Reconstructing the Past

Parsimony, Evolution, and Inference

Elliott Sober

Chapter 2

Pages 37-69

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

method fulfills whatever requirements we might impose. Indeed, I conjecture that any method, regardless of how plausible it is, can be equipped with a process model that renders it the very soul of reason. But this parlor trick promises to offer little insight into the problem of phylogenetic inference.

A more promising approach would be available if we had confidence that a single detailed process model were correct for the set of taxa under investigation. But, as noted earlier, we often are as much in the dark about process as we are about pattern. In this case, it would be natural to investigate how logically weak the assumptions of a process model can be made while still permitting phylogenetic inference to proceed. But like Aesop's farmer, who killed a cow while trying to see how little food is required to keep it alive, we must realize that the strategy of weakening process assumptions must eventually terminate. At some point, pattern will be irretrievable, because our process assumptions will have become too meager.

However, before we descend any further into the details of the biological problem, it will be helpful to develop a more general perspective on the kind of inference problem we here confront. Parsimony, a central idea in phylogenetic inference, has been discussed as a constraint on nondeductive inference in general. Chapter 2 is devoted to this standard philosophical topic. In Chapter 3, I take up another philosophical lead that promises to illuminate the biological problem. The *principle of the common cause* has been recommended as a plausible constraint on how we should reconstruct the causal ancestry of observed effects. Identifying this principle's strengths and limitations will provide additional purchase on the question of how genealogical relationships are to be ascertained. With these two chapters of philosophy under our belts, we shall take a fine-grained look at the systematic issues beginning in chapter 4.

Chapter 2
The Philosophical Problem of Simplicity

2.1. *Local and Global Parsimony*

In the first chapter, I explained the method of phylogenetic inference called cladistic parsimony. That method maintains that synapomorphies (matches with respect to derived characters) are evidence of genealogical relationship, but that symplesiomorphies (matches with respect to ancestral characters) are not. Equivalently, it holds that the best-supported phylogenetic hypothesis is the one that requires the fewest homoplasies. This latter formulation shows why the method has come to be called "parsimony." The idea is that the best explanation of the data is the one that minimizes a particular quantity.

My main purpose in chapter 1 was to raise a question without answering it: What must be true of the evolutionary process if one is to be rational in using parsimony to make phylogenetic inferences? We have seen that the nonexistence of homoplasies is a *sufficient* condition for parsimony to be right in its judgment that synapomorphies have a significance that symplesiomorphies do not possess. However, sufficiency is not necessity. We have yet to see what cladistic parsimony presupposes. The suspicion may present itself that parsimony actually requires that homoplasies are rare, but so far this is only a suspicion.

Discussion among biologists of cladistic parsimony has in some ways recapitulated discussion of parsimony and simplicity among philosophers of science. Philosophers standardly observe that in any research context, there will be many possible alternative hypotheses that are each consistent with the observations. How is one to choose among them? Philosophers and scientists as well have often claimed that it is part of the scientific method to prefer simple, parsimonious hypotheses. A question then arises: What does this preference presuppose? In particular, does appeal to simplicity commit one to thinking that nature is simple?

In this chapter, I shall trace some of the philosophical ideas that have been developed in answer to this question. The idea that the use of simplicity in scientific inference requires substantive assumptions about the physical world has fallen into disrepute. In this century, virtually all writers in the philosophy of science have rejected this thesis as mistaken or confused. It was not always so. An earlier tradition interpreted the use of a simplicity criterion as resting on a substantive thesis about the nature of the world (section 2.3). We must see why that "ontological" position[1] has been replaced by the idea that parsimony is "purely methodological" (section 2.4). I shall argue that all is not well with this now rather standard philosophical position; other work in the philosophy of science, not explicitly addressed to the nature of parsimony, forces us to reopen the question (section 2.5).

The reader will notice that little will be said in this chapter that is specific to the problem of phylogenetic inference. The reason is that philosophers have tended to discuss parsimony as a *global constraint* on scientific reasoning.[2] They usually have thought that there is a single principle of parsimony—one that finds multiple applications in physics, in biology, and, indeed, in all fields of human inquiry.

A powerful impetus behind this assumption about scientific inference is its plausibility in the case of deduction and the attractiveness of the analogy between deduction and induction. In a deductively valid argument, if the premises are true, then the conclusion must be true as well. In an inductively strong argument, the truth of the premises confirms, makes probable, or provides considerable support for the truth of the conclusion. Philosophers often see deduction as a limiting case of induction—one in which the premises provide maximally strong reasons for the conclusion.

This analogy has not gone unchallenged. But even those philosophers who find flaws in it usually maintain that there is another analogy that is persuasive. A valid deductive argument is valid in virtue of its form, not its subject matter. The following two arguments are both valid. Moreover, they are valid for the same reason, even though their subject matters are entirely disjoint:

All fish swim.	All particles have mass.
All sharks are fish.	All electrons are particles.
———————	———————
All sharks swim.	All electrons have mass.

Logicians represent what these arguments have in common by schematizing them in the following way:

1. An "ontological" thesis is one that is about the way the world is—hence my use of this term for the view that the use of parsimony or simplicity presupposes something substantive about the way the world is.

2. Hesse's review article [1967] is a useful guide to the philosophical literature. In addition to the theories she discusses, proposed by Jeffreys [1957], Popper [1959], Kemeny [1953], and Goodman [1958], more recent work includes Quine [1966], Friedman [1972], Sober [1975], Rosenkrantz [1977], and Glymour [1980].

All B's are C.
All A's are B.
———————
All A's are C.

The letters are placeholders for which any term of the appropriate grammatical type may be substituted, thus yielding a deductively valid argument. Rules of valid deduction thus seem to be invariant over change in subject matter. They are *global*, in that they apply to arguments in *all* sciences.

When philosophers, scientists, and statisticians discuss principles of nondeductive inference, they generally assume that those principles will also be global. Open any statistics text and you will find recipes for calculating confidence intervals, likelihood ratios, measures of goodness-of-fit, etc., that are supposed to apply to any empirical subject matter. The text may discuss the example of estimating the average height in a population of giraffes, but everybody knows that this is only an example. The rules are global, applying to the problem of estimating the mean of *any* attribute in *any* population of objects.

Thus, it is entirely natural for those who believe that parsimony is part of the scientific method to think that it is a global principle. Correctly understood, the principle of parsimony is sufficiently abstract to apply to a problem of adjudicating between competing hypotheses, no matter what those hypotheses are about.

This assumption about parsimony may be correct. But I must stress that it is an assumption. In the case of deductive inference, we are on much firmer ground when we claim that there are valid global principles of inference. The reason is that we can actually state rules of deductive inference that are plausible and then point out that they are invariant over changes in subject matter. But no one yet has been able to formalize a global concept of simplicity that is completely plausible as a constraint on all scientific inference. This does not mean that simplicity is not global, only that our understanding of nondeductive inference is far more rudimentary than our grasp of deduction.

Whatever opinion the reader may have on this matter, the status of cladistic parsimony is much clearer. Cladistic parsimony is stated in terms of the ideas of synapomorphy, symplesiomorphy, and homoplasy. It applies to hypotheses about genealogical relationship, not to hypotheses concerning other subjects. Cladistic parsimony is a *local* principle of nondeductive inference.

How, then, is the global notion of parsimony, which is discussed in the philosophy of science, related to the local notion of parsimony used in phylogenetic inference? That is an important question, which I shall leave open for now. We shall see in chapter 4 that some biologists have suggested that cladistic parsimony is a consequence of global parsimony. The idea is that since the scientific method says that simple hypotheses are preferable to complex ones, systematists should use cladistic parsimony to infer phylogenetic relationships. At this juncture, I take no stand on the correctness of this claim. My present point is to note that it requires an argument: we must begin with *two* notions of parsimony—one local, the other global—and then try to see what connection there is between them.

For now, I shall concentrate on the global notion. Why think that the scientific method[3] includes a principle of parsimony? And does this principle assume anything substantive about the physical world?

This chapter will proceed genealogically. In the next section, I describe the historical origins of the way philosophers in this century have come to think about parsimony and simplicity. Here the focus will be on what such methodological criteria say, not on what they presuppose about the way the world is: I shall try to describe those characteristic forms of inference that have been thought to involve appeals to parsimony and simplicity. In section 2.3, I take up the question about presupposition by describing an older tradition of thought, not now much in vogue, according to which parsimony is a reasonable device to use in scientific inference only because the world has certain contingent properties. The idea that the principle of parsimony is substantive, not purely methodological, has met with criticism, which will be detailed in section 2.4. So the result to this point may seem to favor the idea that parsimony is purely methodological rather than substantive. However, in section 2.5, I develop reasons for reversing this verdict. In the end, I reach a conclusion that parallels the hunch introduced in the previous chapter: *Whenever a scientist appeals to parsimony to justify the conclusion that one hypothesis is more reasonable than another in the light of observational data, substantive assumptions about the world must be involved. In practice, parsimony cannot be "purely methodological."*

2.2. *Two Kinds of Nondeductive Inference*

Discussion of simplicity in twentieth-century philosophy of science traces back to two main sources: one philosophical, the other scientific. This dual ancestry is important, in that the issues stemming from the problem's

3. In referring occasionally to "the" scientific method, I may perhaps give the impression that I hold that there is a single corpus of methods that all scientific disciplines have used at all times. I hold no such thing. Indeed, the conclusion I shall reach about the status of simplicity will point to one way in which a science's methodology must be informed by its substantive picture of the world. At this point, however, I shall assume, just for the sake of argument, that there are global and invariant canons of method.

philosophical roots differ markedly from the ones that emerged from the problem's scientific context.

The philosophical provenance of current thought about simplicity goes back to Hume's problem of induction. In discussing whether induction could be rationally justified, Hume gave prominent place to an idea he called the Principle of the Uniformity of Nature. Twentieth-century philosophers who differed over many points of detail often agreed that nondeductive inference exploits a principle of simplicity. Hume's Principle of the Uniformity of Nature was an ancestor of this idea; the thought that what we observe in the local places and times to which we have access also applies to other regions of space and time—perhaps even to the universe as a whole—is the idea that the world is simple in a certain way. Uniformity—the idea of homogeneity through space and changelessness through time—is a kind of simplicity. In thinking that what we locally observe has a more global application, we are making a simple extrapolation from the observed to the unobserved.

Although a genealogical connection is to be found between Hume and twentieth-century discussion, the descent did not proceed without modification. Contemporary philosophers typically thought that induction is a kind of argument in which premises and conclusion are linked by a particular rule of inference. In trying to characterize this form of argument, they took themselves to be describing a pattern of reasoning that scientists and everyday people frequently follow in the process of attempting to learn about the world.

Hume, on the other hand, thought of induction as a habit, not as an inferential process at all. In the *Inquiry* he says that "the most ignorant and stupid peasants, nay infants, nay even brute beasts, improve by experience and learn the qualities of natural objects by observing the effects which result from them" (Hume [1748, p. 52]). In the cases Hume cites, induction does not involve drawing conclusions from premises that include the assumption that nature is uniform. For Hume, inductions of this sort are not reasoned inferences at all. Hume held that the habit of expecting the future to resemble the past in particular respects is like a knee-jerk reflex. Your leg kicks when the physician hits your knee with a hammer, but no inference and no assumption about the uniformity of nature mediates the transition from stimulus to response.

Although Hume thought that induction often proceeds without the aid of inferential argument, I take it that he did not want to claim that we *never* self-consciously reason from past to future by appeal to the uniformity principle. The examples that Hume cites notwithstanding, it is arguable that we do this in scientific controversy and when our beliefs are challenged in everyday life. Hume's point was that it overintellectualizes what we do to interpret expectations about the future as *always* arising psychologically by a process of inference.

At the same time that Hume derided the idea that our expectations about the future are typically mediated by inference, he also claimed that induction assumes that nature is uniform. Indeed, this latter idea is the core of what is now regarded as Hume's skeptical argument about the rational justifiability of induction. How do we know that the sun will rise tomorrow, based on our observations that it has regularly risen in the past? Hume argues that any attempt to justify this expectation by way of reasoned argument must appeal to the assumption that the future will be like the past. But the thesis that it *will* be—the Principle of the Uniformity of Nature—is not something for which we have empirical evidence. Any attempt to use past observations to support this principle must, thought Hume, beg the question. Nor can the uniformity principle be known by reasoning alone. The principle, Hume concludes, is neither *a posteriori* nor *a priori* justifiable. Hume's question about induction led him to a skeptical conclusion: Our inductive practices rest on habit and custom, and cannot be justified by rational argument.[4]

Thus, Hume held that inference about the future presupposes a uniformity principle, even though he did not think that this is a premise assumed by all who have inductive expectations about the future. But how can *induction* assume something that many, if not all, *inducers* do not even believe? Here we must appeal to a distinction that became quite standard in twentieth-century philosophy of science, between the *context of discovery* and the *context of justification.*

The psychological processes by which people come to hold the beliefs they do about the future are part of the context of discovery. Once those beliefs are formulated, we may inquire into their rational foundations. We may ask what the best arguments are that could be offered on their behalf. When people reason badly or not at all, it will emerge that the arguments proposed in answer to this latter question do not describe the psychological processes of discovery. But if our interest is in the question of justification, this need not matter. We want to find the best case that can be made for the belief in question, whether or not this was what drove the agent to that belief in the first place.

4. Hume did not suggest that we abandon this habit. This we cannot do; the making of inductions, Hume thought, is as much a part of human nature as breathing. Nor did he espouse the anarchistic position that all empirical beliefs are equally irrational; Hume criticized superstition (e.g., he derides belief in miracles) and set forth principles for distinguishing strong inductive inferences from weak ones. However, this rejection of "anything goes" occurs within a context. If we accept as legitimate and unobjectionable the idea that we have inductive expectations, we then can separate strong inductions from weak ones. But once we step back from this inevitable part of life and ask what justifies it, we are bound to find that no rational argument can be produced in its behalf.

Hume's skepticism was a conclusion he reached about the context of justification. Our beliefs about the future are not rationally justifiable. The best argument that can be offered in their favor, Hume thought, was as follows. The premises include observations on the one hand and the Principle of the Uniformity of Nature on the other. Taken together, these license our belief that the sun will rise tomorrow or that tomorrow's bread will nourish just as today's did. That this argument rests on a premise that cannot itself be rationally defended shows that the demand for justifying arguments has gone too far. It does not show that anything is radically amiss in our inductive practices as a whole. This, I take it, is how Hume's "naturalism" combines harmoniously with his "skepticism." The skepticism concerns the possibility of a certain sort of reasoned justification, not the legitimacy of this unavoidable aspect of human life.

In what follows, I shall set aside two elements in Hume's views on induction, important though they may be. First there is Hume's psychological opinion that individuals do not standardly formulate arguments when they form inductive expectations. This claim about the context of discovery will not occupy me further. Nor shall I be much concerned with Hume's skepticism about the rational defensibility of inductive inference. Never mind whether his analysis of inductive inference leads to skepticism. My interest will be in a logical structure I shall call Hume's *rational reconstruction of induction.* This is what Hume took to be the "best case" that could be made for the thesis that our expectations about the future are rationally justified. Hume held that the best rational reconstruction of an inductive inference rests on an assumption—the Principle of the Uniformity of Nature. He held that observations by themselves do not suffice to make one inductive conclusion more reasonable than any other that also is consistent with the observations. An extra premise is needed. Hume's suggestion was that this additional ingredient is a simplicity (uniformity) criterion; once this is added, present observations can be brought to bear on what our future expectations ought to be.

The way Hume formulates his uniformity principle is worth noting. He says that inductive arguments assume something about *the world:* they presuppose that *nature* is uniform or that *the future* will resemble *the past.* Twentieth-century formulations have usually preferred to leave the world out of it. Now a principle of simplicity would probably be formulated by saying that we prefer simple *hypotheses* over complex ones. Given that the sun has risen each day in the past, it is simpler to think that the sun will rise tomorrow than that the sun will fail to rise. It is hypotheses, not nature itself, that now are said to be simple. Eventually we must consider whether this switch from world to words makes any difference. Does using the principle about *hypotheses* require us to believe that *nature* is itself simple?

I have presented Hume's rational reconstruction of induction in terms of an inference from observations of past events to a conjecture about the future. But this is not entirely accurate, nor is it sufficiently general. The past is as much beyond our observations as the future is. What we may rely on, according to this Humean point of view, is our *present* observations and mental states. We seem to remember previous sunrises; we use these present memory traces to ground our expectation about tomorrow. But one might just as well ask why our present memories are reliable indicators of what actually happened in the past. After all, memory is fallible. Our beliefs about the past, if they are rationally justified at all, must be justified in virtue of our present mental states. This is what grounds our belief that the sun has risen each day that we have bothered to make an observation.

And, of course, there is the *distant* past as well: there is the past that happened before any of us made any observations at all. Rather than infer tomorrow's sunrise from ones in the present and the recent past, one could equally ask how we infer a prehistoric sunrise from this same data. Hume's problem embraces knowledge of the past just as much as it does knowledge of the future. And equally caught in its grip are our beliefs about generalizations. If we believe, to change the example, that all emeralds are green, based on the fact that all the emeralds we have examined are green,[5] we might ask why the observations justify belief in the generalization.

According to this Humean picture, we begin with observations of the here and now and our present memory traces as data and then try to infer beyond them. There are three sorts of hypotheses (at least) that we might come to believe on this basis. In each case, our conviction that the available data provide evidence for the conclusion we reach depends on the assumption that nature is uniform. This structure is shown below.

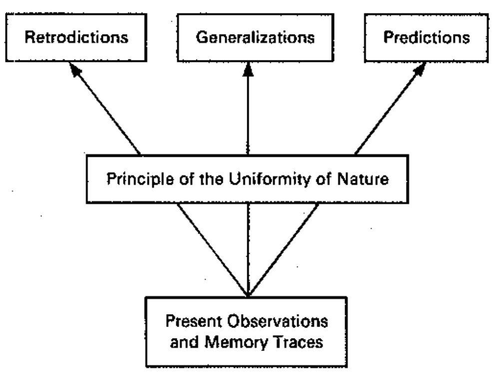

5. Here and in what follows, use of the example of inferring that all emeralds are green will assume that it is not a definitional truth that all emeralds are green.

Many modern philosophers of science have thought that an analog of Hume's insertion of a simplicity (uniformity) criterion into the rational reconstruction of induction is to be found in the so-called curve-fitting problem. Suppose we want to infer the general relationship between two empirical quantities—the pressure and temperature in a chamber of gas, say. To begin with a simple experiment, we might place a sealed pot on the stove and insert a thermometer and a pressure gauge. We then could heat the pot to various temperatures and record the corresponding pressures. The resulting data could be recorded on Cartesian coordinates, each data point representing a single observation. The problem of guessing the general relationship between pressure and temperature in this system then takes the form of deciding which curve to draw. If we are certain that our measurements are perfectly precise, we might demand that a curve pass exactly through each data point. If we are less confident, we instead might require that a curve minimize how much it departs from the data points (this being spelled out by some goodness-of-fit measure). We saw that Hume's problem can be formulated for three sorts of hypotheses—retrodictions, predictions, and generalizations. A parallel multiplicity applies to the curve-fitting problem. We may consider the problem of inferring a general curve from data points or we may focus on issues of interpolation or extrapolation. Instead of asking for the general relationship, I might simply want to know how much pressure the pot would contain if it were heated to some temperature different from the ones I have already produced experimentally.

Philosophers have often used this inference problem to argue for the importance of simplicity in science. Any curve that passes through the data points is consistent with everything we have observed. But there are infinitely many such curves. To choose between them, one must invoke a reason beyond consistency with the evidence. Here it has been customary to invoke simplicity. Smooth curves are simple. Scientists, in preferring the smoothest curve, exhibit their preference for simplicity. The structure of the curve-fitting problem is shown in figure 7.

In both Hume's problem and in the curve-fitting problem, the philosophical thesis has been advanced that a principle of simplicity bridges the gap between observations and hypotheses. However, it might be suggested that a quite different kind of answer should be given to the problems just formulated. When asked how we know that the sun will rise tomorrow, why not appeal to our well-confirmed theory of planetary motion? When asked about the pressure a closed chamber would exhibit if it were heated to a given temperature, why not answer by appealing to our theory of gases? These theories allow us to make predictions, but without any mention of simplicity. If this sort of answer is plausible, our reason for

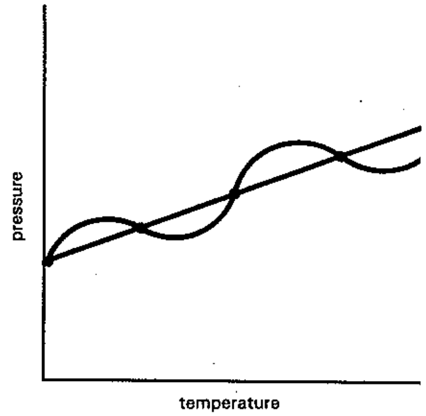

Figure 7.
The curve-fitting problem: Infinitely many curves can be drawn through the data points. Preference for a smooth curve is said to reflect scientists' implicit use of a simplicity criterion in hypothesis evaluation.

thinking that inductive inference must involve an appeal to simplicity seems to have disappeared.

The standard philosophical reply to this suggestion is to ask what justifies our theories of planetary motion and of the kinetics of gases. These ultimately must rest on observations, so Hume's problem must eventually be faced. By appealing to empirical theories, we just postpone having to recognize the use of simplicity in inductive inference. We begin with observations; all our theories, predictions, and retrodictions ultimately trace back for their justification to them and to them *alone*. This conception of how inductive inference works I shall call the *Principle of Empiricism*.[6]

So far I have considered the Humean provenance of modern discussions of simplicity. But there is a second reason that philosophers have had for thinking that simplicity is a central criterion for evaluating hypotheses. It is to be found in that other main source of our present understanding of scientific inference—namely, the arguments found in science itself. Einstein's theories of relativity and the scientific and philosophical discussions of the geometry of physical space that preceded Einstein's work (for example, the writings of Riemann, Gauss, Mach, and Poincaré) exerted a powerful influence on twentieth-century philosophy of science as well.

Rather than attempting to describe the way parsimony figured in Einstein's development of the special and general theories of relativity, I

6. Empiricism is not the truism that the senses are an indispensable source of information about the world. Rather, it is the nontrivial claim that, roughly speaking, observation is a sufficient rational basis for the beliefs we have. See, for example, Popper [1963, p. 54].

shall describe an influential thought experiment of the sort that convinced many writers in the modern period that parsimony crucially influences the kind of picture of the world we construct. Hans Reichenbach [1949, 1951] describes an experiment that Gauss proposed for determining which of the various mathematically consistent geometries is true of the physical space we inhabit. In Euclidean geometry, a triangle always has an angle sum of 180°. However, in Riemannian geometry, the angle sum is greater than 180°. whereas in Lobachevskian geometry the angle sum is less than 180°; in these last two cases, the amount of departure from the Euclidean value is a function of how large the triangle is.

According to Reichenbach, Gauss had tried to solve the question about physical space by measuring the angles of a triangle he set up between the tops of three mountains. The sides of his triangle were light rays. Gauss could detect no departure in his experiment from the Euclidean prediction. On the face of it, this may have been because space really is Euclidean, or because the departure from the Euclidean value was too small to detect, given the size triangle considered and the measurement devices used.

But Reichenbach argued that the problem of testing a geometry is more subtle. Even if Gauss had found a significant departure from the Euclidean value, the possibility would still remain that the light rays did not move in straight lines. For example, if they were acted on by some force that bent them from rectilinear motion, the measurement would not have the implication that Gauss envisioned.

This, too, might seem like a straightforwardly testable matter. After all, in principle we could take a meter rod and see whether the path followed by each light ray is the shortest possible path between the vertices they join—the path traversed by the smallest number of meter rods laid end to end would count as straight. But this suggestion, also, is open to challenge: suppose there were a strange force that not only perturbs light from straight line motion, but also affects the length of our measuring rods. After all, in checking to see if the light moved in straight lines, we assumed that the rods remain the same length as we move them about in the process of measuring.

Suppose this experiment is carried out for a very large triangle, and that the perturbations of all known physical forces are taken into account. We correct for the known physical ways that the signal ray may be affected, and find that the angle sum still departs significantly from 180°. According to Reichenbach, we would then face a choice: we could accept the physical theory we now have at hand, and conclude that space is non-Euclidean. Or we could maintain that space is Euclidean by supplementing our physical theory with a hypothetical force, one that is carefully described so that it cannot be detected by any empirical procedure. This newly postulated force would have no independent confirmation; it is introduced simply to save Euclidean geometry from refutation.

Consider the two total theories we now confront. Each conjoins a geometric claim and a physical one. Reichenbach held that the two theories are *observationally equivalent*; any observation consistent with one is consistent with the other. They may differ, he conceded, in their simplicity. Arguably, it is unparsimonious to postulate an undetectable force. But this difference in parsimony, Reichenbach claimed, is merely aesthetic, since no observational test could ever decide which of the two total theories is true.[7]

Reichenbach's argument might appear to be basically the same as Descartes's puzzle about the evil demon. I take it you believe that you now are seeing a printed page before you. Descartes might ask what justifies this belief as opposed to the hypothesis that your senses are now being misled by an evil demon. If we set up the "normal" and the "evil demon" hypotheses carefully enough, we can ensure that they will be *experientially equivalent*: any experience consistent with one will be consistent with the other. Paralleling Reichenbach, we might remark that it is unparsimonious to think that there are evil demons and then ask why this should count as a reason for thinking one hypothesis true and the other false.

If there were nothing more to Reichenbach's puzzle than this, it would show us nothing *special* about the status of geometry. We would not have discerned a *special* sense in which geometric hypotheses are "conventional" or "untestable" that does not apply with equal force to any hypothesis (such as "there is a printed page before me").

However, for better or worse, twentieth-century philosophy of science did accord such arguments a special significance. It was widely held that choice between empirically equivalent theories on the basis of parsimony was at the core of Einstein's reasoning to the special and general theories of relativity. Although Descartes' problem was a purely philosophical one, it was thought that the family of problems of which the Reichenbach puzzle is an example bore scientific fruit.

The correctness of this argument and its pertinence to scientific questions about the two theories of relativity need not detain us. Whether questions like Reichenbach's show something special about physical theories of space, time, and geometry, or merely show that Descartes' puzzle can be applied to anything, is an important issue I shall not address. I mention this line of argument because it is the second factor that shaped twentieth-century thinking about simplicity. Philosophers and scientists had for a very long time acknowledged the idea we now call Ockham's razor—that "entities should not be postulated without necessity." This methodological maxim took on heightened importance for modern philosophy of science when it was seen as a crucial element in scientific work of the first importance.

From Hume, philosophy of science in our century learned the importance of simplicity in induction;[8] from Mach, Poincaré, and Einstein, it learned the importance of parsimony in theoretical explanation. Reichenbach's [1938, 1949] treatment of these two themes is rather typical in this regard. For hypotheses that are not observationally equivalent—which disagree over some possible observation, whether we shall ever make it or not—a difference in simplicity is a reason for thinking one hypothesis true and the other false. It is simpler to expect the sun to rise tomorrow just as it has in days past, and this counts as a reason for expecting it to do so. However, when the hypotheses differ over no possible observation—as in the case of Descartes' evil demon or, if Reichenbach was right, in the case of total theories of geometry plus physics—we have here a conventional, rather than a substantive, difference. According to Reichenbach and many others, there is no question of deciding which of these hypotheses is true or more plausible, but only of saying which is more convenient.[9]

7. The hypothesis that the sun will rise tomorrow and the hypothesis that it will not are both consistent with my *past* observations. But this does not make them observationally equivalent in the intended sense. Though they agree over what has been observed to date, they do not agree over all *possible* observations. Likewise, hypotheses that agree over the present but disagree about the distant past may be observationally nonequivalent, if a suitably situated observer back then could have gathered pertinent data. The idea of its being "possible" to test two competing hypotheses observationally certainly requires elucidation; philosophies that depend on this idea (like logical positivism) have been heavily attacked for the unclarity of this idea. However, since our purpose here is just to grasp the background of philosophical discussion of simplicity, I shall not pursue this matter further.

8. Even those philosophers who saw themselves as opposed fundamentally to Hume's views about scientific inference were very much influenced by his ideas about uniformity. Thus, Popper [1959] rejects the idea of "induction," but sees simplicity as an indispensable device for comparing competing hypotheses each of which is consistent with the observations. Instead of thinking of induction as a process leading from observations to generalizations, Popper thinks of some generalizations as being better "corroborated." When two generalizations are both consistent with the observations, Popper holds that the simpler is better corroborated in his sense. It was not just "inductivists" who thought simplicity important. Popper's views will be discussed in chapter 4.

9. If simplicity is a reason for thinking one hypothesis true and another false in the case of observationally nonequivalent hypotheses, why does it suddenly become a merely aesthetic consideration when the hypotheses are observationally equivalent? The logical positivists had an answer in their verification theory of meaning: Observationally equivalent hypotheses are synonymous, so assigning different truth values to them would be absurd. However, for those who reject this theory of meaning, the question is a pressing one; I see no reason to think that simplicity is a reason in one kind of problem but not in the other (Sober [1975]).

Puzzles in the philosophy of physics about observationally equivalent theories have no ready analogs in the case of phylogenetic inference. Competing genealogical hypotheses of the kind discussed in chapter 1 are not observationally equivalent, in the sense in which philosophers use that term. So problems peculiar to the case of observationally equivalent hypotheses need not detain us. Nevertheless, the two sources of modern thought about parsimony and simplicity are still both relevant to our inquiry.

Humean ideas about a principle of uniformity pertain to *inductive* arguments. We observe a sample of objects and describe their properties; we then wish to extend that description to objects not in the sample. So, for example, Hume would picture us as noting that each observed *emerald* is *green*, and then asking whether this makes it reasonable to hold that all *emeralds* are *green*. Note that the vocabulary present in an inductive conclusion is already present in the premises. Inductive arguments have this characteristic feature; simplicity as uniformity has been thought to be a maxim that guides this kind of inference.

Parsimony, on the other hand, is not on the face of it a principle concerning smooth extrapolation. It concerns what we should postulate. It applies, not to induction, but to a variety of inference that C. S. Pierce called abduction—to inference to the best explanation. Here we imagine ourselves to confront some observations and to ask which of sereral competing hypotheses best explains them. The hypotheses considered may be stated in a vocabulary that is not already present in the observations. Parsimony in abduction says that one hypothesis is *ceteris paribus* preferable to another when it postulates fewer entities or processes.

Both these simplicity principles have been sketched in only their vaguest outlines. I have not said with any precision how uniformity is to be measured, or how paucity of explanatory machinery is to be gauged. This question of what simplicity is has yet to be answered successfully in any philosophical work. However, it will emerge in what follows that our pressing problems about this philosophical concept can be addressed without a detailed and precise description of what makes one theory, hypothesis, or explanation simpler than another. This is because the main problem for understanding simplicity is not to give a theory that measures this concept with precision, but to describe how simplicity functions within the broader context of hypothesis evaluation.

So, to use a biological formulation, our interest in simplicity will be more in its function than its structure; we want to know what simplicity *does*, not what it *is*. And a point of the first importance has already emerged. Despite the fact that philosophers have not successfully defined uniformity or parsimoniousness, they have in their discussions ascribed to simplicity a very definite methodological role. Whether the problem is inductive or abductive, the role of simplicity is inevitably described as follows: Two hypotheses are both consistent with the observations. Simplicity is then cited as a reason for preferring one to the other. Simplicity plus consistency with the observations is cited as a dual criterion. This description of the role played by a principle of simplicity in scientific inference, we shall see, is pregnant with implications as to whether simplicity is "purely methodological" or demands substantive assumptions about the way the world is.

2.3. *A Lapsed Ontological Tradition*

If one is attracted to the analogy between induction and deduction, and thinks of simplicity as a crucial ingredient in inductive inference, then the idea that simplicity is "purely methodological" may be almost irresistible. It is part of the scientific method, which, like deduction, is a method that is reasonably used to investigate the empirical world, no matter what that world in fact is like. We do not have to understand empirical facts about the subject matter we are investigating before we can use deductive rules of inference. Goodman [1967, pp. 348–349] nicely describes the parallel view about simplicity as holding that "the uniformity required is not in nature's activities but in our account of them. . . . In this version, the Principle of Uniformity does not tell nature how to behave but tells us how to behave if we are to be scientific."

Perhaps the idea that simplicity is "purely methodological" is now the majority view. Be that as it may, matters were not always thus. Hume, I have noted, formulated his principle in terms of the way the world is. The proposition that nature is uniform he regarded as contingent—there being no contradiction, he thought, in its denial. Hume went on to argue that we have no rational justification for this basic ontological assumption. But before Hume, philosophers frequently articulated ontological grounds for the simplicity principles they espoused.

Rather than attempting a full historical account of this ontological tradition, I shall describe one of its most influential practitioners. Newton's views on scientific method assign fundamental importance to maxims of parsimony and simplicity. As we shall see, Newton did not hesitate to justify the use of such principles by appealing to structural features of the world.[10]

10. Another, perhaps more idiosyncratic, representative of this ontological tradition was Leibniz, who held that God created the world so as to maximize the diversity of its phenomena and the simplicity of its laws. The actual world is the best of all possible worlds in this sense. Leibniz extracted detailed methodological advice from this ontological thesis: science was to represent natural phenomena as minimizing or maximizing some relevant quantity, since such "extremal" laws are simplest.

In the *Principia*, Newton lists four "Rules of Reasoning in Philosophy"—the first two emphasizing parsimony, the second two uniformity (Newton [1953, pp. 3–5]):

> 1. *We are to admit no more causes of natural things than such as are both true and sufficient to explain their appearances.* To this purpose the philosophers say that Nature does nothing in vain, and more is in vain when less will serve; for Nature is pleased with simplicity and affects not the pomp of superfluous causes.
>
> 2. *Therefore to the same natural effects we must, as far as possible, assign the same causes.* As to respiration in a man and in a beast, the descent of stones in Europe and in America, the light of our culinary fire and of the sun, the reflection of light in the earth and in the planets.
>
> 3. *The qualities of bodies, which admit neither intensification nor remission of degrees, and which are found to belong to all bodies within the reach of our experiments, are to be esteemed the universal qualities of all bodies whatsoever.* For since the qualities of bodies are only known to us by experiments, we are to hold for universal all such as universally agree with experiments, and such as are not liable to diminution can never be quite taken away. We are certainly not to relinquish evidence of experiments for the sake of dreams and vain fictions of our own devising; nor are we to recede from the analogy of Nature, which is wont to be simple and always consonant to itself . . . .
>
> 4. *In experimental philosophy we are to look upon propositions inferred by general induction from phenomena as accurately or very nearly true, notwithstanding any contrary hypothesis that may be imagined, till such time as other phenomena occur by which they may either be made more accurate or liable to exceptions.* This rule we must follow, that the argument of induction may not be evaded by hypotheses.

The biological reader may find especially interesting Newton's remark in the second rule about "respiration in a man and in a beast." It would be an absurd anachronism to ask whether Newton meant that all similarities are evidence of common ancestry or that only some of them (synapomorphies) are. But we can see here a general principle finding immediate application in the task of explaining organic similarity. Of course, when Newton proposed that respiration in man and beast be traced back to a common cause, he had God in mind, not an ancestral species. But this passage perhaps shows how natural it can be to think that principles of phylogenetic inference flow directly from fundamental principles of scientific method. The idea that matching characteristics cry out for explanation in terms of a common cause will be discussed in chapter 3. For now, I only wish to note that Newton's idea implements an Ockhamite principle of parsimony: Similarities are better explained as stemming from a single (common) cause than as the result of multiple (separate) causes.

Newton, as I said, had God in mind, not descent with modification, as the proper explanation of organic similarity and adaptedness. Like many intellectuals of the time, he saw the design argument as a powerful proof of God's existence (Newton [1953, pp. 65–66]):

> Can it be by accident that all birds, beast, and men have their right side and left side alike shaped (except in their bowels); and just two eyes, and no more, on either side of the face; and just two ears on either side [of] the head; and a nose with two holes; and either two forelegs or two wings or two arms on the shoulders, and two legs on the hips, and no more? Whence arises this uniformity in all their outward shapes but from the counsel and contrivance of an Author? Whence is it that the eyes of all sorts of living creatures are transparent to the very bottom, and the only transparent members in the body, having on the outside a hard transparent skin and within transparent humors, with a crystalline lens in the middle and a pupil before the lens, all of them so finely shaped and fitted for vision that no artists can mend them? Did blind chance know that there was light and what was its refraction, and fit the eyes of all creatures after the most curious manner to make use of it? These and suchlike considerations always have and ever will prevail with mankind to believe that there is a Being who made all things and has all things in his power, and who is therefore to be feared. . . .

I have placed two passages from Newton side by side. In the first, we find that methodology flows from fundamental facts about Nature; in the second, we find that perfection of organic adaptation flows from God. In the following exemplary passage from the *Optics*, Newton traces the perfection (simplicity) of nature in general and the perfection of organic adaptation in particular back to a common source (quoted in Burtt [1932, p. 284]):

> The main business of natural philosophy is . . . not only to unfold the mechanism of the world, but chiefly to resolve these and such like questions. What is there in places almost empty of matter, and whence is it that the sun and planets gravitate towards one and another, without dense matter between them? Whence is it that nature doth nothing in vain; and whence arises all that order and beauty which we see in the world? To what end are comets, and whence is it that planets move all manner of ways in orbs very eccentric, and what hinders the fixed stars from falling upon one another? How came the bodies of animals to be contrived with so much art, and for

what ends were their several parts? Was the eye contrived without skill in optics, or the ear without knowledge of sounds? . . . And these things being rightly dispatched, does it not appear from phenomena that there is a being incorporeal, living, intelligent, omnipresent . . . ?

Here we see Newton claiming that God is the univocal explanation of the mechanism of gravitation, the principle of parsimony, and the adaptedness of organisms. For Newton the perfection of organisms is only part of a larger perfection, one that forms the foundation of the scientific enterprise as a whole.

Newton's view was that the existence of God *suffices* to make science possible; because God created a simple and parsimonious Nature—one that is uniform in space and time and that "does nothing in vain"—methodological maxims of simplicity and parsimony are to be followed. To my knowledge, though, Newton never considered how science would proceed in a Godless universe. Because of this, I abstain from attributing to Newton the view that these methodological principles *presuppose* that God structured the universe in the way He did.

In chapter 1, we saw that a process in which homoplasies are impossible suffices to justify cladistic parsimony. However, from this it does not follow that cladistic parsimony presupposes that homoplasies are rare or nonexistent. The vital distinction was between sufficient conditions and necessary ones. An analogous point arises in the case of Newton's God; Newton asserts a sufficient condition for the use of a global concept of parsimony in scientific method. *If* God created a simple world, then it would seem that science will succeed in uncovering the truth about that world by appeal to a criterion of simplicity. This sufficient condition can be generalized a bit, in that there is no need for *God* to be part of this theory. *Any* process shaping the phenomena we investigate that has the property of making those phenomena simple would suffice. But just as in the case of cladistic parsimony, we must recognize that sufficient conditions need not be necessary. We have no reason, as of yet, to think that global parsimony depends on this assumption about contingent underlying processes.

When Descartes asked how we can know about the world external to the mind on the basis of our experience, he found it necessary to argue for the existence of a God who is no deceiver. Those who reject Descartes' answer must solve his problem in some other way. In similar fashion, we must take seriously the problem that Newton posed, even if we do not accept his theological solution. It is hard to see why we should treat the simplicity of a theory as any indication of its truth unless the processes whereby the phenomena were produced were somehow inclined to make them simple.

The challenge, then, for those who wish to defend the view that parsimony and simplicity are "purely methodological" is to show that these constraints on inference make sense, no matter how the world happens to be structured. How modern philosophy has risen to this challenge we now shall see.

2.4. *The Methodological Critique*

No contemporary philosopher of science would be satisfied with Hume's formulation of the Principle of the Uniformity of Nature. It has become part of the received wisdom in the subject to think that the idea that "nature is uniform" is too vague and amorphous as it stands. This objection, moreover, has accompanied a widespread doubt about the ontological form that Hume gave to his principle. It is not just that the sentence "nature is uniform" is too vague to be worth much; part of the problem, so this consensus concludes, is that Hume tried to describe the presuppositions of induction in terms of some structural feature of the world at large.

Even a first pass over this three-word slogan shows that it is incapable of doing the work in induction that Hume thought it would do. It is quite clear that we do not believe that nature is uniform *in all respects*.[11] The uniformity principle must be refined: how are we to express the expectation of uniformity that supposedly underlies the very activity of trying to learn about the world?

We cannot solve this problem by dividing properties into two classes—the ones we expect to change and the ones we expect to stay the same. We believe that color is constant for emeralds and variable for leaves. On the other hand, if we describe our expectations in a more fine-grained manner, we run the risk of merely restating the inductive beliefs we happen to have. In the case of our belief that all emeralds are green, we are trying to identify the assumption that leads us to think that this hypothesis is well confirmed by the observation of many green emeralds. The "assumption" cannot be that all emeralds are green, at least not if we think that observation has played a nonredundant role in the grounds we have for the belief. Perhaps, then, the presupposition of induction in this case is that emeralds are uniform in color. Given this, even the observation of a single green emerald will lead us to the conclusion that all emeralds are green.

Is it plausible to think that the entire enterprise of inductive inference presupposes that emeralds are uniform in color? This is scarcely credible. It is not hard to imagine how we might come to believe that emeralds are

11. Nor *could* the laws of nature imply that nature is uniform in all respects, a point brought home by Goodman's [1965] grue paradox.

heterogeneous in color. Given the right background beliefs, the observation of green emeralds might lead us to think that green is one among several colors—not the only color—these stones possess.

Even if we withdraw from the hyperbole that the whole inductive enterprise presupposes that emeralds are uniform in color, we still might suspect, more modestly, that if green emeralds are to confirm "all emeralds are green," then one must assume that emeralds are uniform in color. But this also is implausible. For example, the observations would confirm the generalization if one held, instead, that emeralds probably vary in color, but that if one of them is green, the rest probably are too.

The difference between sufficiency and necessity again comes into play. If the vagueness of the idea that nature is uniform could be dispelled, it might be true that an assumption of uniformity would *suffice* for inductive inference to go forward. So, for example, if we were prepared to assume that emeralds are all the same color, then we could see why observing a green emerald would support the generalization that all emeralds are green. But sufficiency is not necessity. We have yet to turn up a shred of evidence for thinking that induction *presupposes* any of uniformity assumptions formulated so far.[12]

This first challenge to Hume's claim that induction presupposes the uniformity of nature I shall call *the respects problem.* We do not even believe that nature is uniform or simple in all respects, so it is difficult to see how all inductive inference presupposes any such thing. If it is replied that the enterprise of induction presupposes that nature is uniform in certain specific respects, this would seem to mean that we cannot use induction to show that nature fails to be uniform in those respects. But this looks quite implausible as well: name any hypothesis that says that nature fails to be uniform in some respect, and it will be possible to show how we could come to have empirical evidence that it is true. And even when we move from trying to formulate a presupposition of *all* inductive inference to the more modest project of trying to specify the presuppositions of a *single* inductive inference, it is hard to see how uniformity is forced on us as an assumption. How, then, can Hume's formulation be even an approximation of anything that deserves to be called a presupposition of all inductive inference?

To the respects argument, I add another, which also makes it implausible to think that Hume was on the right track in saying that inductive arguments always assume that nature is uniform. I shall call this the *no upper bound argument*. A standard picture of inductive inference is that we prefer the hypothesis that is the simplest one consistent with the observations. This suggests that in any inductive inference problem, we formulate a list of hypotheses ordered in terms of their complexity. We then run down the list beginning with the simplest and discard hypotheses inconsistent with the observations until we reach one that is not refuted. This hypothesis is then judged to be more reasonable than the other, more complex, items on the list that also are consistent with the observations.

12. Here I merely restate a long-standing line of argument, endorsed by many writers—for example, by Mill [1859], by Cohen and Nagel [1934, p. 268], and by Salmon [1953, p. 44], who puts the point well when he says that "every formulation of the principle of the uniformity of nature is either too strong to be true or else too weak to be useful."

The important feature of this crude model of how induction proceeds is that it places no upper bound on the complexity of the hypotheses we may have to consider before we find one that is consistent with the observations. It may be the tenth, the hundredth, or the thousandth entry on the list that is the first to go unrefuted. So the hypothesis we prefer by appeal to simplicity may be very complex indeed. In fact, there need be no last entry on the list we are prepared to inspect; for each hypothesis we may consider, we can always construct another that is more complex, yet not beyond the pale.

The idea that nature is simple, if it could be clarified, presumably would imply that there is an upper bound on how complex nature is. But if simplicity functions in hypothesis evaluation in the way just sketched, then its use involves no assumption of an upper bound. This casts further suspicion on Hume's thesis that inductive inference assumes that nature is simple.[13]

The respects argument and the no upper bound argument each suggest that the role of simplicity in hypothesis choice is not happily described by attributing to all inductive inferences the assumption that "nature is simple." But a suspicion may linger: perhaps the idea that the use of simplicity presupposes something substantive about the way the world is can be defended, once this three-word slogan is elaborated.

This brings me to the third argument we need to consider—one that turns on the fact that our inductive practices are enormously *flexible*. If we learn that nature is not uniform in some respect, we incorporate that fact into our set of beliefs and then make inferences in accordance with that knowledge. The idea that simplicity is a criterion in hypothesis choice is quite consistent with the fact that simple hypotheses are often rejected because they clash with observations or with theoretical background assumptions. Inference can be sensitive to background knowledge that the world is a complex place, yet proceed simply and parsimoniously nonetheless.

The flexibility of our inductive practices seems to provide a recipe for refuting the suggestion that inductive inference presupposes that nature is uniform in this or that respect. Once the alleged presupposition is stated

13. The notion of presupposition deployed in this argument will be explored more carefully in section 4.4.

with reasonable precision, we can (i) conceive how empirical evidence might be mustered to show that that uniformity fails to obtain and (ii) show how the use of simplicity and parsimony in nondeductive inference will not only survive that discovery but also play an essential role in inferring that the uniformity in question fails to obtain. The idea is that no discovered complexity can undermine the use of the simplicity criterion, properly understood.

This fact about flexibility suggests that Hume's principle of uniformity is flawed, not just because it is short on details, but because its fundamental outlook is misguided. The source of the problem seems to lie in the fact that Hume's formulation is ontological in character. If so, the way out of the problems we have just surveyed may come with seeing that simplicity involves no assumptions at all about the way the world is. The alternative idea is that simplicity is "purely methodological"; it guides the way we allow observations to shape our judgments about the plausibility of hypotheses, but makes no substantive assumptions about the world those hypotheses purport to describe. The alternative to Hume's ontological formulation, then, is that simplicity, whatever its detailed character turns out to be, is an *a priori* constraint on rational investigation. Its use is consistent with any possible observation and depends on no particular way the world might be.

We have come full circle from the ontological viewpoint epitomized by Newton. That older view, detached from its inessential theological development, saw the principle of parsimony as depending for its justification on the kinds of processes that make the world the way it is. The alternative is to think of simplicity as depending on no such process story at all. This, I believe, is the status that many philosophers now assign the simplicity concept. As plausible as it may seem in the light of the respects argument, the no upper bound argument, and the fact of induction's flexibility, I shall argue in the next section that it is profoundly mistaken. The lapsed ontological tradition described in the previous section contains the germ of an important insight.

2.5. *The Raven Paradox*

Hume investigated the nature of inductive inference at a very high level of generality. Although he discussed various examples—the sun's rising tomorrow, bread's providing nourishment tomorrow, etc.—Hume tried to find one assumption about the world that underlies *all* inductive inference. It was this search for a universal principle that led him to the idea that all induction assumes that nature is uniform.

In the previous section, we found that the substantive principle Hume proposed is seriously flawed. This engendered the suspicion that the principle of simplicity implies nothing about the way the world is, but only constrains how we must reason if we are to be rational.[14]

I now propose to criticize this idea—that simplicity is "purely methodological." However, my line of argument will be strategically different from Hume's. Rather than showing that there is one substantive assumption underlying all inductive inference, I shall suggest that every nondeductive inference from observations to hypothesis must involve substantive assumptions about the world. When we infer tomorrow's sunrise from the ones observed in the past, an additional assumption is involved; the same is true when we infer that tomorrow's bread will nourish from the fact that the bread we ate in the past has done so. But the assumption about sunrises will almost certainly be quite different from the one about bread. Hume went wrong in thinking that there is a *single* uniformity principle that must link premise to conclusion in every inductive inference.

My argument will not focus specifically on simplicity, but will apply to any principle that is said to involve nondeductive inference from observations to hypotheses. Simplicity may be thought to mediate our extrapolation from data to generalization (or to prediction or retrodiction). But even if a connecting principle is described without reference to simplicity, my point remains: *A set of observations confirms, disconfirms, or is irrelevant to a hypothesis only relative to a set of empirical background assumptions.* Confirmation is a three-place relationship between hypothesis, observations, and background assumptions. The same lesson applies to the concept of differential support: *A set of observations supports one hypothesis better than another only relative to a set of empirical background assumptions.*

The bearing of these general claims about confirmation and support on the concept of simplicity is this. Regardless of how simplicity is formulated in detail, it has usually been understood from Hume down to the present as playing a certain epistemological role. *Simplicity has been understood as a principle that takes us from observations to hypotheses.* Given a set of competing hypotheses, simplicity and consistency with the evidence determine which of these hypotheses is to count as "best." My claim is that whenever simplicity performs this function, it embodies empirical assumptions about the way the world is. Explicit mention of empirical background assumptions is often suppressed when an argument appeals to simplicity

14. Not only has this view been widely endorsed as a specificity about simplicity; it also has found support as a claim about induction as a whole. For example, Strawson [1952, pp. 261–262] has argued that "the rationality of induction, unlike its 'successfulness', is not a fact about the constitution of the world. It is a matter of what we *mean* by the word 'rational' in its application to any procedure for forming opinions about what lies outside of our observations. . . ." This is not a universally accepted view, but it is a common and influential one.

or parsimony, but substantive background assumptions there must be nonetheless.[15]

Since confirmation is a three-place relation, there is nothing intrinsic to a set of observations that settles whether they favor one hypothesis rather than another. A background theory *T* may be constructed that implies that the first hypothesis is preferable to the second in the light of the observations, but a different background theory *T'* can be described that implies just the reverse. In the absense of any background theory at all, the observations are powerless to say which competing hypothesis is to be preferred. If "simplicity" or "parsimony" is invoked as the reason for preferring one hypothesis to the other, this must be understood as implicitly assuming something about the underlying background theory. *The fact that confirmation is a three-place relation tells us that simplicity cannot be "purely methodological" when it forges an evidential connection between observations and hypotheses.*

A few words of clarification are needed before I argue for this thesis about confirmation. First, it is important to be clear that the issue here is induction, not deduction. It is obvious that empirical assumptions must be added to observations about emeralds if one wishes to *deduce* that all emeralds are green. But I am not talking about deduction; my claim is that empirical assumptions must be made if we are to claim that "all emeralds are green" is confirmed by green emeralds, or that such observations better support this generalization than one that says that emeralds will change color at the year 2000. Second, the empirical assumptions I have in mind are subject matter specific. Even if "the future resembles the past" were empirical, this sort of assumption is too vague to allow observations to bear on hypotheses. What is required, in each confirmational context, is assumptions about the investigative situation and the subject matter under scrutiny.

The thesis that confirmation and disconfirmation are three-place relations has an exception, which we must identify and set aside. If hypothesis *H* implies an observation statement *O*, then the falsehood of *O* will refute *H*, there being no need to invoke a background theory that connects hypothesis to observations. Likewise, if *O* implies *H* and *O* is true, then so is *H*, no mediating background theory being required here either. But these are enormously special cases, utterly atypical of the relationship between hypothesis and observation. First, it is entirely standard that a hypothesis under test must be conjoined with auxiliary assumptions, if it is to deductively imply anything observable (Duhem [1914]; Quine [1960]). Second, often hypotheses under test, even when embedded in a background theory, do not deductively imply any observation statement. This second point is especially pertinent to testing probabilistic hypotheses, a point to which we shall return in chapter 4.

15. An "empirical assumption" is a proposition that could, in principle, be supported or infirmed by observations, but that, in the context of inquiry, is assumed without argument.

How might one attempt to show that confirmation is a three-place relation between observations, hypotheses, and background assumptions? One possibility would be to describe a fully adequate theory of confirmation and show that a consequence of that theory is that confirmation has the characteristic here attributed to it. If the one true confirmation theory were evident to all, this strategy might be useful. However, there now are a number of conflicting approaches, each confronting its own serious problems. The subject is too much in flux for this strategy to be promising. A second possibility—the one I shall pursue here—is more indirect and perhaps less compelling in its results. It involves arguing from intuitive assessments of examples. As in all such argumentation, even if my analysis of the examples in plausible, this will not force one to the conclusion I wish to draw. Rather, my analysis will take the form of a plausibility argument for the thesis I have in mind.

Earlier in this chapter, I discussed how simplicity has been thought to constrain inductive inference and how parsimony has been thought to constrain inference to the best explanation (abduction). In this section, I shall defend my thesis about the essential role of background assumptions by focusing on a problem about induction; in chapter 3, I shall defend the thesis in the context of a problem concerning inference to the best explanation.

The inductive problem I want to examine is Hempel's [1965b] much discussed paradox of the ravens. Why is it that observations of black ravens confirm the hypothesis that all ravens are black, but that observations of white shoes seem entirely irrelevant?[16] Hempel's question was a request for the specification of the rules we follow in scientific confirmation. Perhaps such rules can withstand Hume's skeptical challenge; perhaps they cannot. But quite independent of this question about justification, one would like to understand how the scientific method works.

Hempel argued that a few simple principles seem to lead to a paradoxical conclusion. Suppose that an hypothesis of the form "All *A*'s are *B*" is confirmed by observing anything that is both *A* and *B*. Suppose next that if an observation confirms a generalization, then it confirms any logically equivalent generalization. So if an observation confirms "All *A*'s are *B*," it also must confirm that conditional's contrapositive, namely, "All non-*B*'s are non-*A*." Applying this reasoning to the raven example, we obtain the following result: If black ravens confirm "All ravens are black," they also

16. Hempel assumes for the purposes of his example that the definition of what it is to be a raven does not settle whether all ravens are black.

confirm "All nonblack things are nonravens." But this latter hypothesis is confirmed by observing anything that is neither black nor a raven. So white shoes also confirm it. It follows, finally, that white shoes confirm "All ravens are black." But this is paradoxical. Hempel set the problem of determining whether black ravens and white shoes really to differ in their confirmational significance.

Hempel's solution was to argue that both sorts of objects confirm the generalization. Appearances to the contrary he dismissed as misguided. But of greater interest to us here is the way Hempel set his problem. We are to consider the hypothesis and its relationship to the observations without assuming anything in the way of empirical background assumptions. This, Hempel [1965b, p. 20] declared, is a "required methodological fiction." Now it is patent that in science and in everyday life, we exploit empirical assumptions in evaluating how observations bear on generalizations. Why did Hempel impose this extreme idealization?

Hempel's problem took the form it did because of the Principle of Empiricism noted in section 2.2. All our knowledge must ultimately trace back to observations and to observations alone. Even if we came to know that all ravens are black for reasons far more complex than the observation of a number of "positive instances," according to the empiricist we could have learned this truth by observation proceeding from a *tabula rasa*. Armed with concepts like "raven" and "black" and with rules of scientific method, we could formulate the generalization and then confirm it by making the relevant observations. The crucial Empiricist Principle is that this is possible even if we begin with no substantive beliefs about the world.

Hempel thought that black ravens and white shoes both confirm "All ravens are black." This is perfectly consistent with saying that black ravens have *greater* confirmational value. Hempel expresses sympathy with the efforts of various workers to show why black ravens are worth more. For example, we know that the world contains fewer ravens than it does nonblack things. This might suggest that looking at a raven and seeing if it is black offers more confirmation than looking at a nonblack thing and seeing if it is a nonraven. Formal theories of confirmation have been developed to ground this intuition.

Hempel is untroubled by the suggestion that the *degree of confirmation* afforded by black ravens as opposed to white shoes depends on an empirical fact about ravens—e.g., that ravens are rarer than nonblack things. But Hempel insisted on his methodological fiction when it comes to the question of whether black ravens and white shoes confirm at all. For him, both sorts of observations confirm the generalization, without the need of any background assumptions. For Hempel, this is a matter of the "logic of confirmation."

Good [1967] argued that whether black ravens confirm, disconfirm, or are neutral depends on what else one believes. He posed the following thought experiment. Suppose we believed that either there are lots of ravens, of which 99% are black, or there are very few ravens, of which 100% are black. This background assumption entails that the more black ravens we observe, the *less* confidence we should have in the hypothesis that all ravens are black. Good thought this showed how a black raven could *dis*confirm the generalization; it does so by reducing the amount of confidence we are entitled to have in the truth of the hypothesis. The title of Good's article was his conclusion: the white shoe is a red herring.

Good also argued that things that are neither A nor B can disconfirm "All A's are B," provided that the problem is embedded in the right background context. His example showed how white crows can *dis*confirm the hypothesis that all ravens are black. Suppose we believe that crows and ravens are biologically related, so that polymorphisms found in one are evidence that they also are present in the other. If we see that crows are sometimes white, this disconfirms the hypothesis that ravens are always black. Good intended a general lesson here: No observation has confirmational meaning, save in the context of a background theory.

Hempel's [1967] reply was that Good's examples fail to address the problem as posed. In the first, Good assumes that we know something about ravens and their color and then shows why black ravens will disconfirm "All ravens are black." In the second, Good assumes we know something about the relationship of ravens and crows and then shows that white crows will disconfirm the generalization that all ravens are black. But Hempel demands that we indulge in his methodological fiction: knowing *nothing* empirical, we are to say whether black ravens and white shoes are both confirmatory.

What would it mean for confirmation to be a relationship between observations and hypothesis alone? In his reply to Hempel, Good [1968] confesses that

> [t]he closest I can get to giving [a two-place confirmation relation] a practical significance is to imagine an infinitely intelligent newborn baby having built-in neural circuits enabling him to deal with formal logic, English syntax, and subjective probability. He might now argue, after defining a crow in detail, that it is initially extremely likely that all crows are black, that is, that H is true. "On the other hand," he goes on to argue, "if there are crows, then there is a reasonable chance that they are of a variety of colors. Therefore, if I were to discover that even a black crow exists I would consider H to be less probable than it was initially." I conclude from this that the herring is a fairly deep shade of pink.

This exchange does not definitively settle the matter of whether confirmation is necessarily a three-place relationship. One might suggest that even though confirmation typically proceeds in the presence of substantive background knowledge, it also must be able to proceed in the absence of such. This claim, central to Hempel's formulation of his problem, is highly conjectural. What is more, it forces the empiricist to defend a peculiarly asymmetrical position. When it is pointed out that observing black ravens can favor the hypothesis that ravens are *heterogeneous* in color, the empiricist will demand to be shown what background assumptions underwrite this surprising extrapolation. Yet, when the extrapolation is from black ravens to the hypothesis that ravens are *homogeneous* in color, the empiricist will see no need to be shown why this "natural" extrapolation makes sense. But why should this be so? Why think that there are extrapolations that are reasonable without there needing to be any background beliefs that make them reasonable? This conviction is based largely on the philosophical faith that empiricism must be true.

It is noteworthy that reasonably developed theories of confirmation that in one way or another exploit the notion of probability all imply that confirmation is a three-place relation.[17] It is perhaps even more noteworthy that theories of confirmation that treat the relation of observation and hypothesis as purely "logical" and presuppositionless have had a dismal track record indeed. I conjecture (following Good [1967] and Rosenkrantz [1977]) that the indispensability of a background theory is a basic fact about confirmation.

If it is a fact, then it is highly significant for our inquiry into the notion of parsimony, in both its global and its local forms. Whenever observations are said to support a hypothesis, or are said to support one hypothesis better than another, there must be an empirical background theory that mediates this connection. It is important to see that this principle does not evaporate when a scientist cites simplicity as the ground for preferring one hypothesis over another in the light of the data. *Appeal to simplicity is a surrogate for stating an empirical background theory.*

How should we understand the idea of a "background theory" in the above principle? We can begin by saying something negative: it is not a set of observation statements. The positive characterization that will suffice for now is vague: a background theory describes how the hypotheses under test and the possible observations are "related," given the design of the experiment. As noted before, we assume that observations do not by themselves deductively ensure the truth or falsity of the hypotheses under test. In the same way, the background theory, in describing the

17. Eells [1982, pp. 58–59] uses a Bayesian analysis of confirmation to isolate a simple empirical assumption that suffices for a positive instance to confirm a generalization.

"relationship" of hypotheses to observations, does not have its truth or falsehood definitively settled by observations either.

In Good's first example, the background theory describes the probability of finding something that is black and a raven, if all ravens are black; in the second, it describes the probability that all ravens are black, if not all crows are black. The "relationships" between observations and hypotheses described by these background theories are probability relationships; these are not given to us *a priori*, nor are they properly treated as mere summaries of the observations we have made. I use the term "background *theory*" to mark the idea that the assumptions that allow observations to bear on hypotheses are not themselves mere observations.

Hume thought that the Principle of the Uniformity of Nature suffices to connect observations to hypotheses inferentially. But this principle, even if its vagueness could be dispelled, is at far too lofty a level of generality to be of much use in inductive inference. If black ravens confirm "All ravens are black," it is not because nature is uniform, but because of some far more specific assumptions about ravens and about the process of sampling.

Hume's problem, like Hempel's, is formulated in such a way that this important relativity to background assumptions is suppressed. This empiricist formulation has encouraged philosophers to believe that confirmation is consistency with the evidence plus simplicity, where "simplicity" is given a purely methodological construal. But once we see the implications of the idea that confirmation is a three-place relation between hypothesis, observations, and a background theory, we see that simplicity must do the work of a background theory; and like a background theory, the use of simplicity in a given context of inference must carry with it substantive implications about the way the world is.

The redescription of inductive inference offered here has consequences for the skeptical argument that Hume advanced. Hume's argument depends on the idea that there is a principle that cannot be justified by reason alone, which is required in all inductive inferences. Because it is required by all inductive inferences, it cannot be justified by any of them. And so we reach Hume's skepticism: the principle must be used in induction, though it cannot be justified either *a priori* or empirically.

I have questioned Hume's claim to have found such a principle. What we do find in any articulated inductive argument is a set of empirical assumptions that allow observations to have an evidential bearing on competing hypotheses. These background assumptions may themselves be scrutinized, and further observations and background theory may be offered in their support. When asked to say why we take past observations to support the belief that the sun will rise tomorrow, we answer by citing our well-confirmed theory of planetary motion, not Hume's

Principle of the Uniformity of Nature. If challenged to say why we take this scientific theory seriously, we would reply by citing *other* observations and *other* background theories as well.

As we pursue these questions of justification—pushing farther and farther back for the "ultimate" assumptions that underlie our empirical beliefs—will we eventually reach a stage where an empirical belief that is not strictly about the here and now is sufficiently supported by current observations, taken all by themselves? This is what the Principle of Empiricism demands. But here we see empiricism in conflict with the thesis about confirmation: If hypothesis H and an observation statement O are not deductively related (they are logically independent), then O confirms or disconfirms H only relative to a background theory T. The third term *never* disappears; there is no room for a simple "principle of induction" (or "principle of simplicity") that takes us directly from observations to reasonable generalizations, retrodictions, or predictions (Rosenkrantz [1977]).

To assess the bearing of this thesis about confirmation on Hume's position, we must take care to separate Hume's argument from the conclusion he reached. First, it should be clear that the *form* of Hume's skeptical argument cannot be reinstated: I see no sign of a premise common to all inductive arguments that cannot be rationally defended. However, this does not mean that Hume's skeptical conclusion is off the mark. If Hume required that we show how present observations all by themselves provide reasonable support for our predictions, retrodictions, and generalizations, he was right to conclude that they do not. The thesis that confirmation is a three-place relation sustains Hume's skeptical thesis, but not the argument he constructed on its behalf.

Although my arguments based on the raven paradox concern induction, I believe they apply with a vengeance to many abductive arguments as well. If empirical background assumptions are needed for observations to have confirmational significance for *observational* generalizations, still more are they needed to connect observations with statements whose vocabulary is not wholly observational.

In the next chapter, we shall consider a principle that has considerable plausibility as a guide to inference to the best explanation: If two events are correlated, then it is better to explain them by postulating a single common cause than by postulating two separate causes. This principle deserves to be viewed as a version of the parsimony idea: One cause is preferable to two. We shall see that this principle depends for its plausibility on empirical background assumptions. This result will help substantiate my claim that the conclusions reached in the present chapter about induction apply to abduction as well.

And, of course, phylogenetic inference is a kind of inference to the best explanation. The general thesis that I have defended about simplicity and parsimony will receive further support if I can show that the power of observations to discriminate among competing genealogical hypotheses depends on the assumptions about the evolutionary process one is prepared to make, and that this is no less true when parsimony is cited as the principle that connects observations to hypotheses. This biological issue will occupy our attention in chapters 4–6.

2.6. *Hume Was Half-Right*

Hume did not err when he thought that an inductive inference from observations to predictions, retrodictions, or generalizations must be supplied with a "missing premise." One cannot cite one, two, or many green emeralds as reason enough for thinking that all emeralds are green, or for thinking that this hypothesis is better supported than one that says that emeralds will remain green only until the year 2000.

In saying this, I am not making the trivial and obvious point that one cannot *deduce* the generalization, prediction, or retrodiction from the observations. A number of commentators have seen Hume as making a point no more important than this, and have gone on to block his skeptical argument as follows. Hume wants to conclude that observations do not rationally justify our beliefs about the future. If Hume's argument for this conclusion is just that one cannot deduce such beliefs from present observations alone, then he must assume that only deductive arguments provide reasons. But this is both question-begging and radically implausible.

My view is that Hume was right in a more significant way. Even granting that nondeductive arguments provide reasons, I still maintain that observations alone are not enough to ground inductive conclusions. Hume was on the right track when he thought some extra ingredient is required.[18]

However, Hume's description of what this additional element must be was defective in both its details and its general character. I see no way of making sense of the idea that "nature is uniform" or that "the future will resemble the past" in such a way that these slogans deserve to be viewed as presuppositions of all inductive inference. What is more, I very much

18. Stove [1973, p. 43] adopts a "deductivist" reading of Hume's argument, according to which induction is said to presuppose the uniformity principle on the grounds that one cannot deduce the conclusion from the observations without it. Beauchamp and Rosenberg [1981] defend the related exegetical thesis that Hume's chief point is that ampliative arguments cannot be demonstrative; however, they also maintain that Hume was not a skeptic about induction. Stroud [1977, pp. 53–77] argues that Hume is better seen as claiming that observations cannot count as reasons—whether deductive or nondeductive—unless the uniformity principle is correct. My gloss of Hume's argument owes a substantial debt to Stroud.

doubt that one can fill in the extra assumptions needed for observations to have an evidential bearing on hypotheses at the level of generality that Hume sought. Hume's mistake was to think that since each inductive inference requires assumptions additional to observations, there must be an additional assumption that every inductive inference requires (see Edidin [1984, p. 286]). Logicians describe this fallacy in terms of the order of quantifiers. "Every person has a birthday" should not be confused with "there is a day on which every person was born." Yet, something like this confusion gives rise to the conviction that inductions about the rising of the sun, the nourishment of bread, and the color of emeralds must share a common premise.

As mentioned earlier, I do not maintain that each inductive inference is utterly unique—that there are no general patterns that the philosopher or statistician can hope to codify. This would be mystery mongering and also contrary to what one can see merely by opening a statistics textbook. A sampling problem about emerald color and one about the nutritional value of bread may have a common structure. But I do not detect in this common structure a general empirical assumption that must be made by all inductive inferences, which therefore cannot be defended by empirical evidence.[19]

Hume's skeptical conclusion concerning the justifiability of induction by a reasoned argument can be reformulated in the light of this criticism of his reconstruction of how inductive inference proceeds. I have argued that each inductive argument must rely on some premise or other whose truth is not guaranteed by present experience and memory traces (or deductions therefrom). The background assumptions required to show how observations have an evidential bearing on hypotheses go beyond what is observed in the here and now. Hume challenged us to show why present observations play the evidential role they do without our assuming additional propositions that require an inductive justification. But this, quite simply, cannot be done. Present experience is no guide to the future, except when it is augmented with contingent assumptions about the connection of past to future. To hold true to the Principle of Empiricism—that beliefs about the future must be justified in terms of present observation *alone*—is, as Hume rightly saw, to be led straight to skepticism.

The bearing of this conclusion about induction on the use of simplicity in scientific inference is indirect but important. Scientists have at different times and in different disciplines used simplicity as a reason for discriminating between competing hypotheses each consistent with the observations. Nothing said here shows that this appeal to simplicity is illegitimate. However, I do think that such appeals to simplicity must be seen for what they are. Whenever a nondeductive inference goes from observations to an evaluation of the plausibility of competing hypotheses, empirical assumptions must be involved. This thesis does not dissolve when simplicity or parsimony is given as the principle that brings observations to bear on hypotheses. Appeals to simplicity must count as highly abstract and abbreviated summaries of background assumptions about the empirical subject and inference problem one faces. Such appeals should not be viewed as unmediated applications of some perfectly general and *a priori* principle of scientific reason.[20]

19. I am not advancing here the stronger thesis that every element in our inductive practices is justifiable in some interesting way. Perhaps there are "primitive postulates" of a logical or mathematical variety that are so ultimate that they elude nontrivial justification. This position concerning the likelihood concept will be discussed in section 5.4.

This conclusion does not contradict the earlier observation that our inductive practices are highly flexible. If it is suggested that induction presupposes that nature is uniform in some specific respect, we can usually show how the negation of that "assumption" is something we can gather evidence on and eventually assimilate into our stock of beliefs; if a principle of simplicity is part of inductive practice, it can facilitate our learning this alleged nonuniformity as well. This suggests that there is no substantive assumption about the world that *all* applications of a principle of parsimony require. But it hardly vindicates the idea that a given application of parsimony to an inductive problem proceeds without substantive assumptions. Again, the order of the quantifiers is at the heart of the resolution I am proposing of the dispute as to whether simplicity is substantive or "purely methodological."

I began this chapter by asking whether the use of parsimony in science makes substantive assumptions about the way the world is. With other philosophers, I have been skeptical of the idea that the use of simplicity assumes that nature is simple. But this does not mean that the use of simplicity requires *no* assumptions—only that the simplicity of nature is not one of them. We have moved away from an overly simple version of the thesis that the use of simplicity has substantive presuppositions; we no longer need formulate this idea as holding that every use of simplicity or parsimomy must make the same assumption. To pursue this matter further, we must descend to a less lofty level of generality; we must see how the use of simplicity in specific scientific inference problems involves nontrivial assumptions. The phylogenetic inference problem is the case study that will flesh out this more general inquiry into the nature of scientific inference.

20. Lyell's defense of uniformitarianism within geology is a nice example of how simplicity can be used to make contingent subject matter specific assumptions look like they are direct consequences of *a priori* methodological principles. See Hooykaas [1959], Rudwick [1970], and Gould [1985] for discussion.