it as a rival to polymer-based mechanisms of inheritance. They even go so far as to call the makeup of a cell's cytoplasm its "compositional genome" and talk about the "compositional genome replication mechanism" (Segre and Lancet, 2000 p. 218). The question, then, is whether this process of compositional

inheritance is a biological mechanism or merely a mechanism in the more permissive sense discussed in Section 2.1, according to which any identifiable causal process is a mechanism.

As discussed above, transfer of cytoplasmic contents from parent to offspring is a hereditary activity, so its status as a mechanism comes down to the characteristics of the entities that carry out this activity. Cytoplasmic contents of parental cells are, of course, part of the parent, so cytoplasmic inheritence meets the criterion that the mechanism be a part of the life it is supporting. However, transfer of cytoplasmic contents does not appear to meet the non-aggregativity requirement detailed in Chapter Two, section 2.2. Since cytoplasm is liquid, solutes continually float around and mix without sacrificing functionality. Concentrations of the various solutes are kept within fairly strict ranges, but a wide variety of radical rearrangements of cytoplasmic contents will have little effect on the success of cytoplasmic inheritance, so any cytoplasmic functions are achieved aggregatively (Wimsatt, 1974, 1997; Craver, 2001). Therefore, cytoplasmic inheritence is not, in itself, a biological mechanism of inheritance. Perhaps cytoplasmic inheritance qualifies as a physical mechanism in the much more permissive sense in which a mechanism is just a describable physical process, but it is not a biological mechanism in the restricted sense associated with the mechanistic approach.

This leaves cytoplasmic inheritance in a somewhat unclear position. It is true that offspring must receive cyplasmic contents from their parents in order for offspring to inherit parental traits, but the process does not seem to be a mechanism. Instead, I believe that cytoplasmic inheritance is best viewed as a side effect of other mechanisms. Mechanisms in the plasma membrane such as ion channels and pumps work to maintain proper concentrations of various solutes. Cytokenisis mechanisms pinch the cell in two, and since the cytoplasm is roughly homogenous, each resulting new cell receives roughly equivalent cytoplasmic contents.

So it seems that cyoplasmic inheritance, while extremely important, is nonetheless achieved as a sort of side-effect of other mechanisms that are geared toward other, non-hereditary activities. This does not mean that research into the effects of cytoplasmic contents on heredity are uninteresting or unfruitful. Rather, it means that study of cyoplasm in and of itself, and its roles in heredity are studied from the view of chemistry more than from the mechanistic approach in biology.

Developmental Systems Theory

Rejection of the gene-centric view in favor of a more inclusive account of heredity and development has led some authors to embrace Developmental Systems Theory (DST). Supporters of DST argue that development of an organism involves a complex web of interactions between a wide variety of what they call "developmental resources" (Gray, 1992; Griffiths and Gray, 1994; Oyama, 2000). Developmental resources include anything that is needed for proper development of offspring. Genes, of course, are developmental resources, but so are epigenetic control systems, cytoplasmic contents, organelles, cell walls, and cell membranes. More boldly, DST even includes factors external to the organism, such as local nutrients, energy sources, and perhaps even ambient temperature and pressure. All of these things affect development such that changing any one of them could change an offspring's phenotypic traits, and so all of them are developmental resources according to DST.

Supporters of DST rightly criticize gene-centrism for picking just one part (the genes) of the complex web of interacting developmental resources as causally primary. However, I believe that DST suffers from a fate similar to that of the holists of the early 20th century. Recall from Chapter One, Section 1.2, that holism was a naturalistic (rather than vitalistic) theory of

living systems that focused on whole organisms in addition to the individual parts of organisms. Holists tried to resist the trend toward reductionism, but they found it difficult to develop holistic research programs that could rival the success of the reductionist research programs (Allen, 2005). Reductionists found it easier to develop clear and useful explanations, predictions, and counterfactuals by isolating and analyzing individual parts of organisms, while holists faced the seemingly intractable problem of accounting for all parts and processes of an organism all at once. Similarly, DST, in its attempt to account for all factors relevant to development all at once, faces difficulty formulating clear and useful explanations, predictions and counterfactuals. This is not to say that DST is wrong or misguided, but rather that it faces significant hurdles as a research program.

The mechanistic approach helps in diagnosing the challenges faced by DST. One problem is that DST seeks to include all mechanisms of heredity and development into one theory. This is a laudable goal, but the mechanisms involved are very diverse and are realized in different ways in different species, and so a single, over-arching theory may not be possible. As a further difficulty, DST includes processes that are not biological mechanisms, and such processes are less open to the style of mechanism-based research that has been so successful in biology over the last several decades (Woodward, 1997, 2002b; Machamer, Darden and Craver, 2000; Darden, 2002). For example, DST includes cytoplasmic contents as developmental resources, and while that seems correct, cytoplasmic inheritance is not achieved via a biological mechanism. Further, and even more problematically, some of the developmental resources DST includes exist out in the environment. Recall from Chapter Two that in order for a process to be a biological mechanism, the entities that realize the proposed mechanism must be part of an organism. External nutrients and energy sources are not part of the developing organism, and

neither are general environmental conditions such as temperature and pressure. Therefore, external developmental resources are not part of any biological mechanisms and so are not open to the kind of investigations available under the mechanistic approach (Woodward, 1997, 2002b; Darden, 2002). Again, this does not mean that DST advocates are incorrect in their belief that external resources are crucial for biological development to produce phenotypic traits similar to parental traits. Rather, it means that DST faces special research challenges that can hamper its success relative to programs that focus on individual biological mechanisms.

Summary

In this Chapter I showed how and why the gene-centric view was established, and argued that it should now be abandoned in favor of a less restrictive account of biological heredity that includes non-genetic hereditary processes, some of which are biological mechanisms. I also showed that Developmental Systems Theory, while capturing important truths about biological development, includes a variety of processes that are not mechanisms and so faces special difficulties in providing explanations, predictions and counterfactuals expected of a strong research program.

In Chapter Four I shift to generation of new variations, which is the second of the three processes needed for ongoing, adaptive evolution. I use the mechanistic approach to address how variations are generated and how they contribute to biological evolution.

CHAPTER FOUR

GENERATION OF NEW VARIATIONS

4.1 Introduction

Before the rise of modern biology, common wisdom was that all species represent unchanging types of organisms, and any variation was mere deviation from an established standard. An evolutionary view, by contrast, requires a source of new variations that can accumulate indefinitely and eventually produce new species. Once evolution of species finally began to replace typological thinking, debate shifted to how variations generated, and which sorts of variations are inheritable and which are not.

In this chapter I use the mechanistic approach defended in Chapters One and Two to address current debates over generation of biological variation and its role in evolution. My focus will be on two processes commonly associated with Lamarckian evolution: acquired characteristics and adaptively directed changes. Lamarck proposed that individual organisms could acquire new characteristics that were adaptive improvements and that those newly acquired characteristics could be inherited by offspring. Although Lamarck was not the first to suggest such possibilities, he was the first to work them into a full theory of biological evolution, and so they have come to be most associated with his name. By the middle of the 20th century, directed generation of variation and inheritance of acquired characteristics had lost their place in evolutionary theory, but I argue in favor of their limited reintroduction.

Section 4.2 is an overview of the history of biological theorizing on generation of new variations. I show how the relevant concepts and terms have changed over time, and explain why Lamarck's ideas fell out of favor. In Section 4.3 I work to resolve some of the confusion left over from the historical account presented in Section 4.2, especially when 19th century

Lamarckian ideas are re-considered in light of contemporary cellular and molecular biology. In Section 4.3 I present a variety process for producing inheritable variations, and I assess 1) whether they are biological mechanisms, 2) whether they generate acquired or nonacquired characteristics, and 3) whether the changes are blind or directed toward increased adaptedness. I show that many of the processes that generate changes are not biological mechanisms, but among those that are mechanisms, some are directed toward increased adaptedness and allow for inheritance of acquired characteristics.

4.2 Historical Development of Concepts and Terms

As already mentioned several times, generation of new variations is only evolutionarily relevant if the change can be inherited by future offspring. As a result, much of the history presented in Section 3.2 also will be relevant here. For example, as was noted in Chapter Three, a dominant assumption in biology up into the 19th century was that species are natural *types* of organisms, and individual members of a species never stray too far from that type. Based on that assumption, generation of new variations is of relatively little concern because they will not be preserved over multiple generations. The occasional birth defect or monstrosity may be observed, but such deviations do not persist in any natural species.

Such typological thinking was tempered by plant and animal breeders who recognized that novel traits could emerge in a population and be passed on to offspring, and that selective breeding could change a species quite radically. However, typologists were able to dismiss the plant and animal breeders' track record by pointing out that domesticated animals, when returned to the wild, tended to revert to their wild forms and to lose any special traits that had been promoted through artificial breeding programs. Thus, the effects of artificial breeding programs,

though sometimes quite dramatic, seemed superficial and readily reversible, and there was little reason to think that the very large, permanent changes needed for general evolution of one species into a new species was at all possible (Mayr, 1982).

By the early 19th century, enough evidence from the fossil record had been gathered to indicate very strongly that some sort of long-term change in species on the planet had occurred, whether through evolution of species or for some other reason. Of course, divine intervention could account for the fossil record (e.g. Paley, 1802), but I am considering naturalistic options only. One naturalistic option for typologists who wanted to avoid any kind of species evolution was to say that old species regularly went extinct and were somehow replaced by new species, perhaps via immigration from other areas (e.g. Cuvier, 1825; see also Mayr, 1982; Hull, 1984). Not all typologists were opposed to evolution of species, but those inclined toward evolutionism faced a problem because the most popular form of typology assumed that an organism's set of traits form a coherent whole, and that changing just one or a few traits would disrupt the essence of the organism to the point that it would not be viable (Mayr, 1982; Dawkins, 1983). Thus, gradual, step-by-step evolution was considered impossible because small changes would result in non-viable offspring. A prominent way of reconciling evolution with typology was to adopt some version of saltationism in which very rare but very radical changes could appear suddenly in a lineage—even over a single generation—thus converting the species into a different viable type. Although no plausible mechanisms were suggested that could perform the activity proposed by saltationsts, it seemed consistent with the fossil record, which was relatively discontinuous at the time, and so saltationism (or something like it) remained entrenched among evolutionists in the mid-19th century (Mayr, 1982; Hull, 1984). The term 'mutation,' introduced

into the biological lexicon by Hugo De Vries in the early 20th century, originally referred to the whole-organism changes proposed by saltationists.

An obvious deviation from 19th century typological thinking was J.B. Lamarck (1809) who presented one of the first full theories of biological evolution. Lamarck's theory was discussed briefly in Chapter Three, Section 3.2, but there the focus was on Lamarck's account of inheritance. Here, I will focus on his account of how new variations are generated. Put briefly, Lamarck believed that organisms contain certain special fluids in their bodies that have a strong tendency to expand and develop the tissues in which they flow. He suggested that this general tendency to expand and develop is the ultimate driver of new characteristics for all organisms. Then, since Lamarck also believed that whatever new complexity an individual achieves during its lifetime can be passed on to its offspring (via inheritance of acquired characteristics), lineages tend to become more complex over time.

Organisms are not merely complex, but also very well adapted to their particular environmental circumstances. Lamarck explained adaptedness by proposing that the rate at which a tissue grows and the manner of its expansion are affected by the behavior of the organism. That is, increased use of a particular tissue will promote growth of that tissue, while decreased use will retard or even reverse growth of that tissue. This is Lamarck's famous "use and disuse" hypothesis about the source of most new variations. He believed that use and disuse can lead to greater adaptedness because all creatures tend to form particular survival habits in their environments, and such habits inevitably require use of some tissues more than others. For example, an organism may come to rely on a certain way of obtaining food makes use of certain parts of its anatomy more than others. Since, according to Lamarck, any repeated use of a tissue promotes that tissue's growth and development, formation of a habit will result in growth and

development of any tissues associated with that habit. In this way, an individual organism gradually becomes better equipped to accomplish whatever goals its habits were intended to accomplish. In other words, the organism will become increasingly better adapted to its environment. Then, since Lamarck allowed for inheritance of acquired characteristics, the trend toward increased adaptedness could continue indefinitely. Finally, since each species faces a different set of challenges, different species will develop different habits, eventually leading to wide diversity between species. Thus, Lamarck was able to account for both adaptedness and diversity of species.

Lamarck's (1802) theory of biology was far more complete and detailed than any rival naturalistic theories of the time. However, Lamarck's completeness came at a cost because his theorizing often went far beyond what the available evidence could support. Prominent biologists of the time were trying to establish their field as a serious science in which theorizing is tightly constrained by evidence, so Lamarck's penchant for speculation earned him an unfavorable reputation, leading most scientists to distance themselves from Lamarckian theory (Burkhardt, 1984; Hull, 1984). It didn't help that Lamarck's inclusion of an innate tendency for living tissues to expand and increase complexity can seem like a supernatural, teleological force of the sort that was being expunged from biology. Thus, even though many biologists were sympathetic to some of Lamarck's proposals, they actively distanced themselves from his overall theory lest they be perceived as sympathetic to his speculative methods (Stocking, 1962; Mayr, 1982; Burkhardt, 1984; Hull, 1984).

Charles Darwin was much more careful about theorizing beyond the available evidence than was Lamarck, even though some of his views on generation of new variations were similar to Lamarck's. For example, as described in Chapter Three, Darwin (1868) hypothesized that

inheritance was achieved by way of what he called "gemmules" that allowed for inheritance of acquired characteristics (Winther, 2000). Gemmules were supposed to be small particles generated in every tissue of the body and that carried information about the characteristics of the tissues in which they were generated. Over an organism's lifetime, gemmules from all of its tissues would collect in its gonads where they could then be used in production of offspring. Importantly, since new gemmules were being produced all the time, any new characteristics that an individual acquired during its lifetime could give rise to gemmules carrying information about those newly acquired characteristics. If the new gemmules got passed on to offspring, then the offspring would inherit characteristics the parent had acquired. Hence, Darwin advocated a kind of inheritance of acquired characteristics that is typically associated with Lamarck. At the same time, however, Darwin also proposed that new variations could be generated within the gonads themselves, especially during times of environmental change (Darwin, 1859, 1868). It is important to note that any changes to hereditary material happening in the gonads would not produce any phenotypic changes to the possessor of the gonads, and so would not result in an acquired characteristics. If the change is to produce any phenotypic effects, they will only show up in future offspring. Another difference between changes happening in the gonads versus those happening to somatic cells is that changes occurring in the gonads were unlikely to be directed toward increased adaptedness because of their relative isolation from the environment. While somatic cells could adapt to environmental challenges via direct interaction with the environment, gonadal cells cannot. In his earlier versions of the *Origin of Species*, Darwin preferred to emphasize the gonadal source of variation that produced so-called blind variations—those not correlated with adaptedness in any way—but felt forced to shift emphasis in later editions. Estimates of the age of the earth were far too small for Darwin's preferred processes to

produce the wide diversity of organisms, and so Darwin, in later editions of Origin of Species, downplayed blind variation and emphasized directed generation of variation and inheritance of acquired characteristics, despite their association with Lamarck.

In the latter 19th century, however, evidence began to accumulate against the possibility of inheritance of acquired characteristics (IAC) and directed generation of new variations. As described in Section 3.2, August Weismann's distinction between germ cells and somatic cells seemed to count against such possibilities, but so too did work by people like Francis Galton. Galton transferred blood from rabbits of one color into rabbits of a different color, with no effect on the color of the latter rabbits' offspring. If acquired characteristics could be inherited, and if hereditary information about an acquired characteristic was transferred to germ cells via the bloodstream (as with Darwin's gemmules), then transferring blood from one kind of rabbit to another should cause the latter rabbit to produce offspring similar to the rabbit from which the blood was taken. Since offspring continued to resemble parents regardless of whether the parents had received blood from a differently-colored rabbit, Galton concluded that the bloodstream did not carry hereditary information, thus depriving supporters of IAC the main route by which hereditary information about acquired characteristics was thought to be delivered to germ cells. Weismann himself conducted somewhat more blunt experiments in which he cut the tails off of a large number of rats to see if the newly acquired "no-tail" characteristic could be passed on to offspring. Despite the crudeness of such experiments, August Weismann used the results to argue in favor of what came to be called "neo-Darwinism," which combined Darwin's natural selection and common decent with Weismann's assumptions that variation is blind and acquired characteristics are not inherited (Mayr, 1982; Jablonka and Lamb, 2005; see also Section 3.2).

Since little was known about the entities involved in inheritance, there was not enough evidence on any side to drive consensus. Plus, typology still infected the thinking of most biologists, resulting in problematic ways of framing their debates over production and inheritance of new variations. For example, as described by Mayr (1980, 1982), many biologists of the time distinguished between so-called continuous and discontinuous variations. Continuous variations were the normal, though somewhat minor, variations seen in any species. There always will be some small amount of variation in exact size, shape, etc., and such variability tends to be continuous in the sense that body size or limb length may fall anywhere within a range of values. However, such variability was assumed merely to be variation around a standard type, with species never drifting too far away from their standard. On the other hand, discontinuous variations were the occasional, dramatic differences observed in some individuals that did not conform to the standard type. That is, the new trait was truly different rather than just at the extreme end of a continuous range of possible variations on a standard. By and large, discontinuous variations were either not inherited at all, or existed only for a few generations, thus supporting typologists' view that species remain fixed. Even the rising number of evolutionists in the latter 19th century still held some amount of typological assumptions and so were inclined toward saltationism rather than toward gradual evolution (Mayr, 1982; see also section 3.2).

Another distinction with roots in typology was that between soft inheritance and hard inheritance. Soft inheritance, if it existed, allowed for changes to accumulate over multiple generations, while hard inheritance did not. Typologists predictably preferred hard inheritance since they believed species could not change in any significant way. On the other hand, those inclined toward gradual evolution preferred inheritance that was at least somewhat soft so as to

preserve small evolutionary changes over time. This distinction became central to the debate over neo-Lamarckism because directed variation coupled to inheritance of acquired characteristics (IAC) was the most prominent example of soft inheritance. Neo-Lamarckians of the late 19th and early 20th centuries acknowledged that Lamarck's original ideas were unlikely to be correct, but still maintained that at least *some* new, inheritable variations could arise as adaptive responses to environmental challenges. However, since it was assumed that any adaptively-directed changes would have to be acquired by a parent before they could be passed on to offspring,[39] debate over the existence of directed variation was collapsed into the debate over IAC. Evidence was mounting against IAC, so biologists—even those sympathetic to gradual evolution—were turning against neo-Lamarckism (Mayr, 1982; Hull, 1984; Jablonka and Lamb, 2005). When combined with the unorganized and even dangerous nature of some prominent neo-Lamarckian research programs,[40] the case for directed generation of new variations eventually was lost (Stocking, 1962; Hull, 1984; Burkhardt, 1984; Mayr, 1982).

Typology eventually was undermined by a mature field of genetics, but when Mendelian genetics was first re-discovered at the beginning of the 20th century, it seemed to support typology and count against long-term evolution by natural selection. The reliable transmission of traits and consistent patterns of heredity looked a lot like hard inheritance, thus supporting the typologists' views that either evolution did not happen or, if it did happen, it occurred in discontinuous jumps (saltations) as described by saltationists. The sustained popularity of saltationism indicates just how little was known about the entities that carry out heredity. Given our contemporary understanding of the mechanisms of inheritance and development, it is

[39] The alternative is that adaptively-directed generation of new variations occurred inside germ cells in the gonads, a possibility that was even more difficult to defend because of the relative isolation of gonads from the environment.

[40] The most famous example has come to be known as "the Lysenko affair" in which Soviet biologists used neo-Lamarckian assumptions to drive agricultural policies. The result was mass starvation. (Medvedev, 1969)

difficult to see how saltationism could work without divine intervention. Nevertheless, saltationism was compatible with the assumptions of Mendelian inheritance and seemed to explain glaring discontinuities in the existing fossil record, so it remained popular (Mayr, 1982).

Things finally began to change when early geneticists noticed some so-called discontinuous changes, such as a change in eye color, would occasionally appear and be inherited over multiple generations according to Mendelian principles. By incorporating new variations that could be preserved over multiple generations, the young, though highly respected field of genetics was inching toward compatibility with gradual evolution. Thomas Hunt Morgan, who first recorded this effect in his famous fruit fly breeding experiments, used the term 'mutation' to refer to these inheritable, single-trait changes even though they were far more modest than the whole-organism "mutations" proposed by Hugo de Vries and other saltationists of the previous decades. Nevertheless, Morgan and his followers still distinguished their newly discovered mutations from what they called the "wild type," where the wild type was the standard or normal version of the trait, and so some typological thinking still persisted.

By the 1920's, basic Mendelian genetics had given rise to population genetics. Population genetics had the conceptual and mathematical tools for describing and tracking changes in gene frequencies in a population over many generations. When combined with Morgan's new account of mutations as inheritable changes in single phenotypic traits, population genetics was eventually shown to be compatible with gradual evolution, resulting in the Modern Synthesis of the 1930's and '40's (see Section 3.2 for more discussion). Only when genetics was shown to be compatible with gradual evolution by natural selection did neo-Darwinian evolution seem like a better explanation for biological adaptedness and diversity than the alternatives (Mayr, 1982; Kitcher, 1985)

According to the Modern Synthesis, acquired characteristics could not be inherited so it was assumed that any adaptively-directed variations were not evolutionarily relevant. However, it was still an open possibility that adaptively-directed changes could occur within germ-line cells rather than being acquired by parents. Evidence concerning adaptively-directed generation of new, inheritable variations in germ cells is difficult to obtain because one would have to examine every change made to any germ cell's hereditary material in order to determine its potential impact on an offspring's adaptedness, which would be extremely expensive and time consuming (Jablonka and Lamb, 2005; Sarkar, 2005) especially before the rise of molecular biology. Consequently, such studies typically relied on a statistical approach (Sarkar, 2005). For example, Luria and Delbrük (1943) studied multiple generations of microbes grown under stressful conditions and concluded that new variations were not generated as adaptive responses to environmental challenges. Even though the work had some serious flaws that throws their results into doubt (Jablonka and Lamb, 2005; Sarkar, 2005), and even though the limited results could not be used to conclude that directed generation of adaptive changes never happens, most biologists were coming around to the neo-Darwinian view that inheritable variations were only generated blindly.

Work in molecular biology in the mid- to late 20th century changed the field dramatically, though the Modern Synthesis remained as the basis for evolutionary theory. One major change was that the term 'gene' came to be associated with transcribable section of DNA (see Section 3.2) and, consequently, the term 'mutation' came to be associated with an inheritable change in DNA nucleotide sequence. In addition, the rise of molecular biology provided yet stronger reasons to doubt the existence of directed generation of new, inheritable variations. When the roles of DNA and proteins became clear, it became *extremely* difficult to conceive of a natural

process that could introduce any adaptive changes into germ line DNA in such a way that the changes would lead to the development of the new, adaptive characteristics in offspring. Any such process would somehow need to take into account relevant changes in environmental conditions and then change the nucleotide sequence of relevant germ-line genes in just the right way so as to result in development of properly adaptive phenotypic traits. Even with the high level of understanding of contemporary genetic engineers, it still is not possible, in most cases, to predict the phenotypic result of a given change in nucleotide sequence, much less to know precisely what nucleotide sequence changes would result in properly adaptive phenotypic changes. As a result, the suggestion that organisms are able to engineer themselves in such precise ways smacks of divine intervention rather than natural causation. Therefore, by the late 20th century, the two key Lamarckian mechanisms—adaptively directed generation of new variations and inheritance of acquired characteristics—seemed all but impossible.

By the latter 20th century, the consensus view in evolutionary biology consisted of the Modern Synthesis updated with a much better understanding of molecular biology. As for generation of new variations, since most of the recent advances in evolutionary biology (and in biology generally) had been driven by the success of genetics—both population genetics and molecular genetics—the understanding of how new variations are generated continued to be heavily gene-centric (see Section 3.2). Chromosome shuffling and recombination (crossing over and karyogamy) during sexual reproduction were understood to be the sources of most differences between individuals of the same species. However, since chromosome shuffling and recombination only work with existing genes, they are limited in their potential to generate completely new traits. Plus, asexual organisms do not include such processes at all, yet they must have some way of generating new variations. Therefore, the ultimate source of all new

variations was assumed to be blind genetic mutations, now understood as inheritable changes in DNA nucleotide sequence due to replication errors or direct damage to nucleic acid molecules. Mutations most often are detrimental because randomly changing one component of any biological mechanism is likely to disrupt that mechanism's function (Wimsatt, 1974, 1997; Woodward, 1997, 2002b; Craver, 2001), and so most mutations decrease adaptedness. However, some mutations may have no phenotypic effect at all (so-called neutral mutations), while still others may actually increase adaptedness in subtle or dramatic ways. Thus, genetic mutations were assumed to be the ultimate source of new, inheritable variations on which differential natural selection could act.

By the end of the 20th century, some other significant sources of variation (besides basic mutations) had gained recognition, though all still concerned genes one way or another. For example, a kind of genetic recombination called transposition occurs when a section of DNA is moved from one area of a chromosome to another. Active genes commonly are associated with a regulatory region of DNA that controls the rate at which the genes are accessed by transcription machinery, and transposition can separate an active gene from its regulatory region, thus rendering it inactive. Alternately, transposition may move an inactive gene to a position that associates it with a regulatory region and allows the gene to become active. Either way, the result usually is a change in phenotype. If the section of DNA being transposed is a regulatory region rather than a transcribable gene, then the phenotypic effects can be quite dramatic because one set of genes may lose their regulatory region and cease being active, while another set of genes may gain a regulatory region and thereby become active contributors to phenotype. When transposition occurs between microbial cells, it is called horizontal gene transfer (HGT). In HGT, one cell copies and transfers a section of its genome to another cell that incorporates the

transferred segment into its own genome. If the transferred segment of DNA can be transcribed, then it can become an active contributor to the new cell's phenotype. As with mutations, any new variations generated as a result of transposition or HGT may be detrimental, neutral, or beneficial to an organism's adaptedness, and so transposition is generally considered a blind source of variation.

In recent decades, some authors have argued that not all sources of biological variation are blind and that acquired characteristics very often can be inherited (e.g. Landman, 1991; Jablonka and Lamb, 1989, 2005; Steele, et.al, 1998; Stolovicki, et.al., 2006). I expand upon those efforts in the rest of this chapter. In Section 4.3 I clarify the distinction between acquired versus non-acquired traits, and between blind versus directed changes. In Section 4.4 I show that many processes that generate changes are not mechanisms, and that among those that are mechanisms, some are not entirely blind. I also show that inheritance of acquired characteristics is commonplace, especially among bacteria.

4.3 Change-Producing Activities

As discussed in Sections 3.1 and 4.1, it is important to maintain the distinction between heredity and generation of new variations (Sarkar, 2005; Kronfeldner, 2007). When characterizing generation of new variations in isolation, it is not appropriate to incorporate a particular variety of inheritance, genetic or otherwise, just as it was inappropriate to include an account of change-production in the characterization of heredity defended in Section 3.3. Of course, only inheritable changes are evolutionarily relevant, so it quickly will become necessary to restrict the account only to inheritable changes.

When the requirement for inheritability is absent, characterizing the relevant change-producing activity is fairly uncomplicated. It is just any causal production of a trait that was not present in the recent lineage. This allows for changes to living organisms as well as changes to hereditary material that result in new traits appearing in offspring. Of course, this account includes a wide variety of changes that are evolutionarily uninteresting, such as superficial injuries, and so it is time to restrict the account to changes that are inheritable. So, the relevant change-producing activity is causal production of a trait that was not present in the recent lineage and that is subsequently inheritable.

By restricting the account to changes that were not present in the recent lineage, we leave out traits that appear as a result of normal developmental mechanisms inherited from parents. Embryological development, sexual maturation, and aging are all developmental processes that can produce very dramatic changes in an organism, but they are pre-arranged changes that play out according to regulatory mechanisms inherited from an organism's parents. Therefore, developmental changes that roughly mimic the developmental changes undergone by parents are not evolutionary changes because they are not changing the lineage. Note, however, that my account only excludes traits that appeared in the recent lineage. This allows for important phenotypic traits that can reappear in a lineage after being absent for many generations. For example, consider the blind cave fish *Amblyopsis spelaea* (Jeffery, 2001; Romero and Green, 2005) whose lineage has lacked eyes for thousands of years. If a population is moved out of the caves and into the light, functional eyes could reappear in the lineage as a result of a surprisingly simple change in developmental regulation. Reappearance of eyes and vision in a lineage certainly is evolutionarily significant, so those sorts of changes should be allowed by our account of change-inducing activity.

To reiterate, the activity we are interested in is production of new, inheritable variations that were not present in an organism's recent lineage. By itself, this account says nothing about whether the inheritable changes are acquired or are generated blindly As discussed in Sections 3.2 and 4.2, the most recent debates over Lamarckian evolution focused on the possibility of changes that are directed toward increased adaptedness and on the possibility of inheritance of acquired characteristics (IAC). Since genetic heredity seemed unlikely to support IAC, most biologists assumed that acquired characteristics could not be inherited and so were not evolutionarily relevant. In addition, they also assumed that any processes for generating new, inheritable variations are blind rather than directed toward greater adaptedness. Since both will be an important part of the discussion in Section 4.4, I will take some time to characterize directed generation of new variations and acquired characteristics more carefully.

First I consider acquired characteristics. In order for a new characteristic to be an acquired characteristic, the change must occur to an individual organism after the organism has been produced as a viable offspring. Acquired changes can be contrasted with non-acquired changes that occur during production of cells that will be used to produce an offspring. For example, replication errors inside germ cells can result in offspring with some new phenotypic trait, but nobody acquires that change. The parent does not change, so the parent has not acquired any new characteristics, and the offspring does not yet exist so it cannot acquire anything at all. Therefore, any new traits introduced during production of cells that will be used to produce an offspring do not qualify as *acquired* changes. We also want to exclude any changes that occur to an organism as part of its normal developmental processes because, as explained above, such changes are not new to the lineage and so are not evolutionarily relevant.

Historically, inheritance of acquired characteristics was not included in the Modern Synthesis because it seemed that genetic inheritance and the germ cell / somatic cell distinction did not allow for it. Thus, it may be tempting to assume that IAC, if it exists, must use a non-genetic mechanism of inheritance. However, despite IAC's historical conflict with genetic inheritance, we should not reject an obvious instance of IAC just because it happens to involve genetic inheritance. It seems clear that if an existing organism acquires a new phenotypic trait because of a change to its genotype, and if that change is genetically inheritable, then we have inheritance of an acquired characteristic.

This brings us to the other main tenet of neo-Lamarckism, adaptively directed generation of new variations. Before discussing directed variation itself, however, I will take a moment to separate directed variation from IAC. As explained in Section 4.2, since it was difficult to imagine how adaptively directed changes could be encoded directly into germ-line hereditary material, it was assumed that any directed changes would have to somehow involve somatic cells and their interaction with the environment. In other words, it was assumed that adaptively-directed changes would inevitably be acquired changes. Thus, the debate over directed generation of variation was bound up with the debate over IAC. However, other than the historical association of IAC with directed variation, it is possible for an acquired, inheritable characteristic to be generated blindly. It also is possible for there to be a process that generates adaptively-directed changes that are not acquired changes (Sarkar, 2005). Therefore, evaluation of IAC and of directed variation can and should be done separately.

Sarkar (2005) explores the concept of directed variation in the context of genetic mutations, but his discussion can be generalized to include inheritable changes of any sort rather than just genetic changes. Sarkar begins by characterizing a very strong notion of directed

generation of new variations—what he calls "instructed" changes—as a single change that is beneficial on the first try, sort of the way a human civil engineer would direct a specific change to a bridge design in order to make it stronger or more efficient. In no way is this a blind process because, as Sarkar explains, information concerning what changes would be beneficial (rather than neutral or detrimental) is playing the main role in directing which changes are generated. Sarkar (2005) then goes on to describe a more relaxed kind of directedness in which an adaptive change "occurs (or occurs more often) in an environment where it enhances the fitness of the organism than in an environment in which it does not do so" (2005: 290). While I believe he is on the right track, the problem I see with Sarkar's more relaxed account is that this sort of directedness can come too cheaply. When faced with some small environmental stress, a population could simply increase its overall rate of change in order to increase the chances of producing an adaptive change. The problem is that all kinds of changes—good, bad, and neutral—would become more likely, and would do so during a time when the beneficial ones would be beneficial, and so this would meet Sarkar's criteria for being directed even though it really just involved increasing the rate of blind changes. Increasing the rate of blind changes, while interesting, would not seem to be very directed in the relevant sense. Rather, in order to count as directed (while still falling short of being "instructed" in Sarkar's sense), the process would at least need to boost the chances of a beneficial change more than it boosts the chances of neutral or detrimental changes. I call such an activity "semi-directed" because it is directed only insofar it uses a blind change-producing process aimed at producing an adaptive change via trial and error. In Section 4.4, I argue that a semi-directed process for producing inheritable variations is at work in bacterial colonies.

Producing Changes As a Biological Function

The activity of a biological mechanism must be a biological function, so when assessing whether any change-producing processes are biological mechanisms, we must determine whether the activity counts as a biological function. The activity of a process is a biological function when the activity promotes continuation of the life in which the process is participating. Among the varieties of change-generating activities discussed above, directed variation is the best candidate because a process that reliably generates changes that increase adaptedness clearly is promoting survival. Semi-directed generation of variation is substantially riskier than fully directed variation because a semi-directed process would be expected to produce some detrimental changes as well as adaptive ones. However, since a semi-directed process increases the chances of an adaptive change more than it increases the chances of a detrimental change, a semi-directed variation still qualifies as a biological function, as long as the risk of occasional detrimental changes is outweighed by the risk of doing nothing. Thus, semi-directed variation is especially beneficial when faced with an environmental challenge that could end up being fatal if left unsolved.

Truly blind variation can also qualify as a biological function as long as it is tightly controlled. For example, sexual reproduction is beneficial mainly because of the way it blindly mixes hereditary contributions from both parents to produce viable offspring with a unique set of traits (see Section 4.4 further discussion). That is, sexual reproduction involves steps that rearrange and recombine the hereditary endowment from each parent in ways that ensure offspring receive all the genes and other cellular components they need, but that each offspring receives a slightly different combination of components, leading to a highly restricted degree of variation among the offspring. Although it is far more difficult and risky than asexual reproduction, sexual reproduction has evolved as a common strategy because its change-

generating features are thought to promote long term survival of the lineage by preventing buildup of detrimental genetic mutations (e.g. see Agrawal, 2001) and by providing the lineage with some flexibility to adapt to changing conditions. The change-generating parts of sexual reproduction are considered blind because the process does not distinguish between combinations of traits that will improve adaptedness from those that are neutral or detrimental. That is, the change-generating processes of sexual reproduction are just as likely to result in offspring that are less adapted than their parents as it is to result in offspring that are better adapted than their parents. Nevertheless, since it usually produces viable offspring, the risks of allowing this small amount of blind variation are outweighed by the long term survival benefits, and so highly controlled blind mutation during sexual reproduction qualifies as a biological function.

4.4 Varieties of Change-producing Processes

In this section I draw upon my account of change-producing activities in Section 4.3 to analyze and characterize the varieties of actual change-producing processes. My main goals are to show that the two processes traditionally associated with Lamarck—directed variation and inheritance of acquired characteristics—are both operative in the biological world and play significant roles in biological evolution, particularly among microbes. In doing so I also will determine which change-producing activities qualify as biological functions carried out by biological mechanisms. I will start with the more familiar variation that results from sexual reproduction and then move on to more contentious cases.

Meiosis and Sexual Reproduction

Most of the cells in sexually reproducing organisms have two full sets of chromosomes. Every chromosome is paired with another, very similar chromosome, and the two together are called a homologous pair. One member of each homologous pair came from the organism's mother while the other came from its father. Germ cells, also known as gametes, have only one set of chromosomes (one member of each homologous pair) so that when a gamete from one parent fuses with a gamete from the other parent, the result is a new cell with two sets of chromosomes, one set from each parent. Gametes are produced through a specialized kind of cell division called Meiosis. During Meiosis, homologous pairs of chromosomes line up along the center of the proto-gamete, and when the cell divides, homologous pairs are separated, with those that happened to line up on the left going to the cell on the left, and those that happened to line up on the right going to the cell on the right. The result is that an individual's gametes receive some chromosomes that were originally inherited from that individual's mother and some that were originally inherited from that individual's father (i.e. from the eventual offspring's grandparents).[41] Cells do not regulate which member of each homologous pair goes left and which one goes right because there apparently is no way to regulate which side each homologue lines up on before the proto-gamete divides. Thus, each gamete receives a randomly-determined combination of maternal and paternal chromosomes, potentially resulting in offspring with new phenotypic variations.

Meiosis is achieved by biological mechanisms whose function is to produce gametes with one full set of chromosomes each (see section 3.4 for more discussion of hereditary mechanisms). Any diversity resulting from this step is achieved passively (and blindly) because

[41] The actual chromosomes being used are not the very same ones physically handed down from the organism's parents. Rather, these are high-fidelity copies of those chromosomes.

there is no mechanism for determining which side each homologue lines up on. Meiosis ensures that each gamete receives a full set of chromosomes, but does not regulate how many of a resulting gamete's chromosomes are maternal and how may are paternal. It is entirely possible that a given gamete could end up with exclusively maternal or exclusively paternal chromosomes. Therefore, generation of new variations at this step is not a biological function, but rather a side-effect of mechanisms whose actual function is to produce cells with one full set of chromosomes.

Contrast the above process with crossing over, which also happens during Meiosis, but involves highly regulated exchange of genetic fragments between homologous chromosomes before they separate. Any phenotypic changes that show up in offspring as a result of crossing over could be adaptively beneficial, detrimental, or neutral, and so crossing over is also a blind process. However, crossing over requires precise coordination of a large number of enzymes, cofactors, and other cellular components, and does not appear to serve any other role than to increase genetic variety among gametes. That is, Meiosis without crossing over would still produce viable offspring, only with less diversity. As explained in Section 4.3, generation of variation is beneficial for long-term continuation of the lineage, and therefore a potential biological function, as long as it does not significantly threaten viability. Since crossing over produces exactly that kind of highly constrained variability, it qualifies as a biological function. Then, since the entities that manage crossing over are part of the organism and achieve their function non-aggregatively, they meet the criteria for biological mechanisms presented in Chapter Two. Therefore, crossing over is indeed a mechanism for producing new variations.

Genetic Mutation

Crossing over and gamete fusion can generate a fair amount of variety in a population, but their potential is limited because they only work with existing chromosomes. The ultimate source of entirely new genetic variations is genetic mutation. As explained in Section 4.2, a genetic mutation is an inheritable change in nucleotide sequence, and if the mutated section of DNA contributes to an individual's phenotype, then the mutation can change one or more phenotypic characteristics. Most mutations are a result of uncorrected replication errors, though some are caused by chemical or radiological damage. In the former case, mutations are malfunctions of replication mechanisms whose function is to copy DNA precisely, and in the latter, they are completely unregulated insults coming from the environment. Therefore, most mutations are blind and unregulated. Although an occasional beneficial mutation may show up, completely unregulated blind changes, such as externally-induced damage, usually are detrimental and all organisms go to great lengths to minimize the occurrence or impact.[42] Therefore, blind, unregulated genetic mutation cannot be a biological function under most circumstances, and consequently cannot be the activity of a biological mechanism. Nevertheless, replication and DNA repair mechanisms are not perfect, so a steady rate of mutation—called the background mutation rate—is to be expected.

Mutations introduce novel variety in a population that can help the population adapt to unpredictable future circumstances. Organisms that reproduce with perfect fidelity may not be able to adapt to changing conditions, so allowing some amount of mutation is good. This suggests that as replication mechanisms evolved, natural selection may have struck a balance between perfect replication fidelity and allowance for mutations (Sniegowski, et. al., 2000).

[42] Organisms have a wide variety of ways to minimize unregulated blind changes. Examples include outer coverings that block ultraviolet radiation, enzymes that neutralize toxins and free radicals, and mechanism that proofread and repair DNA.

However, even if the background mutation rate has evolved to its present value as a result of natural selection, replication errors or damage they are still blind and are achieved passively. Therefore, genetic mutation as a result of replication errors is, at best, a side effect of replication mechanisms rather than a diversity generating mechanism in itself.

However, if genetic mutation due to replication errors was regulated in the right way, rather than just being a passive side-effect of very high fidelity replication mechanisms, then genetic mutation could potentially qualify as a biological function. Since genetic mutations typically are detrimental, and since a biological function must promote continued survival, a mutational process that serves a biological function would need to somehow favor beneficial changes over detrimental ones, or only generate mutations when the dangers of doing so are outweighed by the potential benefits.

Sarkar's (2005) description of "instructed" mutations (see Section 4.3) would fit the bill very well, but there are no known biological processes that can precisely alter genes in an exclusively adaptive way, so I will set that possibility aside. A more plausible possibility is the semi-directed use of blind genetic mutation in a way that is more beneficial than harmful. Although they are not discussing biological functions in quite the same way I am, Jablonka and Lamb's (2005) description of four different "interpretive" mutational strategies will be helpful in understanding how blind mutation might become a biological function. Their first mutational strategy is "induced global mutation" in which the entire genome is subjected to increased mutation rates, but only during periods of time determined by the cell and usually in response to some environmental trigger. Jablonka and Lamb's induced global mutations are the same as Sarkar's (2005) personal notion of "directed" mutation discussed in Section 4.3. Jablonka and Lamb's (2005) second mutational strategy is "local hypermutation" in which a specific area of

the genome is particularly susceptible to mutations at all times. The third strategy is "induced local mutations" in which a specific area of the genome is caused to be particularly susceptible to mutations at regulated times. The fourth strategy is "induced regional mutation" in which a specific set of genes (which may or may not be near each other) are caused to be particularly susceptible to mutations at regulated times. All of these have potential to be biological activities, and the third and fourth strategies even have the potential to be semi-directed. Recall that a semi-directed process of generating variations is one that uses a blind mutational process in a way that promotes beneficial mutations more than it promotes detrimental mutations.

An interesting candidate for a semi-directed process occurs in bacterial colonies that have been subjected to an environmental stress harmful enough to cause a stress reaction in the colony but not so terrible as to kill the constituent cells too quickly. For example, a common experimental design involves growing cells whose ability to metabolize lactose has been disrupted on a medium with lactose as the only carbon source (e.g. Cairns et al. 1988; Cairns and Foster 1991; Oda et al. 2001). Under such starvation conditions, some of the cells will enter a state of stress called an SOS response. This triggers (among other things) replication of the cell's DNA using an error-prone method, thereby introducing random mutations (Torkelson et al. 1997; Foster 2000, 2004; Godoy et al. 2000; Bjedov et al. 2003; Kivisaar 2003; Aertson and Michiels 2005). This cellular state is referred to as a transient mutator state or hypermutational state, and it results in duplication of many genes, most of which will have some mutations. If the hypermutational process hits upon a mutation that allows a cell to solve the environmental problem—in this case, the ability to harvest and metabolize lactose—then the cell turns off its SOS response and acquires enough energy to reproduce and to activate its conjugation (gene

sharing) mechanisms. In this way it can confer its newfound adaptive advantage to its offspring and to its neighbors.

Even though entering a hypermutational state increases the risk of detrimental mutations, it is done only during times when doing nothing would be lethal. Since it at least increases the chances of finding a solution, inducing a hypermutational state as part of the SOS response can promote the continued survival of the bacterial colony. Therefore, inducing a hypermutational state during a time of stress qualifies as a biological function. To be clear, the suggestion at hand is that an ability to enter a hypermutational state during times of stress is itself a response that has evolved, via evolution by natural selection, because it is advantageous to enter such a state during times of stress, and it is beneficial because it increases the chances of hitting on an adaptive response to the environmental stress. This interpretation of the hypermutable state is criticized by Brisson (2003) who argues instead that the hypermutable state is not an adaptive response, but rather the result of malfunctioning replication systems during a time of low nutrient supply (see also Lenski and Mittler 1993; Sniegowski and Lenski 1995; and Torkelson et al. 1997). In other words, the appearance of occasional beneficial mutations is just a lucky side effect of a malfunctioning system and is not a biological function that has evolved as a response to stressful conditions.

Several authors, such as Foster (2000, 2004), Kivisaar (2003), and Aertson and Michiels (2005) argue against Brisson's view on empirical grounds. Before discussing their arguments, however, I will counter Brisson's view on theoretical grounds. Brisson (2003) argues that the hypermutable state that cells enter during times of stress is not an adaptation but is rather a side effect of a malfunctioning system. He writes, "The breakdown of cellular mechanisms resulting in a fortuitous adaptive mutation does not indicate that the system has evolved to fail under

stressful conditions. Because of myriad mutation generating mechanisms and the differential response at various loci, it is unlikely that they evolved for the purpose of producing adaptive mutations" (Brisson 2003: 33). The problem with this argument is that it ignores the very real problem of distinguishing between a natural function and a mere side effect or malfunction of a biological system (see, for example, Wright 1973, 1976; Cummins 1975; Brandon 1981; Falk 1981; Neander 1991; and Millikan 1998). Brisson (2003) seems to be employing a historical, selectionist account of what determines a biological function. Under such an account, an adaptive function of a biological structure or mechanism is the feature or activity of the mechanism that resulted in its being selected in previous generations. That is, in order to be considered an adaptation, the feature must have aided in the survival and reproduction of the organisms that possessed it, resulting in its being naturally selected and passed down to present day offspring. On this account, a side effect is any feature that plays no positive role in the system being selected (Wright 1976; Millikan 1998; see also Section 2.3). For example, a function of the mammalian heart is to pump blood while its beating sound is merely a side effect. The former is a feature of hearts that has caused past generations of animals to survive and reproduce very well, while the beating sound is a feature that, in and of itself, is irrelevant to survival or reproduction. Brisson (2003) considers the hypermutable state to be a side effect rather than an adaptive function, but the ability to enter the hypermutable state does in fact lead to increased survival and reproduction because doing nothing would eventually be fatal, and so entering the hypermutable state should qualify as an adaptive trait according to the criteria used by Brisson himself. Therefore, Brisson's general theoretical strategy does not work against the view that the hypermutable state is an evolved biological function.

A separate concern is the possibility that the hypermutable state is applied to the entire genome, that is, to all of the cell's DNA rather than to specific genes or types of genes. Such a case would amount to Jablonka and Lamb's (2005) induced global mutation and Sarkar's overly permissive sense of directed mutations, but would not be semi-directed because it would not increase the chances of a beneficial mutation more than it boosts the chances of detrimental mutations.

The research team that first reported this effect (Cairns et al. 1988; Hall 1998) argued that the hypermutational state was precisely directed toward whatever genes were most likely to produce a relevant change in phenotype, perhaps even something akin to Sarkar's (2005) instructed mutations. This possibility was later abandoned when it was discovered that under a particular stressful regime, such as the lactose-only conditions described above, there occurred other, unrelated mutations at a higher rate as well. For example, Torkelson, et.al. (1997) detected changes to genes involved in amino acid metabolism during the lactose-only conditions. Such results were interpreted by some to indicate that the hypermutational state is applied genome-wide and therefore not directed (Lenski and Mittler 1993; Sniegowski and Lenski 1995). However, I believe that the results do not rule out semi-directed mutagenesis because the other genes that were subjected to hypermutation were all peripheral genes rather than core genes. Core genes are absolutely essential for cell functioning, such as those that support transcription, translation, and basic metabolism, and changes to these genes are almost always lethal. Peripheral genes are genes that may be important, but are not absolutely required under all circumstances. For example, genes associated with the acquisition of a specific nutrient resource or with antibiotic resistance are peripheral genes because even when such genes are damaged or absent, the cell may still continue to survive with its core functions left intact. Notice that the

conditions under discussion involve the presence or absence of such peripheral genes, and it may be that during those conditions, the hypermutable state is applied only to peripheral genes, leaving the core genes unaffected. This corresponds to Jablonka and Lamb's (2005) induced regional mutations, and would raise the chances of getting a beneficial mutation without aiming the hypermutational mechanisms at the core genes, which would likely be fatal. Therefore, this process would achieve semi-directed generation of new variations. To find out if this is indeed the case, we would need to determine whether the core genes are also experiencing hypermutation. Unfortunately, this is a difficult task because the vast majority of cells that sustain mutations to their core genes are unable to survive and therefore are not detected by standard assay techniques. Research involving metagenomic analysis—in which a representative sample of all the DNA in a given sample of cells is sequenced and identified—may help to answer this question, but such research has yet be conducted.

Before continuing, it should be noted here that any phenotypic changes resulting from the hypermutable state are *acquired* changes because they are occurring to organisms that already exist (see Section 4.3). Further, any changes acquired in this way are inheritable because unicellular organisms reproduce via simple cell division rather than using isolated germ cells. Recall that Weismann's distinction between germ cells and somatic cells in multicellular organisms was an important basis for the assumption that acquired characteristics are not inherited. Since unicellular organisms lack that distinction, any characteristics they acquire, whether through regulated hypermutation, unregulated genetic mutation, horizontal gene transfer, or any other process, are inheritable. Therefore, inheritance of acquired characteristics occurs in unicellular organisms all the time.

Another way for the hypermutation process to be semi-directed in the appropriate sense is for it to occur in a subpopulation of cells in a colony of closely related individuals. For example, Foster (2004) reports that approximately 0.1% of the *E. coli* cells in the colonies under study entered a hypermutational state during times of stress. In so doing, only the hypermutating subpopulation faces the potentially lethal danger of such a strategy, although it also faces the increased possibility of hitting on a beneficial mutation. Due to the dangers of hypermutation, it is quite possible for a hypermutating cell that has hit on a beneficial mutation to also have hit upon some lethal mutation that kills it. However, since healthy cells are able to pick up DNA from dead cells, any beneficial results can still be obtained by the colony (Redfield, 1988). The result is that the colony is able to subject only a small number of its constituent cells to the dangers of hypermutation (even genome-wide hypermutation), while the entire colony gets to benefit from any advantageous results. This sort of kamikaze strategy is especially effective when more than one mutation is required for the beneficial phenotype to be produced (Hall 1998; Foster 2004) and can work with regulatory genes as well as transcribable genes (Stolovicki et al. 2006; David et al. 2010), resulting in even greater potential for diversity.

We see an analogous kind of semi-directed generation of new variations in the mammalian immune system and its ability to proliferate a large variety of antibodies to combat an unknown pathogen. The immune system uses a blind process of recombination to search for antibodies that can bind to the unknown invader, and since the genetic manipulation is restricted to genes associated with antibody production, there is relatively little chance of producing seriously detrimental variations. Therefore, the response of the mammalian immune system is semi-directed toward increased adaptedness. Of course, for any new traits to be evolutionarily relevant, the newly acquired adaptive characteristics (the new antibodies) would have to be

inheritable. Some authors have argued that genes associated with acquired immunity can be incorporated into germ line DNA, allowing it to be inherited via normal genetic inheritance (Steele, 1981; Steele, et.al., 1998; Blanden, et.al. 2006). Any process resulting in reliable genetic inheritance of acquired, adaptive characteristics in mammals would be evolutionarily very significant, but the possibility has not yet received enough research attention to determine how often it actually happens. However, there *is* reliable evidence that at least some of a mother's acquired immunity can be passed on to offspring via breast milk (see, e.g., Low, et.al. 2010). Although it is not genetic inheritance, such a process nevertheless is a process of inheritance according to the criteria established in Section 3.3. Therefore, mammals employ a semi-directed mechanism for acquiring adaptive characteristics that are potentially inheritable, and so mammals engage in a very limited kind of Lamarckian evolution in their defense against pathogens.

This leads us back to microbes and to O'Malley and Dupré (2007) who argue that mature bacterial colonies should be categorized as multicellular organisms. They start their argument by pointing out that the original account of what it takes to be a multicellular organism was formulated when attention was focused exclusively on macrobes (plants and animals mostly, but also fungi), and so the account is unfairly skewed in favor of macrobes. A multicellular organism, according to the traditional account, is composed of cells that all have the same genome but are specialized to perform particular functions. Proper coordination of the various specialized cells allows the multicellular organism to accomplish tasks that no individual cell would be able to accomplish alone. O'Malley and Dupré (2007) point out that bacterial colonies known as "biofilms" also display cell specialization and coordination that enable whole biofilms to perform tasks (such as certain sorts of food collection) that no individual cell would be able to

accomplish alone (see also Shapiro, 1998), making them little different from uncontroversially multicellular organisms such as algae, fungi, or even jellyfish and hydras.

The main impediment to bacterial multicellularity, then, is the relatedness of cells in a biofilm versus those in standard multicellular organisms. Bacteria draw themselves together into colonies via chemical signaling, but the signaling is not species specific, so a single biofilm could be composed of many different taxa, and the characteristics of the colony depend on the characteristics of the kinds of cells that end up in the colony. Thus, bacteria appear to engage in a sort of highly flexible, mix and match development that stands in contrast to the relatively rigid development of most macrobes. Therefore, a critic of O'Malley and Dupré might argue that biofilms do not qualify as multicellular organisms, but merely as collections of cells. O'Malley and Dupré counter this line of reasoning by pointing out that the cells composing a standard multicellular organism often are not as genetically identical as the critic assumes. First of all, mitosis of somatic cells during normal development inevitably results in some replication errors, and so there is some diversity in somatic cell genomes. Secondly, most multicellular organisms live symbiotically with other organisms, usually microbes. Symbiotic cells and their host organisms perform essential roles for each other such that neither would be viable without the other. Therefore, argue O'Malley and Dupré (2007), the genetic diversity of cells in a biofilm is not such a radical departure from the diversity of cells that compose standard multicellular organisms. O'Malley and Dupré conclude that microbial multicellularity is just a different multicellular strategy than that displayed by macrobes.

I am sympathetic to O'Malley and Dupré's conclusion, but even if the 'multicellular organism' label is rejected for biofilms, they nevertheless are coherent, coordinated biological systems capable of reproducing themselves, and so they are subject to evolution. Distinct

biofilm lineages may be difficult or even impossible to individuate over long time periods, but the same can be said for the difficulty of individuating bacterial lineages due to their penchant for horizontal gene transfer (Woese, 1998, 2002; Vestigian, et. al., 2006). With this view of biofilms as multicellular organisms (or, at least, multicellular *somethings*), the hypermutable state of bacterial colonies described above looks even more similar to the semi-directed generation of variation happening in the mammalian immune system. In mammalian acquired immunity, the increased rate of genetic changes is restricted to B cells and also to genes associated with antibody binding features. In bacterial colonies, the increased rate of genetic changes is restricted to about 0.1% of individuals (Foster, 2004) and may be restricted to peripheral genes. The evolutionary impediment faced by mammals was that acquired immunity is, at best, only partly inheritable. But with bacteria, since there is no distinction between somatic cells and germ cells, any adaptive characteristics a colony acquires as a result of the hypermutable state can be inherited by a new colony, even if the hypermutated cells all die in the kamikaze strategy described above. Therefore, semi-directed hypermutation in bacterial colonies closely resembles the process advocated by neo-Lamarckians at the turn of the 20th century (see Sections 3.3 and 4.3), though, of course, updated by modern molecular biology.

Non-genetic Changes

In Chapter Three, I argued that genetic inheritance (DNA replication and transmission to offspring) is not the only hereditary process. All other parts of a living cell, including any cell walls, membranes, organelles and cytoplasmic contents must also be transferred to offspring in order for offspring to resemble their parents. Changes to these other, non-genetic parts of cells are also potentially inheritable because formation of new cell structures is determined in part by

existing cell structures. Some changes to cell walls and membranes are potentially inheritable (Landman, 1991), though such changes typically are a blind result of damage or malfunction. As such, they are similar to blind genetic mutation due to replication errors. That is, they are not the result of any change-producing mechanisms, but nevertheless are potentially inheritable.

Similarly, changes tracked by evolutionary developmental biology (evo-devo) and Developmental Systems Theory typically are not the result of any change-producing mechanisms, but instead are the result of blind changes to factors that contribute to some stage of development (see the cave fish example in Section 4.3). The change may be a genetic mutation to a regulatory section of DNA or a change in some environmental factor, such as temperature or nutrient availability, that blindly alters development of phenotypic characteristics. Since the risks associated with introducing changes into a developmental pathway tend to outweigh any possible benefits, mechanisms of heredity and development tend to be very tightly controlled in order to ensure that offspring are viable and resemble their parents. Therefore, introduction of developmental changes is not a biological function, but rather is a malfunction or side effect of hereditary and developmental mechanisms.

Summary

In this Chapter I showed how the gene-centric view that emerged from the Modern Synthesis led most biologists to assume that directed variation and inheritance of acquired characteristics did not occur. I then showed that more recent research shows that both processes really do happen in significant portions of the biological world. The most significant examples occur in bacterial colonies where the two processes work in tandem to accelerate the pace of bacterial evolution. The historical figure that came to be most closely associated with directed

generation of variations and IAC is Lamarck, and when those processes were assumed to be absent in the biological world, Lamarck ended up as somewhat of a pariah in the scientific community. Therefore, recognizing that both processes are operative in nature is somewhat of a vindication for Lamarck, although the underlying mechanisms turned out to be much different than what he had in mind.

I also showed that blind variation typically is not accomplished by a biological mechanism. Rather, blind generation of new variations typically occurs as a side effect of mechanisms whose activity involves heredity and development. That is, many of the biological mechanisms for heredity reproduction allow some amount of imprecision (e.g. DNA replication errors) or randomness (e.g. orientation of homologous pairs of chromosomes in preparation for Meiotic division) that result in blind variations.

In chapter Five, I shift to natural selection, which is the third process needed for ongoing, adaptive evolution. I use the mechanistic approach and the preceding accounts of heredity and generation of new variations to understand natural selection as a process separate from the other two.

CHAPTER FIVE

NATURAL SELECTION

5.1 Introduction

Natural selection was the key element needed to develop a successful theory of biological evolution and origin of species. The most important task of a theory of biological evolution is to explain why organisms are so well adapted to their environments, and before evolution by natural selection was proposed, adaptedness of organisms could only be explained through divine intervention, unobservable vital forces, or pure luck, none of which are scientifically satisfying. Creationism, which was the dominant view until less than a century ago (and which still has many contemporary adherents), explains that organisms are the result of God's plan and creative abilities. That explanation is scientifically unsatisfying because God's plans and causal methods (if they are supposed to be distinct from natural causes) are mostly mysterious and opaque to investigation. The theory just does not make enough testable claims about how and why the creator worked. Plus, any significant testable claims made by creationists have turned out to be false. Creationism's more recent variant, intelligent design, attempts to shed the religious component (Behe, 1996, 2007; Monton, 2009) but even if it is successful in doing so, it so far has not been any less opaque to research than creationism (Young and Edis, 2004; Matzke, 2007). Other naturalistic attempts, such as Lamarckism or vitalism, proposed special animating forces or fluids that automatically worked toward greater complexity and adaptedness. But those proposals required goal-oriented forces that were not much better than the divine intervention they were supposed to replace. As understanding of molecular biology increased, Lamarckism and vitalism had trouble describing their animating forces or fluids in a way that opened them up to empirical investigation, and so they seemed utterly mysterious. Only when natural selection

finally was accepted into evolutionary theory was it possible to explain the adaptive features of organisms without appealing to any mysterious, indescribable forces.

As discussed in Chapter One, modern-day scientists claim to be interested in finding mechanisms that can explain how and why phenomena occur, rather than just describing the phenomena themselves (Casadevall and Fang, 2009). Perhaps it is no surprise, then, that once natural selection replaced the other, more mysterious options for explaining adaptedness, biologists started to refer to natural selection as a mechanism (Skipper and Millstein, 2005). However, the process of natural selection is unlike other proposed biological mechanisms such as hearts, neurons, or ion pumps. Standard biological mechanisms are, roughly speaking, discrete units organized to accomplish a specific biological function. Natural selection, by contrast, involves ever-changing interactions between individuals and their environments, and so the entities involved are not all part of a single system. In section 5.4 I discuss recent debates over whether natural selection is a mechanism and I draw upon my account of biological mechanisms in Chapter Two to argue that natural selection is not a biological mechanism.

Before assessing whether natural selection is a mechanism, however, I must make clear what I believe to be the activity of natural selection. There continues to be a great deal of debate over the correct characterization of natural selection and over what actually gets selected (the so-called units of selection). In Section 5.3 I join the debate by providing my own account of natural selection and defending it against various other prominent accounts. My account differs from most others in its strict separation of the three distinct processes that constitute biological evolution: hereditary reproduction, generation of new variations, and natural selection. Most other accounts of natural selection include reproduction or generation of new variations (or both), and so are not proper descriptions of natural selection in and of itself. In order to frame

the debate over the correct characterization of natural selection, I first provide an overview of the history of attempts to explain the existence and adaptedness of organisms. Some of the history described in section 5.2 overlaps with the historical overviews in Sections 3.2 and 4.2, but here the emphasis is on how natural selection came to be part of the explanation for adaptedness of organisms.

5.2 Historical Development of Concepts and Terms

In order to trace the history of natural selection, it is necessary to understand the two most salient phenomena that natural selection is intended to help explain. Any legitimate theory of biological origins must explain the adaptedness of organisms as well as the great diversity of organisms (see also Sections 3.2 and 4.2). Before any acceptable naturalistic theories of evolution had been proposed, the best available explanation for diversity and adaptedness was that all organisms had been designed and created by a supernatural being (e.g. Paley, 1802; Plantinga, 1991). As intentional creations by a supernatural being, each species was also viewed as an immutable type of organism. Typological thinking has been addressed in previous chapters (particularly Sections 3.2 and 4.2), but the issues are relevant again here so I will repeat some of the most important points.

Even though some superficial variability could be observed in any species, most biologists in the 19th century assumed there to be a standard form or essence for every species and that individual members of a species never stray too far from the standard form. Such typological thinking about species was tempered by plant and animal breeders who recognized that novel traits could emerge in a population and be passed on to offspring, and that selective breeding could change a species quite radically. However, despite the plant and animal breeders'

track record of changing species in dramatic ways, typologists were able to dismiss their results by pointing out that domesticated animals, when returned to the wild, tended to revert to their wild forms and to lose any special traits that had been promoted through artificial breeding programs (Mayr, 1982, 2001b). Thus, the effects of artificial breeding programs, though sometimes quite dramatic, seemed superficial and readily reversible, and there was little reason to think that the very large-scale, permanent evolution of one species into a new species was possible.

By the early 19th century, enough evidence from the fossil record had been gathered to indicate very strongly that some sort of long-term change in species had occurred, and so typologists needed to come up with an explanation. For those who wished to avoid any large-scale changes in species, evidence from the fossil record could be interpreted as showing that older species occasionally disappeared and were replaced by new ones. One possibility was that a supernatural creator that designed and instantiated each type of organism occasionally decided to eliminate some and replace them with others, but such speculations about divine intervention are difficult if not impossible to test. A naturalistic option proposed by the influential anatomist and paleontologist George Cuvier was that old species were replaced by new species via immigration from other areas (Mayr, 1982; Hull, 1984). Although naturalistic, Cuvier's proposal was not much of an explanation because it left mysterious where all those new species were coming from and why the fossil record shows an overall trend toward increased complexity. Nevertheless, absent a better theory, suggestions such as Cuvier's were tentatively accepted by typologists opposed to evolution.

There were typologists who were not opposed to evolution of species, but they faced a problem because a standard typological assumption was that an organism's set of traits form a

coherent whole, and that changing just one or a few traits would disrupt the essence of the organism to the point that it would not be viable (Mayr, 1982). For this reason, gradual, step-by-step evolution was considered to be impossible because any small changes would result in non-viable offspring. A prominent way of reconciling evolution with typology was to adopt some version of "saltationism" (Mayr, 1982; Dawkins, 1983) in which very rare but very radical changes (saltations) could appear suddenly in a lineage—even over a single generation—thus converting the species into a different viable type. Although no plausible mechanisms were suggested that could perform the activity proposed by saltationsts, it actually seemed consistent with the fossil record, which was relatively discontinuous at the time, and so saltationism (or something like it) remained entrenched among evolutionists in the mid-19th century (Mayr, 1982; Hull, 1984).

Saltationism was a common view up until the early 20th century, though the seeds of its replacement were sown in the mid- to late 19th century, particularly by Darwin (1859).[43] One of Darwin's (1859) great contributions to evolutionary thinking, apparently inspired by 19th century economist Thomas Malthus, was his reconception of species as populations rather than as types. One of the most important differences between population thinking and typological thinking concerns the nature of variability. Typology assumes that each individual in a species is an imperfect instantiation of an abstract universal, and so it allowed for variability in the form of small deviations from a standard, unchanging body plan. This was labeled "continuous" variation (Mayr, 1982) since the variability was allowed to drift anywhere within a small range. Although dramatic differences might appear on occasion, according to typology, such "discontinuous" variations could not persist, thus drawing the species back toward the standard

[43] Alfred Russell Wallace, roughly simultaneously, made many of the same contributions as Darwin, though Darwin has ended up with most of the credit.

type (see Section 4.2 for more discussion of continuous and discontinuous variations). In populations, by contrast, there is no standard body plan to which each individual is compared, and so variability is considered normal rather than deviant. As a consequence, there is no distinction between so-called continuous variation and discontinuous variation in populations—it is all just variability. It should be noted also that typology, even among scientists of the 19th century, was rooted in religious ideology. There remained an assumption among many intellectuals of the time that God had created a perfectly balanced and harmonious natural world, with every species perfectly designed to fulfill its own natural purpose (Mayr, 1982). This stands in stark contrast to population thinking, which allows for dramatic imbalances between population size and resource availability leading to a great deal of suffering and death. Thus, in addition to any empirical or theoretical objections, certain ideological commitments stood in the way of Darwin's ideas.

If the constraints of typology are shed, then species are allowed to change very dramatically over time, and can do so gradually. Plant and animal breeders had already provided examples of dramatic changes in captive lineages, and so Darwin (1859) used their examples to argue that similar kind of change occurs in the rest of nature. Plant and animal breeders intentionally select which individuals survive and reproduce, and their decisions are based on pre-determined sets of desirable characteristics. The result is a population of individuals that, by and large, only have the desirable characteristics. Darwin argued that an analogous process is at work in wild populations. All organisms tend to produce more offspring than can be supported by available resources. As long as resources are scarce, there inevitably will be competition for those resources, with some individuals winning and others losing.[44] Individuals that are better

[44] This is where Darwin apparently was inspired by Malthus' description of human populations and economic boom and bust cycles.

equipped to obtain the scarce resources will, more often than not, out compete those less able to obtain resources. As a result, by and large, only the best competitors end up surviving and reproducing offspring like themselves, while the lesser competitors die, leaving the population overall better adapted to its environment. The results are so similar to the results from plant and animal breeders, that Darwin considered them both to be processes of selection. Plant and animal breeders are engaged in "artificial selection" because they use a pre-conceived set of characteristics as the basis for their intentional, intelligently guided selection. In wild populations, survivors are determined by environmental conditions rather than by a pre-set plan, and so Darwin dubbed the process "natural selection."

Combined with hereditary reproduction and a source of new variations, natural selection explains how organisms have come to be so well adapted to their environments. Since individuals that happen to be better equipped to survive and reproduce in their environments are the ones most likely to actually survive and reproduce, natural selection results in a population with more well-adapted individuals than poorly adapted ones. Evolution by natural selection also helps explains diversity[45] because each population faces a different set of environmental challenges, and so different sets of characteristics will be selected in different environments. Darwin went so far as to suggest that all life may be descended from a single ancestral population from which many sub-populations splintered off and migrated all over the word, each gradually adapting in unique ways to each new environment they encountered.

Darwin's theory of evolution by natural selection combined with his assumption of common decent is what I'm calling "Darwinian" evolution.[46] Although it provided a plausible naturalistic explanation for diversity and adaptedness of organisms and was widely read and

[45] Diversity also is partially explained by the sources of new variations (see Chapter 4).

[46] Contemporary evolutionary theory is still basically Darwinian because it includes these two core concepts.

debated, Darwinian evolution was not accepted by most of the scientific community in the mid-19th century. Darwin's theory, with its emphasis on harsh, messy competition and unguided change clashed too much with the utilitarian and typological views of nature that still were prominent at the time, and it seemed unsupported by the discontinuous nature of the fossil record.

In contrast to most others in the scientific community, plant and animal breeders were more open to Darwinian evolution because they understood the power of selective breeding. Breeders were joined by biometricians and naturalists who studied trait variations in wild populations and so were less inclined toward typological assumptions. Biometricians understood that although casual observation seemed to support the typological belief in species conformity and immutability, more careful and systematic measurements revealed a great deal of variation and gradual change. Perhaps the best example of biometric work in support of Darwinian evolution is the famous set of experiments on peppered moths conducted by Bernard Kettlewell in the mid 19th century. Kettlewell showed that peppered moths changed from predominantly light colored to predominantly dark colored when their environment (English woodlands) became darker as the result of coal burning. Kettlewell not only showed a correlation between a changing environment and a changing phenotype, but he showed that the change was a result of Darwinian natural selection (combined, of course, with hereditary reproduction). Through painstaking observations of moth behavior in the wild, Kettlewell showed that when trees had light colored bark, dark colored moths were being eaten by birds more often than light colored moths because the dark variety was not sufficiently camouflaged. As a result, moth populations in areas with light colored trees were dominated by light colored moths. Then, when the trees

had become dark due to accumulated soot, the light colored moths became the prey of choice for local birds, and over just a few generations, the dark colored moths predominated.

Although work like Kettlewell's seems in retrospect to be very persuasive, it failed to move many scientists who were skeptical of natural selection. They acknowledged the work of biometricians and naturalists, but complained that their work showed only very minor, superficial changes that could quickly change back.[47] To be swayed by Darwinian evolution, critics wanted to see *speciation* in which one species gradually evolves into a different species. That is, typologists could countenance some microevolution, but if they observed one type gradually change into some other type, then they would be forced to get rid of types altogether and accept something like Darwin's gradual evolution of species.

Another problem for Darwin was the postulated age of the earth, which put severe constraints on the pace of evolution. As discussed in Sections 3.2 and 4.2, Darwin originally preferred that new variations were generated blindly with respect to their adaptedness, and that differential natural selection be allowed to sort the adaptive from the maladaptive. By the late 19th century, however, the earth was thought to be no more than several million years old, which did not allow enough time for natural selection acting on blind variation to produce the huge diversity of organisms we see today. In light of this time constraint, Darwin enlarged the roles of directed generation of new variations and inheritance of acquired characteristics in order to help explain the complexity and adaptedness of organisms (Mayr, 1982). While this helped Darwinian evolution conform to the time constraint, it did so at the expense of downplaying the role of natural selection in explaining the adaptedness of organisms, and made his theory more closely resemble Lamarck's. As discussed in previous chapters, many prominent and influential scientists had rejected Lamarckian processes, and so Darwin faced an uphill battle.

[47] Today this is called "microevolution," where large-scale evolution of species is called "macroevolution."

Another setback for evolution by natural selection came from the re-discovery of Mendelian genetics at the beginning of the 20th century. Given the understanding of basic Mendelian inheritance patterns at the time, even if a new, adaptive trait appeared in one or a few individuals, it was not clear that such rare traits could be preserved and proliferated in future generations. Darwin's theory required that new traits occasionally arise and spread in a population, but according to the geneticists's formulae, such rare traits would quickly disappear from any sexually reproducing species, thus providing a barrier to any gradual process of evolution such as Darwin's. This was an important obstacle because the field of genetics commanded greater scientific respect than the work of biometricians and naturalists whose work was inherently more qualitative and less experimental than the geneticists'. Thus, Darwinian evolution by natural selection still remained a minority view in the scientific community (Mayr, 1982).

The path toward acceptance of natural selection finally opened up in the late 1920s and early '30s when basic Mendelian genetics gave way to population genetics. Population genetics had the conceptual and mathematical tools needed to show how new, rare traits could be preserved and proliferated in a population over many generations, which made genetics compatible with gradual evolution by natural selection. Since genes are importantly responsible for an individual's phenotypic characteristics, tracking any changes in the makeup of a population's gene pool can allow researchers to track the evolution of that population. Since natural selection was no longer seen as incompatible with genetics, the scientific community finally felt comfortable accepting gradual evolution by natural selection (Mayr, 1982). Plus, improved knowledge of radioactive isotopes in geology resulted in a revised age estimate for the earth at around 4.5 billion years, which is plenty of time for gradual evolution to produce the

adaptedness and diversity of life on earth from a single common ancestral population. All of these advancements finally overwhelmed the entrenched resistance and allowed for establishment of the Modern Synthesis, which combined population genetics with neo-Darwinism[48] to create the first comprehensive theory of gradual evolution to be accepted by the scientific community.

Fitness

Natural selection has often been expressed as "survival of the fittest," a phrase coined by Herbert Spencer (1864) and later adopted by Darwin, but the meaning of 'fitness' changed significantly after adoption of the Modern Synthesis. Darwin used fitness as shorthand for an individual's degree of adaptedness to local environmental conditions, and so fitter individuals are more likely to survive and reproduce than less fit individuals. The meaning changed when the Modern Synthesis placed population genetics at the center of evolutionary theory. Population genetics is primarily concerned with tracking alleles[49] over multiple generations. By tracking how allele frequencies change over time, population genetics can quantify evolution of a population. In addition to mere tracking of alleles, population genetics can provide a numerical estimate—called a fitness value—of how likely each allele is to be passed on to the next generation, and can use those values in probability calculations to generate predictions and explanations concerning how allele frequencies change over time. The success of this approach began shifting the focus of natural selection away from individuals and onto genes, and 'fitness'

[48] Neo-Darwinism, promoted by August Weismann, combined Darwinian evolution by natural selection with blind generation of new variations and no inheritance of acquired characteristics (see Section 3.2 and 4.2).

[49] Alleles are gene variants and the source of much of the phenotypic variability observed in a population.

came to refer to the potential of individual alleles to be passed on to the next generation rather than referring to the overall adaptedness of whole organisms.

The mathematical grounding and the predictive abilities of population genetics helped give evolutionary biology its scientific credibility and was the original basis of the gene-centrism of evolutionary theory described in Chapters Three and Four. Since population genetics was central to the Modern Synthesis, any hypotheses concerning the details of evolutionary theory had to be compatible with genetics, an approach that proved to be very successful. One prominent example concerns altruism in which one individual appears to sacrifice its fitness in order to boost the fitness of some other individual. "Survival of the fittest" implies a zero-sum game in which the strongest survive while weaker individuals are left to die. Fitter individuals are much more likely to have their alleles represented in the next generation, so altruistic sacrifice of one's own fitness seemed evolutionarily self defeating. Any altruists would be less likely than their selfish conspecifics to pass on their altruistic traits, and evolution of selfish behavior would quickly overwhelm any altruistic behavior. It seemed, therefore, that altruism could not evolve, so the existence of actual altruism in the biological world was rather puzzling.

One proposed solution was to include *group* selection in addition to organismal selection or gene selection. According to group selectionists of the early 20th century, individual altruism could be evolutionarily advantageous because it is good for the group to which an individual belongs. In other words, groups (populations) with altruists may be selected over groups without altruists. This solution was shot down by Williams (1966) who showed that such thinking is based on a mistaken understanding of how biological evolution was understood within the Modern Synthesis. Williams explained that a trait can only make it into the next generation if individuals that possess the trait pass it on to their offspring. Groups cannot reproduce

themselves, only individual organisms can, so even if a trait somehow benefits a group, the trait cannot be passed on unless the individuals that possess the trait pass it on. And since it was assumed that altruism inevitably decreases an individual's chances of reproducing while selfishness increases those chances, in the long run, selfish traits will always win out over altruistic traits. So, according to this reasoning, even if altruism is in some sense good for the group, it still would be an evolutionarily self defeating strategy in the face of selfish competition, Therefore, Williams (1966) concluded, there can be no traits that have evolved for the good of the group at the expense of the individual.

Later mathematical work in population genetics to reveal how altruism, under certain circumstances, could be an evolutionarily stable trait without running afoul of Williams's criticisms. For example, if altruism is directed toward close relatives, then it can promote reappearance of one's own genes in the next generation by promoting the survival and reproduction of close relatives with whom one shares many genes (Hamilton, 1972). Thus, an individual can actually increase its own fitness (the population genetics sense) by aiding family members, even at the expense of its own individual survival.

The above example is significant because it demonstrates the effectiveness of evolutionary explanations based on population genetics relative to explanations coming from other subfields of evolutionary biology. In contrast to population genetics, ecological field studies are much more difficult to control than are laboratory breeding studies, and the multiple interacting species in an ever-changing environment that ecologists observe are so complex that they do not lend themselves to mathematical analysis. As for the fossil record, it is only a static record, and, despite major advancements, still is seriously incomplete. Therefore, any evolutionary explanations or predictions from fields like ecology or paleontology have

traditionally relied much more on creative interpretation rather than more objective statistical analyses, at least as compared to population genetics. This has left researchers in those other fields open to the charge of generating "just-so" stories in which creative scientists weave a narrative linking their disparate pieces of evidence in an "explanation" of how various phenotypic traits have evolved. At their worst, such narratives ignored important facts from other fields that rendered the explanation implausible, which is what happened with the "group selection" explanation for altruism described above. Even when they did not conflict with well-established knowledge from other fields, such narratives typically provide only one possible way that a phenotypic trait may have evolved, which falls short of scientifically rigorous explanation.

The relative success of explanations in terms of genes and population genetics had the effect of shifting the focus of evolutionary theory, and natural selection in particular, ever further away from organisms and even more onto genes. Not everyone has gone along with that shift, however, so there is a distinct lack of consensus on the correct characterization of natural selection and its role in adaptive evolution. In Section 5.3 I show how the shift toward gene-centrism has led to a conflation in the contemporary literature between natural selection and reproduction, and I propose my own account of natural selection to resolve the problem.

5.3 Activity of Natural Selection

Recall from Chapter Two that a biological mechanism must have both an activity and entities that carry out that activity. So, in order to determine whether natural selection is a mechanism, we must have a precise characterization of the activity of natural selection. Unfortunately, there continues to be a great deal of debate over what natural selection is, how it works, and what it selects. Much of that debate, I believe, results from a failure to consistently

distinguish natural selection from the other two processes at work in adaptive evolution. In this section I present my own account of natural selection as a process separate from reproduction or generation of new variation and defend my characterization against other prominent accounts. As explained in the Introduction, it is possible for each process to occur without either of the others, and only when all three are acting in concert can there be long-term, adaptive evolution. Of course, all three are closely intertwined in most biological systems, but they still can be distinguished from one another. Therefore, my account of natural selection on its own does not include variation between individuals or reproduction. This puts me at odds with many prominent accounts of natural selection and so I must clarify exactly where my account differs from others, and explain why I believe mine is preferable.

Perhaps my most significant deviation from standard accounts is my insistence that natural selection, by itself, does not involve reproduction. First of all, recall from the Introduction that there is a difference between viability selection and so-called reproductive selection, and that viability selection helps explain adaptedness while reproductive selection does not. Reproductive selection certainly requires viability, but by itself, reproductive selection does not promote evolution of environmental adaptations. In fact, reproductive potential can be increased by diverting energy or other resources toward reproduction at the expense of personal survival. Therefore, if there is no viability selection favoring individuals that are better adapted to surviving in their environments, evolution will not lead to increasing adaptedness even if it leads to increasing reproductive potential.

In addition to viability selection and reproductive selection, sexually reproducing species may also experience sexual selection in which individuals select mates based on their possession of certain attractive traits. Like viability selection, sexual selection involves environmental

interactions, at least insofar as other members of one's own species are part of an individual's environment, but does not include reproduction. Sexual selection often includes mating, but it does not necessarily involve reproduction because even sterile individuals may be sexually selected despite their inability to reproduce. Of course, in order for sexual selection to lead to evolution of a sexually attractive trait, hereditary reproduction must follow sexual selection, but reproduction itself remains a separate process from the selection process, whether it is sexual selection or viability selection. Like reproductive selection, sexual selection does not necessarily have anything to do with environmental adaptedness. In fact, many famous examples of sexual selection involve selection of traits that are detrimental individual survival. Therefore, while evolution by sexual selection helps explain some prominent features of some species, it does not explain environmental adaptedness of species.

Here, then, is my own account of natural selection so that I can contrast it with other views in the literature. In broad terms, natural selection (as it occurs in biology) is the constant causal interaction between a living individual and its environment, where the interactions cause the individual to continue surviving and functioning or cause it to die. Natural selection is causal because it necessarily involves physical interaction between an individual living system and its environment. It is constant because an individual always exists in some environment or other, and so is always interacting in some way with an environment. As highly complex, low-entropy, dynamic systems, living organisms are continuously working to counteract their own senescence and decay, even when not facing any special challenges like pathogens or predators. Every living system, therefore, is exposed to natural selection every moment that it lives. If the individual continues to survive, then it is selected (or selected for), but if it does not, then it is selected against.

Now, to fill in the details. Natural selection requires at least one living individual existing in an environment, and the physical characteristics of that environment must causally interact with the individual in ways that depend on the individual's own physical characteristics as well as those of the environment. That is, the interaction between individual and environment cannot be such that it would be exactly the same no matter what characteristics either possessed. Further, it must be the case that some potential interactions between environment and individual could be disruptive to the individual (e.g. may cause it to break apart or stop functioning) while other potential interactions are beneficial the individual (e.g. may support its structural integrity or allow its continued functioning), while still others may be roughly neutral with regard to survival or functioning of the individual. If the individual has characteristics that lead to beneficial interactions with the environment, then it will have greater chances of surviving and functioning, and so has been *selected for*. An individual that possesses characteristics that lead to detrimental interactions with the environment will have decreased chances of continued survival and functioning, and so has been *selected against*. Whenever these sorts of causal interactions occur, an individual is being naturally selected for (or against) continued survival and functioning. Although natural selection often is associated with population genetics, according to my account, natural selection is more closely associated with physiology, ecology, biochemistry, embryology, and the like because those fields study how individual organisms work to survive and function in their environments. Those fields, rather than population genetics, are central to understanding natural selection because they get at the causes—the mechanisms, in many cases—that underlie evolutionary theory's ability to explain why some individuals survive while others do not, and thus its ability to explain how adaptations continue to evolve.

Although I have been speaking in terms of biology and living organisms, it should be noted that natural selection can be generalized to include non-living individuals as well (Lewontin, 1970; Dawkins, 1980; Dennet, 1995; Hull, et. al., 2001). Whenever interaction between an individual thing and its environment affects the continuing existence or activity of that individual, then the individual is being subjected to natural selection. One example is pebbles on a beach that are sorted by wave action, where some pebbles are removed while others remain. Other examples include "survival" and "death" of computer-based entities that exist in a digital environment in which the entities compete for computer resources they need to avoid deletion (e.g. Grabowski, et. al., 2011), and cosmic evolution involving natural selection working on a diverse set of universes (Smolin, 1992, 2006).

There are a number of features of my account of natural selection that set it apart from most others. For example, contrary to some prominent accounts (e.g. Millstein, 2006; Matthen and Ariew, 2002), my account of natural selection requires only a single individual rather than a population of phenotypically diverse individuals. This is consistent with my strict distinction between natural selection by itself and the larger process of adaptive evolution. Adaptive evolution is a process that accumulates adaptive changes, but selection by itself can only result in an individual succeeding or failing. An individual just continuing to survive is not change, but merely continuation. So selection cannot produce evolutionary changes by itself. When natural selection happens within a population of individuals with diverse characteristics, then there is the possibility of *differential* natural selection, where some individuals tend to have more successful interactions with the environment than others, resulting in a change in the population when some die while others live and go on to reproduce.

As further support for my view that natural selection by itself does not require a population, artificial selection, which is closely analogous to natural selection, can also happen to just one individual. In artificial selection, where humans decide which plants or animals to keep, if the human caretaker is left with just a single individual, then she can decide to keep the individual alive (and maybe get it to reproduce, if it can do so asexually) or she can decide to let it die. Thus, she can *select* that single individual. Similarly, if a wild population is down to a single individual, that individual can still be subjected to natural selection within its environment. In fact, even in populations, especially ones in which individuals are scattered across a wide area, each organism faces its environmental pressures as an individual. Therefore, natural selection not only *can* operate on individuals, but *always* acts on individuals.

Another consequence of my characterization of natural selection is that it does not require diversity. This puts me at odds with Ernst Mayr (2001a, 2001b), one of the most influential founders of the Modern Synthesis. I believe that natural selection does not require diversity because, as I argued above, natural selection acts on single individuals, and a lone individual cannot be diverse in the way required. Plus, if we consider a population in which all individuals are identical, each individual may still be subjected to natural selection despite the lack of diversity (see the Introduction for more discussion of natural selection in a population of with zero diversity).

Another issue concerning the activity of natural selection is whether natural selection necessarily involves competition, as many authors have claimed (e.g. Dawkins, 1976, 1980; Hull, 1980; Mayr, 1976, 1982; DeDuve, 2005). Competition is a very good example of natural selection in action, but natural selection does not require competition. First of all, if it is true that natural selection can act on single individuals living apart from any larger population, then

competition is not required for natural selection because a lone individual has nobody with which to compete. Similarly, if there is a great surplus of resources, then even in a crowded population, individuals will not need to compete with each other at all, yet natural selection may nevertheless be operative because each individual must still be viable in that environment, and some individuals may still be better than others at obtaining the abundant resources. Perhaps such a situation could still be thought of as a kind of competition, analogous to a pie eating contest (whoever gorges most wins), but it is a very watered down notion of competition when compared to the old notions of survival of the fittest[50] that were used to popularize Darwin's theory.

Another reason to think that natural selection does not require competition is that natural selection can involve cooperation. There are many examples of mutualistic symbioses in which members of different species promote each other's survival because they provide each other with valuable resources or services. This is not mere altruism. Altruism involves giving up resources without receiving a comparable benefit in return, while mutualistic symbiosis occurs when each individual receives an important benefit by helping its co-symbiont. Such symbiotic cooperation is the antithesis of competition, and yet is an interaction that promotes continued survival and functioning. Therefore, natural selection does not necessarily require any kind of competition.

Contrasts with other prominent accounts

I maintain that most prominent accounts of natural selection in the literature end up including too much of one or more of the other two processes in evolution. That is, in their accounts of natural selection, they inappropriately include hereditary reproduction or generation of new variations even though those are separate processes that need not occur in order for

[50] This characterization is attributed to Herbert Spencer in the mid to late 19th century, and was adopted by Darwin as a shorthand way of describing natural selection.

natural selection to occur. To reiterate, all three processes are needed for ongoing adaptive evolution, but all three are separable and need not occur together.

Much of the literature on natural selection is centers on the units of selection, and debate over the units of selection was prompted by the need to explain altruistic behavior in biological populations. Recall from Section 5.2 that evolution of altruism was puzzling because it involves one individual sacrificing its own ability to survive or reproduce in favor of another individual, and that seems maladaptive. However, since altruistic behavior appears to have evolved and persisted in many different lineages, it must have some adaptive benefit. This led some to propose "group selection" in which groups of altruists were naturally selected over groups of non-altruists, thus leading to a prevalence of altruists in creatures such as ants, bees, and many species of birds and mammals. The problem with this explanation, as pointed out most famously by Williams (1966) was that groups cannot reproduce, only individuals can. Consequently, even if a group benefits from altruistic behavior (because the group persists longer than it would if all of its members were strictly selfish), altruistic individuals still would be less likely to reproduce than selfish individuals, and so altruism would be eliminated in favor of selfishness. Therefore, the proposed explanation for altruism coming from group selectionists failed, and it was concluded that group selection did not occur. In other words, groups were not thought to be units of selection.

A problem with Williams's (1966) conclusion is that his argument was against the possibility of group traits evolving by natural selection, not against natural selection of groups, per se. Williams pointed out correctly that groups typically cannot reproduce in the way required for adaptive evolution by natural selection, but then he used that fact to conclude that groups cannot be units of selection. The problem is that an inability to reproduce does not imply an

inability to be subjected to natural selection. Instead, it only means that any natural selection cannot contribute to ongoing evolution. Therefore, Williams (1966) presents an argument against adaptive evolution of groups, not against group selection by itself.

Unfortunately, Williams's approach persisted in the literature, including Lewontin's (1970) extremely influential essay "The Units of Selection." Lewontin argues that adaptive evolution by natural selection requires just three basic elements. First, there must be a population of individuals with variable phenotypes. Second, different phenotypes must have different rates of survival and reproduction. And third, there must be a correlation between parents and offspring in the contribution of each to future generations. As long as these three elements are present together, explains Lewontin, then the population will evolve by natural selection. Lewontin (1970) explicitly left open the possibility that the "individuals" in his account could be genes, organisms, or even groups. As long as the entities in question meet the criteria, then they can be units of selection.

Although I largely agree with Lewontin (1970) regarding the three elements required for adaptive evolution, he makes the same mistake as Williams (1966) by drawing conclusions about units of selection based on what it takes for something to undergo adaptive evolve. Natural selection may be a necessary element of adaptive evolution, but something can be a unit of selection—it can still be a target of natural selection—yet fail to participate in the larger process of adaptive evolution due, for example, to an inability to reproduce. Another problem with Lewontin's (1970) account that should be mentioned here is that, as explained by Mitchell (1987), Lewontin's second element—that different phenotypes must have different rates of survival and reproduction—only *implies* that the phenotypes themselves are the reason for the differential rates of survival and reproduction. Unless that is made explicit, then the different

rates of survival and reproduction could be due to chance (sometimes called "drift") rather than natural selection. Therefore, Lewontin's (1970) three elements, as stated, will produce evolution by whatever means, rather than just evolution by natural selection. Mitchell (1987) suggests adding a fourth element stating that the differential survival and reproduction must be due to differential adaptedness of the different phenotypes. That way, differential survival and reproduction will be the causal result of differential interactions between individuals and their environments, where some individuals have phenotypes that allow for more beneficial interactions than do others. Whether it is considered a fourth element or a modification of Lewontin's second element, Mitchell (1987) clearly is correct that beneficial or detrimental interactions between individuals and their environments is necessary for any natural selection.

Williams (1966) and Lewontin's (1970) failure to distinguish natural selection from the larger process of adaptive evolution shows up throughout the literature on units of selection. Lloyd (2005) offers a similar critique of the debate over units of selection, but rather than framing the issues in terms of three separate process that make up adaptive evolution, she points to four separate questions with which debaters appear to be concerned. First is the interactor question, which concerns what is interacting with the environment. Second is the replicator question, which is less concerned with any environmental interaction and more concerned with what gets copied and passed on through multiple generations. Third is the beneficiary question, which explores what actually benefits from the various adaptations that have evolved. Fourth is the manifestor of adaptation question, which concerns what thing actually possesses or displays adaptive traits.

Before discussing how Lloyd's (2005) four questions help us organize the literature on natural selection, I will take a moment to describe my ways of addressing each question, and

how they relate to natural selection. First of all, I believe that the interactor question is the only one directly relevant to natural selection. As explained above in Mitchell's (1987) critique of Lewontin (1970), natural selection happens when an organism interacts with its environment. The replicator question, by contrast, concerns reproduction rather than natural selection. As for the beneficiary of adaptations question, much depends on what it means to benefit in this context. If we assume that "to benefit" is to be aided in survival, then any living system that is helped to continue living by the presence of an adaptive trait is a beneficiary of that adaptation. This is not directly relevant to units of selection because an item does not need to be subjected to selection or contribute to adaptedness in any way in order to be a beneficiary of adaptation. For example, every organism's genome contains so-called "junk DNA" that does not seem to have any functionality, but since DNA maintenance and replication machinery cannot distinguish functional DNA from junk DNA, it all survives and is passed on to the next generation together. Thus, junk DNA benefits from an organism's adaptedness even though it does not contribute to adaptedness and is not itself subjected to any environmental interactions. The manifestor of adaptations question is more closely related to natural selection than is the replicator question or the beneficiary question because whatever manifests an adaptation is a target of natural selection. An adaptation is a trait possessed by an individual (and typically acquired via evolution) that helps the individual overcome some challenge it faces in its interactions with its environment. Therefore, in order for an individual to manifest an adaptation, the adaptation must be part of the individual, the individual must be interacting with its environment, and the adaptation must be helping in those interactions. Those individuals that manifest more and better adaptations are naturally selected, while those that manifest fewer or less effective adaptations may be selected against. The upshot of this paragraph is that of Lloyd's (2005) four questions, the interactor and

manifestor of adaptation questions are most relevant to natural selection, while the replicator and beneficiary questions are less relevant.

When Lloyd (2005) discusses how her four questions relate to the units of selection debate, we can see how the debate often ends up answering either the replicator question or the beneficiary of adaptations question rather than the interactor question. For example, in the debate over altruism that spawned the units of selection debate, group selectionists (at least those in the 1960's) seem mainly to be considering the beneficiary question. As Sober and Wilson (1994) explain in their review of the debate over units of selection, "The problem of the units of selection…concerns whether traits evolve because they benefit individual organisms or because they are good for the group in which they occur" (534). The idea is that individuals act altruistically even though individual altruists do not benefit from that trait. Instead, it is the group that benefits because individual acts of altruism lead to a larger, longer-lasting group. Group selectionists used this observation to argue that groups could be units of selection. They also concluded that natural selection of groups could lead to adaptive evolution of groups.

By contrast, Williams's (1966) critique of group selection addressed the replicator question. Groups cannot replicate (reproduce) in the required way, so groups cannot undergo adaptive evolution. With this argument, Williams successfully refuted the conclusion that groups could undergo adaptive evolution, but he did not refute the conclusion that groups could be subjected to selection. Nevertheless, Williams's success against group selectionists helped shift the focus of the debate onto replicators, leading some authors to conclude that replicators were actually the units of selection. The perceived success of Williams's arguments, combined with the success of population genetics (as discussed in Section 5.2) made the gene-selectionist view very popular (Robert, 2004; Jablonka and Lamb, 2005).

Perhaps the most influential arguments in favor of gene selection came from Dawkins (1976, 1980) who actually addressed all four of Lloyd's (2005) questions, but placed the greatest weight on the replicator question. Dawkins's arguments and their serious problems were discussed in Chapters Three and Four, but they are relevant here as well, so I must discuss them again, though with greater emphasis on his account of gene selection rather than on heredity or variation. Most relevant to the units of selection debate is Dawkins's distinction between replicators (nucleic acid-based genes), and vehicles (organisms). In Dawkins's terminology, vehicles interact with their environments in ways that affect whether the replicators get replicated and passed on to the next generation. Although he describes the vehicles as the things that interact with the environment, Dawkins argues that since genes are the only things that persist across generations, genes are the beneficiaries of any adaptations. Therefore, he concludes, genes are the units of selection.

One of the most serious problems for Dawkins' argument is that selection requires interaction between selector and selectee, which means that in order to be a unit of natural selection, it is necessary that the selectee interact with its environment. Genes interact with other intracellular components, but by themselves they do not interact with the larger environment in the required way, and so genes cannot be units of selection. Therefore, whole, integrated organisms (vehicles) are units of selection because whole organisms are what interact with the environment and either survive or die.

Dawkins understood this concern, and responded by arguing that genes interact with their environments via their construction of organisms. That is, phenotypes interact with the environment, but what get selected are genes responsible for those phenotypes. As discussed in greater detail in Section 3.3, this gene-centric argument falsely assumes that a living organism is

a mere support system for genes. While it may be an interesting and occasionally useful heuristic to think of an organism as nothing more than the genes' way to make more genes, a useful heuristic should not be confused with reality (Gould, 2002; Robert, 2004). The reality is that it is incorrect to pick out genes (or any other proper part of an organism) as the causal starting place or sole beneficiary of a complex, dynamic, and interconnected living system. Organisms may not be able to live without their genes, but they also cannot live without their membranes or their precisely regulated cytoplasmic contents or any other major cellular components. In other words, genes are just one part of an integrated, living system in which every part works to sustain every other part, and so it is misleading to claim that genes make organisms (Robert, 2004; Jablonka and Lamb, 2005). This also means that genes are not the sole beneficiaries of adaptations as Dawkins claimed. Genes do not persist across generations any more or less than the rest of an organism, and so they do not have the special status that Dawkins awards them (see also Section 3.4).

S.J. Gould (2002) diagnosed the problems with Dawkins's gene selectionism as a tendency to focus on certain measurable *results* of evolution by natural selection instead of focusing on natural selection itself. For example, one measurable result of natural selection acting on reproducing organisms is a change in allele frequencies in the population to which the organisms belong. Population genetics is in the business of tracking such changes, and since population genetics formed the basis of the Modern Synthesis, the concepts and tools used in population genetics ended up being the concepts and tools used for understanding biological evolution more generally (Gould, 2002; Robert, 2004). For example, "fitness," which is defined by population genetics as an individual's genetic contribution to the next generation, has often been used as a synonym for natural selection, and a common definition of biological evolution is

change in allele (or gene) frequencies in a population over time (e.g. Dobzhansky, 1951). As Gould (2002) points out, tracking individual fitness and populational changes in allele frequencies are useful techniques, but they do not provide a sufficient account of adaptive evolution by natural selection. In particular, such measurements do not distinguish between changes that result from differential natural selection versus those that are mere accidents having nothing to do with differences in phenotypic characteristics between individuals. In other words, population genetics focuses on what happens after natural selection has already occurred—reproduction—rather than on the causal interactions between individuals and their environment that constitute natural selection. While population genetics tracks differential contributions to the next generation, natural selection is more concerned with *how* individuals managed to survive long enough to contribute (Gould, 2002). Therefore, fitness is concerned with hereditary reproduction and the replicator question rather than natural selection and the interactor question.

Gould (2002) was not the only modern author to focus on causal interaction in characterizing natural selection, although he was one of the few to do so consistently. One example of inconsistent focus on causal interaction is Hull (1980). Hull largely followed Dawkins's (1976) approach, but sought to emphasize the importance of causal interaction in natural selection by referring to organisms as "interactors" rather than "vehicles." According to Hull's (1980) formulation, a replicator is "an entity that passes on its structure, directly in replication," while an interactor is "an entity that directly interacts as a cohesive whole with its environment in such a way that replication is differential." With these definitions in place, Hull (1980) defines natural selection as "a process in which the differential extinction and

proliferation of interactors cause the differential perpetuation of the replicators that produced them" (p. 318).[51]

While Hull (1980) is correct to emphasize causal interaction in his account of natural selection, he does not distinguish fully between the process of natural selection and the process of reproduction with heredity. While it is true that long-term biological evolution can only proceed when natural selection is followed by reproduction with heredity, natural selection and hereditary reproduction nevertheless are separate processes and make separate contributions to biological evolution. To put the problem another way, recall that evolution by natural selection is supposed to explain why organisms appear so well adapted to their environments. However, Hull's characterization of natural selection as "differential extinction and proliferation of interactors" does not explicitly tell us that the individuals that survived usually were the ones that had traits that helped them survive in their environments. In this way, Hull (1980) makes a mistake similar to Lewontin (1970) in assuming that differential survival and reproduction is sufficient for natural selection. As Mitchell (1987) explained in her critique of Lewontin (1970), natural selection is not just differential survival, but rather differential survival due to differential interactions between individuals and their environment. Without this added element, the differential survival and eventual reproduction (if any) could be due entirely to chance rather than adaptedness, and if survival and reproduction are due only to chance, then the amazing adaptedness of organisms is left unexplained.

Hull's (1980) mistake of conflating two separate processes is already there in his definition of an interactor. It seems fine for Hull to say that interactors interact with their

[51] Notice that Hull describes natural selection as a two-step process and that each step works with a different entity (or set of entities). The first step involves interactors while his second step involves replicators. Since there is no single entity involved in both steps, Hull concludes that there actually is no unit of selection, that is, no single kind of entity on which selection acts.

environments, but when he adds that they do so "in such a way that replication is differential," he is not saying something general about interactors, but rather something contingent about interactors that are also reproducers. If a reproducer (e.g. a typical biological organism) has beneficial interactions with its environment, then the reproducer is more likely to successfully reproduce when it tries to reproduce. However, it is perfectly fine for a non-reproducing, living individual to also have beneficial interactions with its environment despite its inability to reproduce. Therefore, if we want to have a proper account of natural selection itself, then we must not include its contingent (though evolutionarily very interesting and important) connection to the separate process of hereditary reproduction.

Hull, et. al. (2001) try to update Hull (1980) and to provide a general account of selection by drawing on a variety of selection processes observed in nature. However, their account still ends up conflating natural selection with other processes. In fact, Hull, et.al. (2001) include all three processes when they define selection as "repeated cycles of replication, variation, and environmental interaction so structured that environmental interaction causes replication to be differential" (p. 513). In other words, Hull, et. al. (2001) collapse *all* of evolution by natural selection into just natural selection. They admit as much in the very next sentence when they write, "The net effect is the evolution of the lineages produced by this process." While it is true that all three processes working in tandem will result in adaptive evolution, it is a mistake to pack all three processes into natural selection.

A similar mistake shows up in the later of work Ernst Mayr. For example, Mayr (2001a) writes that natural selection (rather than the larger process of adaptive evolution) is a two-step process, the first of which includes "all the processes leading up to the production of a new zygote" that are responsible for generating variation. Then, "At the second step, that of selection

(elimination), the 'goodness' of the new individual is constantly tested" (p. 119). The problem here is that step one—all of the processes leading up to production of a new zygote—includes hereditary reproduction and generation of new variation, both of which are separate processes from natural selection. Natural selection only really occurs during the second step, which Mayr himself refers to as "selection." It is curious that Mayr chose to include the first step in his account of natural selection when he seemed to understand that the second step was where selection actually happens. In an earlier work, Mayr (1976) makes the far more reasonable claim that biological *evolution* (rather than just natural selection) is a two-step process, where the first step is the production of variation and the second is natural selection.[52] Perhaps this means that Mayr's later collapse of all three processes into natural selection was an honest mistake. Unfortunately, the problematic formulation is repeated in Mayr (2001b) when he writes, "Now that it is realized that chance in evolution is part one of the two-step nature of the process of natural selection, the processes of selection or elimination during the second step of natural selection can make use of the positive contribution made by random variation at the first step" (p.493). Again we see Mayr explicitly separating natural selection into two steps, with selection happening only in the second step, and so the change seems to be intentional.

If the problematic change in Mayr's later thinking was intentional, then I am unsure what prompted it. Perhaps Mayr (2001a, 2001b) intended to use the source of variation (blind mutations and recombination) to characterize natural selection as natural, rather than some other, non-natural kind of selection. If so, the move seems unjustified. The 'natural' part of natural selection refers not to the source of variation, but to the interactions that all organisms have with their environments in their struggle to survive. Organisms whose variability was generated

[52] This is better for his account of natural selection, but I disagree with the way he combines reproduction and generation of variation into a single process.

intentionally by human genetic engineers—and therefore is unnatural in the relevant sense—could then be placed in the wild and subjected to natural selection. As Mayr (1976) himself reported in his history of evolutionary thinking, "Darwin was satisfied…to take for granted the existence of genetic variability and the replenishment of genetic variability as the source material for natural selection" (p. 273). And in Mayr's (1982) chapter on Darwin's conception of natural selection he writes, "Variation and its inheritance belong to the subject matter of genetics, and Darwin's assumptions and theories [on variation and inheritance] will be analyzed in detail [in a separate chapter]" (p. 491). Darwin did not invoke a specific source of variation to clarify the naturalness of natural selection and to distinguish it from artificial selection. Artificial selection occurs when humans (or some other intelligent beings) precisely control which individuals in a population are allowed to survive and to breed, all in order to promote certain characteristics preferred by the selectors. Natural selection, by contrast, does not have an intelligent chooser doing the selecting. It is just what tends to happen when organisms struggle to survive in their environments. Therefore, a proper account of natural selection by itself need not and should not include a source of variation. An account of ongoing, adaptive evolution should incorporate a source of variation, but an account of natural selection by itself should not.

A different approach to characterizing natural selection (one that does not seek to identify units of selection) comes from Matthen and Ariew (2002). They argue that natural selection is not actually a causal process at all, but rather a mere statistical phenomenon indistinguishable from drift. They start their argument by distinguishing two concepts of fitness, which they call vernacular fitness and predictive fitness, and then argue that neither can be a cause of evolution. As Matthen and Ariew (2002) explain, predictive fitness is "a statistical measure of evolutionary change, the *expected* rate of increase (normalized relative to others) of a gene, a trait, or an

organism's representation in future generations, or, on another interpretation, its *propensity* to be represented in future generations, suitably quantified and normalized" (56, emphases in original). Since predictive fitness is a statistical measure rather than a physical entity or force, it is not the sort of thing that has causal power, so Matthen and Ariew correctly conclude that predictive fitness cannot cause of evolutionary change.

Vernacular fitness, on the other hand, is "an organism's overall competitive advantage traceable to heritable traits," and "According to one standard way of understanding natural selection, vernacular fitness—or rather the variation thereof—is a *cause* of evolutionary change" (Matthen and Ariew, 2002, p. 56, emphasis in original). If that way of understanding natural selection is correct, then natural selection cannot be a causal contributor to adaptive evolution because variation itself cannot be a cause of anything. However, that way of understanding natural selection is mistaken. First of all, Matthen and Ariew's account of vernacular fitness, which they borrow from Kitcher (1984), states that competitive advantage must be "traceable to heritable traits." That addition is necessary for ongoing adaptive evolution by natural selection, but is not necessary for natural selection by itself. If the goal is to understand whether natural selection can be invoked as part of a causal explanation for adaptive evolution, then we must be careful to isolate the role natural selection is supposed to play in our explanations, and we must not expect natural selection to play any roles reserved for hereditary reproduction or generation of variation.

That brings us to Matthen and Ariew's (2002) other reason for rejecting natural selection as a causal process. Even if we remove the part about being traceable to inheritable traits, vernacular fitness (and adaptedness) is inherently qualitative because it cannot quantify exactly how much of an individual's survival (or reproduction) is due to its adaptedness and how much

comes from pure luck. Matthen and Ariew argue that if natural selection is a causal process, then we should be able to separate its causal contributions to survival and reproduction from all other contributions, much like physical forces can be isolated in Newtonian Mechanics (e.g. Sober, 1984) Since this does not appear to be possible, they conclude that natural selection is not a causal process, and it cannot be used as part of a causal explanation of adaptive evolution. By contrast, predictive fitness, which simply tracks reproductive rates and marks no distinction between adaptive success and drift, is rigorously quantitative. Because of this, Matthen and Ariew (2002) believe that our conception of natural selection should shed any particular physical realizations in order to make it more compatible with mathematical modeling (similar arguments are found in Beatty, 1984; and Darden and Cain, 1989). To that end, Matthen and Ariew (2002) rely on equations from population genetics that assign different growth rates to different parts of a population (e.g. genes) to determine how the makeup of the population changes over time. Different growth rates of genes will cause their frequencies in the population to change and, as Matthen and Ariew explain, "The resulting differential growth of gene frequencies is natural selection" (p. 56). This purely statistical approach "*defines* selection as what happens to the parts of a population when these parts have different growth rates…This definition of selection is mathematical in nature, and independent of the particular causal laws that produce growth" (p. 73-74, emphasis in original).

Finally, Matthen and Ariew (2002) argue that the above account not only captures natural selection, but all of biological evolution. As they explain, "In this way of looking at things, the distinction between evolution (the total change of gene frequencies due to all causes), and natural selection (the portion of evolution due to differences in competitive advantage) is unmotivated. Natural selection is…the aggregative result over time of differential growth rates in a

population" (p. 78). The central problem with Matthen and Ariew's (2002) account of evolution is that it does not provide an explanation for adaptive evolution. While it is true that their proposed account of natural selection and overall evolution is more quantitative, it only captures one aspect of adaptive evolution. Their account can tell us that there was differential reproduction, but since it does not include causal interaction between individuals and their environments, their account does not include any explanation for why reproduction was differential. This is similar to Gould's (2002) critique of gene selectionists and to Mitchell's (1987) critique of Lewontin (1970). Both complained about an over-emphasis on the sorting effects of natural selection rather than the causal reasons for the sorting.

To be fair, Matthen and Ariew (2002) do not completely ignore the underlying causal processes, but they downplay the central contribution those processes make to a full, explanatory theory of adaptive evolution. When Matthen and Ariew (2002) provide their standard formulation of evolution by natural selection, they refer to the physical instantiation of the population, including all phenotypic traits and the larger environment, as the "substrate" that instantiates a process of evolution. Every process of adaptive evolution must occur in a substrate, but Matthen and Ariew believe that our theory of adaptive evolution by natural selection is independent of any substrate. I maintain, roughly following the likes of Mitchell (1987), Gould (2002), Millstein (2002), Bouchard and Rosenberg (2004), and Rosenberg and Bouchard (2005), that the complex causal interactions happening in the substrate are absolutely central to evolutionary theory. When biologists talk about natural selection in action, they are talking about what's happening in the substrate, that is, in the daily lives of individual organisms and their complex causal interactions with their environments.

As explained in Chapter One, scientists seek to understand and explain the underlying causal processes that produce an interesting regularity or other phenomenon, not just describe the final results (Mitchell, 2002; Casadevall and Fang, 2009). Population genetics and statistical models play important roles in describing and quantifying some of the changes that result from evolution by natural selection, but the full explanation requires contributions from physiology, ecology, and the other life science sub-fields that study how organisms interact with their environments. Those are the fields that explain why some individuals survive while others do not, and so they are indispensable in understanding how adaptations evolve. As explained in Section 5.2, the Modern Synthesis was a successful theory of evolution because of the way it reconciled population genetics with natural selection and the work done by naturalists,[53] cell biologists and paleontologists. Even though population genetics achieved some successes before the Modern Synthesis was formed, by itself it could not provide a full theory of evolution. Only when the explanations coming from those other disciplines were included was there an acceptable theory of evolution, and this still is true today. It's still true that evolutionary theory successfully explains adaptedness (and diversity) in large part because of the contributions from the areas of biology that explore the complex causal reasons why some individuals are better able to survive than others. And while Matthen and Ariew (2002) may be correct that natural selection is difficult to describe and measure as a force analogous to forces described in Newtonian mechanics, that does not mean natural selection is not a causal process. It just means that we cannot expect quantitatively precise measurements of adaptedness and selection in our studies of natural selection in action. Therefore, accounts such as found in Matthen and Ariew (2002), Beatty (1984), and Darden and Cain (1989) that claim natural selection is best

[53] Naturalists studied and catalogued organisms as they existed in the wild. Today, they likely would be called "ecologists" or "taxonomists."

understood through population genetics, or that subsume all of evolutionary theory into population genetics, do not work.

In summary, my own account of natural selection focuses on interactions between an individual and its environment, and makes no claim that natural selection, by itself, can produce adaptive evolution. Natural selection by itself requires only that at least one living individual exists in an environment, and that the physical characteristics of that environment must causally interact with the individual in ways that depend on the individual's own physical characteristics as well as those of the environment. Further, it must be the case that some potential interactions between environment and individual could be disruptive to the individual (e.g. cause it to break apart or stop functioning) while other potential interactions are beneficial the individual (e.g. support its structural integrity or allow its continued functioning). Whenever these sorts of causal interactions occur, an individual is being naturally selected for (or against) continued survival and functioning. This account of natural selection can now be used to help determine whether there is a biological mechanism that has natural selection as its activity.

5.4 Is Natural Selection a Mechanism?

The biological literature on evolution contains many casual references to natural selection as a mechanism, but only recently has that attribution been seriously examined. Skipper and Millstein (2005) introduce the issue very well:

> "Actually, evolutionary biologists call natural selection, and the other evolutionary mechanisms, many things, which they all seem to think amounts to the same thing. Natural selection is a 'cause,' a 'force,' a 'process,' a 'mechanism,' a 'factor'. Sometimes, natural selection is called a 'principle,' or 'concept,' but when the explication continues, cause, force or mechanism talk is apparent. We think there is no question that contemporary evolutionary biology exemplifies the view that natural selection is a mechanism. The interesting and harder question to answer is philosophical, namely, 'What is the nature of the mechanism of natural selection?'" (p. 329).

Unfortunately, in their analysis of the issue, Skipper and Millstein (2005) end up using an account of natural selection that includes both hereditary reproduction and generation of new variation. Their description of the structure of natural selection (p.330) includes seven steps, only the second of which, "Interaction," is actually natural selection because that is the step at which an individual organism interacts with its environment. Their first step, "Initial Conditions," includes both variation and heredity, while steps three through seven are all "Effects" of what happened after all the interaction at step two. Like Hull, et. al. (2002), Skipper and Millstein (2005) include all of adaptive evolution into their account of natural selection. They then use that account of natural selection (which is really the whole process of evolution by natural selection) to conclude that natural selection is not a mechanism because there are no identifiable entities that carry out all those steps. I agree that there is no single mechanism that could be responsible for the whole process of adaptive evolution, but that does not tell us whether natural selection by itself is a mechanism.

In reply to Skipper and Millstein (2005), Barros (2008) argues that natural selection does, or at least *can*, meet the criteria for mechanism laid out by MDC (2000). Since he is relying on MDC's (2000) account of mechanisms, Barros (2008) must identify the activity and the entities that constitute his proposed mechanism of natural selection. As an example, he describes a predator-prey relationship between crabs and sea snails in which presence of the crab predator seems to have resulted in adaptive changes in snail shells. In arguing that natural selection is a mechanism, Barros (2008) first argues that predation by a crab is accomplished by a mechanism, namely its claw's crushing mechanism. "This characterization of crab predation as a mechanism is consistent with the MDC account of mechanisms. The entities (the crab's claw and the snail's shell) and activities (crushing) are organized in a way that they may be productive of a change (a

crushed shell) from the start of the attack to the end of the attack" (p. 314). With this much I agree. Even when we throw in my own requirements that the entities work non-aggregatively and that the activity serve a biological function (see Chapter Two), the crab's claw and its crushing of snail shells is a good example of a mechanism.

In order to extend his argument to natural selection rather than just predation with crab claws, Barros (2008) believes he must move to the population level rather than the individual level. He does this because he believes that natural selection is supposed to explain how adaptive traits evolve, and adaptive evolution is something that happens to populations rather than to individuals. Barros concludes that "Natural selection thus is a multistage mechanism that tends to increase the prevalence of a beneficial trait within a population over time" (p. 317). This again collapses all of evolution into just natural selection, but as I have argued repeatedly in this chapter, natural selection by itself is *not* supposed to explain long-term adaptation. Rather, the explanation for long-term, ongoing adaptation comes from evolution by natural selection, which includes hereditary reproduction and generation of new variation in addition to natural selection. Barros's (2008) specific problem is that the population-level component of his characterization of natural selection involves variation and hereditary reproduction rather than just natural selection, and so he is considering the wrong activity for his proposed mechanism of natural selection.

Although I strongly disagree with the characterizations of natural selection used by Skipper and Millstein (2005) and Barros (2008), their arguments nevertheless contain the chief reasons I believe that natural selection is not a biological mechanism. In order for there to be a biological mechanism of natural selection, the activity and entities must match the criteria laid out in Chapter Two. The activity must be a biological function that supports the continued life of

which the proposed mechanism is a part, and the entities must be part of the living process they are supporting and must achieve their function nonaggregatively. I believe that even when we are working with the correct activity for natural selection, on both points natural selection fails to meet the criteria for a biological mechanism.

As for the entities, since natural selection always involves an interaction between an individual and its environment, any instance of natural selection involves some entities from the individual and some from its environment, and so the entities involved in the selection process will never all be part of a single organism. Therefore, the entities responsible for natural selection are not all part of any one living individual, and so cannot be considered a biological mechanism. Barros tried to resolve this problem, but I think his efforts do not work. In discussing environmental interactions he writes, "The source of the critical environmental factor…need not be only one entity. It could be the sum of the activities of many entities imposing various types of selection pressure. As I use the term here, selective pressure is an abstract activity that refers to the impact that the critical environmental factor has on the population being studied" (317). Barros then goes on to include a very wide variety of environmental factors, such as rainfall and local air temperature, as potential entities in the mechanism of natural selection. Inclusion of so many different factors working in so many different ways makes it impossible to identify any specific entities that are responsible for the activity of natural selection, and so Barros is forced to set his mechanistic account of natural selection at a more abstract level, independent of any specific physical entities. "Indeed," he writes, "the environment in a particular location can be seen as an abstract entity that is the sum of all other entities in that location" (317). In addition to inappropriately packing all of adaptive evolution into just natural selection, Barros is stretching our notion of entities beyond what MDC

(2000) had in mind for mechanisms, and the entities he describes certainly do not fit into my account of biological mechanisms.

As for the activity of natural selection, it is not a biological function and so cannot be the activity of a biological mechanism. A biological function necessarily promotes continuation of life, while natural selection oftentimes is selection *against* continued survival. Selection against continued survival obviously does not promote the continuation of life, and so cannot be a biological function. Even when there is selection for continued survival, any mechanisms at work are ones that have something other than selection as their function. For example, mechanisms might function to acquire food (as in the crab claw example) or to regulate internal temperature, but those functions, by themselves, do not count as selection. When those mechanisms operate in a particular environment, natural selection helps sort the more effective mechanisms from the less effective ones.

When we look out into the environment, we also do not find anything that has natural selection as its function. Most of the external environment is made up of nonliving entities or climatic conditions, and as discussed in Section 2.4, nonbiological entities in the environment do not have any functions at all. They just are what they are. Of course, other biological entities can be part of an individual's environment, but any biological functions they have still are aimed at their own survival. Natural selection of other organisms is not their function. So, biological mechanisms have functions aimed at the organisms of which they are a part, and natural selection happens when those mechanisms operate in a particular environment. Parental care of offspring may be the closest thing there is to a biological mechanism of natural selection. Parental care is an interaction between an individual (the offspring) and part of its environment (its parent), and its function is to continue the ongoing, multi-generational, living process of

which the parent and the offspring are both parts. However, even this kind of natural selection necessarily involves interaction between an individual and some part of the environment, and a single biological mechanism must be part of a single living individual. Therefore, even parental interaction fails to qualify as a biological mechanism of natural selection.

While it is true that natural selection is carried out by causal interactions between an individual and its environment, the activity of natural selection is not a biological function and, therefore, natural selection is not a biological mechanism according to the criteria laid out in Chapter Two. The same may be said about evolution overall. There is no biological mechanism that carries out adaptive evolution of a lineage because the entities and activities do not meet the criteria for biological mechanisms (Skipper and Millstein, 2005). Recall that in Section 1.2, I discussed some other senses of "mechanism" used by scientists and philosophers. For example, scientists (especially those that are not biologists) often call any specifiable physical process a mechanism (Casadevall and Fang, 2009), and sometimes abstract procedures or algorithms are thought of as mechanisms. Perhaps one of these other senses of "mechanism" might apply to natural selection, but natural selection is not a *biological* mechanism.

Summary

In this chapter I showed that natural selection requires nothing more than causal interaction between an individual and its environment that affect the individual's continued survival. It does not require a diverse population of individuals or any competition, though without those factors, not much evolution is likely to occur. If there is more than one individual, and those individuals differ in their interactions with the environment, then there can be differential natural selection in which some individuals continue to survive while others do not.

Then, when differential natural selection is associated in the right way with hereditary reproduction and generation of new variations, a lineage of organisms will undergo adaptive evolution.

I also showed that many characterizations of natural selection found in the literature do not distinguish properly between natural selection itself and the larger process of adaptive evolution. A satisfactory account of adaptive evolution requires all three, but since each process can happen apart from the others, and since each makes a separate contribution to adaptive evolution, it is misleading to include all three under the single label of natural selection.

Finally, I showed that contrary to many casual assertions, natural selection is not a biological mechanism because the activity of natural selection is not a biological function. There may be a more permissive or abstract sense of 'mechanism' that applies to natural selection, but it is not a biological mechanism. For similar reasons, the overall process of adaptive evolution is also not a biological mechanism. There are many biological mechanisms that contribute to the process of adaptive evolution, but adaptive evolution itself is not a biological mechanism.

BIBLIOGRAPHY

Aertson, A. & Michiels, C. (2005) "Diversify or Die: Generation of Diversity in Response to Stress" *Critical Reviews in Microbiology*, 31, 69-78.

Agrawal, A. (2001) "Sexual selection and the maintenance of sexual reproduction" *Nature*, 411, 692-695.

Allen, G. E. (2005) "Mechanism, Vitalism and Organicism in Late Nineteenth and Twentieth-Century Biology: the Importance of Historical Context" *Studies in the History and Philosophy of Biology and the Biomedical Sciences*, 36, 261-283.

Aristotle. (1984) "De Anima" In *The Complete Works of Aristotle* (pp. 641-692). Princeton, New Jersey: Princeton University Press.

Aristotle. (1993) *Posterior Analytics* Oxford University Press.

Armstrong, D. (1983) *What Is a Law of Nature?* Cambridge University Press.

Barros, B. D. (2008) "Natural Selection as a Mechanism" *Philosophy of Science*, 75, 306-322.

Beatty, J. (1994) "The Evolutionary Contingency Thesis" In E. Sober (Ed.), *Conceptual Issues in Evolutionary Biology* MIT Press.

Bechtel, W. & Abrahamsen, A. (2005) "Explanation: A Mechanist Alternative" *Studies in the History and Philosophy of Biology and the Biomedical Sciences*, 421-441.

Behe, M. (1996) *Darwin's Black Box: the Biochemical Challenge to Darwin's Evolution* New York: Free Press.

Behe, M. (2007) *The Edge of Evolution: the Search for the Limits of Darwinism* New York: Free Press.

Bjedov, I., Tenaillon, O., Gérard, B., Souza, V., Denamur, E., Radman, M., et al. (2003) "Stress-Induced Mutagenesis in Bacteria" *Science*, 300, 1404-1409.

Boas, M. (1952) "The Establishment of the Mechanical Philosophy" *Osiris*, 10, 412-541.

Bouchard, F. & Rosenberg, A. (2004) "Fitness, Probability, and the Principles of Natural Selection" *British Journal for the Philosophy of Science*, 55, 693-712.

Brandon, R. (1981) "Biological Teleology: Questions and Explanations" In C. Allen & M. Bekoff (Eds.), *Nature's Purposes: Analyses of Function and Design in Biology* (pp. 79-98). Cambridge: MIT Press.

Brisson, D. (2003) “The Directed Mutation Controversy in an Evolutionary Context” *Critical Reviews in Microbiology*, 29(1), 25-35.

Burkhardt, R. (1984) “The Zoological Philosophy of J.B. Lamarck” In *Zoological Philosophy* R. Burkhardt (Ed.), (pp. xv-xxxix). Chicago: University of Chicago Press.

Cairns, J. & Foster, P. (1991) “Adaptive Reversion of a Frameshift Mutation in *Escherichia coli*” *Genetics*, 128, 695-701.

Cairns, J., Overbaugh, J., Miller, S. (1988) “The Origin of Mutants” *Nature*, 335, 142-145.

Campbell, N. (Ed.). (2004) *Biology* Redwood City, CA: Benjamin/Cummings.

Carnap, R. (1966) *Philosophical Foundations of Physics* New Yourk: Basic Books.

Cartwright, N. (1983) *How the Laws of Physics Lie* Oxford University Press.

Casadavall, A. & Fang, F. (2009) “Mechanistic Science” *Infection and Immunity*, 77(9), 3517-3519.

Cleland, C. & Chyba, C. (2002) “Defining 'Life'“ *Origins of Life and the Evolution of the Biosphere*, 32, 387-393.

Craver, C. (2001) “Role Functions, Mechanisms, and Hierarchy” *Philosophy of Science*, 68, 53-74.

Craver, C. & Bechtel, W. (2007) “Top-Down Causation Without Top-Down Causes” *Biology and Philosophy*, 22, 547-563.

Craver, C. & Darden, L. (2005) “Introduction to Special Issue on Mechanisms” *Studies in the History and Philosophy of Biology and the Biomedical Sciences*, 36, 233-244.

Cummins, R. (1975) “Functional Analysis” *Journal of Philosophy*, 72(20), 741-765.

Cuvier, G. (1825) “Discourse On the Revolutionary Upheavals On the Surface of the Globe and on the Changes Which They Have Produced in the Animal Kingdom” accessed at http://www.victorianweb.org/science/science_texts/cuvier/cuvier-e.htm.

Darden, L. (2002) “Strategies for Discovering Mechanisms: Schema Instantiation, Modular Subassembly, Forward/Backward Chaining” *Philosophy of Science*, 69, S354-S365.

Darden, L. & Cain, J. (1989) “Selection Type Theories” *Philosophy of Science*, 14(1), 106-129.

Darwin, C. (1859) *The Origin of Species* New York: Bantam Classic.

Darwin, C. (1868) *The Variation of Animals andPplants Under Domestication* London: J. Murray.

Dawkins, R. (1976) *The Selfish Gene* Oxford University Press.

Dawkins, R. (1980) *The Extended Phenotype* Oxford University Press.

Dawkins, R. (1983) "Universal Darwinism" In D. S. Bendall (Ed.), *Evolution From Molecules to Men* (pp. 403-425). Cambridge: Cambridge University Press.

DeDuve, C. (2005) "The Onset of Selection" *Nature*, 433, 581-582.

Dennet, D. (1989) *The Intentional Stance* MIT Press.

Dennet, D. (1995) *Darwin's Dangerous Idea* New York: Simon and Schuster.

Dennet, D. (1996) *Kinds of Minds* New York: Basic Books.

Dobzhansky, T. (1951) *Genetics and the Origin of Species* New York: Columbia University Press.

Downes, S. M. (2010). "Heritability" *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/fall2010/entries/heredity/

Falk, A. E. (1981) "Purpose, Feedback, and Evolution" *Philosophy of Science*, 48, 198-217.

Fodor, J. (1991) "You Can Fool Some of The People All of The Time, Everything Else Being Equal; Hedged Laws and Psychological Explanations" *Mind*, 100(1), 19-34.

Foster, P. (2000) "Adaptive Mutation: Implications for Evolution" *BioEssays*, 22, 1067-1074.

Foster, P. (2004) "Adaptive Mutation in *Escherichia coli*" *Journal of Bacteriology*, 186(15), 4846-4852.

Glennan, S. (2002) "Rethinking Mechanistic Explanation" *Philosophy of Science*, 69, S342-S353.

Godoy, V., Gizatullin, F. S., Fox, M. S. (2000) "Some Features of the Mutability of Bacteria During Nonlethal Selection" *Genetics*, 154, 49-59.

Gould, S. J. (2002) *The Structure of Evolutionary Theory* Cambridge: Belknap Harvard University Press.

Gray, R. (1992) "Death of the Gene: Developmental Systems Strike Back" In *Trees of Life* P. Griffiths (Ed.), Dordrecht: Kluwer Academic Publishers.

Griffiths, P. & Gray, R. (1994) "Developmental Systems and Evolutionary Explanation" *Journal of Philosophy*, 91(6), 277-304.

Haldane, J. B. S. (1884) "Life and Mechanism" *Mind*, 9(33), 27-47.

Haldane, J. B. S. (1930) *The Philosophical Basis of Biology* London: Hodder and Stoughton.

Hall, B. (1998) "Adaptive Mutagenesis: a Process That Generates Almost Exclusively Beneficial Mutations" *Genetica*, 102/103, 109-125.

Hamilton, W. D. (1972) "Altruism and Related Phenomena, Mainly in Social Insects" *Annual Review of Ecology and Systematics*, 3, 193-232.

Hempel, C. (1962) "Two Basic Types of Scientific Explanation" In *Philosophy of Science: the Central Issues* M. Curd & J. A. Cover (Eds.), (pp. 685-694). New york: Norton.

Hempel, C. (1965) "The Thesis of Structural Identity" In *Philosophy of Science: the Central Issues* M. Curd & J. A. Cover (Eds.), (pp. 695-705). New York: Norton.

Hull, D. (1980) "Individuality and Selection" *Annual Review of Systematics*, 11, 311-332.

Hull, D. (1984) "Lamarck Among the Anglos" In *Zoological Philosophy* R. Burkhardt (Ed.), University of Chicago Press.

Hull, D., Langman, R., Glenn, S. (2001) "A General Account of Selection: Biology, Immunology, and Behavior" *Behavioral and Brain Sciences*, 24, 511-573.

Jablonka, E. & Lamb, M. (1989) "The Inheritance of Acquired Epigenetic Variations" *Journal of Theoretical Biology*, 139, 69-83.

Jablonka, E. & Lamb, M. (2005) *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life* Cambridge: MIT Press.

Jeffery, W. (2001) "Cavefish as a Model System in Evolutionary Developmental Biology" *Developmental Biology*, 231(1), 1-12.

Kant, I. (2000) *Critique of the Power of Judgment* Cambridge University Press.

Kauffman, S. (2000) *Investigations* Oxford University Press.

Kitcher, P. (1984) *Vaulting Ambition: Sociobiology and the Quest for Human Nature Cambridge* MIT Press.

Kivisaar, M. (2003) "Stationary Phase Mutagenesis: Mechanisms that Accelerate Adaptation of Microbial Populations Under Environmental Stress" *Environmental Microbiology*, 5(10), 814-827.

Korzeniewski, B. (2001) “Cybernetic Formulation of the Definition of Life” *Journal of Theoretical Biology*, 209, 275-286.

Kronfeldner, M. (2007) “Is Cultural Evolution Lamarckian?” *Biology and Philosophy*, 22, 493-512.

Lamarck, J. B. (1809) *Zoological Philosophy* traslated by Burkhardt, R.,1984, University of Chicago Press.

Landman, O. (1991) “The Inheritance of Acquired Characteristics” *Annual Review of Genetics*, 25, 1-20.

Laura M. Grabowski, L., Bryson, D., Dyer, F., Pennock, R. T., Ofria, C. (2011) “Clever Creatures: Case Studies of Evolved Digital Organisms” *Advances in Artificial Life: ECAL 2011*, 276-283.

Lenski, R. & Mittler, J. (1993) “The Directed Mutation Controversy and Neo-Darwinism” *Science*, 259(5092), 188-195.

Lewis, D. (1973) *Counterfactuals* Oxford: Blackwell.

Lewontin, R. C. (1970) “The Units of Selection” *Annual Review of Ecological Systematics*, 1, 1-18.

Lloyd, E. (1986) “Evaluation of Evidence in Group Selection Debates” *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1986(1), 483-493.

Lloyd, E. (2005). “Units and Levels of Selection” *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition). http://plato.stanford.edu/entries/selection-units/ 2005

Low, E. N., Zagieboylo, L., Martino, B., Wilson, E. (2010) “IgA ASC Accumulation to the Lactating Mammary Gland is Dependent On VCAM-1 and alpha4 Integrins” *Molecular Immunology*, 47(7-8), 1608-1612.

Luria, S. E. & Delbrük, M. (1943) “Mutations of Bacteria from Virus Sensitivity to Virus Resistance” *Genetics*, 28, 492.

Machamer, P., Darden, L., Craver, C. (2000) “Thinking About Mechanisms” *Philosophy of Science*, 67, 1-25.

Matthen, M. & Ariew, A. (2002) “Two Ways of Thinking About Fitness and Natural Selection” *The Journal of Philosophy*, 99, 55-83.

Matzke, N. (2007) “The Edge of Creationism” *Trends In Ecology and Evolution*, 22(11), 566-567.

Mayr, E. (1976) *Evolution and the Diversity of Life: Selected Essays* Harvard University Press.

Mayr, E. (1980) "Prologue" In *The Evolutionary Synthesis* E. Mayr & W. Provine (Eds.), (pp. 1-48). Harvard University Press.

Mayr, E. (1982) *The Growth of Biological Thought* Harvard University Press.

Mayr, E. (1998) *This Is Biology* Harvard University Press.

Mayr, E. (2001) "The Philosophical Foundations of Darwinism" *Proceedings of the American Philosophical Society*, 145(4), 488-495.

Mayr, E. (2001) *What Evolution Is* New York: Basic Books.

Medvedev, Z. (1969) *The Rise and Fall of T.D. Lysenko* Columbia University Press.

Millikan, R. (1998) "In Defense of Proper Functions" In *Nature's Purposes: Analyses of Function and Design in Biology* C. Allen, M. Bekoff, G. Lauder (Eds.), (pp. 295-312). Cambridge: MIT Press.

Millstein, R. (2002) "Are Random Drift and Natural Selection Conceptually Distinct?" *Biology and Philosophy*, 17(1), 33-53.

Millstein, R. (2006) "Natural Selection as a Population-Level Causal Process" *British Journal of the Philosophy of Science*, 57, 627-653.

Mitchell, S. (1987) "Competing Units of Selection? A Case of Symbiosis" *Philosophy of Science*, 43(3), 351-367.

Mitchell, S. (1997) "Pragmatic Laws" *Philosophy of Science*, 4, S467-S479.

Mitchell, S. (2000) "Dimensions of Scientific Law" *Philosophy of Science*, 67(2), 242-265.

Mitchell, S. (2002) "*Ceteris Paribus* – An Inadequate Representation for Biological Contingency" *Erkenntnis*, 57, 329-350.

Monton, B. (2009) *Seeking God in Science* Broadview Press.

Muenzinger, C. F. (1935) "Mechanism, Vitalism, and the Organismic Hypothesis" *Philosophy of Science*, 2(4), 518-520.

Nagel, E. (1977) "Teleology Revisited" *Journal of Philosophy*, 84, 261-301.

Neander, K. (1991) "Functions as Selected Effects: The Conceptual Analyst's Defense" *Philosophy of Science*, 58, 168-184.

O'Malley, M. & Dupré, J. (2007) "Size Doesn't Matter: Towards a More Inclusive Philosophy of Biology" *Biology and Philosophy*, 22, 155-191.

Oda, Y., DeVires, Y., Forney, L., Gottschal, J. (2001) "Acquisition of the Ability for *Rhodopseudomonas palustris* to Degrade Chlorinated Benzoic Acids as the Sole Carbon Source" *Microbiology Ecology*, 38, 133-139.

Oparin, A. (1964) *Life: Its Nature,Oorigin, and Development* Academic Press.

Oyama, S. (2000) *The Ontogeny of Information: Developmental Systems and Evolution* Duke University Press.

Paley, W. (1802) *Natural Theology, or Evidences of the Existence and Attributes of the Deity Collected From the Appearances of Nature* Daniel and Samuel Whiting.

Plantinga, A. (1991) "When Faith and Reason Clash: Evolution and the Bible" *Christian Scholar's Review*, 21, 8-32.

Railton, P. (1978) "A Deductive-Nomological Model of Probabilitistic Explanation" In *Philosophy of Science: the Central Issues* M. Curd & J. A. Cover (Eds.), (pp. 746-765). New York: Norton.

Redfield, R. (1988) "Evolution of Bacterial Transformation: Is Sex With Dead Cells Ever Better Than No Sex at All?" *Genetics*, 119, 213-221.

Rheinberger, H. (2000) "Gene Concepts: Fragments From the Perspective of Molecular Biology" In *The Concept of the Gene in Development and Evolution. Historical and Epistemological* P. Beurton, R. Falk, H. Rheinberger (Eds.), (pp. 219-239). Cambridge University Press.

Rheinberger, H. & Müller-Wille, S. (2010). "Gene" *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/spr2010/entries/gene/

Robert, J. (2004) *Embryology, Epigenesis, and Evolution* Cambridge University Press.

Romero, A. & Green, S. M. (2005) "The End of Regressive Evolution: Examining and Interpreting the Evidence from Cave Fishes" *Journal of Fish Biology*, 67, 3-32.

Rosenberg, A. & Bouchard, F. (2005) "Matthen and Ariew's Obituary for Fitness: Reports of its Death Have Been Greatly Exaggerated" *Biology and Philosophy*, 20, 343-353.

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

Ruben, D. (1990) "Arguments, Laws, and Explanation" In *Philosophy of Science: the Central Issues* M. Curd & J. A. Cover (Eds.), (pp. 720-745). New York: Norton.

Rupert, R. (2008) "*Ceteris Paribus* Laws, Component Forces, and the Nature of Special-Science Properties" *Noûs*, 42(3), 349-380.

Salisbury, F. & Ross, C. (1992) *Plant Physiology* Wadsworth.

Salmon, W. (1984) *Scientific Explanation and the Causal Structure of the World* Princeton University Press.

Salmon, W. (1997) "Causality and Explanation: A Reply to Two Critiques" *Philosophy of Science*, 64, 461-477.

Sarkar, S. (2005) "Neo-Darwinism and the Problem of Directed Mutations" In *Molecular Models of Life: Philosophical papers on molecular biology* (pp. 287-301). Cambridge: MIT Press.

Schrödinger, E. (1945) *What Is Life?* Cambridge University Press.

Segré, D. & Lancet, D. (2000) "Composing Life" *EMBO Reports*, 1(3), 217-222.

Segré, D., Lancet, D., Kedem, O., Pilpel, Y. (1998) "Graded Autocatalysis replication Domain (GARD): Kinetic Analysis of Self-Replication in Mutually Catalytic Sets" *Origins of Life and Evolution of the Biosphere*, 28, 501-514.

Segré, D., Shenhav, B., Lancet, D. (2001) "The Molecular Roots of Compositional Inheritance" *Journal of Theoretical Biology*, 213, 481-491.

Shapiro, J. A. (1998) "Thinking About Bacterial Populations as Multicellular Organisms" *Annual Review of Microbiology*, 52, 81-104.

Skipper, R. & Millstein, R. (2005) "Thinking About Evolutionary Mechanisms" *Studies in the History and Philosophy of Biology and the Biomedical Sciences*, 36, 327-347.

Sniegowski, P. D., Gerrish, P., Johnson, T., Shaver, A. (2000) "The Evolution of Mutation Rates: Separating Causes from Consequences" *Bioessays*, 22, 1057-1066.

Sniegowski, P. D. & Lenski, R. (1995) "Mutation and Adaptation: The Directed Mutation Controversy in Evolutionary Perspective" *Annual Review of Ecology and Systematics*, 26, 553-578.

Sober, E. (1984) *The Nature of Selection: Evolutionary Theory in Philosophical Focus* Cambridge: MIT Press.

Sober, E. (1997) "Two Outbreaks of Lawlessness in Philosophy of Biology" *Philosophy of Science*, 64, S458-S467.

Sober, E. (2003) "Philosophy of Biology" In *A Blackwell Companion to Philosophy* (pp. 317-343). Oxford: Blackwell.

Sober, E. & Wilson, D. S. (1994) "A Critical Review of Philosophical Work on the Units of Selection Problem" *Philosophy of Science*, 61(4), 534-555.

Steele, E. J. (1981) *Somatic Selection and Adaptive Evolution* University of Chicago Press.

Steele, E. J.; Lindley, R.; Blanden, R. (1998) *Lamarck's Signature* Perseus Books.

Sterelney, K. (2001) *Dawkins vs. Gould: Survival of the Fittest* Cambridge: Totem Books.

Stocking, G. W. (1962) "Lamarckism in American Social Science: 1890-1915" *Journal of the History of Ideas*, 23(2), 239-259.

Stolovicki, E., Dror, T., Brenner, N., Braun, E. (2006) "Synthetic Gene Recruitment Reveals Adaptive Reprogramming of Gene Regulation in Yeast" *Genetics*, 173, 75-85.

Szathmáry, E. & Smith, J. M. (1997) "From Replicators to Reproducers: the First Major Transitions Leading to Life" *Journal of Theoretical Biology*, 187(4), 555-571.

Tabery, J. G. (2003) "Synthesizing Activities and Interactions in the Concept of a Mechanism" *Philosophy of Science*, 71, 1-15.

Tooley, M. (1977) "The Nature of Laws" *Canadian Journal of Philosophy*, 7, 667-698.

Torkelson, J., Harris, R., Lombardo, M., Nagendran, J., Thulin, C., Rosenberg, S. M. (1997) "Genome-wide Hypermutation in a Subpopulation of Stationary-phase Cells Underlies Recombination-dependent Adaptive Mutation" *The EMBO Journal*, 16, 3303-3311.

Turner, J. E. (1940) "The Distinction Between 'Mechanics' and 'Mechanism'" *Philosophy of Science*, 7(1), 49-55.

Williams, G. (1966) "Adaptation and Natural Selection" In *Conceptual Issues in Evolutionary Biology* E. Sober (Ed.), Cambridge, MA: MIT Press.

Wimsatt, W. (1974) "Complexity and Organization" *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1972, 67-86.

Wimsatt, W. (1986) "Forms of Aggregativity" In *Human Nature and Natural Knowledge* A. Donagan, A. Perovich, M. Wedin (Eds.), (pp. 259-291). Dordrecht: Reidel.

Wimsatt, W. (1997) "Aggregativity: Reductive Heuristics for Finding Emergence" *Philosophy of Science*, 64, S372-S384.

Wimsatt, W. & Griesemer, J. R. (2007) "Reproducing Entrenchments to Scaffold Culture: The Central Role of Development in Cultural Evolution" In *Integrating Evolution and Development: From Theory to Practice* R. Sansom & R. Brandon (Eds.), (pp. 227-323). Cambridge: MIT Press.

Winther, R. (2000) “Darwin On Variation and Heredity” *Journal of the History of Biology*, 33, 425-455.

Woodward, J. (1997) “Explanation, Invariance and Intervention” *Philosophy of Science*, 64, S26-S41.

Woodward, J. (2002) “There Is No Such Thing As a *Ceteris Parbus* Law” *Erkenntnis*, 57, 303-328.

Woodward, J. (2002) “What Is a Mechanism? A Counterfactual Account” *Philosophy of Science*, 69, S366-S377.

Woodward, J. (2003) *Making Things Happen: a Theory of Causal Explanation* Oxford University Press.

Woodward, J. & Hitchcock, C. (2003) “Explanatory Generalizations” *Noûs*, 37(1), 1-24.

Wright, L. (1973) “Functions” *Philosophical Review*, 82, 139-168.

Wright, L. (1976) *Teleological Explanations: An Etiological Analysis of Goals and Functions* Berkeley: University of California Press.

Young, M. & Edis, T. (Eds.). (2004) Why Intelligent Design Fails: A Scientific Critique of the New Creationism Rutgers University Press.