Introduction

What is the philosophy of biology?

Philosophy asks two kinds of questions

Philosophy, Aristotle wrote, begins with wonder. And, for a long time, philosophy meant the same thing as science. Indeed, in some universities, physics is still called "natural philosophy," and philosophy is taught in the department of "moral science." The reason is not hard to see. The history of Western philosophy is the history of a discipline that has been "spinning off" sciences since about 300 BC when Euclid wrote the *Elements* and established the separate discipline of mathematics. It was only much later, in the seventeenth century, that physics finally established itself as a discipline distinct from philosophy, followed in the late eighteenth century by chemistry, and, as we will argue in the next chapter, by biology as late as 1859, when Darwin published *On the Origin of Species*. This process continues, for there are other disciplines, still in the process of spinning themselves off from philosophy. As the sciences establish their separate existences, two questions arise: Do the sciences leave anything to philosophy when they "spin off," and, if so, why do they leave unfinished business to philosophy? The answer to the first question is obvious. Each of the sciences leaves to philosophy issues that they might be expected to answer but have not. Consider the question of what a number is. A number is not after all a numeral, which is just the symbol we use to name a number. For "2," "II," "two," "dos," and "dho" all name the very same number, in Arabic, Roman, English, Spanish, and Hindi notation. We may hold, as many followers of Plato still do, that numbers are "abstract objects," or that there are no such things and that numbers are mental constructs. But it will be in vain to look to mathematics for an answer to the question of what a number is. That question has remained one for philosophical inquiry since Plato. Or consider the question of what time is. Time is a variable in many of the most important physical laws. Newton's second law, for example, tells us that force equals mass×acceleration, $F = ma$, where acceleration is defined as the rate of change of velocity with respect to time, $a = \mathrm{d}v/\mathrm{d}t$. But the question of what is time, t in the equation, has remained unanswered in physics and left to philosophers.

Biology too has left questions that philosophy addresses. In fact, the questions biology leaves to philosophy are hard to avoid and of great interest beyond biology (and beyond philosophy for that matter). This is part of the reason that the philosophy of biology has become one of the liveliest and most publicly visible of philosophy's subdisciplines. Another is that the questions biology leaves to philosophy are the most immediately relevant to many distinctively human concerns. For example, it is to biology that many look for insight into "human nature." It is biology that appears to address the question of what is "life" and whether things have a meaning or purpose beyond the merely physical and chemical processes that constitute them. Now, biological science itself does not tell us whether it has the power to answer these questions. And for that reason there are lively debates about biology's scope and limits, its authority to answer such perennial questions of deep human concern. These questions about biology's scope and limits are clearly philosophical ones.

Like the other natural sciences, biology is an experimental discipline, and, as such, it is a fallible one. For experiments, observation, and collecting data can never establish the truth of a theory with perfect certainty. Like other scientists, however, biologists have the confidence that though their findings are always subject to revision and improvement, their method—the scientific method—is the right one, indeed, the only way to assure the increasing reliability of their results. But there are disputes within biology, and between biologists and other scientists, both about what the "scientific method" is and about whether various research programs and their results honor that method. Then there are disputes about whether and why the application of the scientific method in biology differs from that in the physical sciences. And, finally, there are disputes about whether there is any such thing as *the* scientific method, with the emphasis on the uniqueness suggested by the definite article. All of these issues are well and truly part of the agenda of the philosophy of biology and, of course, the philosophy of science generally. For the sciences cannot themselves answer questions about the warrant of their own methods, the justification of their modes of research, and the adequacy of each discipline's distinctive approaches to its own and other disciplines' domains. A physicist's argument that biology should be more like physics, or a chemist's claim that biological facts need to be explained by chemistry, cannot be settled by experiment and observation, if they can be settled at all. These questions are the purview of the philosophy of these sciences and the philosophy of science in general. This does not mean that scientists have no right to express views about these matters or that only philosophers of science are qualified to do so. It means merely that when informed participants debate these issues, they are engaged in a philosophical dispute.

Recall now our second question. If there are questions that the sciences cannot answer, why do such questions exist? This can be construed as a question about the limits of science. It is well known that many people reject the findings and theories of natural science in favor of other beliefs, often religious ones, and often with the accompanying claim that some facts of the world are forever beyond the reach of science. No science is more often met with claims of this sort than biology. Questions about the meaning of life are often said to lie in this unreachable domain. Some go even further, arguing that questions about the origin of life, or of the human species, lie there also. Further, there are social and behavioral scientists, and scholars in the humanities too, who deny the relevance of biology to their research questions, for example questions having to do with the causes of human behavior or the foundation of ethics. Now it would seem that those who hold that biology, or other natural sciences, cannot answer certain questions owe an account of why not, as of course do those who argue that science *can* answer them. And these accounts of the limits of science, or of the absence of limits, will be philosophical arguments, as traditionally understood.

Like biology, philosophy is divided into subdisciplines: metaphysics studies the basic kinds of things, processes, and properties in the universe, and addresses questions about them such as: What are numbers? Does God exist? Are all events governed by physical law, and, if so, is there such a thing as human freedom? Epistemology, or the theory of knowledge, treats the nature, extent, and grounds of knowledge: What distinguishes knowledge from mere opinion? Why are mathematical truths more certain than scientific theory? Can we reliably infer the future from the past? The philosophy of science, of course, overlaps these two subdisciplines considerably. It also intersects with logic, the subdivision that seeks to identify the principles of valid reasoning, and that therefore is of the greatest importance in science and mathematics. Beyond these three subdivisions of philosophy, there are those of ethics, aesthetics, and political philosophy. These last subdivisions might seem most clearly to be addressing questions beyond the limits of scientific inquiry, questions about what ought to be the case, and not just what, as a matter of fact, is the case. But it is a more than curious fact about biology that it is the only scientific discipline that anyone has ever supposed might be able to answer the questions of moral and political philosophy. Evolutionary biology in particular has often, at least since Darwin's day, inspired a hope of putting ethics on a "scientific" footing. We will address this hope in the last chapter of this book. Meanwhile, let us draw a working definition of philosophy from this section: it is the discipline that addresses those questions that the sciences cannot (yet, or perhaps ever) answer and the questions about why the sciences cannot answer these questions. Thus, the philosophy of biology addresses those questions that arise from biology but that biology cannot answer, at least not yet, and the further questions about why biology may be unable to answer these questions.

Philosophy and language

So, what are these questions biology raises but cannot address? Here are some candidates:

1 Is life a purely physical process? Are biological processes "nothing but" complex physical and chemical ones? If so, what does this mean for the science of biology as an independent discipline?

2 Does evolution have any goal or purpose, perhaps one that might give our existence meaning or intelligibility?
3 Is there any such a thing as evolutionary progress? Is complexity increasing in evolution? If so, is that increase inevitable? And what, if anything, does increasing complexity say about values? Are more complex organisms somehow better than, or higher forms of life than, less complex ones?
4 Does the theory of natural selection conflict with theism, and, if so, how can we rationally choose between them?
5 What is human nature? Which of our traits are essential to us? Are some traits innate? Do any determine our characters more than others? Are they fixed or not? Are socially important human traits more the result of heredity, nature, and our genetic programs than the result of learning, nurture, and our environments?
6 To what extent are humans adapted in the biological sense? To what environmental conditions are we adapted, and at what level does this adaptation occur—the individual human, the family or the lineage, the whole population, or perhaps the species?

If we ask any one of these questions, almost inevitably the right initial response turns out to be: "It depends." And what it depends on is the meaning of key words in each of the questions. How we eventually answer these questions will turn on what meaning we agree to confer on terms such as "life," "purpose," "progress," "complexity," "theism," "genetic program," "adaptation," and so on. For this reason a great deal of the philosophy of science, and analytical philosophy generally, is given over to the clarification of the meaning of the concepts in which questions are framed. Philosophy is not itself an experimental, observational discipline. It does not have its own domain of data about the world. Rather, philosophy addresses the questions raised by the sciences—at least in part—by clarifying the concepts on which these questions hinge.

Sometimes, the result of such a philosophical analysis is to show that a question is ambiguous and that the difficulty or debate about its answer reflects the failure to see the ambiguity. It might reveal that a crucial concept such as "life," "program," or "adaptation" has two or more alternative meanings. Armed with this insight, we can then decide which alternative meaning is relevant and appropriate. This may not settle the matter. The focal question may remain unanswered. But at least we will have a clearer idea of what the question means. And we will also have a clearer idea of what would count as a satisfactory answer.

How do we go about deciding on the meaning of a crucial concept? Only rarely will looking up the word in a dictionary help, for dictionaries usually provide many alternative meanings and our problem is to decide which among the alternative meanings is the one relevant to our inquiry. Just try to answer the question whether life is wholly a matter of physical and chemical processes by looking up the word "life" in a dictionary. Moreover, many of the concepts with which the philosophy of science is concerned are discussed in technical terms, neologisms, the meaning of which are given in large part by the scientific theories in which they figure. Consider the term "positive charge" in physics. Suppose someone asked what it is that positively charged protons have and that negatively charged electrons lack (the word positive implying that something is present or added and negative implying an absence or loss). This silly question simply reflects ignorance of the relevant theory and a reliance on the dictionary meanings of "positive" or "negative." To be clear on the meaning of the concepts with which the philosophy of biology deals, we need to understand the scientific theories in which these concepts figure. This of course makes the biologist who understands these theories at least as much of an expert on questions in the philosophy of biology as the philosopher!

So deciding on the meaning of a scientific concept requires that we understand the theory in which it figures. Further, understanding a scientific theory requires that we be able to identify the domains in which it explains and predicts phenomena, and the experimental techniques and instruments that can be employed to test the theory. And, indeed, many of the questions the philosophy of biology considers are questions about the domain of a theory and the domain's appropriate methods of investigation. Consider, for example, question 6 above, about whether biological theory can explain human social phenomena. Does the domain of the notion of adaptation by natural selection include human behavior? In other words, is human behavior the sort of phenomenon that the theory could in principle explain? Does the theory's domain extend to human societies? Just what is the range of entities to which the notion is applicable?

What all this means is that the process of identifying the meanings of the scientific terms we need to make our philosophical questions unambiguous is not really separable from the development of scientific theory itself. It also means that the difference between philosophy and theoretical science is not a matter of kind but of degree. Of course there will be differences between laboratory and field science on the one hand and theory and the more abstract inquiries of the philosopher on the other, but these differences lie on a continuum. Because philosophers' interests are abstract, they do not require laboratories. Instead, they often proceed by undertaking "thought experiments." Philosophers will often have to create "science fiction" scenarios, to explore scientifically impossible scenarios, in order to extract the logical relations of implication, exclusion, and compatibility between scientific theories and data—and among theories themselves. Scientists are advised not to lose patience with such explorations. For one important aspect of scientific progress is—beyond the increasing precision of tests that confirm or falsify scientific theories—the broadening of the domain of those theories. And such advancement requires the same kind of thought experiment, albeit more tightly constrained by immediately available data than the philosopher needs to worry about.

Once the key terms in a question have been made clear, we can turn to considering how it may be answered. Of course, it may be that, once made clear, a question no longer troubles us. Perhaps the answer to the question is obvious, or perhaps the question rests on a false presupposition, or is otherwise "defective" in a way that is obvious. Not every interrogative sentence expresses a bona fide question. Some are what philosophers call "pseudo-questions." Some obvious examples include the following: "Do green ideas sleep furiously?" "What time is it on the sun when it is noon at Greenwich, England?" or "Did you phone your wife?" asked of a 10-year old girl. The first of these "questions" looks grammatically like one, but once we know the meanings of the terms that express it, we see that it is a pseudo-question, one that has the right syntax but really has no coherent content. The second question can be disposed of once we recognize that local time at a point on the Earth depends on the Earth's position with respect to the sun, and it makes no sense to ask what the sun's position is with respect to itself. The last question makes syntactic and semantic sense but is based on several false presuppositions: that the pronoun "you" refers to a married person, and a married male person to boot. None of these questions can be answered, but they can be disposed of as not needing answers. Some philosophers have held that many or all philosophical questions are like these pseudo-questions. On their view philosophical problems are dissolved, not solved. They are disposed of, not answered.

Suppose that one held, as some scientists who have no patience with philosophy do, that there are no real philosophical questions, no questions in the philosophy of science. One might hold, for example, that all real questions can, at least ultimately, be answered by science, given enough genius, enough time, and enough money, leaving nothing to philosophy. On this view, questions such as "What is time?" or "Is abortion morally wrong?" will turn out to be either questions to which empirical inquiry broadly considered can give definitive answers or pseudo-questions expressing pseudo-problems that need dissolution, not solution. If all real questions can be answered by science, then there is no such subject as philosophy, defined as the discipline addressing questions not answered by science and questions about why science cannot answer these questions.

The view that science will ultimately answer all real questions and that the remainder will turn out to be pseudo-questions, faces a serious problem, however. For it must be granted that there are many questions raised by science that it cannot *yet* answer. And in that case, why be so confident that all these questions are either answerable by science or pseudo-questions? There are only two ways to respond. The first is quite tedious. It is to take on each and every apparently unanswerable question and show what is the matter with it, show why we need not take it seriously, or else show that it is in principle answerable. The second is to show that in principle there can be no real questions beyond the reach of science. But notice that either of these two endeavors is properly and recognizably a philosophical project! We have a right to conclude, therefore, that even those who assert that science alone will eventually answer every real question owe us an argument for this claim, and that any such an argument will be a philosophical one. That makes philosophy pretty much unavoidable, even for those who deny that there are any real questions for philosophy to address.

In any case, in the absence of such an argument, we can safely assume that the sciences really do raise questions that they cannot answer and that once we have identified these questions, the philosophy of science should address them.

The agenda of the philosophy of biology

Darwinian theory is central to the philosophy of biology. One reason is its relevance to the questions listed at the beginning of the previous section, questions that interest almost all thinking people. Another is the very large amount of evidence that the theory is correct, a claim that cannot be made by other theories—coming mainly from the social and behavioral sciences—relevant to those same issues. In the physical sciences, there are other theories that are more strongly confirmed by scientific experiment. For example, quantum electrodynamic theory makes predictions that have been confirmed to 12 decimal places. That is an accuracy roughly equivalent to measuring the distance from the tip of the spire of the Empire State Building in New York City to the point of the Space Needle in Seattle to within the breadth of a single hair. But, for all its accuracy, the theory appears to have little explanatory relevance to human life. The atomic theory that stands behind the Periodic Table of the Elements is also a very well-established theory with ever-increasing application in technology and engineering. But its account of the chemical relations among the atoms that compose our bodies, for all its completeness, will not answer questions about human nature, human behavior, human institutions, and human history. Darwin's theory does not attain the standards of accurate prediction and detailed explanation that theories in physics and chemistry do, but it is potentially far more relevant to questions about ourselves.

On the other hand, there are theories in the social and behavioral sciences that, unlike Darwinian theory, were developed explicitly to explain and (more recently) to predict human behavior, human action, and the large-scale social processes, i.e. culture and history. Indeed, social and behavioral scientists have been offering such theories at least since the late nineteenth century. Most of them should be familiar: Freud's psychodynamic theory, Skinner's behavioral learning theory, the competing theories of social structure and function attributed to Durkheim and Weber, Marxist economic theory, classical, Keynesian, and neoclassical economic theory, and their successors. One reason that there are so many such theories, and that we could go on listing others, is that none has secured anything like the scientific confirmation required for general acceptance in science, social or natural, and therefore we continue to seek more such theories. Were any of these theories well enough confirmed, we might be able to rely on them to explain human affairs, or at least to do so to a greater extent than a theory such as Darwin's, which may have significant implications for the human sciences but secures its considerable scientific support in other domains. Alas, none of these theories has secured general acceptance in its discipline to match the well-established role of Darwinian theory in biology.

Darwin's theory of natural selection and its subsequent scientific elaboration more fully combines explanatory relevance to human affairs with independent scientific confirmation than any other theory in science. And this is what makes the theory a potential lightning rod for public controversy. Exploring its implications for humans, some see in it the gravest threat to religion generally or theism in particular. Others find in it the rationalizations for the worst excesses of capitalism. Some treat it as destructive of the very essence of our humanity, on which our values and the very meaning of life depend. Still others see Darwinian theory and the biological understanding it inspires as finally providing the basis for an enduring moral concern for all living things and the planet on which we and other living things find ourselves.

Whether or not Darwinian theory has any such implications is a question that biology certainly cannot yet answer. It may turn out to be a question that biology can never answer. And that of course is what makes the question a philosophical one. And it explains why the philosophy of biology has become so consequential a subject, so consequential that among all the technical subdisciplines of philosophy it is about the only one to find itself represented on bestseller lists, to be expounded in courts of law examining constitutional issues of church and state, and to be the subject of debate in popular culture generally.

The aim of this book is to shed light on at least some of these human questions, but to do so we will need to guide the reader through the narrower scientific and philosophical issues on which answers to the big questions may turn. Thus, a great deal of our concern will be with matters the relevance of which for the lively public debates—the nature–nurture debate, the intelligent design debate, and so on—may not be obvious until understood. To get to the big questions, we will need to travel through issues that may look technical, complicated, and even out of touch with the target questions. We think, we hope, that the pay-off is worth the journey, and also that the journey itself will prove valuable in its own right.

Decades ago, the famous evolutionary biologist Theodosius Dobzhansky wrote, "Nothing in biology makes sense except in the light of evolution." This statement needs some explanation and qualification. First, evolution is descent with modification, the notion that all organisms are modified descendants of a common ancestor. It is broader than Darwin's theory of natural selection, which is a mechanism of change, an explanation for how modification occurs. (And as will be seen, selection is not the whole story.) Second, the statement overreaches somewhat. Biological questions can be posed the answers to which involve evolution only very indirectly (for example, questions relating to the physical properties—the biomechanics—of biological materials). Nevertheless, understood as a claim about shared ancestry, as well as natural selection, we think it is close to true. And that is why, as will be seen, evolution emerges as central in every chapter and virtually every section of this book. Biology is inescapably historical.

We begin in the first chapter by discussing the theory of natural selection, its structure, the scientific problems it raises, common misunderstandings of the theory, and its major metaphysical consequence, the extension of the mechanistic worldview of the physical sciences to the life sciences. This extension raises an epistemological problem about the kind of knowledge that biological theory provides. For Darwin's theory does not look much like the sorts of theories familiar in physics and chemistry, the explanatory and predictive powers of which have vindicated mechanism as a metaphysical worldview for these disciplines. Differences between biology and the physical sciences, and indeed between it and the human sciences, must be reflected in the epistemology of biological science, in the kinds of knowledge it provides. For this reason, philosophers of biology have been as interested in the grounds of the theory of natural selection as in its structure. In Chapter 2 we consider how and why scientific theory should turn out to look so different in biology from the way it looks in physics. We do so by examining the question of why there seem to be no scientific laws in biology, or none to rival those of physical science in scope, simplicity, and power. Answering this question will reveal a great deal about the nature of biological theory and also shed light on the human sciences too, as we shall see in the last chapter.

Chapter 3 continues the examination of epistemic issues raised by Darwinian theory, in particular three "technical issues" about evolution that vex biology but that are often invisible to nonspecialists. One is the nature and extent of biological adaptation and the role of constraints of various kinds in shaping organismal design. It will turn out that adaptation and constraint—often considered to be alternatives in evolutionary explanation—are for certain kinds of questions jointly essential to explanation. The second is the role of statistics and probability in biology. It will be seen that the notion of objective chance—so essential to Darwinian thinking—is only imperfectly understood and remains problematic. The third is the foundation of functional explanation and description. We will show that two very different conceptions of function survive in biology, and that the imperfect overlap between them has consequences for how questions about function are posed and answered, both in biology and in the social sciences. In general, we try to show how these apparently abstract matters bear on the larger questions that drive interest in the philosophy of biology. For example, we show in this chapter how the problem of reconciling the theism of the Abrahamic religions with biology's commitment to natural selection turns in part on how we are to understand "probability" and "drift."

Chapter 4 examines the relationship between molecular biology and the other subdisciplines of biology, from cell biology to paleontology. It raises the question of whether all biological processes can or must eventually be explained by theories about their macromolecular constituent processes. The issue is reductionism. Biologists and philosophers have argued mainly against reductionism, yet it persists both among many physical scientists and even a few prominent biologists. It is clear that answers to the reductionism question will drive a good deal of future scientific research in the discipline. Further, the reductionism question is relevant to a number of important philosophical issues such as the mind–body problem and determinism versus free will. All of this makes reductionism a threat or a promise that few philosophers or biologists will be neutral about. Reductionism is a very old issue in biology. But in addressing it we cover some new territory, issues that have arisen or become especially problematic only in recent decades on account of new discoveries. One is the problem of what is a gene. The modern understanding of genetic mechanisms makes the concept of a gene problematic, varying as it does from one research context to another. The gene of molecular biology seems not to refer to the same concept as the gene in population genetics. If population biology is reducible to genetics, in what sense of the word "gene" is it so reducible? Another issue has to do with the dynamics of complex systems of interacting components, such as the gene networks in an organism are said to be. Such networks seem, from an antireductionist standpoint, to have higher-level properties and to be affected by higher-level controls, that raise new challenges to the reductionist view. Finally, the principle of natural selection seems to present a barrier to the reduction of biology to physical science. In particular, it seems to create an unbridgeable gap between explanation at the level of chemistry and physics and that at the level of macromolecules. If so, then the scope of reduction will be limited, necessarily coming to an end at the level of molecular biology.

In the last three chapters of the book, we turn to some more specific issues. The question of whether evolution is progressive—raised briefly in Chapter 1—is addressed at length in Chapter 5, along with the further issue of the evolution of complexity. Progress has an evaluative component, which raises the question of whether it is even a proper subject for putatively value-neutral science. If it is, if progress can be understood in a way that makes it suitable for scientific study, what does evolutionary theory predict about progress? Is it an expectation or merely a possibility? And then, what is the relationship between progress and complexity? If they are related, what does the history of life tell us about complexity and how it changes? The discussion reveals how advances in empirical science sometimes can hinge critically on advances in conceptual clarification.

In Chapter 6, we return to the connected questions of metaphysics and epistemology that biology raises. The metaphysical ones are those about whether, along with genes, cells, and organisms, biology must recognize "higher levels" of organization—for example groups or societies of organisms—and questions about whether there is something causally unique about genes and the genome that should accord them a special explanatory role in biology. Finally, in Chapter 7, we consider the relationship between biology and the social sciences and, more narrowly, between biology and human nature. Humans are members of a biological species, and therefore arguably human adaptations are not exempt from the operation of natural selection. But the degree to which human psychology and behavior is molded by selection, and the mechanism by which it is molded—for example by selection at the level of the individual versus the level of the group—are open questions. And then there is a pressing further question: if biology *is* relevant to human affairs, what are the implications for distinctively human concerns such as ethics?

Our outline of the agenda of the philosophy of biology is not aimed at settling any of its debates. Indeed, the authors of this book have divergent views about almost all of the unavoidable questions biology raises and cannot (yet) answer. Our aim is to provide the reader with the resources to see how serious the questions are and what would count as good answers to them.

1 Darwin makes a science

Overview

There is an important sense in which biology as a science began only when Darwin hit upon the theory of natural selection in the late 1830s, although he did not publish the theory until 1859 (after A.R. Wallace hit upon it too, and threatened to scoop him). Of course there had been scientists making important discoveries about the biological world at least since Aristotle in the third century BC. In the 200 years prior to Darwin's birth, Harvey and van Leeuwenhoek stand out for their discoveries that, respectively, the heart beats to circulate the blood and all living things are composed of cells. And there was Linnaeus' system of classification of living things and his naming system for genus and species, the binomial nomenclature. But it can be argued that until Darwin's achievement, none of these findings, explanations, or classifications could be organized into anything with a right to call itself a science. Darwin's evolutionary theory explains more than just common descent, the shared ancestry of all organisms on Earth. It identifies a causal process that produces the adaptations we see everywhere in nature, one that replaces other accounts of the adaptation, other accounts that could not be causal or even in principle scientifically testable.

In this chapter we consider this argument that biology did not really exist as a science at all until Darwin's discovery of the mechanism of natural selection. We also discuss some controversies. Natural selection has been controversial from the very first time the idea was publicly expounded. Some of these controversies are based on misunderstandings, but some are real. In this chapter we separate common misunderstandings about Darwin's theory from the real issues that any defender of the theory must come to grips with.

Teleology and theology

Before Darwin, philosophers such as Immanuel Kant had despaired of our ever creating a science of biology on a par with sciences such as physics and chemistry. "There will never be," Kant (1790) wrote, "a Newton for the blade of grass." What Kant meant by this claim was that biological processes could not be understood or explained by the operation of the sort of mindless causal properties of mass and velocity, position and momentum, force and acceleration that promised to suffice in Newton's mechanics to explain everything physical. By the end of the nineteenth century, electric charge and electromagnetic fields were added to the list of causes, enabling science to explain almost all physical processes, including heat, flight, electricity, and magnetism. And, soon after, most of chemistry could similarly be explained on the basis of atomic theory.

But until Darwin the biological seemed permanently, and logically, conceptually, necessarily, out of the explanatory reach of merely physical causes. Take a cotton plant: it moves its leaves throughout the day to track the sun, and it does so *in order to* maximize the amount of sunlight that falls on its petals. Even more impressively purposeful or goal directed is the cowpea plant. When well-watered plants of this species move in a way that maximizes the amount of sunlight to fall on their leaves, they do so apparently *in order to* produce starch from water and CO_2 through a chemical reaction catalyzed by chlorophyll. And the plant produces starch *in order to* grow. But when the surrounding soil is dry, these same plants move their leaves *in order to minimize* their exposure to sunlight so that they retain water that would otherwise evaporate. It looks like explanation in biology connects events, states, processes, and things with their *future* goals, ends, and purposes, not with the *prior* causes that bring them about. It was Aristotle who distinguished the prior physical causes we are familiar with in physical explanations, from the purposes, goals, or ends with which biological processes are explained. The former he called "efficient causes" and the latter "final causes." The Greek word for "end" or "goal" is *telos* from which comes the English word "teleological." A teleological explanation shows why something happened by identifying the end, purpose, or goal that it brought about. Why does the heart pump? Kant would have answered that it does so in order to circulate the blood. Circulating the blood is an effect of the heart pumping, and this effect explains it, even though circulation happens afterward as a result of the pumping. Things have not changed much in three centuries. Ask a molecular biologist why the DNA molecule contains thymine whereas the RNA molecule transcribed from the same DNA molecule contains uracil (even though both would appear to perform nearly the same function). The answer is teleological: Although the two molecules are otherwise the same in nucleotide composition, DNA is made of thymine *in order to* minimize mutation (in particular, what are called point mutations arising from deamination), whereas RNA contains uracil *in order to* minimize the costs of protein synthesis.

And of course it is not just biological explanations that are "teleological," i.e. that cite future ends, goals, or purposes to explain past structures, processes, and events. The whole vocabulary of biology is teleological. Consider some of the most basic nouns in biology: codon, gene, promoter, repressor, organelle, cell, tissue, organ, fin, wing, eye, coat, stem, chloroplast, membrane. Almost all of these terms are defined—at least conventionally—by what the thing *does*, or what it does when working normally. And not just anything it does, for each of these does many things. Take a shark's fin, for example: it provides stability while swimming, but it also reflects light, makes turbulence behind it in the water, adds weight and surface area to the body, signals to humans the presence of a predator near the surface, attracts the interest of connoisseurs of shark fin soup, and so on. But only one (or maybe a couple) of these things a fin does is its *function*. The function of a fin is the only one among these effects that define what is to be a fin: a fin is an appendage of a fish or whale, one of whose functions is to provide stability. In other words, it is something the animal has "in order to" provide stability while swimming. Well, if fish have fins in order to swim stably, one may ask, who arranged this neat trick for them? And the same question arises for practically every other feature of organisms that has biological interest. For almost everything biological is ordinarily described in terms of its function. So almost everything biological raises a teleology problem. In contrast, a question such as "What is the function of the electron?" is not one physicists ordinarily consider.

Teleological explanations, which explain by citing goals, ends, or purposes, are troublesome. For they explain events, states, and processes, not by showing how they came about from prior causes but by identifying the future effects they will lead to. The trouble is we know that future events cannot bring about past ones. For one thing, it is hard to see how something that does not yet exist (because it is in the future) could bring about something that does already exist and may have existed for some time in the past. For another, we seem to be allowing the behavior of something seeking a goal to be explained by the goal even when it fails to achieve the goal. A sperm cell moves up the uterus "in order to" fertilize the ovum, even when, as in almost every case, it fails to do so.

Aristotle may in fact have recognized the first of these problems, the impossibility of future causation. For he argued that final causes had to be "immanent," meaning somehow embodied or represented in the prior states of the organism's life, directing its course towards some goal.

Of course some immanent teleological explanations seem unproblematic. These are the "in order to" explanations we employ to explain our own behavior. "Why are you taking organic chemistry?" "In order to get into medical school." Or "Why do you want to go to medical school?" "In order to please my parents." In these cases, the "in order to" relation reflects our desires, and our beliefs about the means to bring them about. So, we can "unpack" the explanation of why I am taking organic chemistry into: (i) the desire to get into medical school; and (ii) the belief that taking organic chemistry is necessary for getting into medical school. The beliefs and desires that underwrite the "in order to" explanations of our actions are almost never made explicit. But making them explicit turns the apparent teleological explanation of why I am taking organic chemistry into a nonteleological explanation in terms of prior causes for later effects. I am taking organic chemistry (now), because at sometime in the past I came to desire to go to medical school, and I came to believe that taking organic chemistry is necessary for going to medical school.

But in biology there does not seem to be a similar strategy available for turning statements about purposes, goals, ends, and the means to achieve them into causal relations between earlier events and later ones they bring about. Because, to a first approximation, science seeks to explain by uncovering prior causes, biology before Darwin was arguably not a science. Of course, before Darwin, one could explain all the "in order to" explanations in biology on the model of explanations of human action, simply by appealing to the "desires" and "beliefs" of God. Why does the heart beat? The explanation that it does so in order to circulate the blood turns out to be shorthand for something like: it was God's will (i.e. God wanted) that blood circulate through vertebrate bodies, and he knew (i.e. believed correctly) that making a heart that beats would be a good way to do so. Of course, as God is omnipotent (all powerful), he can cause the object with the desired future effect to exist. For each "in order to," there is a set of statements about God's knowledge and his will (God's infallible and always benevolent versions of our beliefs and desires) that show the underlying causal basis of the teleological explanation.

Now there are several problems about this way of saving teleological explanation. To begin with, invoking God to explain natural phenomena is, in the view of many, simply to change the subject from science to theology. Now the acceptability of teleological explanations will hinge on the soundness of arguments for and against God's existence. Second, invoking God's will and his omnipotence to explain biological events and processes seems to be too easy. As far back as the eighteenth century, Voltaire was ridiculing "in order to" explanations for this reason. In his book *Candide*, Voltaire has Dr. Pangloss explain why the nose has a bridge by pointing out that noses bear bridges in order to support eye glasses. We detect adaptations everywhere in nature—the exoskeleton of insects adapted to prevent dehydration; the intricate complexity of the mammalian eye so perfectly suited to the available sources of light, reflectance, luminosity, etc.; even perhaps morning sickness in early pregnancy, seemingly exquisitely arranged to protect the fetus from foods the mother might eat that are even slightly harmful. In each of these cases, the explanation turns out to be exactly the same. God's good will, her complete knowledge, and her omnipotence, together account for the arrangement.

But surely an omniscient, omnipotent God could have chosen some different arrangement of things to attain the very same outcome. God could have made water less evaporative so that insects would not dehydrate so quickly, or arranged the digestive systems of pregnant females to digest all poisons instead of becoming more sensitive to them. Why didn't God do so? Notice that an attempt to answer this question by identifying the constraints imposed by the physical and chemical laws and the local conditions in which

God operated to realize her will immediately raises questions about why God should be constrained in any way. She can create, arrange, suspend any chemical or physical law, or local conditions she chooses. There is, of course, no answer to the question, why did God choose the course she did, and not some other one, at least none that is open to testing by data, experiment, observation, etc. This question is pretty clearly a matter of theology, not science.

Making teleology safe for science

So appeals to God will not bail out teleological explanations for science, will not turn them into causal ones. This of course is where Darwin's theory of adaptation by natural selection comes in. According to the most widely known contemporary statement of Darwin's theory (Lewontin 1978), adaptation results if three facts obtain:

1 There is reproduction with some inheritance of traits in the next generation.
2 In each generation, among the inherited traits there is always some variation.
3 The inherited variants differ in their fitness, in their adaptedness to the environment.

The simplicity of these statements hides their tremendous explanatory power, and also leaves unspoken some important implications and fosters several potentially serious misunderstandings. Before discussing these implications, and forestalling these misunderstandings, it is as well to give a simple illustration of the explanatory power. Why do giraffes have long necks? The short answer could be "in order to reach the tasty leaves at the tops of the trees that other animals can't reach." A slightly more scientific way of expressing the same explanatory facts is to say, "Having a long neck is an adaptation for the giraffe" (or "The function of the giraffe's neck is to reach leaves that other savannah mammals cannot"). But the fuller version of the explanation goes something like this: The length of a giraffe's neck is a somewhat inherited trait. Long-necked giraffes have long-necked offspring, not invariably so and not always as long, but usually and sometimes longer. Never mind for the moment the details of why such traits are inherited in this pattern. Observation and measurement are sufficient to convince us that they are. Observation also reveals that, as with all inherited traits, there is always variation in the length of necks in each generation of giraffes. This variation is never in just one direction, say only toward longer necks; some long-necked giraffes have offspring with shorter necks, and vice versa. This will be true no matter whether trees get taller or shorter or other animals, say some insect species, come along who can compete with giraffes for the highest leaves in the trees. This point about variation in heritable traits is sometimes expressed by calling them "blind," though this expression is plainly metaphorical. More often the independence of variation from features of the environment that might make a variation useful or not is expressed by calling it "random" (and this is the source of a possible misunderstanding that we will forestall later). Now, let us say that at some time in the distant past, a long-necked variant appeared among a small number of giraffes, just as average and short necks appeared as well. And it appeared not because a long neck would be advantageous but just because variation is the rule. Further, let us say that this longer-necked giraffe did better at feeding off the high leaves than shorter-necked ones and did better than other mammals competing with giraffes for resources in the same environment. That is to say the hereditary trait of having a longer neck was "fitter" in the giraffe's environment. So giraffes with the long necks survived longer and had more longer-necked offspring. As the total giraffe population that could be supported by their environment was limited, the proportion of longer-necked giraffes in the whole population increased from generation to generation. This was because in each generation they out-competed the shorter-necked giraffes for limited resources (leaves high enough up on trees that only giraffes could reach them) and, therefore—owing to their longer life, greater strength, etc.—had more offspring. After a sufficiently large number of generations, the population of giraffes came to consist only of long-necked ones. Thus, Darwin's theory explains why giraffes have long necks by identifying a causal process that in the long run would produce long necks without any person or force acting "in order to" provide for the nourishment of giraffes. Having a long neck is an adaptation for giraffes. That is to say they have it because in the past there was hereditary variation in neck length and the longer variants just happened to be fitter in the environment where giraffes found themselves. (We feel constrained to note that the point of this story is only to illustrate how adaptation arises in principle, using a well-worn example that many find easy to grasp. In fact, however, giraffes may have evolved long necks for very different reasons. They could have been an adaptation for intimidation of predators or of other giraffes, perhaps in male–male competition. Or it could be that in giraffes, neck length, and body size are connected in growth in such a way that animals with larger bodies grow disproportionately longer necks. If so, then selection for large body size might have produced a long neck as a side effect. The treetop leaves possibility is an example of an evolutionary "just so" story. We will discuss such stories, and alternative nonadaptive modes of evolutionary explanation in Chapter 3.)

Darwin called this process natural selection and the name has stuck. The theory of natural selection explains the traits of extant flora and fauna by tracing their evolution back through successive rounds of natural selection by the environment operating on the variation in hereditary traits each generation presents. "Natural selection" is not an entirely apt name for the process, as it misleadingly suggests the notions of choice, desire, and belief built into the theological account of adaptations. It evokes an agent doing the choosing, if not God then perhaps Mother Nature, actively picking the best of the litter. But the selection process is more passive than that. Perhaps "environmental filtration" is a better label than "natural selection." The environment does not "select," but rather it filters, preventing the less fit from passing through. Moreover, it is particularly important to recognize that environments change over time, and that what is adaptive in one environment can be maladaptive in another. For instance, as global warming accelerates, the grizzly bear's thick warm coat may become maladaptive. This fact has important implications for the notion that natural selection generates continued improvement in absolute terms, that later organisms are better, in some important sense, than earlier ones. In fact, arguably, Darwinian theory demands no such thing. The theory implies only that there will be adaptation to local environments. But as environments change, and improvement tracks only local environments, there is no commitment in the theory to long-term "progress." Indeed, extinction is a fate not restricted to the dinosaurs. We shall discuss progress further in Chapter 5.

Another potential source of misunderstanding has already been mentioned. The theory requires that in every generation heritable traits vary to some degree, and that this variation is "random." The theory requires inheritance of traits and it requires variation in these traits across generations. It is entirely silent on the mechanism of inheritance and the source of variation. Darwin had theories about both inheritance and variation but they were mistaken. The later independent discovery of the right theory of heredity and the source of variation greatly strengthened biology's confidence in Darwin's theory of natural selection. But the theory would have worked with many different hereditary mechanisms and sources of variation, and it did not imply or require any particular one. At most it required that there is one or more mechanisms of heredity and one or more sources of variation in heritable traits for each generation in every evolving lineage. The theory of natural selection does however rule out one cause of variation in heritable traits, namely a future cause in which new variation is guided by the needs of the individual who bears it. Indeed that is the major thrust of the word "random" in the phrase "random variation" in Darwin's theory. It is not that the appearance of a new trait is undetermined, that it is not fixed by prior causes. It is rather that the causes that fix it are independent of, unconnected with, the factors that determine its adaptedness. We say that variation is random "with respect to" adaptation. To put it another way, the usefulness of a trait in the environment in which it appears—its goal, purpose, or end—is not among the causes responsible for its appearance. Philosophers, theologians, and others noticed almost immediately after the appearance of *On the Origin of Species* that Darwin's theory made goals, purposes, ends, and future causes of any sort completely superfluous to biology.

Nowadays a great deal is known about the mechanism of heredity and the source of variation. Hereditary transmission proceeds mainly via genes composed of nucleic acids, and variation results from recombination of genes and mutation. As some of this mutation is caused by quantum processes (such as radioactive decay), at least some of the variation that the environment filters many would call random in the sense of lacking a deterministic cause. But nothing in the mechanism of natural selection requires indeterminism. Indeed, the term “blind” may be more apt, less misleading, than “random” for the sort of variation required by natural selection. Variation can be said to be blind with respect to need, or to the environment. In that case, the process of natural selection as a whole could be described—in the sociologist Donald Campbell’s apt phrasing—as “blind variation and selective retention” (Campbell 1974). The phrasing is especially apt in that it emphasizes that, according to Darwin’s theory, nature has no foresight.

Notice that the three requirements listed above for the operation of natural selection—reproduction, heredity, and differential fitness or adaptedness—do not mention organisms. They do not mention actual animals and plants and their traits, which are assumed to be the subjects or the “domain” of the theory of natural selection. One reason is that the theory is supposed to explain not just the origin of adaptation in these organisms but also the evolution of higher units—colonies or societies of multicellular units—and of lower ones—single-celled organisms. Also, the theory is intended to apply more broadly to explain the evolution of genes and other molecules within organisms, units that are not living at all. And finally, many believe it is supposed to apply to the origin of life, that is to the evolution of single cells from large macromolecules, which again are not organisms at all. Thus, formally the theory cannot be expressed solely as a claim about giraffes, or about mammals generally, or even about animals generally, or, for that matter, about organisms. Rather, it must be expressed as a general claim about the evolution of reproducing things with heritable variation and differential fitness or adaptedness.

To express the generality of natural selection as a mechanism, David Hull (and, independently, Richard Dawkins) introduced the terms “replicator” and “interactor” (or, for Dawkins, “vehicle”). In Hull’s definition, a replicator is anything that passes on its structure largely intact through successive replications. An interactor or a vehicle is anything that acts as a cohesive unit in its environment in such a way as to make a difference for the replicators that generate it. These terms have taken on a life of their own in evolutionary theory and other biologists and philosophers have modified them in various ways. Nevertheless, it is easy to see how the two concepts provide the theory with the generality it requires. And here is one way in which such generality is useful. We can begin to paint a picture of how life originated on Earth, even without knowing the details of the process. Perhaps the first evolving entities were simple macromolecules that functioned simultaneously as replicators and interactors. Then variation arising in these macromolecules could in some cases have produced associations of them, which if better adapted than their predecessors would have preferentially survived. Further random variation and filtration of the better adapted might have eventually produced a separation of the replication and interaction functions, as well as further buildup, eventually generating larger and more complex entities of a sort that we are willing to call living. Long continued, this process could in principle produce the entire range of adaptation we know today, in other words, the entire explanatory domain of Darwinian theory. The point is that the generality of the replicator–interactor concepts enables us to tell this story—to develop hypotheses about the origin and diversification of life on Earth—without knowing any of the actual details of the actual process: which macromolecules, combining in what way, under what environmental conditions. Indeed, as we will discuss later, the theory is sufficiently general that we can use it to speculate about the origin and evolution of life not just on Earth but anywhere in the universe.

Misunderstandings about natural selection

Skeptics, detractors, and students learning Darwin's theory of natural selection for the first time are often incredulous. How could a theory based on such a simple mechanism as blind variation and selective retention actually explain all of the adaptation we see in biology? Most of the hereditary variations we see in nature are either slight differences that appear and reappear irregularly, or they are larger but extremely maladaptive hereditary defects. How could the environment selecting on extremely slight differences from generation to generation produce a structure such as, for example, the eye, a structure that—whether in an insect, octopus, or human—consists of many intricate parts, all of them highly adapted to the particular environments of insects, octopi, and humans. Darwin recognized this problem for his theory in one of the most famous passages of *On the Origin of Species*:

> To suppose that the eye, with all its inimitable contrivances for adjusting the focus to different distances, for admitting different amounts of light, and for the correction of spherical and chromatic aberration, could have been formed by natural selection, seems, I freely confess, absurd in the highest possible degree.
>
> (Darwin 1859: 186)

The skeptic allows that Darwin's mechanism can bring about slight changes, for instance in a laboratory where the experimenter can manipulate the environment of some rapidly reproducing organism such as a bacterium or a fruit fly. And the skeptic can easily see how animal breeders can modify their stocks in ways advantageous to farmers or fanciers over thousands of years. But the changes produced in both cases are quite small, compared with the change from, say, single-celled protist to mammal, from an amoeba-like ancestor to a modern goat. What is more, the experimenter or the breeder begins with a highly adapted creature and is able to carefully control the environment—the probabilities of reproduction for each organism—to bring about the desired effect. What the skeptic really wants, one suspects, to confirm Darwin's theory, is an experiment that begins with a random collection of early-Earth molecules and produces complex, intelligent species like us over a period of about 3.5 billion years of unmanipulated evolution by natural selection. The biologist must admit that no such experiment, or even one close to it, is in the offing. So, why are biologists so strongly convinced that Darwinian natural selection underlies evolution? The reason is that natural selection—random variation and environmental filtration—is the only mechanism known in nature that can produce adaptation, that can produce the "in order to" that characterizes so many of the features of organisms. In fact, as we explain in the next section, it is hard even to think of an alternative mechanism.

For some, Darwin's theory presents puzzles having to do with complexity, randomness, and directionality. They ask how complex functional designs can arise by a random process such as natural selection. For example, the evolution of a complex structure like a wing capable of sustained flight (as in a bird) from a fin (as in a primitive fish) might seem to be impossible, given the randomness of the process of natural selection and the enormous number of modifications necessary. The first part of the answer is that natural selection is not random. It is a process that requires some randomness in its "input." Variations arising are not targeted toward solving problems posed by the environment. But the "output" of natural selection is decidedly nonrandom, the differential survival and reproduction of the variants that are better adapted. The second part of the answer is that natural selection can act cumulatively, and that is what makes complex adaptations possible. Selection first transformed a fin into a walking limb, strong enough to support a large animal on land, and then later transformed a walking limb into a wing, capable of producing sustained powered flight. Complexity is possible because later adaptations build on earlier adaptations. In other words, complex adaptations are not produced in big leaps but in smaller steps, each one of which is adaptive, and function can change from one step to the next. A wing is not a better fin or even a better leg. It is something entirely different, serving a different function. Looking only at the endpoints, the gap covered might seem impossibly large, and the reason is that natural selection to some extent covers its tracks. Looking at a fin or wing, the intermediate walking limb stage is not evident, at least not superficially evident. Though natural selection covers its tracks, a great deal of biology has been devoted to uncovering them. One of Darwin's earliest arguments for natural selection was based on the close similarity in parts, their numbers, and spatial relations to one another (their "homology") of the bones in fins, legs, and wings. Two hundred years later, molecular biologists can trace the genealogy of the bird's wing back through the reptile's leg to the fish's fin in the similarities and differences of the gene sequences that control the development of each. They can show how the DNA sequence differences and similarities between the genes involved in limb development in birds, reptiles, and fish enable us to date their common ancestors and say something about how the homologies and the differences among the bones of their different limbs are due to differences in DNA sequences. In the fruit fly *Drosophila*, it is known that a small mutation in the right gene is all it takes to turn its antennae into a pair of legs. Increasingly, molecular biology is able to uncover the tracks evolution has hidden, so that adaptations begin to look expectable instead of miraculous.

What about directionality? The notion of cumulative change might seem to suggest a kind of directedness to the process of adaptation, a drive toward greater complexity. In fact, however, it is an open question whether there is in evolution any preferred tendency for complexity to build up. What is clear, however, is that nothing in the current understanding of natural selection predicts a drive toward greater complexity. Increases occur, but in our fascination with them we tend to forget the frequent decreases. Winged animals become flightless, as in the evolution of penguins. Animals with walking limbs lose them when they return to the water, as in the evolution of whales (from a common ancestor with hippopotamuses!). Complexity is reversible, and selection is expected to favor decreases whenever opportunities for adaptive simplicity arise . . . which could be often!

Thus, the randomness of variation is not a problem for selection theory. Nor is the buildup that seems to underlie complex adaptation. Given random variation and environmental filtration, plus at least occasional accumulation, the evolution of structures such as wings and eyes is not surprising. (As we shall see, the big problem for the theory of natural selection is not an adaptation like the eye but an adaptation like sex. What is it about the environment of living things that makes sex adaptive?) That said, it must also be pointed out that, as a cause of adaptation, natural selection has its limitations.

Evolutionary biologists sometimes describe the challenges that the environment presents to organisms as "design problems," though they recognize that the expression is even more misleading than "natural selection." (If Darwin is right, there is no designer who sets the "design problems" or solves them. So, the expression is a metaphorical way of identifying a dimension of the organism's environment that poses a challenge to its survival and reproduction.) Being fittest is a matter of a line of descent solving these "design problems" better than its competing lineages. But the best among competing solutions to a design problem does not have to be, and rarely is, a complete, or elegant, or even a very good solution. Variations that arise will often be "quick and dirty" solutions to design problems, advantageous for the moment but perhaps not in the long run—and perhaps not as advantageous as other possible variants that have not arisen yet but that might "solve" the design problem better. A "better but slower to emerge" solution may yet appear, but if it does, it will have to compete with the quick and dirty one. Often too the quick and dirty solution just makes the better one unreachable.

Examples of such satisfactory but imperfect solutions to design problems are not hard to find. The giraffe's long neck is adapted to browsing at treetop levels, but it makes drinking difficult; it would be easy to dream up an anatomical structure that would solve both problems at once. Or consider the relatively poor design reflected in the frequency with which we choke on food or drink. Did the alimentary canal have to intersect the respiratory system?

The classic example of an imperfection is the “blind spot” in the human eye. A simple experiment reveals it: hold a piece of paper with a black dot on it in front of one eye and cover the other; move it until the dot disappears from your visual field. Near the center of one’s visual field, where one would suppose vision is and needs to be most acute, there is no vision at all. This is owing to the fact that the optic nerve is connected to the retina, not from the back but from the front, and then bends around 180 degrees to connect up to the brain behind, passing right through the visual field. This strikingly bad piece of “design” is presumably a vestige of a much earlier quick and dirty solution to the problem of vision in vertebrates. It is certainly not essential to high-resolution vision as it did not arise in the largely independent evolution of eyes in molluscs—squids, octopi, and their relatives. Why was the attachment of the optic nerve not later reversed in vertebrate evolution? Perhaps the necessary variation did not arise. Alternatively, it could be that the many parts of the vertebrate eye have been selected for compatibility with one another, and their mutual dependencies are now too deeply entrenched to permit major rearrangements, even when big rearrangements would offer significant improvements. The quick and dirty solution excluded the slow and elegant one. The lesson is that while natural selection explains adaptations, apparent perfections of design, it also explains some of the imperfections of design we see in organisms.

Is Darwinism the only game in town?

So natural selection explains the appearance of purposiveness of adaptation. And this is not a process involving an active, literal process of “selection.” Rather it is the passive filtering out of the maladapted and the less well adapted. Also, natural selection operates on random variation, but selection is not itself a random process. It produces adaptation to a local environment or, speaking metaphorically again, solutions to “design problems.” In other words, it filters available variations for the best quick and often dirty solution to the organism’s present design problems, a solution in which the good, or the merely “good enough,” is sometimes the enemy of the best.

The scientific evidence that supports the theory of natural selection is diverse and immense, direct and indirect, from laboratories and from the field. But in addition to all the evidence biologists have amassed in favor of the theory, there is another powerful argument for it. It is one that biologists are reluctant to rely on and philosophers hesitant to articulate. It is, however, important to state this almost a priori argument in favor of the theory of natural selection’s explanation of adaptation. Physical science provides an account of the origin and development of the universe, from the Big Bang onward, in which true teleology, purposes, and goals have no role, at least outside of human purposes and goals. And future causation has no role at all. The methods of physical science excluded explanation by future causes long before it was shown to be physically impossible by Einstein's special theory of relativity. The method and the theory place a consistency constraint on the rest of science, including biology. Either biology must honor this prohibition against future causation or it must take a stand against physics and deny the truth of its most well-established theories while rejecting one of its most fundamental methodological rules. Naturally, this is something no biologist is willing to do.

Are there alternatives to natural selection as the cause of adaptation? Two hundred years ago, the French biologist Jean-Baptiste Lamarck (1809) offered a theory based on use and disuse, accepted by some before Darwin (and to some extent by Darwin himself). Consider again the neck of the giraffe. Lamarck's theory was that from an early age, each giraffe was stretching its neck to reach the tasty leaves at the tops of trees, and that a lifetime of stretching not only lengthened the giraffe's neck but the longer neck was inherited by its offspring. Enough generations of stretching and its transmission and, *voila*, long-necked giraffes. This theory has the virtue, like Darwin's, of offering a straightforward account of adaptation based entirely on past causation. The disadvantages are: there is no evidence in its favor and plenty of evidence against it. Obvious evidence against Lamarck is all around us in human history. For example, in China girls' feet were bound for millennia, without any effect on their size at birth or their size if left unbound. Equally important, Lamarck's theory requires that there be a causal chain from the act of neck stretching in some giraffe parent to that parent's hereditary material (in modern terms, its DNA), so that the parent can produce an offspring with a longer neck. But there is no evidence whatever that use or disuse of any part of the body has any effect on the nucleotide sequence in the body's germ cells. Accordingly, in addition to lacking any empirical support whatever, Lamarckism is incompatible with the modern genetic theory of heredity, which is itself strongly confirmed. So, it seems we can rule out any sort of environmental fit as the source of heritable variation, and hence we can rule out Lamarckism.

There are other in-principle alternatives. For example, the adaptedness of organisms to their environment here on Earth could have been engineered by a species of technologically sophisticated aliens. However ludicrous it sounds, we cannot rule such an alternative on purely physical grounds. That is to say this alternative does not contradict any known physical laws. But as these aliens are themselves highly enough adapted to engineer adapted organisms here on Earth, the question immediately arises of how *their* adaptations—including superior intelligence and sophisticated technology—arose on whatever world those aliens inhabit. One possibility is the operation of a Lamarckian mechanism on their world. Suppose it were to turn out that intelligent life emerged elsewhere in the universe owing to Lamarckian mechanisms in which use and disuse of traits to solve "design problems" caused changes in hereditary material that controlled the character of traits in the next generation. Such a mechanism would of course greatly expedite evolution and explain why the aliens were so far advanced as to be able to engineer adaptation on our planet. A Lamarckian genetic mechanism would of course be an extremely elegant adaptation. Instead of quick and dirty temporary or merely satisfactory solutions to design problems, it would swiftly provide extremely efficient ones.

But the existence of so perfect an adaptation as a Lamarckian genetic mechanism would inevitably raise the question of how *it* could have emerged in a world without purposes, goals, ends, or future causes or, of course, a designer to put it in place. Some philosophers and biologists will argue that once we exclude future causation and God, the only causal process that could put such an elegant adaptation in place is the very same one that biologists believe put the quick and dirty adaptations in place in our world: Darwinian natural selection. More generally, it will be argued, once we exclude God and future purposes, the explanation for the existence of any particular adaptation must invoke causal processes operating on prior traits that are less well adapted (or perhaps not adaptive at all) than the one the existence of which is to be explained. Otherwise, our explanation will beg the question of how adaptation is possible at all. The great appeal of Darwinian natural selection is that it honors this requirement on explanations of adaptation. The question is whether any other mechanism could do so.

Suppose we impose this requirement on explanations. Then even if Lamarckian mechanisms actually did operate on Earth or, indeed, even if the Earth's flora and fauna were the result of extraterrestrial aliens' gardening and zoo-keeping, the ultimate source of adaptation would still have to be blind variation and selective retention! For if the extraterrestrials' traits (like the ability to cultivate terrestrial flora and fauna) are adaptive for them, we will still need to explain their traits causally. Further, such a Lamarckian mechanism is itself so wonderfully adaptive that its emergence requires nonadaptive, purely causal explanation as well. What, other than Darwinism, could produce it? It's not hard to imagine a hereditary mechanism arising without natural selection in which use and disuse have some arbitrary effect on the hereditary material. But it is much harder to see how a mechanism could arise in which the effect is *adaptive*, again without natural selection. In sum, Darwinism seems to be "the only game in town," not just the best explanation of adaptation but the only physically possible purely causal explanation, the only one consistent with what we already know about the physical laws (the laws of special and general relativity, quantum mechanics, thermodynamics) that govern the universe. Of course biologists have no need to adopt so strong an argument for Darwinian theory. Indeed, some will want to treat the theory as a much more limited one, one that makes claims only about the Earth over the last 3.5 billion years. As we shall see, this limitation of the theory's domain to Earth has some advantages and also presents some difficulties. If the theory is understood as a specific claim about the natural history of the Earth, then the evidence needed to test it would be finite, and we could, at least in principle, establish its truth (or falsehood). On the other hand, limiting the theory in this way saps it of some of its explanatory power, even for local adaptation on Earth.

Of course, arguing that the only way adaptation could have arisen is by natural selection is not the same thing as claiming that adaptive evolution was inevitable on Earth or anywhere else. Adaptive evolution by natural selection might never have happened, if the environment never stayed constant long enough to give blind variation a chance to generate adaptations that could be "selected for" or, on the other hand, if the environment changed too little or too slowly. Nor would much adaptive evolution have happened if the rate of introduction of new variation was much lower, or much higher, than it is. So in this sense adaptive evolution is not an inevitable feature of any universe that obeys our physical and chemical laws. And of course evolution here on Earth might slow down or even come to a halt if the tempo of variation changed radically. So the claim that adaptive evolution by natural selection is "the only game in town" is not a claim that adaptive evolution was inevitable. Rather it is a claim that *when it occurred*, it happened only through random variation and environmental filtration.

Though it is tempting to hold that natural selection is the only possible way that adaptation could have arisen, there are reasons to be a little wary of this claim. Recall that Kant, the eighteenth century philosopher of science who denied that there would be a Newton for the blade of grass, believed that he could show that Newtonian mechanics was "the only game in town." That is to say Kant sought to explain the universality of Newton's laws by showing that they are a body of necessary truths, that the universe could not behave in accordance with any laws other than the ones Newton discovered. To show this Kant tried to derive Newton's laws about the universe from principles of logic that govern human thought. Kant's derivations are invalid. What is worse (he would have been embarrassed to learn), about 100 years after he thought he had proved their necessity, Newton's laws were shown to be false—good approximations, but strictly speaking false in light of the theory of relativity and quantum mechanics. The lesson of Kant's failure is cautionary. We cannot imagine, and have no evidence for, an Earthly or extra-terrestrial alternative to the process of adaptation Darwin hypothesized, but perhaps we should not assert categorically that there is none.

The reigning confidence in Darwin's mechanism for the explanation of the appearance of purpose has led psychologists, sociologists, anthropologists, economists, and other students of human behavior, action, and institutions to invoke his theory to explain apparently purposive phenomena in all their disciplines. The twentieth century psychologist B.F. Skinner suggested that operant conditioning in animals and humans, which is known to be able to produce highly purposive behavior, is just blind variation and selective retention applied *ontogenetically* instead of *phylogenetically*. That is to say it is a mechanism that builds up learned behavior in an animal's lifetime and especially its early development (ontogeny), not just in the course of its evolution (phylogeny). Similarly, neuroscientists (such as Edelman and Kandel) have proposed that, in development, the neural connections in the brain are a result of random generation of connections and selective retention of functionally appropriate ones, the connections that allow the brain to work properly. And a "blind variation and selective retention" mechanism has also been used in the development of artificial intelligence programs (e.g. genetic algorithms) to write programs to solve problems in computer science in cases in which the computer scientist does not know how to proceed.

Darwinian thinking is not obligatory is these realms. Human behavior has many sources other than conditioning. The connections of the brain might instead have formed in a highly directed way (as a result of selection in the past, of course), rather than by blind variation and selective retention. Computer scientists can apply their own native purposiveness (itself the product of selection in the past, of course) to solve problems directly, without any Darwinian computer intermediate. But Darwinism—while clearly not the only game in town in these fields—has nevertheless generated some fascinating possibilities, such is the power of the principle. Small wonder that Darwin's great defender, Thomas Huxley, when he first read *On the Origin of Species* in 1860 is supposed to have exclaimed, "How stupid of me not to have thought of that!"

Philosophical problems of Darwinism

The extension of the theory of natural selection beyond biology to all of the behavioral and social sciences has been controversial. When added to the controversies that Darwinism has provoked at home in biology, it should not be surprising that the theory has been subjected to close scrutiny, both scientific and philosophical. What are these controversies and why should they concern philosophers?

One is about how much of biology the theory can explain. Biologists all agree that it explains adaptation. But the question remains whether every feature of every plant and animal that has emerged in the natural history of this planet is an adaptation, or even whether most are. Some biologists think that the explanatory power of the theory of natural selection is quite limited. It might explain appendages in animals, but it is not so obvious that it explains bilateral symmetry, the similarity between left and right sides of many animals. Does it explain why some rhinoceros species have two horns and others just one? Does it explain the Cambrian explosion, a geologically short period of time 500 million years ago in which most of the major modern groups of animals arose? Does it explain why there are both sexual and asexual species? Does it explain human behavior and human social organization? The debate about how many of the characteristics of organisms are biological adaptations is implicitly one about the range of the theory's application. Equivalently, it is a debate about the limits of selective explanation.

Further, there is a universally acknowledged chance element in biology, often called "drift." The role of drift has been a vexed question among biologists for almost a century, and one to which philosophers have devoted themselves as well. And yet another group of critics is concerned with the domain of the theory, those who question its application to human beings. In particular, some of those who favor radical social change think that adaptationist explanations of human traits such as gender and sex roles, intelligence, violence, and criminality somehow undermine their preferred social programs. These are all disputes about the domain of the theory's application. To settle them, we first need a statement of the theory itself that everyone will accept. There can be no agreement on a theory's range of application so long as there is no agreement on precisely what the theory says. And here again philosophers can help, disambiguating the theory of natural selection, identifying its logical implications for the various domains to which the theory's application is disputed.

So there are those who deny that the theory explains everything in biology. And then there are those who deny that it explains anything, an eclectic group including mainly religious thinkers but also some respected scientists. Some of these critics have strong motives to undermine the theory, many of the devout because it threatens the view that life was created by God to fulfill a divine plan. For them as we shall see, natural selection has a feature that makes it completely unsuitable as a tool for God to use in creating the Earth's flora and fauna. And some contend that the theory is somehow logically or conceptually defective, and therefore unacceptable.

Perhaps the most famous charge is that the theory is a trivial tautology, vitiated by a circular definition which deprives it of all explanatory force. These opponents claim that the theory has no domain of explanatory application, that it can be refuted and dismissed even before any evidence has been examined, owing to purely logical or methodological defects. The charge focuses on the concept of fitness and a bumper-sticker-length summary phrase for natural selection, coined by the nineteenth century philosopher, Herbert Spencer, "survival of the fittest." The theory explains evolution as the result of the increase in the population of "fitter" variants, whose fitter ancestors survived longer in competition with less fit conspecifics and therefore had more opportunities to reproduce than the less fit. Or, more briefly, the theory says that those who are fitter are the ones that survive and reproduce more. But, the critic says, when evolutionary biologists define fitness, they can do so only in terms of survival and reproduction, and therefore we can substitute "survive and reproduce" for "are fitter." And, in that case, the theory says merely that those who survive and reproduce are the ones that survive and reproduce, a completely untestable claim. And therefore the theory of natural selection can no more explain any actual case of evolution than the statement that all bachelors are unmarried adults can explain why Sir Elton John was unmarried (at the time he was still a bachelor). Presumably the fact that he was then a bachelor simply redescribes this fact about him in fewer words than that he was then an unmarried male, and because it is simply another description of the same fact about Sir Elton, it cannot explain that fact. Explanations must adduce facts that go beyond the facts they explain. Otherwise "self-explanation" would suffice. Why X? Because X. Accordingly, unless evolutionary biologists or philosophers of biology can provide a noncircular definition of fitness, the theory of natural selection explains nothing biological. As we shall see, it is by no means easy to provide such a definition.

Finally, we need to understand the explanatory and evidential relations between the theory of natural selection and other theories in the natural sciences, in particular those of physics and chemistry. To see why, consider the treatment accorded the theory of natural selection by nineteenth century physics. In the 1880s Lord Kelvin argued that Darwin's theory of natural selection and the chronology of life's appearance based on it must be false, because the evolution of modern organisms would have required—by the estimates of that time—at least hundreds of millions of years. Could the Earth be that old? Given the best available account of combustion, Kelvin calculated, the sun—and therefore the Earth—could be no more than 40 million years old. Ergo, evolution by Darwinian mechanisms could be excluded by their incompatibility with the best available physics. Kelvin's argument posed a serious threat to the theory, not removed finally until the late 1940s by Hans Bethe's Nobel Prize-winning account of the thermonuclear reactions that power the sun's combustion, which increased the calculated age of the solar system, and therefore the Earth, by a factor of 100. The point is that Darwinian theory needs to be logically and empirically consistent with theory in the physical sciences. Further, some would argue that if it can be grounded in physical science, explained by or derived from physical theories, then this would provide the strongest possible support for Darwinism. For then all of the evidence for the best current theories in physics and chemistry would count as evidence for the theory of natural selection too.

But establishing whether the theory is consistent with, or derivable from, physical theories again requires that we have an agreed upon, clear, and explicit understanding of exactly what the theory claims. Providing such an account of the exact content of the theory of natural selection is a central concern of the philosophy of biology.

So quite a number of questions hover around the theory of natural selection. There are questions about whether the theory is true and nontautologous. There are questions about scope, about how much of the biology of organisms it explains. And then there are questions about the relationship of the theory to theories in other fields, especially in physics and chemistry. Finally, there are questions about whether the theory, even if well supported by data and theory in the natural sciences, has implications for the social and behavioral sciences. All of these require that we get straight about the meaning and structure of the theory of natural selection. And this is a job for philosophers.

Summary

Organisms show such pervasive adaptation to their environments that in explaining them it is hard to escape the notion of purpose: plants have chlorophyll *in order to* produce starch. But we have seen that a world of physical causes such as Newton described has no room for real ends, goals, or purposes, at least none that cause the events that lead up to their attainment. The future cannot cause the past. Accordingly, biology needs an alternative account of how its "teleology" is possible, or it needs to banish teleology from its descriptions and explanations.

Before Darwin this problem was solved by appeal to the existence of an omnipotent and benevolent designer, God. The weaknesses of this appeal were known: lack of predictive power, absence of independent evidence for God's existence, incoherences in the theist's definition of God. But there was no alternative explanation of the adaptedness of living things.

Darwin's theory provides a far more scientifically attractive theory, one free from the specific weaknesses of theism, and easy to link up with the rest of science. It provides a purely causal, nonteleological explanation of biological processes and structures that exploits what physical science tells us about them. But Darwin's theory has its own conceptual problems, to which must be added the common misunderstandings of the theory. Solving these problems and dissipating these misunderstandings is crucial to understanding the nature of biology.

Among the misunderstandings of Darwin's theory are the ideas that evolution enables us to identify organisms such as humans as "higher" and others such as yeast as "lower," that there has been persistent progress in evolution from more primitive to more sophisticated, that natural selection is powerful enough to provide perfect solutions to "design problems," and that evolution by natural selection is a wholly random process that miraculously and improbably produces order from disorder.

Among the criticisms it faces are the claims that it is no more empirically testable than the theory of God's design that it replaces, that it leaves unexplained various imperfections and defects among biological creatures, and that it is hard to reconcile with other more established parts of science such as physics. These are serious problems but they are ones that biologists typically ignore. We will make it our business to either answer these criticisms or show why they must be taken seriously.

Suggestions for further reading

There is no substitute for reading Darwin's own words. *On the Origin of Species* is, as he put it, "one long argument," and readers always find insights and evidence that no summary or textbook version of the theory can provide. His complete works are available online at http://darwin-online.org.uk/contents.html. The first edition of the *Origin* is "pure" Darwin; later editions added qualifications, and even an invocation of God in the book's last paragraph that was not there in the first edition. Modern "ultra-Darwinist" expositions of the theory of natural selection include Dawkins's *Blind Watchmaker*, and Dennett's *Darwin's Dangerous Idea*. These works may be contrasted with the many scientific and popular works of Stephen J. Gould, in whose view Darwin is far less univocal and more qualified in his commitment to natural selection as the agency shaping evolution. This matter is taken up at length in the next chapter. Gould's magnum opus is *The Structure of Evolutionary Theory*. More popular expositions of his view can be found in *Ever Since Darwin* and *The Panda's Thumb*.

Ernst Mayr, a renowned evolutionary biologist, wrote a number of histories of biology, and especially evolution, and participated in many of the controversies among philosophers and biologists over his 100-year lifetime. *The Growth of Biological Thought* develops many of his views about Darwinism while tracing its prehistory as well as its post-Darwinian fate. Other important works on the reception and interpretation of the theory of natural selection by philosophers of biology include Michael Ghiselin's *The Triumph of the Darwinian Method*, and Michael Ruse's *The Darwinian Revolution*.

Jonathan Hodge and Gregory Radick's *The Cambridge Companion to Darwin* contains a collection of papers by leading contributors to the philosophy of biology identifying the major implications of Darwin's theory for a range of philosophical issues.

William Paley's *Natural Theology: Or Evidences of the Existence and Attributes of the Deity Collected from the Appearances of Nature* represented the best explanation of adaptation prior to Darwin, an inductive argument for the existence of a Designer. Darwin studied these volumes carefully as a university student. Excerpts from this work are widely available in introductory anthologies of philosophy.

Many academic journals publish papers on Darwin, his theory, and its philosophical significance. The ones focusing steadily on these subjects include *Biology and Philosophy* and *Studies in the History and Philosophy of the Biological and Biomedical Sciences*.