

of reasoning. The fallibilist response to skepticism defended in Chapter 6 clearly relies on the assumption that inductive reasoning can yield justification. The most widely discussed general pattern of inductive inferences is

Argument 7.1: Inductive Pattern

- 1-1. All As examined up till now have been Bs.
 1-2. The next A to be examined will be B.

A standard example used to illustrate this sort of reasoning is the inference from the observation that the sun has risen every day in the past to the conclusion that it will rise tomorrow. To make this example fit the displayed pattern, we must take the stated premise—the sun has risen every day in the past—to mean that the sun has risen on every day that has been observed up till now.

Argument 7.2: The Sun Rise Argument

- 2-1. All days examined up till now have been days on which the sun has risen.
 2-2. The next day (tomorrow) will be a day on which the sun rises.

For present purposes, it is acceptable to take this premise to be true. This seems to be a good example of inductive reasoning.

Not all inductive inferences follow exactly this pattern. The stated pattern draws a conclusion about only the next case. But sometimes from the same premise people draw a general conclusion:

Argument 7.3: The Sun Rise Argument (II)

- 3-1. All days examined up till now have been days on which the sun has risen.
 3-2. All days will be days on which the sun rises.

Although the conclusion here is different, the reasoning is similar. In both Arguments 7.2 and 7.3, observed patterns in previous cases are used to predict the future.¹

In the cases just mentioned, the premise is about *all* the observed As. But some very similar inferences are not. Suppose that every fall a gardener plants some lily bulbs in his garden. Some of the bulbs sprout and some do not. Suppose that over a period of many years the gardener has observed that about 80 percent of the bulbs have sprouted each year. If the gardener is not overly

Skepticism (II)

The arguments for skepticism discussed in Chapter 6 were all arguments for "high-standards skepticism." They relied on the assumption that the standards for knowledge are extremely high, and that we do not, or cannot, satisfy them. Fallibilism and modest foundationalism provide a sensible reply to all these arguments. This, however, is not the end of the discussion of *The Skeptical View*. Another kind of skeptical argument challenges the claim that our reasons for our ordinary beliefs are as good as fallibilists and modest foundationalists think they are. In other words, this kind of argument denies that we meet ordinary standards for justification. In this chapter we will examine two such arguments.

I. THE PROBLEM OF INDUCTION

A. Inductive Inferences

Inductive reasoning is at the heart of science and is crucial to common-sense reasoning as well. Very roughly, inductive reasoning is reasoning that relies on observed patterns to draw conclusions about what occurs in other cases. If you have gone to a restaurant several times and found the food to be very good, you are likely to believe that you will find it to be good the next time you go there as well. If a researcher finds that the patients he has seen with a particular disease always or usually recover when treated in a certain way, then the researcher may conclude that this same pattern will apply to future patients. These are simple examples of inductive reasoning. It is clear that *The Standard View* rests on the assumption that we can learn about the world through this sort

optimistic, he is likely to believe that the same will happen this year. The inference this gardener makes is

Argument 7.4: The Lilies Argument

4.1. 80 percent of the lily bulbs I've planted each year in the past have sprouted.

4.2. 80 percent of the lily bulbs I will plant this year will sprout.

Obviously, additional information might undermine such an inference. For example, if the gardener knows that the weather is predicted to be unusual in the coming months, or if he has purchased bulbs from a different and unknown source, then he might be less likely to draw the conclusion. Still, the general pattern of inference seems correct.

People sometimes think that all inductive inferences are inferences in which one draws a conclusion about the future from premises about the past. But not all inferences that rely on the same style of reasoning are exactly like that. Consider a modification of the lily example. Suppose the gardener is too busy to look in the garden all spring during the season they sprout. At the end of the spring, the gardener might make essentially the same inference, concluding that 80 percent of the bulbs he planted *have* sprouted. So now the inference is entirely about the past, but the reasoning is the same.

The central feature of inductive inferences is thus that they involve inferences from observed cases to unobserved cases. It is sometimes said that the principle upon which inductive inferences rest is that the future will be like the past. But the real principle is that unobserved cases are like observed cases.

It is clear that *The Standard View* and modest foundationalism rely on the epistemic merit of inductive reasoning. It is not just our predictions about what will happen in the garden that are at stake. The justification of your belief that your favorite chair will support you rather than eject you when you sit in it depends upon induction. Much of what we commonly take ourselves to know similarly depends upon the legitimacy of inductive reasoning. There is, however, a long-standing philosophical question about the merits of such reasoning. We turn next to it.

B. Hume's Problem

David Hume raised a question about the merits of inductive inferences that has long troubled philosophers. Stated most simply, Hume's problem (or question) is: (Do we have any good reason to accept the conclusions of inductive arguments? Are these arguments any good?)

One classic statement of Hume's problem is in the following passage:

All reasonings may be divided into two kinds, namely demonstrative reasoning, or that concerning relations of ideas, and moral reasoning, or that concerning matters of fact

and existence. That there are no demonstrative arguments in the case seems evident; since it implies no contradiction that the course of nature may change, and that an object, seemingly like those which we have experienced, may be attended with different or contrary effects. May I not clearly and distinctly conceive that a body, falling from the clouds, which, in all other respects, resembles snow, has yet the taste of salt or feeling of fire?²²

Here Hume says that the reasoning in inductive inferences is not *demonstrative*. That is to say, the conclusions could be false even though the premises are true. This is surely right. In the next passage he goes on to consider the possibility that inductive arguments involve "moral reasoning." By this he does not mean that they involve questions about morality, but rather that "these arguments must be probable only."²³ He writes:

... all our experimental conclusions proceed upon the supposition that the future will be conformable to the past. To endeavor, therefore, the proof of this last supposition by probable arguments, or arguments regarding existence, must be evidently going in a circle, and taking that for granted, which is the very point in question.²⁴

The idea here seems to be that if you think that inductive inferences are good inferences because they have worked, then you are in this very argument relying on the supposition that the future will be like the past. You are thereby assuming in this argument the very thing at issue. The question was: Why think inductive inferences are any good? Why think that the future will be like the past? To make this assumption in arguing that it will is to assume the very thing at issue.

As noted earlier, inductive inferences are really inferences from the observed to the unobserved, and inferences from the past to the future are just a particular case of this. But they are clear and interesting instances of inductive reasoning, and it will be harmless to follow Hume's lead and discuss induction as if it always involved past-to-future inferences. Hume's idea, then, seems to be that inductive inferences turn on some version of a principle such as

PF. The future will be like the past. (Or somewhat more precisely, if x percent of the observed As have been Bs, then X percent of the unobserved As are Bs.)

We could also have formulated this as a uniformity of nature principle, since it says that patterns found to hold in nature will continue to hold. There are details about this principle that need attention. Obviously, a specific inference of this sort is stronger when many As in many different circumstances have been observed. Furthermore, the future will surely not be like the past in all respects. A 49-year-old approaching his next birthday might use (PF) to argue that because on all his birthdays up till now he has been under 50, he will be under 50 on his next birthday. Clearly, something goes wrong here. However, since Hume is challenging the idea that anything at all like (PF) is justified, we will set these details aside.

One way to interpret Hume's remarks is as follows. Inductive inferences depend upon principle (PF) or some variant of it. But (PF) is not a necessary truth; it cannot be proven by a demonstrative argument. And if we attempt to establish (PF) by means of any nondemonstrative (or moral) argument, we rely on (PF) itself. We thereby argue "in a circle" and fail to establish the principle. And there is no other argument for (PF) available. Hume's own view seems to be that the mind is such that we just make these inferences as a result of habit, but there is no real justification for them. That is a disappointing, and skeptical, conclusion. If it is true that science relies essentially on inductive reasoning, it follows that there is no good justification for scientific reasoning. If *The Standard View's* contention that we do know a lot about the world depends upon the adequacy of inductive reasoning, then Hume's argument casts doubt on *The Standard View*. If modest foundationalism implies that inductive reasoning yields justified conclusions, then Hume's problem casts doubt on modest foundationalism. It is not surprising that many philosophers have tried to find an answer to Hume's problem.

We may set out a precise form of *Hume's Argument* as follows:

Argument 7.5: Hume's Argument

- 5-1. If (PF) can be justified at all, then it can be justified either by a "demonstrative" argument or by a "moral" argument (an inference from observed facts).
- 5-2. Only necessary truths can be justified by demonstrative arguments.
- 5-3. (PF) is not a necessary truth.
- 5-4. (PF) cannot be justified by a demonstrative argument. (5-2), (5-3)
- 5-5. All moral arguments assume the truth of (PF).
- 5-6. Any moral argument for (PF) would assume the truth of (PF). (5-5)
- 5-7. Any argument for a principle that assumes the truth of that principle fails to justify that principle.
- 5-8. (PF) cannot be justified by a moral argument. (5-6), (5-7)
- 5-9. (PF) cannot be justified. (5-1), (5-4), (5-8)

This argument is valid. Lines (5-4), (5-6), (5-8), and (5-9) follow from previous steps. So the only legitimate responses will involve the rejection of one of the other premises: (5-1), (5-2), (5-3), (5-5), (5-7). Each of these premises seems quite reasonable. Perhaps we are stuck with Hume's surprising conclusion.

It deserves emphasis that Hume's problem is not dependent, at least not explicitly, on high-standards skepticism. He is not asking how we can be certain that (PF) is true. Instead, he is denying that we have any good reason at all to believe it.

C. Three Responses to Hume's Problem

C1. **Inductive Defenses of Induction** One might be tempted to respond to Hume's problem by pointing out that induction works. That is, we have done well making inductive inferences in the past, so it is reasonable to conclude that it will continue to work. As an argument for (PF), this thought might be formulated as follows:

Argument 7.6: The Inductive Argument for (PF)

- 6-1. (PF) has been true in the past.
- 6-2. (PF) will be true in the future. (6-1)

Perhaps from (6-2) we could go on to infer that (PF) is simply true, and thus that we are justified in making use of it. If Argument 7.6 does justify (PF), it must also expose some flaw in *Hume's Argument*. We will return to this point shortly. Hume, of course, would think that Argument 7.6 is a version of the sort of argument that "goes in a circle" and takes "for granted . . . the very point in question." One way an argument could assume the point in question is by taking that very point as a premise. In this case, since the truth of (PF) is the point in question, an objectionable argument would have (PF) as a premise. But the premise of *The Inductive Argument for (PF)* is not (PF) itself. In addition, (6-1) seems to be quite well justified. Thus, *The Inductive Argument* seems not to take the point in question as a premise. (One response to Hume, then, is that premise (5-5) of his argument is false. Argument 7.6 is a moral argument for (PF), but it does not assume the truth of (PF). The reason Argument 7.6 does not assume the truth of (PF) is that (PF) is not among its premises. And once (5-5) is rejected, the remaining steps of the argument are left without support. *Hume's Argument* seems to be undermined.)

There is, however, a second way in which an argument can assume the point in question. One way is to have the rule as an explicit premise. This is what we have discussed so far. The other way is for the rule to be the one that connects the argument's premise(s) to its conclusion. This is what happens in the case of Argument 7.6. (PF) is a not a premise of the argument, but it is the rule of inference or the principle that connects the premise of Argument 7.6 to its conclusion. If some rule of inference is in question—if we are wondering whether we are justified in using it—then an argument that uses that very rule assumes the truth of that rule. (PF) is the very rule required to get from the premise to the conclusion of Argument 7.6, so that argument does assume the truth of (PF). Hence, premise (5-5) is true after all.⁵ The argument fails to answer Hume's challenge.

The fact that Argument 7.6 does not provide an adequate answer to Hume does not show that the premise of Argument 7.6 is false or that inductive arguments are in general defective. The problem is that this argument does not

establish the legitimacy of using (PF) in a circumstance in which the legitimacy of using it is in doubt.

C2. Pragmatic Defenses of (PF) Some philosophers have noted that inductive reasoning has certain advantages over any competing strategies for forming beliefs about unobserved things.⁶ Two related analogies will bring out the idea. First, consider a situation in which a doctor is going to perform an operation on a patient. The doctor is not certain that the operation will work, but she does know that

A. If anything will work to save the patient, the operation will.

To take a second case, suppose that you are put in the following unfortunate situation:

... you were forcefully taken into a locked room and told that whether or not you will be allowed to live depends on whether you win or lose a wager. The object of the wager is a box with red, blue, yellow, and orange lights on it. You know nothing about the construction of the box but are told that either all of the lights, some of them, or none of them will come on. If the colored light you choose comes on, you live; if not, you die. But before you make your choice, you are also told that neither the blue, nor the yellow, nor the orange light can come on without the red light also coming on. If this is the only information you have, then you will surely bet on red.⁷

In this situation it is true that

B. If any bet will be successful, then a bet on red will be successful.

Advocates of the pragmatic justification of induction contend that something similar is true of induction; They say

C. If anything will work to form accurate beliefs about unobserved things, induction will.

The reason (C) is true has to do with the self-correcting nature of induction. Suppose some rival to induction were being considered. Perhaps reading tea leaves provides an alternative way to form true beliefs about the future. If so, then this pattern will be discovered over time. And inductive reasoning will eventually sanction it. That is, induction will license an argument for the conclusion that predictions based on tea leaf readings are true. If any general policy for forming beliefs proves to work correctly, induction will eventually approve of it. This may take time, so the case for (C) is not exactly parallel to the case for (A) and (B), but it still does provide some sort of vindication of (PF).

(The pragmatic response to *Hume's Argument*, then, is that there is another way to justify (PF) in addition to demonstrative and moral arguments. There is the pragmatic argument just given. Thus, premise (5-1) is false.)

This justification of inductive reasoning may provide less than some would want. First, reconsider the analogies. Even though (A) is true, it does not follow that the operation has much chance to succeed. Even though (B) is true,

there is no reason to think that a bet on red will be successful or is even likely to be successful. The operation and the bet may be nothing more than the best of some very bad options. It is not clear that this defense of induction implies that induction is any better than the best of a set of bad options.

Furthermore, the point about the self-correcting nature of induction is a bit misleading. If nature is uniform, then induction eventually will lead to good principles (presumably), but if it is not, then induction need not. There is no guarantee that induction will yield good principles for forming beliefs about unobserved things.

Finally, if what was sought is a case for the epistemic rationality of (PF), the defense seems to fall short. It does not show that we have good reason to believe that (PF) is true. At most, it shows that we are at least as well off using (PF) as we are using any alternative to it. And that is less than what was sought. Therefore, these considerations suggest that there is no pragmatic justification of induction in the sense of "justification" intended by Hume. *Hume's Argument* has not been refuted.

C3. An *a priori* Defense of Induction *Hume's Argument* is about the past-to-future principle, (PF).

PF. The future will be like the past.

Hume says, correctly, that no "demonstrative" argument establishes (PF). It is not true by definition, and demonstrative arguments are the sort that prove that sort of thing. Hume says that any argument for (PF) based on experience is "circular," or takes for granted the very thing in question. And that seems right as well. In a chapter of *The Problems of Philosophy*, Bertrand Russell attempted to set out the same problem Hume had discussed, but, interestingly, he formulated the principle under discussion somewhat differently. A somewhat simplified version of the claim Russell discussed is

PFR. Knowing that things have been a certain way in the past gives you a good reason to believe that they will be that way in the future.⁸

The key difference between (PF) and (PFR) is that the latter is about what we have reason to believe. If (PFR) is true, then the premises of inductive arguments can provide us with good reasons to believe their conclusions. These reasons, of course, are not conclusive. One can have other reasons that defeat, or undermine, the conclusion of an otherwise good inductive argument. (The 49-year-old who thinks he will still be under 50 at his next birthday has such defeating reasons.) (PF) and (PFR) also differ in a way that is directly relevant to *Hume's Argument*.

Hume is surely right to say that (PF) is not true by definition and cannot be established by means of a "demonstrative" argument. Things could change tomorrow, as he says. (PF) is not a necessary truth. But that same fact does not establish that (PFR) is not true by definition.⁹ One imperfect analogy illustrates why. Suppose there is a jar with 1,000 marbles in it. You know that 999 of the

marbles are black and one is white. You have randomly selected one but not looked at it. You will think you have picked a black one, and that is a reasonable belief. Now compare:

- M1. If there are 1,000 marbles in a jar, 999 are black, 1 is white, and I has been randomly selected, then the one selected is black.
 M2. If you know that there are 1,000 marbles in a jar, 999 are black, 1 is white, and I has been randomly selected (and you have no other relevant information), then it is reasonable for you to believe that the one selected is black.

(M1) and (M2) are related to one another in much the same way (PF) and (PFR) are related to one another. (M1), like (PF), says that if one condition obtains, then another will obtain. (In the case of (PF), the first condition is that some regularity has held in the past, and the second is that it will hold in the future.) (M2), like (PFR), says that if you know that the first condition obtains, then you have good reason to think that the second obtains. (M1) is not necessarily true. Indeed, there are situations in which (M1) is false. It is false in the case in which you select the white marble. In contrast, (M2) may well be necessarily true. Quite plausibly, the definition or nature of the concept of being reasonable makes it true. There is no possible situation in which the situation in the antecedent could be true, yet you would not be reasonable in believing that the selected marble is black.¹⁰ (M2) is something we can know to be true *a priori*; that is, we can know it simply by understanding the concepts involved. We do not have to observe cases and infer its truth.¹¹

Analogously, according to the present response to Hume's problem, (PFR) is true by definition and thus knowable *a priori*. It is part of the concept of being reasonable to use past cases as one's guide to the future. There is no possible situation in which the condition it mentions—knowledge that things have been a certain way in the past—could fail to give you a good reason to think that they will be that way in the future. There may be cases in which that belief is false, and there may be cases in which that good reason is overridden by other reasons (as in the example about the 49-year-old predicting his age on his next birthday). But there are no cases in which information about past regularities fails to provide some reason for beliefs about the future. That is just how being reasonable works.

The response to *Hume's Argument*, then, is that the argument as formulated is sound. (PF) cannot be proven. However, the epistemic merit of inductive reasoning does not depend upon the truth of (PF). Instead, inductive reasoning depends on the truth of (PFR). And, according to the present response, (PFR) is a necessary truth. If *Hume's Argument* were reformulated to be about (PFR), it would have a version of premise (5-3), modified to be about (PFR). That premise would say

- 5-3*. (PFR) is not a necessary truth.

This premise is false. *Hume's Argument* is not sound when modified to be about (PFR).

Notice that this response to Hume does not rely on the assumption that everyone knows that some principle along the lines of (PFR) is true. The claim is not that, in order to know things by inductive inference, people have to study epistemology and come to see that (PFR) is true. Rather, (PFR) is true, and, because it is true, everyone (including those who haven't thought about it) is justified in believing the conclusions of good inductive arguments. In other words, if you are justified in believing the premise of an inductive argument, and have no defeating evidence for its conclusion, then you are justified in believing its conclusion.¹²

D. Induction and Tea Leaves

Critics may think that the solution to Hume's problem just advanced is no better than merely stipulating that induction is reasonable. It might be charged that defenders of any other practice for forming beliefs about unobserved objects could defend their practice in a similar way. For example, if a person's practice is to believe the first thing that pops into his head about unobserved objects, that person might argue that some analogue of (PFR) concerning that method of belief formation is reasonable. Or, to take a more colorful case, consider a tea leaf reader, Madam Malarkey.

Example 7.1: Madam Malarkey, The Tea Leaf Reader

Madam Malarkey uses the configuration of tea leaves to form beliefs about unseen objects. If you want to know something about some object, Madam Malarkey will look into the tea leaves, and by means of some secret formula, use what she sees in them to form a belief about the unseen object. Critics object that her methods are irrational nonsense. Some attempt to challenge her, asking if she has found that her beliefs have proved to be correct. She replies, of course, that to worry about track records and past performance is a sheer prejudice of irrational inductivists. The tea leaves tell her that the tea leaves are the way to proceed. And, when further challenged, she offers an *a priori* defense of her approach. She says that there is a principle that is true by definition:

TLR. Knowing that the tea leaves predict that p will be true provides good reason to believe that p will be true.

Madam Malarkey's defense is, surely, pure malarkey. But is the *a priori* defense of induction any better? Is Madam Malarkey on as good grounds as we are? Perhaps not. There are a few things to be said in behalf of induction.

First, suppose it turns out that we cannot offer a proof that (PFR) is true. It is a mistake to infer that it is false or that our inductive reasoning is not reasonable. Suppose you have some premises (or evidence) and believe something on that basis. It is one thing to say that the conclusion is reasonable only

if the evidence supports the conclusion. It is another to say that the conclusion is reasonable only if you can "show" or "prove" that the evidence supports the conclusion. It is difficult to see why the latter, more demanding, requirement is correct.

Second, we must admit that Madam Malarkey may be unimpressed by our response. But we should be careful to distinguish two projects: convincing intransigent skeptics or fools and seeing if there is a sensible view according to which our ordinary beliefs are reasonable ones. Hume was neither intransigent nor a fool. But it is plausible to think he conflated two principles, (PF) and (PFR). Perhaps he would find some merit in the *a priori* response, even if Madam Malarkey would not.

Third, and most important, there is a good reason to think that the *a priori* defense of induction is superior to Madam Malarkey's defense of (TLR). The reason is based on a distinction between fundamental principles and derivative principles. Some principles are, if true, true only derivatively or as a result of something more fundamental. If someone proposed, as a fundamental principle, that it is reasonable to believe the things reported in a specific newspaper, that claim should surely be rejected. Even if the newspaper is in fact worthy of trust, any principle specifically about the newspaper is a derivative principle. The same is true of (TLR). One can imagine, just barely, situations in which it would be true. Perhaps there are some possible, though unrealistic, situations in which something observable about tea leaves is regularly connected to properties of unobserved objects one is inquiring about. Were such patterns discovered, it would be reasonable to accept (TLR), or some variant of it. But that is not the case in the actual world. In the actual world we have good reason to think tea leaves are not reliable predictors. In any case, (TLR) is the sort of thing that, if true, is at best contingently true. And, in fact, we don't have evidence that it is true. In contrast, (PFR) is not in the same way derivative or contingent. When properly understood, there is no situation in which it is not reasonable to use past patterns (of the appropriate sort) as guides to future results.

Finally, it possible that the idea of inference to the best explanation, to be discussed later in this chapter, can be of some help here. We will return to this point at the end of the chapter.

E. Conclusion

The *a priori* defense of induction provides a plausible response to the problem Hume set for inductive reasoning. The key to the response requires seeing the problem not as one of proving that the future will be like the past, but rather as one of defending the idea that past (or observed) cases are reasonably used as evidence about future (or unobserved) cases. The response relies on the idea that it is an *a priori* fact about the nature of evidence, not a contingent fact about how things are in the actual world, that it is reasonable to use observed cases as evidence.

This defense leaves open many hard questions about inductive reasoning. As noted, it is not true that the future will be like the past in all respects, nor is it reasonable to believe that it will. We know that we will be older in the future than we were any time in the past. Determining exactly which observed patterns it is reasonable to believe will continue to obtain is an exceedingly difficult problem.¹⁵ Nevertheless, the *a priori* defense of inductive reasoning does at least provide a suitable response to Hume's problem. It is safe to conclude that *Hume's Argument* does not undermine scientific reasoning and *The Standard View*.¹⁴

II. ORDINARY-STANDARDS SKEPTICISM AND BEST EXPLANATIONS

A. Alternative Hypotheses and Skepticism

Advocates of *The Skeptical View* have another argument for their view. The argument can be brought to light by raising a simple but difficult question:

What, exactly, is the feature of your evidence that gives you such good reason to think, for example, that you really do see a book rather than that you are dreaming, hallucinating, a BIV, etc.?

Here, the question is not one of certainty. Skeptics who raise this question admit that we need not have certainty in order to have knowledge. However, they contend that if our evidence is good enough to give us knowledge, then our evidence must be good enough to provide good reason to think that our ordinary beliefs are true and the skeptical alternatives are false. However, they claim, when one looks at one's evidence, it is not so clear that our reasons are that good.

The issue can be formulated somewhat more precisely as follows. At any moment, one's current observations are one's present experiences and apparent memories. I now seem to see a computer on a desk, seem to remember seeing the same sort of desk yesterday, and so on. More generally, as I now seem to remember and experience things, my experiences follow patterns. The objects I do not simply appear stay still or move around in relatively smooth ways. Objects do not simply appear and disappear in a random or disorganized way. Furthermore, places look similar over time, or they change in regular ways. My office looks today approximately the way it looked yesterday. When I go home, my house will look similar to the way it looked when I left. The plants in my garden change gradually in regular ways. Things appear in just about the way relatively stable and persisting objects would appear to a perceiver with a relatively stable perceptual system. We can sum this up as follows:

O. I have memory and perceptual experience that are regular and orderly.

The commonsense explanation of (O) is

CS. There is a world of enduring and relatively stable physical objects. My experiences are typically caused by these objects stimulating my sense organs.