EVIDENCE AND EVOLUTION

The logic behind the science

ELLIOTT SOBER

CAMBRIDGE
UNIVERSITY PRESS

Contents

Preface

Biologists study living things, but what do philosophers of biology study? A cynic might say "their own navels," but I am no cynic. A better answer is that philosophers of biology, and philosophers of science generally, study science. Ours is a second-order, not a first-order, subject. In this respect, philosophy of science is similar to history and sociology of science. A difference may be found in the fact that historians and sociologists study science as it is, whereas philosophers of science study science as it ought to be. Philosophy of science is a *normative* discipline, its goal being to distinguish good science from bad, better scientific practices from worse. This evaluative endeavor may sound like the height of hubris. How dare we tell scientists what they ought to do! Science does not need philosopher kings or philosophical police. The problem with this dismissive comment is that it assumes that normative philosophy of science ignores the practice of science. In fact, philosophers of science recognize that ignoring science is a recipe for disaster. Science itself is a normative enterprise, full of directives concerning how nature ought to be studied. Biologists don't just describe living things; they constantly evaluate each other's work. Normative philosophy of science is continuous with the normative discourse that is ongoing within science itself. Discussions of these normative issues should be judged by their quality, not by the union cards that discussants happen to hold.

Pronouncements on "the scientific method" all too often give the impression that this venerable object is settled and fixed – that it is an Archimedean point from which the whole world of scientific knowledge can be levered forward. The fact of the matter is that a thorough grasp of scientific inference is a goal, not a given. Like our current understanding of nature, our present grasp of the nature of scientific inference is fragmentary and a work in progress. Scientists themselves disagree about the methods of inference that should be used, and so do statisticians and philosophers. For this reason, the first chapter of this book, on the

concept of evidence, is not a report on a complacent consensus. The position I develop on what evidence means in science is controversial. It is an intervention in the long-standing disagreement between frequentists and Bayesians. I wrote this chapter for neophytes, not sophisticates. No prior understanding of probability is presupposed; I try to build from the ground up.

The methods of inference used in science take two forms. Some are entirely general, in the sense that they apply no matter what the subject matter is. These are the sorts of procedures described in texts on deductive logic and statistics. A method for estimating the average blood pressure in a population of robins is also supposed to apply to the problem of estimating the average weight in a pile of rocks. The different sciences also include methods that are narrower in scope; these methods are tailor-made to apply to a specific subject matter. For example, in evolutionary biology, a concept of parsimony has been developed that underwrites inferences about phylogenetic trees; this method is not general in its subject matter, it applies only to hypotheses about genealogies of a certain sort. The usefulness of this concept of parsimony has been controversial in evolutionary biology. When I consider the role of parsimony considerations in evolutionary biology in Chapters 3 and 4, I again will be intervening in a methodological dispute that is alive within science itself.

When scientists disagree about which of several competing inference methods they should use, it often is fairly obvious that there is a philosophical dimension to their dispute. But philosophical questions also can be raised when there is a thoroughgoing scientific consensus. No competent biologist now doubts that human beings and chimps have a common ancestor. The detailed similarities that unite these two species are overwhelming. It takes a philosopher to see a question in the background – why does detailed similarity provide evidence of common ancestry? Philosophers can ask this question without doubting the good judgment of the scientific community. They want to uncover the assumptions that need to be true for this inference from similarity to common ancestry to make sense. Analyzing inferences that seem to be obviously correct has long been a favorite project for philosophers.

Two grand ideas animate the Darwinian theory of evolution, both in the form that Darwin gave it and also in the form that modern Darwinians endorse. These are the ideas of common ancestry and natural selection. In each case, we can think of Darwinian ideas as competing with alternatives. The hypothesis that the species we now observe trace back to a common ancestor competes with the hypothesis that they

originated separately and independently. The hypothesis that a trait in a species – say, the long fur that polar bears now have – evolved by natural selection competes with the hypothesis that it evolved by random genetic drift and with other hypotheses that describe other possible causes of character change and stasis. Most of Chapters 3 and 4 is devoted to understanding how the Darwinian position can be tested against its competitors. But I also spend time exploring how ideas about natural selection and common ancestry interact with each other. Biologists use information about common ancestry to test hypotheses about natural selection. And inferences about ancestry often rely on information about how various traits have evolved. The two parts of the Darwinian picture are *logically independent* of each other, but they are *methodologically* interdependent.

This book is aimed at philosophers of science and evolutionary biologists. Both tend to have little patience with creationism, so I want to explain why I devote Chapter 2 to its evaluation. I do not think that "intelligent design" is a substantive scientific theory, but I am not satisfied with the standard reasons that have been offered to explain why this is so. For example, Karl Popper's ideas on falsifiability are often used in this context, but philosophers of science have long realized that there are serious problems with Popper's solution to the demarcation problem – the problem of separating science from nonscience. In Chapter 2, I try to develop a better account of testability that clarifies what is wrong with the hypothesis of intelligent design. Another standard critique of creationism begins with the fact that many of the adaptations we find in nature are highly imperfect. It is claimed that an intelligent designer would never have produced such arrangements. I explain in Chapter 2 why I find this criticism of creationism problematic. Although it isn't true that every word of Chapter 2 matters to the material in Chapters 3 and 4, there nonetheless is a through-line from Chapter 1 to Chapters 3 and 4 that passes through Chapter 2. The Duhem–Quine thesis about scientific testing is introduced in Chapter 2 and so is the concept of a fitness function; both play important roles in what comes after.

Chapter 3 begins where Chapter 2 leaves off, by asking whether hypotheses about natural selection are in any better shape than hypotheses about intelligent design. It is not fair switching standards – setting the bar impossibly high when evaluating creationism, but lowering the bar when evolutionary hypotheses are assessed. I begin with the apparently simple problem of explaining why polar bears now have (let us assume) fur that is, on average, 10 centimeters long. Which is the more plausible

explanation: that the trait evolved by natural selection or that it evolved by drift? In the first few sections of Chapter 3, I describe what needs to be known if one wishes to test these hypotheses against each other. The result is a catalog of difficulties. I then argue that the situation is transformed if we take up a different problem: Rather than trying to explain why polar bears have an average fur length of 10 centimeters, we might try to explain why bears in cold climates have longer fur than bears in warm ones. This new problem is easier to solve, and the fact that bears have a common ancestor plays a role in solving it. The rest of Chapter 3 discusses some of the methods that biologists have used to test hypotheses about natural selection; for example, they use DNA sequence data and they also infer the chronological order of the novelties that evolve in a phylogenetic tree.

Chapter 4 addresses a question I mentioned before: Why, or in what circumstances, is the similarity of two species evidence that they have a common ancestor? After developing an answer to this question that is based on the concept of evidence described in Chapter 1, I explore Darwin's idea that similarities that are useless to the organisms that have them provide stronger evidence for common ancestry than adaptive similarities do. Although Darwin's suggestion is right for a large class of adaptive similarities, it emerges that that there is a type of adaptive similarity for which the situation is precisely the reverse. I then consider how intermediate fossils and biogeographical distribution provide evidence concerning common ancestry. The chapter concludes with a discussion of two conflicting methods for inferring phylogenetic trees.

The title of this book may be a little misleading, but I hope that the subtitle corrects a misapprehension that the title may encourage. The title perhaps suggests that this is a book that describes the evidence *for* evolution. There are many good books that do this; they are works of *biology*. The book before you is not a member of that species; rather, it is a work of *philosophy*. My goal in what follows is not to pile up facts that support this or that proposition in evolutionary biology. Rather, I want to describe the tools that ought to be used to assess the evidence that bears on evolutionary ideas. Scientists, ever eager to draw conclusions about nature, reach for patterns of reasoning that seem sensible, but they rarely linger over why the procedures they use make sense. Although this book is not a work of science, I hope that scientists will find that some of the thoughts developed here are worth pondering. I also hope that the philosophers who read this book will be intrigued by the evolutionary setting of various epistemological problems.

CHAPTER 2

Intelligent design

2.1 DARWIN AND INTELLIGENT DESIGN

The first edition of Darwin's *On the Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life* (1859) begins with quotations from two philosophers:

> But with regard to the material world, we can at least go so far as this – we can perceive that events are brought about not by insulated interpositions of Divine power, exerted in each particular case, but by the establishment of general laws. (W. Whewell, *Bridgewater Treatise*)

> To conclude, therefore, let no man out of a weak conceit of sobriety, or an ill-spirited moderation, think or maintain, that a man can search too far or be too well studied in the book of God's word, or in the book of God's works; divinity or philosophy; but rather let men endeavour an endless progress or proficience in both. (F. Bacon, *Advancement of Learning*)

William Whewell was Darwin's contemporary and rejected his theory of evolution, a result that Darwin probably anticipated when he wrote *The Origin of Species*.[1] Francis Bacon wrote more than 200 years earlier. The two quotations are interesting because of what they reveal about Darwin's views on the relationship of belief in God and belief in evolution.

Bacon's remark harks back to an old distinction between the Bible (God's word) and nature (God's work). Sacred texts and natural phenomena provide separate pathways for learning about God. This two-pathway picture was important in the formation of the Royal Society in

[1] The *Bridgewater Treatises* were a series of books that developed the argument for the existence of God that we will consider in detail in this chapter – the argument from design. In the 1833 book from which Darwin drew this quotation, Whewell embraced the view that the origin of species and the origin of languages are beyond the reach of present-day science and are likely to remain so; he argued that both require divine intervention. Darwin's quoting from Whewell does not mean that he expected Whewell to like how he used this passage. See Ruse (1979), Hodge (1991), Brooke (2003), and Snyder (2006) for different views of Darwin's relation to Whewell.

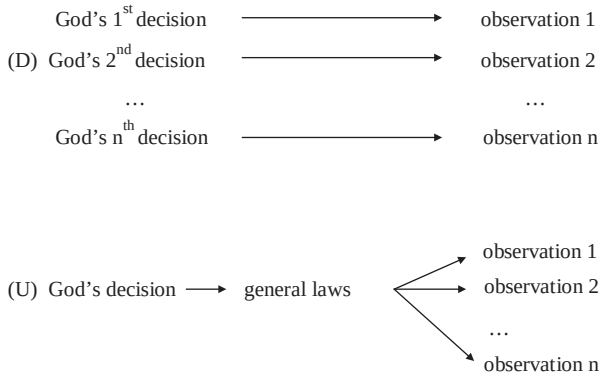

Figure 2.1 Two theistic hypotheses. (D) says that each of our observations traces back to a separate decision made by God; (U) says that God creates a single set of general laws that produces all the details we observe. (D) is a disunified model; (U) is unified.

London and to the philosophy within which the scientific revolution of the seventeenth century developed. The founders of the Royal Society included many clerics who saw "the new science as itself a witness to the Deity's handiwork and therefore to his existence" (Hacking 1975: 169). Darwin quotes Bacon to make a more specific point: that there is no conflict between theism and the theory of evolution. Darwin aimed to describe the processes in nature that account for the features of organisms that we observe; it is logically consistent to add to this biological claim the theological thesis that the evolutionary process occurs because God put it in place. This is the idea that evolution is God's way of making organisms. It now goes by the name *theistic evolutionism*.

The passage from Whewell expresses a different idea. Whewell is discussing the two hypotheses depicted in Figure 2.1. We make a vast number of observations. Should we view each of these observations as the direct result of God's separate decree? Or should we view those observations as knitted together, as flowing from a single cohesive set of laws that God ordained? Whewell's view is that the unified hypothesis is superior to the disunified hypothesis. Darwin used this idea to draw a conclusion that Whewell did not anticipate. The hypothesis that each kind of animal and plant was separately created by an intelligent designer is inferior to the hypothesis that each evolved according to a single set of laws that God created. Darwin's theory of evolution by natural selection was intended to specify the laws that unify the enormous variety of observations we have made and continue to make of the living world (Kitcher 2003).

Darwin went beyond the thesis that the unified hypothesis is *superior* to the hypothesis of disunity; he additionally thought that the disunified hypothesis is *empty*. The "theory" it embodies is easy to state: Whenever you observe something, you simply declare that this is what the designer wanted. In *The Origin of Species*, Darwin ([1859] 1964: 435) puts the point with a touch of irony: "On the ordinary view of the independent creation of each being, we can only say that so it is – that it has so pleased the Creator to construct each animal and plant." If this simple formula were enough to explain the observations in question, there would be no need for science. Not only would Darwin's own theory be unnecessary; there would be no need for theories in any other area of science, either. This does not mean that "God did it" is false, only that it is no substitute for science. Theists regard scientific theories as describing how God brought about various observations, while atheists and agnostics decline to interpret them in this way. Whether this theistic gloss is added or withheld, the practice of science should be the same.[2]

Darwin sometimes failed to live up to his own principles. Consider, for example, the famous sentence that ends *The Origin of Species*:

> There is grandeur in this view of life, with its several powers, having been originally breathed into a few forms or into one; and that, whilst this planet has gone cycling on according to the fixed law of gravity, from so simple a beginning endless forms most beautiful and most wonderful have been, and are being, evolved. (Darwin [1859] 1964: 490)

In saying that life was "breathed into a few forms or into one," Darwin seems to concede that the origin of life is to be understood as the act of a creator. In a letter to J. D. Hooker written four years after *The Origin of Species* appeared, Darwin says what he thinks of his earlier choice of words:

> I have long regretted that I truckled to public opinion, and used the Pentateuchal term of creation, by which I really meant "appeared" by some wholly unknown process. It is mere rubbish, thinking at present of the origin of life; one might as well think of the origin of matter. (Darwin 1887: II, 202–3)

Darwin's considered view was that the origin of life, being an event that occurred *in* nature, needs to be understood in terms of natural processes,

[2] If natural science seeks to answer questions about what happens *in* nature and has nothing to say about *super*natural beings, then "methodological naturalism" is an appropriate scientific research strategy; for discussion of the history of this idea, see Numbers (2003).

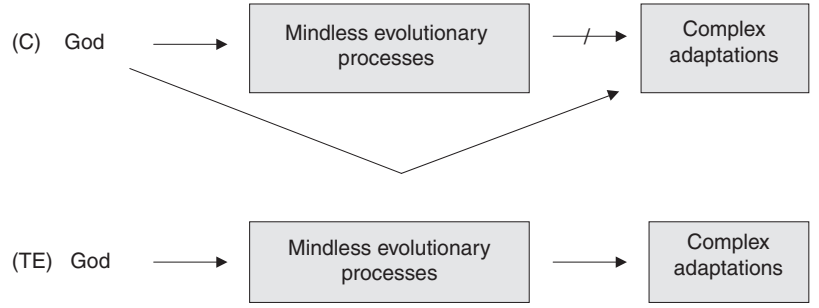

Figure 2.2 Creationism (C) holds that mindless evolutionary processes are incapable of producing the complex adaptations we observe in nature and that God directly produced what we observe. Theistic evolutionism (TE) holds that God indirectly produced complex adaptations in nature by setting the evolutionary process in motion.

not by the facile declaration that it was God's will (Brown 1986). Notice, by the way, how Darwin concludes his book by putting the process of evolution along side "the fixed law of gravity." Newton was a devout theist, but his theism was no substitute for the *Principia*.

It is important not to lose sight of the possibilities that Darwin saw so clearly. Many creationists describe theism and evolutionary theory as if they are incompatible: that the proposition that God exists entails that evolutionary theory is untrue. Some defenders of evolutionary theory agree with creationists on this point (Dawkins 1986; Dennett 1995, 2006; Provine 1989), only they reason from the truth of evolutionary theory to the falsehood of theism. Both parties are invoking a false dichotomy (Ruse 2000). Theistic evolutionism is logically consistent; it also happens to be the viewpoint that many religious people have embraced. My point here is not that theistic evolutionism is *true* or *plausible* (a question to which I'll return in §2.21), but just that it isn't *contradictory*.

Creationism isn't simply the claim that organisms exist and have the features we observe because of a plan that God decreed. This mischaracterization of creationism blurs its difference with theistic evolutionism. The real difference between creationism and theistic evolutionism is depicted in Figure 2.2. Creationists hold that the evolutionary process is fundamentally incapable of producing the complex adaptations we observe; these features require God's *direct* intervention. For theistic evolutionists, God produces complex adaptations *indirectly*, by way of the natural processes he put in place. In addition to creationism and theistic evolutionism, there are other possibilities, such as atheistic evolutionism and agnostic

evolutionism.[3] All these positions are internally consistent; biology does not address the question of whether the universe was created by an intelligent designer. The theory of evolution, which is a theory about *living* things, not about *the origin of the entire universe*, is silent on the question of whether there is a God.

2.2 DESIGN ARGUMENTS AND THE BIRTH OF PROBABILITY THEORY

Bacon's distinction of God's word from God's works means that there are two types of theology: revealed and natural. Revealed theology bases its claims on sacred texts; natural theology seeks to ground its claims on natural observations. When creationism is developed within the context of natural theology, its central organizing concept – both before and after 1859 – is the design argument for the existence of God. Creationists *after* Darwin reject evolutionary theory and claim that the only plausible explanation of the complex adaptations we observe in nature is the hypothesis of intelligent design. Creationists *before* Darwin of course could not have considered Darwin's theory, but they too argued that the complex adaptations we observe in nature provide compelling evidence for the existence of an intelligent designer.

The design argument has evolved, so it might be better to regard it as a family of arguments. I don't propose to give anything like a complete account of the history of this family; rather, in this section I want to mention a few landmarks. My goal in this chapter is to arrive at the strongest, most defensible, version of the argument, and then to say why I think the argument is defective.

The design argument differs from the cosmological argument. The latter argues that the universe as a whole is the result of a first cause (i.e., God). In contrast, the design argument is a claim about what we find *in* nature, not about the existence *of* nature as a whole. The argument from design usually focuses on the complex adaptive features that organisms possess, but other versions of the design argument have been stated. Kepler, for example, believed that the face we see when we look at the moon requires explanation in terms of intelligent design; Newton had the same thought concerning the planets' revolving around the sun in the same direction and in the same plane. And, more recently, the fine-tuning

[3] And, of course, the *rejection* of evolutionary theory is logically consistent with theism, atheism, and agnosticism.

argument maintains that the values of the fundamental physical constants show that the universe was created by an intelligent designer (see Sober 2004b for discussion). These extrabiological examples will not concern us in what follows; our subject will be the *organismic* design argument.

Here is a classic formulation of the design argument developed in the thirteenth century by Thomas Aquinas in his *Summa Theologica* (Part I, Question 2, Article 3):

> We see that things which lack intelligence, such as natural bodies, act for an end, and this is evident from their acting always, or nearly always, in the same way, so as to obtain the best result. Hence it is plain that not fortuitously, but designedly, do they achieve their end. Now whatever lacks intelligence cannot move towards an end, unless it be directed by some being endowed with knowledge and intelligence; as the arrow is shot to its mark by the archer. Therefore some intelligent being exists by whom all natural things are directed to their end; and this being we call God.

To which objects in nature is this argument intended to apply? Aquinas thinks that the adaptive features of plants and animals that lack minds must be explained in terms of intelligent design. Yet, the human mind does not fall within this argument's purview. If the arrow's trajectory requires the archer's mind, doesn't the archer's mind also require an intelligent creator? As we shall see shortly, more recent versions of the design argument usually point to the complexity and functionality of an organ as evidence that it was produced by an intelligent designer, thus providing an opening for the argument to offer an explanation of the human mind. It also is worth noting that Aquinas takes his argument to extend beyond the realm of biology. Being an Aristotelian, Aquinas thinks of lifeless physical objects as goal-directed systems. When we drop an object, it falls towards the center of the earth; for Aristotelians, this is the goal the object has, and it falls in order to achieve that end. The mechanical philosophy that developed within the scientific revolution of the seventeenth century discarded this teleological conception of the behavior of physical objects. Falling objects, planets, and projectiles obey laws, but the idea that they have goals or purposes gradually lapsed from scientific discourse. It might be thought that Darwin did for biology what Newton and others did for physics two centuries earlier – that is, that Darwin demonstrated that it is a mistake to regard organisms as goal-directed systems. I do not agree. Darwin's biology and the evolutionary biology that he inspired seek to understand how features of organisms contribute to their survival and reproduction. These are the

functions those features subserve. For example, the function of the heart is to pump blood; its function is not to make noise. Darwinism does not reject these claims; rather, it provides a framework within which they can be understood. The heart evolved because there was selection for pumping blood; it did not evolve because there was selection for making noise. The concept of function does not require that the heart be capable of conscious striving.[4]

Aquinas's formulation of the design argument concludes that there is one intelligent designer responsible for the goal-directed behavior of all mindless objects. However, what follows is something more modest: that for *each* such object there must be an intelligent designer. It does not follow that *all* such objects trace back to a single intelligent designer.[5] In addition, it is a further step in the argument, requiring further defense, to conclude that this single designer is God. However, the main point to which I want to draw the reader's attention is the connection Aquinas sees between goal-directed behavior and the existence of an intelligent designer. Consider the following two interpretations of his argument:

- If a mindless system exhibits goal-directed behavior, it *must* have been made by an intelligent designer.
- If a mindless system exhibits goal-directed behavior, it *probably* was made by an intelligent designer.

I won't address which of these is the interpretation that Aquinas intended, nor even whether he was aware of the distinction involved here. The point of importance is logical, not biographical: *the distinction between necessity and high probability makes a huge difference for the argument's defensibility*.

The modern mathematical theory of probability began to develop in the seventeenth century. This theory led the design argument to evolve. It became clear to many defenders of the argument that the first of the two versions of the argument just described, which we might

[4] See Wright (1976) for a definition of function according to which function claims are claims about why a feature or organ is present. The present point, that Darwinism does not entail that function talk be discarded, does not require that this definition is correct.

[5] Aquinas's argument commits the *birthday fallacy* (Sober 1990). In the following argument, the premise does not entail the conclusion:

Everyone has a birthday.

There is a single day on which everyone was born.

summarize with the slogan "no design without a designer," is a mistake. The reason is very simple: a mindless random process *can* produce complex and useful devices. It is possible, as we now would say, for monkeys pounding at random on typewriters to eventually produce the works of Shakespeare.[6] The problem is that this outcome, given some fixed number of monkeys and typewriters and a limited amount of time, is very improbable. What is true is that monkeys pounding at random on typewriters *probably* will not produce the works of Shakespeare. For just this reason, it is a mistake for the design argument to claim that complex adaptations *cannot* arise by a mindless random process. The probabilistic formulation of the argument is more defensible.

The birth of probability theory not only transformed the design argument into a probabilistic argument. It also led defenders of the argument who absorbed the point about monkeys and typewriters to think contrastively. Instead of simply declaring that design requires a designer, they were led to consider possible alternatives to intelligent design. Until 1859, the main alternative they considered was *Epicureanism*; here I don't mean the philosophy of eat, drink, and be merry, but the hypothesis due to Epicurus and his followers that physical particles whirling at random in the void eventually combine to produce orderly, stable, and functional arrangements. Design theorists repeatedly held this alternative up to ridicule. For example, Jonathan Swift satirizes Epicureanism in Book 3 of *Gulliver's Travels* (published in 1726) by describing a distinguished professor at the Grand Academy of Lagado who sought to "improve speculative knowledge by practical and mechanical operations"; his innovation was to produce random arrangements of words by twiddling the handles of a device that resembles a foosball game (illustrated in Plate 5 of *Gulliver*). The probability of successfully generating a well-formed sentence of the language – and one that is a new and useful contribution to speculative knowledge as well – is not zero; rather, it is exceedingly tiny. It is not impossible that Chance should produce this result, just very improbable that it should do so.

Swift's satire of Epicureanism may have been inspired by an argument that Richard Bentley, an important figure in the Royal Society, made from "linguistic combinatorics." In his inaugural Boyle lectures of 1692, Bentley asks what the probability would be that a male and

[6] The earliest source I have been able to find for the metaphor of monkeys and typewriters is Borel (1913); Eddington (1928: 72) says that "if an army of monkeys were strumming on typewriters they *might* write all the books in the British Museum."

a female of the same species should each arise by chance. He answers by proposing an analogy, derived from Cicero's *De natura deorum*, between the gigantic number of sequences that can be constructed from the Latin alphabet of twenty-four letters and the still greater number of arrangements there can be of the 1,000 or more parts that comprise the human body (Shoesmith 1987: 136). For both the complex adaptive features of organisms and the orderly pattern of letters in a book, it is absurd to claim that they are due to chance. However, this is not because a random process *cannot* yield the results we observe; rather, the reason is that the *probability* of these results is tiny if a mindless random process is doing the work.

One landmark in the development of probability theory during this period was a version of the design argument published by John Arbuthnot, who was physician to Queen Anne and inventor of the satirical character John Bull. Arbuthnot's "Argument for Divine Providence, Taken from the Constant Regularity Observ'd in the Births of Both Sexes" appeared in the *Philosophical Transactions of the Royal Society* for 1710. The paper provides a tabulation of eighty-two years of London christening records; more boys than girls are listed for each year. Arbuthnot takes this difference at face value; he must have realized that not every birth was recorded, but he nonetheless assumes that the records reflect a real difference in the frequencies of male and female births. The main part of the paper is given over to the task of calculating the probability that this pattern would obtain if the sex ratio were due to chance. By "chance" Arbuthnot means that each birth has a probability of $\frac{1}{2}$ of being a boy and $\frac{1}{2}$ of being a girl. According to this hypothesis, there being more boys than girls in a given year has the same probability as there being more girls than boys in that year; the chance hypothesis also allows for a third possibility, namely, there being exactly as many girls as boys:

$$\begin{aligned} &Pr(\text{more boys than girls born in a given year} \mid \text{Chance}) \\ &\quad = Pr(\text{more girls than boys born in a given year} \mid \text{Chance}) \\ &\quad \gg Pr(\text{exactly as many boys as girls born in a given year} \mid \text{Chance}) = e. \end{aligned}$$

Although Arbuthnot goes to the trouble of explaining how e might be calculated, the details of his calculation don't matter to the argument; the point is just that for each of the years surveyed, e is tiny. Arbuthnot concludes that the probability of there being more boys than girls in a given year, according to the chance hypothesis, is just under $\frac{1}{2}$, and so the probability of there being more boys than girls in each of eighty-two years

is less than $\left(\frac{1}{2}\right)^{82}$. He further asserts that if we were to tabulate births in other years and other cities, we would find the same male bias. So, the probability of all these data – both the data that Arbuthnot presents and the data that he does not have but speculates about – is "near an infinitely small quantity, at least less than any assignable fraction." The conclusion is obvious: "it is Art, not Chance, that governs."

Arbuthnot also notes that males have a higher mortality rate than females, so that the male bias at birth gradually gives way to an even sex ratio at the age of marriage. "We must observe," he says,

> that the external accidents to which males are subject (who must seek their food with danger) do make a great havock of them, and that this loss exceeds far that of the other sex, occasioned by diseases incident to it, as experience convinces us. To repair that loss, provident Nature, by the disposal of its wise creator, brings forth more males than females.

At the end of the paper, Arbuthnot adds, as a *scholium*, that

> polygamy is contrary to the law of nature and justice, and to the propagation of the human race. For where males and females are in equal number, if one man takes twenty wives, nineteen men must live in celibacy, which is repugnant to the design of nature, nor is it probable that twenty women will be so well impregnated by one man as by twenty.

In Arbuthnot's hands, the design argument begins as an explanation of what *is*, but ends as an argument concerning what *ought to be*.[7]

2.3 WILLIAM PALEY: THE STONE, THE WATCH, AND THE EYE

William Paley published his book *Natural Theology, or, Evidences of the Existence and Attributes of the Deity, Collected from the Appearances of Nature* in 1802; it appeared after more than a century of defenses of intelligent design and attacks on Epicureanism. Paley's was neither the first nor the last, but the way he puts the argument is famous:

> In crossing a heath, suppose I pitched my foot against a *stone* and were asked how the stone came to be there, I might possibly answer that for anything I knew to the contrary it had lain there forever; nor would it, perhaps, be very easy to show the absurdity of this answer. But suppose I had found a *watch*

[7] For discussion of the eighteenth-century reaction to Arbuthnot's argument, and of Darwinian theorizing about sex ratio evolution, see Sober (2007b).

upon the ground, and it should be inquired how the watch happened to be in that place, I should hardly think of the answer which I had before given, that for anything I knew the watch might have always been there. Yet why should not this answer serve for the watch as well as for the stone? Why is it not as admissible in the second case as in the first? For this reason, and for no other, namely, that when we come to inspect the watch, we perceive – what we could not discover in the stone – that its several parts are framed and put together for a purpose, e.g., that they are so formed and adjusted as to produce motion, and that motion so regulated as to point out the hour of the day; that if the different parts had been differently shaped from what they are, of a different size from what they are, or placed after any other manner or in any other order than that in which they are placed, either no motion at all would have been carried on in the machine, or none which would have answered the use that is now served by it. To reckon up a few of the plainest of these parts and of their offices, all tending to one result; we see a cylindrical box containing a coiled elastic spring, which, by its endeavor to relax itself, turns round the box. We next observe a flexible chain – artificially wrought for the sake of flexure – communicating the action of the spring from the box to the fusee. We then find a series of wheels, the teeth of which catch in and apply to each other, conducting the motion from the fusee to the balance and from the balance to the pointer, and at the same time, by the size and shape of those wheels, so regulating that motion as to terminate in causing an index, by an equable and measured progression, to pass over a given space in a given time. We take notice that the wheels are made of brass, in order to keep them from rust; the springs of steel, no other metal being so elastic; that over the face of the watch there is placed a glass, a material employed in no other part of the work, but in the room of which, if there had been any other than a transparent substance, the hour could not be seen without opening the case. This mechanism being observed – it requires indeed an examination of the instrument, and perhaps some previous knowledge of the subject, to perceive and understand it; but being once, as we have said, observed and understood – the inference we think is inevitable, that the watch must have had a maker-that there must have existed, at some time and at some place or other, an artificer or artificers who formed it for the purpose which we find it actually to answer, who comprehended its construction and designed its use. (Paley 1809: 1–3)

Four chapters later, Paley connects his discussion of the watch[8] with a claim about the complex adaptations that organisms have. One of his many examples is the eye: ''Every observation which was made in our first chapter concerning the watch may be repeated with strict propriety

[8] Paley did not invent the analogy with a clock. William Derham, an expert on the mechanics of time pieces, gave the Third Boyle lectures, published as *Physico-Theology* in 1711, developing the analogy between watches and watchmakers on the one hand and the universe and God on the other (Hacking 1975: 169–70).

concerning the eye, concerning animals, concerning plants, concerning, indeed, all the organized parts of the works of nature" (Paley 1809: 11).

How did Paley understand the logic of his argument? He often writes as if complex adaptations *must* be the result of intelligent design. For example, in Chapter 2 he says that

> There cannot be design without a designer; contrivance without a contriver; order without choice; arrangement without anything capable of arranging; subserviency and relation to a purpose without that which could intend a purpose [...] Arrangement, disposition of parts, subserviency of means to an end, relation of instruments to a use imply the presence of intelligence and mind. (Paley 1809: 268–9)

Paley's repetitions make his point more than clear. Yet, in other passages, Paley seems well aware of the relevant fact about monkeys and typewriters. For example, in Chapter 15 he considers the fact that "the eyes are so placed as to look in the direction in which the legs move and the hands work" (an example he may have drawn from Plato's *Timaeus* 44D–45B). The obvious explanation, Paley says, is intelligent design. This is because the alternative explanation is chance; if the direction in which our eyes point were "left to chance [...] there were at least three-quarters of the compass out of four to have erred in" (Paley 1809: 269). Paley here grants that it *is* possible for the adaptive arrangement to arise by chance.

2.4 FROM PROBABILITIES TO LIKELIHOODS

The simple point about monkeys and typewriters shows that it is a mistake to claim that complex adaptive features *cannot* be brought into existence by a mindless random process. In response, I suggested that Paley's argument should be formulated as a claim about probability: That the complex adaptive features we observe were *probably* put in place by an intelligent designer. However, a question arises when we look at the details of how this probabilistic argument should be articulated. It concerns the important distinction drawn in Chapter 1 between the probability a hypothesis has in the light of evidence and the probability that the hypothesis confers on the evidence. This is the distinction between $Pr(H \mid O)$ and $Pr(O \mid H)$; the former, recall, is the *posterior probability* of the hypothesis H, whereas the latter is called (unhelpfully) the *likelihood* of H. When Paley talks about the fact that our eyes point in the direction

in which we walk, he considers $Pr(\text{Observations} \mid \text{Chance})$ and says that this has a value of $\frac{1}{4}$. And when Arbuthnot talks about his sex ratio data, he points out that $Pr(\text{Observations} \mid \text{Chance}) < \left(\frac{1}{2}\right)^{82}$. These assessments say nothing about the value of $Pr(\text{Chance} \mid \text{Observations})$. The question we now need to address is whether a suitably probabilistic version of the design argument should describe the probability of intelligent design, or only its likelihood.

We know from Bayes' theorem (§1.2) that the prior and posterior probabilities of the two hypotheses, intelligent design (*ID*) and chance, are related to their likelihoods as follows:

$$\frac{Pr(ID \mid O)}{Pr(\text{Chance} \mid O)} = \frac{Pr(O \mid ID)}{Pr(O \mid \text{Chance})} \times \frac{Pr(ID)}{Pr(\text{Chance})}.$$

The ratio of the posterior probabilities (the "odds") equals the ratio of the likelihoods times the ratio of the priors. Paley's and Arbuthnot's assessments of intelligent design and chance involve a comparison of their likelihoods:

$$Pr(O \mid ID) \gg Pr(O \mid \text{Chance}).$$

However, this does not entail that

$$Pr(ID \mid O) \gg Pr(\text{Chance} \mid O).$$

To reach that further conclusion, we need further assumptions about the prior probabilities (Keynes 1921: 298; Himma 2005).

As I explained in Chapter 1, I am disinclined to appeal to prior probabilities when they reflect only a subjective degree of certainty; however, I have nothing against priors when they can be justified by sampling data or by an empirically well-established theory. It is for this reason that an assessment of Newton's theory of gravitation or of Darwin's theory of evolution should not be formulated so as to depend on assigning them prior probabilities. The same holds, I suggest, for intelligent design. What is the prior probability that the vertebrate eye was the result of intelligent design? I see no way to answer this question in a way that allows that probability to be objective. For this reason, I don't want to formulate the design argument as an argument that seeks to establish that the hypothesis of intelligent design has high probability. Better to think of the argument from design as a likelihood argument. The law of likelihood (§1.3), applied

to the alternatives that Paley and Arbuthnot considered, says that

$$\text{Observation } O \text{ favors } ID \text{ over Chance if and only if } Pr(O \mid ID) > Pr(O \mid \text{Chance}).$$

Understood in this way, the design argument does not seek to establish that an intelligent designer *must* exist, nor even that such a being *probably* exists. The likelihood argument is more modest than these alternatives, and therein lies its strength. I'll consider two nonlikelihood formulations of the design argument in §2.18 and §2.19. But for now, let's go with likelihoods.

2.5 EPICUREANISM AND DARWIN'S THEORY

Arbuthnot, Paley, and the many other defenders of the design argument who wrote before 1859 naturally did not know about the Darwinian theory of evolution. The alternative to intelligent design that they knew about was Epicureanism. Post-Darwinian creationists often write as if Darwin's theory is nothing new – that evolution by natural selection is just like monkeys and typewriters. This is what they intend to convey when they claim that natural selection has the same chance of producing complex adaptations that a hurricane blowing through a junkyard has of assembling scattered pieces of metal into a functioning airplane.[9]

This analogy is fundamentally misleading. In colloquial usage, a random process is one in which all outcomes have the same (or nearly the same) probability. Gambling devices are the paradigm. A fair coin and an unrigged roulette wheel are randomizing devices.[10] However, the essence of the process of natural selection is that some outcomes are far more probable than others. Traits that help an organism survive and reproduce have a higher probability of evolving than traits that hurt. Natural selection is a *biased* process, not a *random* process. When biologists talk about the "random" element in Darwinian evolution, they usually have in mind the origination of novel variants by mutation. The idea is not that a mutation has a probability of $\frac{1}{2}$ of being advantageous and a probability of $\frac{1}{2}$ of being deleterious; the conventional wisdom is that most mutations are

[9] The analogy is due to the astronomer Fred Hoyle.

[10] Probabilists use "random" in a wider sense. They would apply the term to the sequence of heads and tails produced by repeatedly tossing a highly biased coin (probability of heads = 0.9999 on each toss). Ordinary usage is closer to what probabilists call a "uniform" or "flat" distribution.

deleterious. This is because a random change in a complex functioning machine is very unlikely to improve its performance; it is mutation, not natural selection, that resembles the hurricane blowing through the junkyard. The two-part process of mutation + selection contains a random element and a nonrandom element. Do not confuse the part for the whole.

One important difference between a purely random process and a two-step process in which variation is randomly generated and then there is nonrandom retention of favorable variants concerns *time*; the first process will take longer than the second for advantageous features to evolve. Dawkins (1986) provides a nice illustration of this point by using Simon's example (1981) of a combination lock. Suppose that the lock opens when its nineteen windows spell out *METHINKSITISAWEASEL*. Each window has twenty-six alternative states, one for each letter of the alphabet. If the nineteen tumblers are simultaneously spun at random, the chance that this exact sequence of letters will appear on a given spin is $\left(\frac{1}{26}\right)^{19}$. Imagine doing a very large number of experiments with this combination lock; in each experiment, you repeatedly spin all the wheels at random until all of them fall into place and spell the target sentence. Some experiments will hit the target sooner and others will take much longer. But, *on average*, it will take 26^{19} spins to hit the target; this is the *expected value* (§1.2). Now consider a second kind of experiment. You spin the first wheel at random until it hits *M*; after it hits its target, that letter is frozen in place and the second wheel is spun until it hits *E*, and so on. Imagine doing this kind of sequential experiment a large number of times. Some experiments will achieve the target sentence sooner and some will take much longer. But, on average, this sequential process will take $26 \times 19 = 494$ spins to attain the target sentence. The purely random all-at-once process takes much longer on average than the sequential and partly random process; the nonrandom element in this second process involves the selective retention of letters that match the target. There is a third experiment that will probably hit the target even faster. In this experiment, all the wheels are spun simultaneously until one or more target letter is attained, after which those wheels are frozen and the remaining wheels are spun at random. In this third experiment you don't postpone spinning the second wheel until the first wheel hits the target. In all three of these experiments, the target sentence for the combination lock is set by an intelligent being (the designer of the lock), but that isn't relevant to the present point. The point is that a *purely* random process takes longer to evolve adaptive configurations than the *partly* random *partly* nonrandom process of mutation plus selection. This is why analogizing the process of natural

selection with a hurricane blowing through a junkyard is fundamentally misleading.[11]

Another difference between Epicureanism and Darwinism consists in the fact that the evolutionary process involves branching (as lineages split) and the sequential accumulation of different modifications (both those that are adaptive and those that are not) in different lineages. Darwinism includes the idea of common ancestry as well as the idea of natural selection, and the former introduces an element that was not standard in the Epicurean picture. According to Epicureanism, there was a time of random mixing at the end of which all the stable configurations we now observe had come into existence; it is not intrinsic to this picture that the stable configurations that now exist share common ancestors. To see how this marks an important difference between Epicureanism and Darwinism, consider what Paley says about Epicureanism in Chapter 5 of *Natural Theology* (1809: 49–51). He argues that Epicureanism makes the false prediction that we should see unicorns and mermaids. He also says that it mistakenly predicts that organisms should fail to form a nested taxonomic hierarchy (e.g., *Mammalia* within *Vertebrata* within *Animalia*). There is an irony in Paley's second objection, in that Darwin later claimed, correctly, that his own theory *predicts* hierarchy.[12] Nor is it surprising, on Darwin's theory, that some conceivable organisms do not exist. Evolution by natural selection is path-dependent; the traits that evolve earlier in a lineage constrain the traits that will probably evolve later. One lineage leads to fish and another to human beings; there is no reason to expect this process to

[11] Because he had long admired the philosophical writings of John Herschel, Darwin must have been disappointed to read the following footnote in the 1867 edition of Herschel's *Physical Geography of the Globe* (quoted in Hull 2000: 59):

> We can no more accept the principle of arbitrary and casual variation and natural selection as a sufficient account, *per se*, of the past and present organic world, than we can receive the Laputan method of composing books (pushed *a l'outrance*) as a sufficient one of Shakespeare and the *Principia*. Equally, in either case, an intelligence, guided by a purpose, must be continually in action [. . .] We do not believe that Mr. Darwin means to deny the necessity of such intelligent direction. But it does not, so far as we can see, enter into the formula of his law.

[12] In *The Origin of Species*, Darwin ([1859] 1964: 128–9) says that

> it is a truly wonderful fact . . . that all animals and all plants throughout all time and space should be related to each other in group subordinate to group [. . .] On the view that each species has been independently created, I can see no explanation of this great fact in the classification of all organic beings; but to the best of my judgment, it is explained through inheritance and the complex action of natural selection.

produce organisms that are half-fish and half-human. Perhaps Paley was right about Epicureanism, but Darwinism is different.

The misleading analogy between natural selection and a hurricane blowing through a junkyard should be junked. Darwin suggests something better in his book *The Variation of Animals and Plants Under Domestication*:

> Let an architect be compelled to build an edifice with uncut stones, fallen from a precipice. The shape of each fragment may be called accidental; yet the shape of each has been determined by the force of gravity, the nature of the rock, and the slope of the precipice, – events and circumstances all of which depend on natural laws; but there is no relation between these laws and the purpose for which each fragment is used by the builder. In the same manner the variations of each creature are determined by fixed and immutable laws; but these bear no relation to the living structure which is slowly built up through the power of natural selection, whether this be natural or artificial selection. (Darwin 1876: 236)

Natural selection is no more a random process than intelligent design is. As for the "randomness" of variation, the point is that novel variants do not arise because they would be useful; this does not mean that they are uncaused.

2.6 THREE REACTIONS TO PALEY'S DESIGN ARGUMENT

There are three possible reactions one might have to the design argument, once it is formulated as an argument about likelihoods and we acknowledge that Paley was comparing intelligent design with chance and did not consider Darwin's theory of evolution.

One reaction, now common among biologists, is that Paley reasoned correctly given the alternatives he was considering but that the dialectical landscape shifted profoundly when a third hypothesis was formulated. Translated into the language of likelihoods, this reaction consists in the thought that Paley was right in his claim that

$$(\mathrm{L})\ Pr(\text{Observations} \mid \text{Intelligent design}) \gg Pr(\text{Observations} \mid \text{Chance}).$$

What he could not have anticipated is that

$$Pr(\text{Observations} \mid \text{Darwinian evolution}) \gg Pr(\text{Observations} \mid \text{Intelligent design}).$$

The pattern here is familiar. A better theory displaces an inferior one, but then a new theory comes along that is better still.

The second possible reaction is that Paley's argument is flawed and that it doesn't take the development of Darwin's theory to see what is wrong with the theory of intelligent design. I associate this reaction with David Hume's *Dialogues Concerning Natural Religion*, which appeared posthumously in 1779. Hume, of course, didn't know about Darwin any more than Paley did. The point is that the *Dialogues* present a number of serious criticisms of the design argument (some of which I'll consider in §2.11). If any of these criticisms are correct, they show that there are flaws in Paley's argument that we can recognize without knowing anything about Darwin's theory. And even if Hume's criticisms miss the mark, perhaps there are other criticisms of Paley's reasoning that do not depend on the theory of evolution by natural selection.

The third possible reaction to Paley's argument is the one that post-Darwinian creationists have. Whether or not they acknowledge that chance and Darwinian evolution are different hypotheses, their view, expressed in a likelihood framework, is that Paley was not only correct in asserting the inequality (L); essentially the same argument also shows that

$$Pr(\text{Observations} \mid \text{intelligent design}) \gg Pr(\text{Observations} \mid \text{Darwinian evolution}).$$

This is the thought that intelligent design is better supported by what we observe than both the hypothesis of chance and the hypothesis of Darwinian evolution.

So as to leave no doubt in the reader's mind as to which of these reactions I favor, let me say this: *I stand with Hume*. Although I think that some of Hume's criticisms of the design argument are off the mark, I do think there is a devastating objection to Paley's argument that does not depend in any way on Darwin's theory.

2.7 THE NO-DESIGNER-WORTH-HIS-SALT OBJECTION TO THE HYPOTHESIS OF INTELLIGENT DESIGN

Paley's analogy between the watch and the eye is seductive. Surely Paley reasoned correctly when he considered the watch. And since his reasoning about the eye apparently follows the same pattern, it seems irresistible to conclude that he also reasoned well about the eye. Evolutionists who grant Paley this much usually hasten to point out that Paley did not know about Darwin's theory (much less about the evolutionary biology that developed after 1859); their point is that even though Paley demonstrated that

intelligent design is more plausible than chance, it also is true that Darwinian evolution is more plausible than intelligent design. This is the first possible reaction to Paley's argument described in the previous section.

But why think that Darwin's theory is better supported than the hypothesis of intelligent design? A standard way to defend this assessment is to point to the many imperfect adaptations that are found in nature. This style of argument has a long history. Darwin gives voice to it in Chapter 14 of *The Origin of Species* when he says "on the view of each organic being and each separate organ having been specially created, how utterly inexplicable it is that parts [. . .] should so frequently bear the plain stamp of inutility!" More recently, Stephen Jay Gould (1980) used the example of the panda's "thumb" to make the same point.

The name for this feature is misleading since pandas don't have opposable digits. Rather, they have a spur of bone that sticks out from their wrists. The thumb and the paw together form a V through which the panda repeatedly runs branches of bamboo, laboriously stripping the stalks to get them ready to eat. Pandas spend a large portion of their waking lives at this task. The thumb is extremely inefficient. Gould's point is that no designer worth his salt (the phrase is due to Raddick 2005) would have given the panda this device for preparing its food. A truly intelligent designer would have done better. On the other hand, Darwin's theory of evolution by natural selection says that inefficient devices of this kind are not at all surprising. Darwin thought of natural selection as a gradual process that improves adaptedness; natural selection does not necessarily lead to perfect adaptation, whatever that might mean. Selection modifies the traits found in ancestors by small changes; the result is not that the *best* of all *conceivable* adaptations evolves; rather, natural selection causes traits to evolve that do a *better* job than the alternatives that are *actually* present in the evolving lineage.[13]

I hope it is clear that Gould's argument is a likelihood argument. He claims that the hypothesis of intelligent design makes the panda's thumb very improbable, whereas the hypothesis of evolution by natural selection makes the result much more probable. Creationists have a serious objection to Gould's argument. It can be expressed by a rhetorical question: How does Gould know what God (or some unspecified designer) would have wanted to achieve in building the panda? Gould is assuming that an

[13] Darwin was inspired by Lyell's geology, according to which huge changes in the Earth were brought about by a long series of small alterations. See, for example, Darwin ([1859] 1964: 95) and also Darwin's comment about "absolute perfection" (1964: 202).

intelligent designer would have wanted to supply pandas with a super-efficient device (like a stainless-steel can opener) for preparing bamboo and would have had the ability to achieve this objective. But why is it so clear that God would have wanted to do this? Perhaps God realized that if pandas had better tools, they would eat all the bamboo, which would cause the extinction of the bamboo forest and of pandas as well. And maybe these two extinctions would have triggered a cascade of others. Perhaps God realized that these bad consequences would follow if pandas had better tools, and so he decided to slow them down. Creationists don't need to assert that *they* know what God would have had in mind if he had built the panda. All they need to say is that *Gould* does not know this. Gould adopts assumptions about the designer's goals and abilities that help him reach the conclusion he wants – that intelligent design is implausible and Darwinian evolution is plausible as an explanation of the panda's thumb. But it is no good simply *inventing assumptions* that help one defend one's pet theory. Rather, what is needed is *independent evidence* concerning what God (or some other intelligent designer) would have wanted to achieve if he had built the panda. And this is something that Gould does not have. I think creationists are right to object in this way to Gould's argument.[14] We will see in §2.12 that this good point comes back to haunt the theory of intelligent design.

Paley anticipates the no-designer-worth-his-salt objection. After describing the watch found on the heath, he responds to various objections that might be made against the hypothesis of intelligent design. Here is what he says about one of them:

> Neither [. . .] would it invalidate our conclusion, that the watch sometimes went wrong, or that it seldom went exactly right. The purpose of the machinery, the design, and the designer, might be evident, and in the case supposed would be evident, in whatever way we accounted for the irregularity of the movement, or whether we could account for it or not. It is not necessary that a machine be perfect, in order to show with what design it was made: still less necessary, where the only question is, whether it were made with any design at all.

Paley then says that what is true of imperfect adaptations is also true for traits of unknown function. The watch manifestly has the function of telling time and it is a complex machine. If it is imperfect and if it has parts whose functions are unknown, that does not matter.

[14] Behe (1996: 223) says that "another problem with the argument from imperfection is that it critically depends on a psychoanalysis of the unidentified designer." See also Nelson (1996).