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.[13] 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 experience either 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.

Of course, (CS) can be fleshed out in many ways. In fact, one can regard many of the results of scientific investigation as spelling out the details of this barebones "theory." There are alternative explanations one might offer for (O). They include:

BIV. I am a brain in a vat connected to a powerful computer. The computer stimulates my brain, giving me sensory experiences. The computer has been programmed to make my experiences regular and orderly.

DR. All my experiences are dream experiences. My dreams are (usually) relatively systematic and orderly.

EG. My experiences are caused by an evil genius. This genius causes me to have regular and orderly experiences in order to trick me into believing (CS).

These alternative explanations are also incomplete. It seems possible that they too can be filled out in ways that provide more detailed explanations of (O).

The question raised by ordinary-standards skepticism is, "Why believe (CS) when there exist these alternative explanations of our fundamental data?" Behind the question raised by ordinary-standards skepticism is a final argument for skepticism. The main idea in the argument is that the evidence we have does not provide good reason to believe the commonsense propositions we all do believe rather than the skeptical alternatives mentioned. The argument can be formulated as follows:

Argument 7.7: The Alternative Hypotheses Argument

7-1. The evidence people have, (O), does not provide better reason to believe ordinary external-world propositions and (CS), than to believe the rival skeptical hypotheses, such as (DR), (BIV), and (EG).

7-2. If one's evidence does not provide better reason to believe one hypothesis than to believe some rival hypothesis, then one is not justified in believing that hypothesis.

7-3. People are not justified in believing (and thus do not know) ordinary external-world propositions and (CS). (7-1), (7-2)

This is a valid argument and, once again, the conclusion asserts a significant skeptical thesis. So any defender of *The Standard View,* and modest foundationalism, must find a good response to the argument. To find reason to deny (7-1), one must uncover reasons to think that our evidence really does support our commonsense beliefs over the rivals. To find reason to deny (7-2), one must find reason to think that our beliefs can be justified even if they are not better supported than their rivals.[15]

Modest foundationalism, as discussed in Chapter 5, holds that our ordinary perceptual beliefs are "proper" responses to perceptual stimuli. It is plausible to view *The Alternative Hypotheses Argument* as a challenge to that view. The responses to be considered in the next section, then, are all ways of spelling out the modest foundationalist response to the argument.

B. Three Responses

B1. Epistemological Conservatism Epistemological conservatism is the view that one is justified in retaining an existing belief provided one's evidence does not provide better support for some rival belief.[16] This view is an epistemological analogue of what seems to be a fairly reasonable practical principle. Suppose one is considering replacing some material possession, such as a house, a car, or a computer. In general, it would be foolish to purchase the replacement if what one ended up with was exactly as good as what one had before the change. Making a change makes no sense unless it is an improvement, in one way or another, on what one already has. One might sensibly replace an older car by a newer car of the same model, thereby gaining increased reliability and added features. One might replace an existing house by a smaller and less expensive one if one's family or financial circumstances indicate that one will be better off by so doing. So the principle does not say that it always best to buy bigger and more expensive things. It just says that it is sensible to make a change only when something is gained by so doing. Perhaps a reason for this is that there is always some cost—financial or otherwise—in making the change. It is foolish to incur these costs to end up in exactly as good a situation as one is already in.

An implication of this practical principle is worth noting explicitly. Suppose that there are two very similar cars, A and B. You might be in the following situation. Given the similarity of the cars, if you already own A, then it is reasonable to stick with A and unreasonable to switch to B. But if you already owned B, it would be reasonable to stick with B and unreasonable to switch to A. Even though the one you have is not better than the other, sticking with what you already have is more reasonable than switching. It is this preference for sticking with what one already has that makes this principle conservative.

Epistemological conservatism says that a principle somewhat analogous to the one just described holds for believing as well. When your evidence supports a couple of theories equally well, then if you already believe one of those theories, it is reasonable to stick with that belief rather than switch. In effect, the fact that you already believe one of them is a tiebreaker. Applied to *The Alternative Hypotheses Argument,* epistemological conservatism says that (7-2) is false: It is reasonable to retain our commonsense beliefs even though they are not better supported by our evidence than their rivals are.

Epistemological conservatism is subject to an important objection. To begin, there is a significant way in which the analogy between believing and practical actions breaks down. Suppose that you already own a car and are considering

keeping it or replacing it. As we are understanding this example, all your options include owning some car or other. Perhaps you could choose not to have one at all, but that, we may assume, is a terrible option for you. In the belief examples, however, you do have the option of suspending judgment about the propositions in question. To see the significance of this, consider the following example:

> *Example 7.2: Two Suspects*
>
> Detective Jones has definitively narrowed down the suspects in a crime to two individuals, Lefty and Righty. There are good reasons to think that Lefty did it, but there are equally good reasons to think that Righty did it. There is conclusive reason to think that no one other than Lefty or Righty did it.

What should Jones think? Under these circumstances, it clearly would be unreasonable for him to think that Lefty did it and Righty did not. It would be equally unreasonable for him to think that Righty did it and Lefty did not. Clearly, he should suspend judgment about whether Lefty did it and, equally, he should suspend judgment about whether Righty did it. Furthermore, the mere fact that he already believes that one of them, say Lefty, did it, is of no epistemological significance whatsoever. Suppose Jones came upon the evidence about Lefty first and so reasonably came to believe that Lefty did it. Once he learns that there is equally good evidence for the proposition that Righty did it, he should stop believing that Lefty did it. Returning to the practical matter of the car, when a person owning one car learns that there is another equally good car, even one that he could trade for at no cost, it is not true that he should get rid of the already owned car. Part of the reason for this is that there is no analogue of suspending judgment.

It may be that there are some costs involved in changing beliefs. At the very least, it can be cognitively disruptive. There may be other beliefs that have to be changed when the one under consideration is changed. This might lead you to think that there are factors to be weighed against the evidential considerations and that these other considerations might alter the outcome in some cases. But in thinking along these lines one is allowing practical considerations to enter into epistemological evaluations. As we saw in earlier chapters, there can be practical considerations relevant to beliefs, but such considerations do not affect the epistemic evaluation of beliefs.

It is possible that defenders of epistemological conservatism can find a way to modify their theory to avoid the objection raised. Perhaps there is some relevant difference between the beliefs of the detective in the example just given and our ordinary beliefs that a suitably modified conservative principle could appeal to. But until such a difference is established, it is best to look beyond epistemic conservatism in developing a response to *The Alternative Hypotheses Argument.*

B2: Immediate Perceptual Justification One might think that our experiential evidence really is better evidence for the proposition that there are things as we seem to experience them than for the proposition that we are dreaming that there are such things, or that we are being deceived by an evil demon or a computer into thinking that there are such things, or any other specific skeptical alternative. Advocates of this view will reject premise (7-1) of *The Alternative Hypotheses Argument.*

One statement of this view can be found in the work of Roderick Chisholm. In his *Theory of Knowledge,* Chisholm proposes as fundamental epistemological principles such principles as the following:

> If S believes that he perceives something to have a certain property F, then the proposition that he does perceive something to be F, as well as the proposition that there is something that is F, is one that is *reasonable* for S.[17]

Notice that the antecedent of this principle requires that S believe that he perceives something to be F, not merely that it be true that he seems to see something F. So this principle may not fit our situation exactly, because the modest foundationalist view we are considering takes the fundamental evidence to be perceptual experiences, not beliefs about them. Still, Chisholm's view is close enough to be relevant here.

What bears emphasis is that Chisholm does not derive this principle from some other, more fundamental truths. One can imagine a skeptic producing a rival principle, say, that when one has an experience and believes that one is dreaming or that one is being tricked by an evil demon, one is reasonable in those beliefs. Chisholm's view would be that his own principle is correct and these others are wrong. He might defend his view in part by noting that his principle can explain how we have knowledge of the external world. Chisholm's approach bears some similarity to Moore's view discussed earlier.

Another philosopher who defends a view along these lines is James Pryor. He writes:

> My view is that whenever you have an experience as of *p,* you thereby have immediate *prima facie* justification for believing *p.* Your experiences do not, in the same way, give you immediate *prima facie* justification for believing that you are dreaming, or being deceived by an evil demon, or that any of the skeptic's other hypotheses obtain.[18]

Pryor's idea is that our perceptual experiences make it seem as if there are certain things external to us, things such as trees, houses, other people, and so on. He thinks that we are justified in taking things to be as they seem, provided we do not have defeating evidence. That is why he says the justification is merely *prima facie.* This means that the evidential support can be defeated by other evidence. In the typical case, it is not.

Pryor acknowledges that one might like to have a more informative explanation of why our experiences justify our commonsense beliefs. He rejects a number of potential explanations of this. For example, his view is *not* that our commonsense beliefs are justified by our experiences because there is a reliable

connection between them.[19] One might try to get some mileage out of the fact that our perceptual beliefs seem to be *irresistible* in the light of our experiences, but Pryor denies that this carries any epistemic significance.[20] Nor does he give any weight to the "best explanation" account, which will be covered in the next section. The best he gives in defense of his view is that our experiences have a kind of "phenomenal force": When we see a table, it "feels as if" there really is a table there. And this, he thinks, is significant.

The immediate perceptual justification thesis can thus be summed up in the following principle:

> IPJ. Whenever a person has an experience as of p's being the case, the person has immediate *prima facie* justification for believing p.

Since this justification is merely *prima facie,* it is possible for it to be defeated. It is also possible, of course, for a person to have a second source of justification for a proposition that has justification of the sort described in (IPJ). Thus, the proposition that there is a table in front of me might be justified in this way, and it might have additional justification from the fact that I've heard someone else say that there is a table there.

One difficulty for (IPJ) is that it seems somewhat *ad hoc.* This can be brought out by imagining a defender of the evil genius theory. Such a person can equally well assert that our experiences "just are" evidence for propositions about the evil genius rather than for propositions about ordinary objects. Furthermore, it is possible that the idea that our experiences make it "feel as if" there are external objects is a mistake. Perhaps our beliefs result in part from training or indoctrination. Perhaps they result from a built-in bias. It would be good to have some more general defense of our preferred position. Thus, one might agree with Pryor's claim that our experiences do give us reason to believe commonsense propositions and not skeptical alternatives, but one might also think that there must be some more general theoretical account of reasonable belief that explains this fact.

It is possible that defenders of (IPJ) will defend their view by appeal to considerations similar to those used in defense of inductive reasoning. The idea there was that using observed cases as the basis for drawing conclusions about unobserved cases was simply part of the idea of good reasoning. Analogously, defenders of (IPJ) might claim, believing commonsense external-world propositions on the basis of perceptual experiences just is good reasoning. Yet, to many, the cases seem importantly different. There must be, the critics contend, some explanation of why the experiences we have justify the beliefs they do. Unlike (PFR), principles such as (IPJ) and Chisholm's principle about perception are not fundamental principles.

There is a second, and possibly related, difficulty for (IPJ). The principle makes use of the idea of a person having an experience as of p's being the case. This idea may seem reasonably clear. When you look at a table, you have an experience as of its being the case that there is table there. When you see a book, you have an experience as of its being the case that there is a book before you. However, there is a difficult question about this idea, as is brought out by the following example:

> *Example 7.3: Three People in a Garden*
>
> Three people, Expert, Novice, and Ignorant, are standing in a garden looking at a hornbeam tree. They have a clear and unobstructed view of the tree. The visual appearance present to each of the three people in the garden is exactly the same. (Minor differences due to their slightly different positions are irrelevant to the example.) Expert knows a lot about trees and can easily identify most trees, including this one, immediately. Novice knows a little about trees but is unfamiliar with hornbeams. Ignorant does not know anything about trees. He does not know which of the things in the garden is a tree and which is a flower.

A question a defender of (IPJ) must face is this: What do these people experience? That is, do they all have an experience as of its being the case that there is a hornbeam tree before them, or that there is a tree before them, or that there is a green and brown object before them? Or do they have experiences that differ in content?

If defenders of (IPJ) say that they all have an experience as of there being a hornbeam there, then their theory seems to yield the incorrect result that even Novice and Ignorant are justified in believing that there is a hornbeam there. If they say that they all have an experience as of there being a tree there, then the theory seems to yield the incorrect result that even Ignorant is justified in believing that the thing he is looking at is a tree. Perhaps, then, defenders of (IPJ) should say that the experience is just as of its being the case that there is something partially green and partially brown there. While not obviously unacceptable, this makes the content of what's immediately justified more limited than defenders of the theory seemed to have in mind.

It is possible for defenders of (IPJ) to argue that background information and prior experiences can affect what one's experiences are like. Thus, even though Expert, Novice, and Ignorant all have visually identical experiences, Expert has an experience as of its being the case that there is hornbeam tree before him, while neither Novice nor Ignorant have an experience as of that proposition being true. Novice (and, presumably, Expert) have an experience as of its being the case that there is a tree there, but Ignorant does not.

Thus, defenders of (IPJ) can spell out the application of their view to this example in a way that may yield the desired results. But it must be admitted that the idea of what it is for one to have an experience as of its being the case that p is far from clear. Some better account of how all this works is needed. One has the sense that defenders of this view are simply saying whatever they must to achieve the desired results. A more general theoretical basis for their view is desirable.

These considerations are by no means decisive. They do not refute the idea some perceptual beliefs enjoy immediate *prima facie* justification. Perhaps the idea could be extended to memory as well: Memory beliefs also have immediate *prima facie* justification. If so, we have a response to premise (7-1) of *The Alternative Hypotheses Argument* and some hope that *The Standard View*, and modest foundationalism, can be vindicated. However, the two considerations just raised do show that there is some unclarity in the idea of what it is for a person to have an experience as of its being the case that p and that it is at least reasonable to wonder why our experiences provide evidence for our commonsense beliefs rather than their skeptical alternatives. We will turn next to a view that attempts to provide the desired explanation.

B3: Inference to the Best Explanation A third response to *The Alternative Hypotheses Argument* shares with the second the claim that premise (7-1) of the argument is false. But the third response holds that it is not a simple or fundamental fact about perceptual experience that it supports our commonsense beliefs. Rather, according to this view, perceptual beliefs are supported by our experiential evidence in much the way theories in science can be supported by the relevant experimental evidence. Very roughly, the idea is that there can be a number of alternative theoretical explanations of a particular event or pattern of events. That is, each theory provides an explanation of why things happen as they do. But, on this view, there can be theoretical grounds for thinking one explanation is a better explanation than another, and that this better explanation is therefore more reasonable to believe than the other.[21]

The general idea of what counts as a best explanation is reasonably familiar. Nevertheless, it proves extremely difficult to spell out the idea in any precise way. An example will illustrate the idea.

Example 7.4: The Variable Colleague

You work in an office and have a colleague who works with you every day. You notice that your colleague's behavior is quite variable, although not bizarre or extraordinary. Some days he is in a good mood, other days he is not. You can come up with two potential explanations of your observations. Explanation 1 is that your colleague's mood, and consequently his behavior, varies with how well he sleeps. This explanation can be filled out with accounts of why he sleeps better some nights than others and explanations of how sleep affects behavior. Explanation 2 is that your "colleague" is really two different people, identical twins with markedly different personalities. They never go out in public together and never let anyone know that they are twins. They tell each other everything that happens every day, so there are no revealing episodes in which they appear to be ignorant of the things they should remember from previous days.

Each explanation is consistent with your observations. Each does, in some sense, provide an explanation of why you observe the variations in behavior.

There may be something exotic or intriguing about Explanation 2. Nevertheless, it seems highly unreasonable to accept it. Explanation 1 is a much better explanation. One thing that makes Explanation 1 better is its simplicity in comparison to Explanation 2. Explanation 2 is pointlessly complex. It introduces two people, with odd motivations and habits, and a complex scheme for tricking people, when no such complexity is needed to explain the data. It is far more reasonable to believe the mood variation story. Perhaps a second virtue of Explanation 1 is that it better fits with our background information about people. They just do not undertake the complex schemes that Explanation 2 introduces.

A second example may help to clarify the idea. Suppose we see footprints in the sand along the beach. The footprints are in the shape of boots commonly worn by people, though not typically worn by people at the beach. We might wonder why these prints are on the beach. There is an obvious explanation and a host of alternative explanations. The obvious one and one alternative are

T1. People wearing boots recently walked along the beach.
T2. Cows wearing boots and walking on their hind legs recently walked along the beach.

Both (T1) and (T2), when suitably filled out, do explain the prints you observed. But (T1) has the virtue of simplicity. It does not introduce the pointless complexity that (T2) introduces. (T1) is the better explanation.

It must be admitted that it is very hard to spell out in detail just what a best explanation is. Simplicity and conforming to background information are two characteristics mentioned previously. However, as Peter Lipton points out in a book-length discussion of inference to the best explanation, there are cases in which explanations that seem to be simple are nevertheless quite unreasonable.[22] For example, some conspiracy theorists propose unified explanations of numerous salient assassinations and other significant political events. They propose some international organization behind all of them. There is a kind of simplicity to this, in contrast to the diverse and independent explanations that most experts regard as much more plausible. Perhaps the conspiracy theory is complex in that it attributes a complex set of behaviors and motivations to one organization, even though it manages to unify the explanations of many events.

Lipton briefly states the application of inference to the best explanation to skepticism:

> [As] part of an answer to the Cartesian skeptic who asks how we can know that the world is not just a dream or that we are not just brains in vats, the realist may argue that we are entitled to believe in the external world since hypotheses that presuppose it provide the best explanation of our experiences. It is possible that it is all a dream, or that we are really brains in vats, but these are less good explanations of the course of our experiences than the ones we all believe, so we are rationally entitled to our belief in the external world.[23]

The response to *The Alternative Hypotheses Argument,* then, is that (7-1) is false. We can set this out as a formal argument:

Argument 7.8: The Best Explanation Argument

8-1. Our experiential evidence is better explained by (CS) than by (DR), (BIV), (EG), or any other available alternative.

8-2. If one explanation better explains one's evidence than any other available alternative, then one's evidence better supports that explanation than any of those alternatives.

8-3. Our evidence better supports (CS) than it supports (DR), (BIV), (EG), or any other alternative (so (7-1) of *The Alternative Hypotheses Argument* is false). (8-1), (8-2).[24]

In support of (8-1) we can point to the fact that the alternative explanations do bring in a kind of complexity—amazingly sophisticated computers, evil geniuses monitoring our thoughts, implausibly orderly dreams. These explanations seem *ad hoc,* complex, and ridiculous. *The Best Explanation Argument* thus seems promising.

Still, there are difficult problems. There are, as noted, questions about exactly what in general counts as a best explanation. There is, further, some question about exactly why the evil genius hypothesis is such a bad explanation anyway. In a certain sense, it is elegant and simple. Rather than a complex world of enduring objects, it poses one all-purpose cause of everything. There is something simple about it. So (8-1) is not obviously true.

Finally, there is a hard question about (8-2). As formulated, this premise says that a person is justified in believing a proposition when it is a better explanation of the relevant data than its rivals. The hard question is whether it is enough for the explanation in fact to be the best explanation, or does the person have to realize that it is the best explanation? Example 7.3 can be used to illustrate the basis for thinking that (8-2) should be revised to include the requirement that the person realize that the explanation is best. In that example, Expert, Novice, and Ignorant were looking at what was in fact a hornbeam tree. Novice was unable to identify the tree by the way it looked. But suppose Novice wondered why the tree looked the way it did and why it had leaves of the particular shape it had. In one use of the phrase "best explanation," it seems to be true that the best explanation of Novice's experience was that there was a hornbeam tree in front of him. After all, nothing else would look just like that. But then, by (8-2), Novice is justified in believing that he sees a hornbeam. But this is the wrong result.

We might get around this problem by revising (8-2), requiring of a justified explanation not merely that it be the best explanation but that the believer have reason to believe that it is the best explanation. Novice, in the example just considered, does not satisfy this condition. So the new principle would have the right results.

This suggests that *The Best Explanation Argument* needs revision, perhaps along the following lines:

Argument 7.9: The Best Explanation Argument (Revised)

9-1. We are justified in believing that our experiential evidence is better explained by (CS) than by (DR), (BIV), (EG), or any other alternative.

9-2. If we are justified in believing that one explanation better explains one's evidence than any other explanation, then one's evidence better supports that explanation than any of those alternatives.

9-3. Our evidence better supports (CS) than it supports (DR), (BIV), (EG), or any other alternative (so (7-1) of *The Alternative Hypotheses Argument* is false). (9-1), (9-2)

Perhaps (9-2) avoids the problem proposed for (8-2). And perhaps (9-1) is acceptable as well, at least for those of us who have thought about these matters. But there is a question to worry about here. If the best explanation view is adequate to defend *The Standard View,* and modest foundationalism, then it must be that considerations about best explanations justify the beliefs of people who have not given a moment's thought to skeptical issues and the comparative merits of these explanations. Thus, if this approach is to account for everyone's knowledge, then, apparently, even those who have never thought about any of this must be justified in believing that (CS) explains their observations better than its rivals. Perhaps this is true, but it is likely that critics would have reservations about this view. Thus, they might be more inclined to accept (8-1) than (9-1).

There is, then, a dilemma for defenders of the best explanation response to skepticism. Premise (8-1) seems more plausible than (9-1), at least if the "we" in (9-1) includes ordinary people who have not thought about skepticism. But premise (9-2) is considerably more plausible than (8-2), given the objection based on Example 7.3. The best option for defenders of the best explanation view is to claim that (9-1) is true and to claim that typical people, who have not thought about these matters, are nevertheless justified in believing that (CS) is the best explanation of their experiences.[25]

C. Conclusion

The three responses to *The Alternative Hypotheses Argument* covered in this section do not exhaust the possibilities, but they do give a good indication of the range of responses. Epistemic conservatism seems to be subject to decisive objections. The idea that many commonsense beliefs enjoy immediate *prima facie* justification has some plausibility, but there is a crucial unclarity in the view and it leaves unanswered what seems like a perfectly sensible question: Why do our experiences justify our commonsense beliefs rather than their rivals? Best explanationism attempts

to answer this question. Although there are difficult questions about the details of this view, it survives as a plausible response to ordinary-standards skepticism.

One might wonder whether we are justified in believing that best explanations are always true. One can imagine a critic raising questions analogous to those Hume raised about induction. And the response may well be the same: It is not the case that, necessarily, best explanations are always true. But it is the case that, necessarily, one is justified in believing what one knows to be a best explanation. This may even lend some support to the response to Hume's problem advanced earlier in this chapter. One reason to believe that observed regularities will continue to obtain is that this is part of the best explanation of one's experiences.

We have, then, a plausible, though not decisive, response to ordinary-standards skepticism. One might have hoped for a more clear-cut refutation of skepticism. The difficulty of providing that refutation is distressing and reveals the power of the intellectual challenge posed by *The Skeptical View.*

APPENDIX: CONTEXTUALISM

A widely discussed view in recent epistemology is *contextualism.* Some philosophers think that it provides the basis for a good response to skepticism while at the same time giving *The Skeptical View* its due respect. In this section we will briefly examine contextualism.

Contextualism is fundamentally a view about the way the word "knows" functions. The central idea is that the standards for applying the word "knows" vary from one context to another. Sometimes the standards for its application are very high, and in those settings what we express by saying "S knows that p" is usually false. But in other settings the standards are more easily met, and what we express by uttering those same sentences may be true.

The contextualists' idea can best be appreciated by considering a relatively uncontroversial analogy. Suppose you go to the zoo with a young child. The first few animals you see are monkeys and birds. You then go into the elephant area and see a baby elephant. You say to the child, "Look at that elephant. It is big." You then go through the rest of the zoo and see the adult elephants, among other large animals. On your way out, you see the baby elephant again, now standing near the adults. Pointing at the same baby elephant, you say, "Look at the baby elephant. It is so small. It's not big."

An obnoxious onlooker might charge you with contradicting yourself. He might ask you to make up your mind and decide whether the baby elephant is big or not big. "It can't be both," he might say. But you have a sensible reply to this charge. The response draws on a fact about the way the word "big" functions. Whenever we say that something is big, we are comparing the thing in question to the members of some comparison class. We often do not explicitly state what that class is. Instead, the conversational background or the context help determine the comparison. In your first remarks about the elephant, you were comparing the baby elephant to the general group, animals at the zoo. Relative to that group, the baby elephant is big. Later, you were comparing the baby elephant to a more restricted group, elephants. Relative to that group, the baby elephant is not big. You did not contradict yourself. The context in which you used the word "big" changed.

Contextualists in epistemology think that the word "knows" is in some ways like the word "big". In different settings, the word has different standards for application. In ordinary settings, when we talk about the world and make claims about what we know, ordinary standards are in effect. We often can satisfy them. Sometimes, however, we raise the standards for the application of the word "knows". Contextualists typically say that this happens when we discuss arguments about skepticism. Under those circumstances, they say, the skeptics are right. We do not meet the high standards in place in those contexts.

Contextualists claim that their theory has a strength that the theories we have discussed up until now have all lacked. The strength is that it can explain our varying reactions to claims about knowledge. In ordinary settings, we accept a variety of knowledge claims without reservation. They seem plainly correct. And then skeptical arguments come along and many people deny that they know things. Yet, later on, these people will once again confidently claim to have knowledge. Contextualists claim that people are right each time: In the ordinary settings, their claims to knowledge are correct; and in the skeptical contexts, their denials of knowledge are also correct. There is no contradiction involved here, just as there is no contradiction in the things said at the zoo.

There are a variety of ways in which contextualists can work out the details of their view. Evidentialist contextualism accepts an evidentialist view about justification and can even accept much of what modest foundationalism says about knowledge.[26] According to this approach, when we attribute knowledge to someone, how well justified the person has to be for our attribution to be true varies from one context to another. Ordinarily, ordinary standards obtain. We can satisfy those standards. But sometimes the standards are higher. And sometimes, as when we are discussing skepticism, the standards are so high that we do not meet them. In those contexts, attributions of knowledge are typically not true.

Nonevidentialist contextualism typically makes use of an account of knowledge that is closer to some of the nonevidentialist theories discussed in Chapter 5. One such view is the *relevant alternatives* theory.[27] According to this view, a person knows a proposition to be true just in the case the person can "rule out" or "eliminate" all relevant alternatives to that proposition. But what counts as a relevant alternative depends upon the context of the person attributing (or denying) knowledge (and not the context of the person being talked about). In ordinary contexts, only ordinary alternatives are relevant. But in some contexts, such as when skepticism is at issue, a broader range of alternatives count as relevant. With respect to many propositions, we can rule out ordinary alternatives, but not some of the more exotic ones. Hence, in ordinary contexts it is true that we have knowledge, but in other contexts it is not.

An example will make the idea clearer. Suppose Jones sees Smith down the hall. In commenting on this, you say that Jones knows that he sees Smith. This is because Jones can rule out the alternatives to its being Smith that he sees. The alternatives are that he sees, say, Black or White. But he can tell, by the per-

son's size and shape, that it is Smith and not Black or White. So, in this setting, when you say that Jones knows that he sees Smith, what you say is true.

But now suppose the conversation turns to skepticism. Now there are new relevant alternatives. Maybe Jones is having a hallucination. Maybe Smith's previously hidden identical twin is down the hall. Jones cannot rule out these alternatives, because things would appear just the same if they were the case. So now, according to the theory, it would not be correct to say, "Jones knows that he sees Smith."

A question defenders of the relevant alternatives theory must face concerns their use of the phrase "rule out" or "eliminate." They claim that we cannot rule out alternatives such as those just mentioned. This is because these alternatives are consistent with our experience.[28] In other words, Smith cannot rule out the alternative that he sees Jones's twin on the grounds that he has never heard that Jones has a twin and, because he knows Jones well, he would have heard about any twin he had. And he cannot rule out the hallucination alternative on the grounds that it is an inferior explanation of his observations. The relevant alternatives theory is, therefore, committed to extremely high standards for knowledge. It holds that you can rule out an alternative when, and only when, that alternative is inconsistent with your observations.

However exactly the details of contextualism are spelled out, the theory has certain attractions. It implies that a great many of our ordinary attributions of knowledge are correct. If, upon arriving at class one morning you say, "I know that I brought my book with me," contextualists can agree that what you say may well be true. In contrast, skeptics would say that you were wrong. But contextualism also can explain the appeal of *The Skeptical View.* It holds that in the contexts in which we discuss skepticism, skeptical arguments are good arguments. In those contexts, the conclusions skeptics draw are correct. This is because discussion of skepticism causes us to be in conditions in which the standards for knowledge are very high, so high that we cannot meet them.

Contextualism has some liabilities as well.[29] For one thing, it concedes a great deal to skeptics, perhaps more than is correct. Many people, especially fallibilists of all sorts, think that skeptics are wrong when they say that people do not have knowledge. They think that there are defects with the arguments for skepticism. But contextualism, at least in the forms discussed here, implies that in the settings in which skepticism is discussed, the claims of the skeptics are correct. As we have seen, there are some reasonably good replies to skepticism, so it is difficult to see why so much should be conceded to skeptics.

Furthermore, it is far from clear that the word "knows" does shift its standards in the way contextualists claim. In the case of the words that most obviously vary with context, such as "big," it is very easy to see that apparent contradictions need not be genuine contradictions. Thus, when you are told in the preceding example that you first said that the baby elephant is "big" and later said that it is "not big," you probably would not feel as if you had contradicted yourself. If you are moderately sophisticated about these matters, you might simply explain that you meant "big for an animal" and "not big for an elephant." You would not

find yourself moved by the skeptical arguments, you probably think that your ordinary claims are called into question. This suggests that "knows" differs crucially from words like "big" and casts some doubts on the contextualist analysis.

Finally, it is worth noting the similarity of what evidentialist versions of contextualism say about skepticism to what modest foundationalists say about skepticism. Recall that evidentialist versions of contextualism say that we have good reasons to believe many of the things we ordinarily believe, reasons good enough to yield knowledge. Relative to ordinary standards, we also have good reason to deny that we are brains in vats. In light of *The Alternative Explanations Argument,* these contextualists must have some account of why these reasons are good enough to give us knowledge (by ordinary standards). They will, presumably, have to appeal to one of the views along the lines of those we discussed in response to that argument. Contextualism by itself—the mere view that the standards for attributions of knowledge vary—does nothing to explain why it is true that we meet ordinary standards. This is not an objection to contextualism. Rather, the point is important because it reveals that the evidentialist version of contextualism as a partial response to skepticism depends upon the adequacy of one of the other previously discussed responses.

The adequacy of the nonevidentialist version of contextualism previously mentioned, the relevant alternatives theory, as a response to skepticism depends upon the merits of the idea that to have knowledge is to be able to rule out, in the special sense of "rule out" that the theory uses, alternatives. This is, in its own right, a controversial theory. One difficulty it seems to face is that it is hard to see how it can account for knowledge based on inductive reasoning. This is because the falsity of an inductive conclusion is never ruled out by one's evidence, and it is hard to see why the falsity of the conclusion is not a relevant alternative in any case of inductive reasoning.[30] Contextualism could also be developed in ways that draw on the other nonevidentialist theories discussed in Chapter 5. It would, however, inherit the difficulties of those theories.

ENDNOTES

1. This ignores special cases in which you have reason to think that the As will be Bs for the near future but will cease being Bs later on. In such a case, you might have good reason to think that the next A will be a B but that not all As will be Bs.
2. David Hume, *Enquiry Concerning Human Understanding,* 2nd ed., edited by L. A. Selby-Bigge (Oxford: Oxford University Press, 1962), Section IV, Part II, p. 35
3. *Enquiry Concerning Human Understanding,* Section IV, Part II, p. 35.
4. *Enquiry Concerning Human Understanding,* Section IV, Part II, pp. 35–6.
5. It is possible that a defender of the inductive justification of induction would object instead to premise (5-7). The claim would be that the inductive justification does appeal to (PF) (in the way described in the text) but that the argument can still justify the principle. The considerations presented in the present paragraph would seem to undermine this response as well.
6. See Hans Reichenbach, *Experience and Prediction* (Chicago: University of Chicago Press, 1938).
7. Brian Skyrms, "The Pragmatic Justification of Induction," *Choice and Chance,* 2nd ed.

8. Bertrand Russell, *The Problems of Philosophy* (Oxford: Oxford University Press, 1959), p. 65. Russell goes on to formulate explicitly a somewhat different, and more detailed, principle.
9. The idea presented here is based on the proposal made by Peter F. Strawson in *Introduction to Logical Theory* (New York: John Wiley & Sons, 1952).
10. Assuming, of course, that you do not have defeating evidence about its color.
11. The topic of *a priori* knowledge is a complex and controversial one. The claim here is simply that (PFR) is the sort of thing that can be known simply through understanding, in much the way it can be known that all bachelors are male or all mothers are parents. For more on *a priori* knowledge, see Chapter 8.
12. In light of this, one might want to reexamine what was said earlier about the inductive defense of induction.
13. This problem is often called "The New Riddle of Induction." The classic formulation of the puzzle is in Nelson Goodman, *Fact, Fiction, and Forecast* (Cambridge, MA: Harvard University Press, 1955).
14. In Chapter 8 we will consider extending the *a priori* defense of induction to questions about perception and memory.
15. The discussion here focuses on the best explanation of our overall experiences. More narrowly focused versions of the same issues could be developed. These would examine the best explanations of the specific experiences one is having at a particular time. Essentially, the same points would apply.
16. For a critical discussion of this view, see Richard Foley, "Epistemic Conservatism," *Philosophical Studies* 43 (1983): 165–82.
17. *Theory of Knowledge* (Englewood Cliffs, NJ: Prentice Hall, 1966), p. 45.
18. "The Skeptic and the Dogmatist," *Nous* 34 (2000): 517–49. The quotation is from p. 536.
19. "The Skeptic and the Dogmatist," endnote 6.
20. "The Skeptic and the Dogmatist," endnote 37.
21. An interesting defense of this response to skepticism can be found in Jonathan Vogel, "Cartesian Skepticism and Inference to the Best Explanation," *Journal of Philosophy* 87 (1990): 658–66. For critical discussion, see Richard Fumerton, *Metaepistemology and Skepticism* (Lanaham, MA: Rowman and Littlefield, 1995), pp. 207–14.
22. *Inference to the Best Explanation* (London: Routledge, 1993). See especially Chapter 4.
23. *Inference to the Best Explanation*, p. 72.
24. (8-3) denies premise (7-1) of *The Alternative Explanations Argument*. But (8-3) does not imply that we are justified in believing (CS). (8-3) leaves open the possibility that (CS) is merely the best of a bad set of explanations, and thus not justified. To use the kind of view under consideration here to support the claim that our commonsense beliefs are justified, rather than just to refute (7-1) of the argument, it would be necessary to defend a stronger claim: (CS) is not only better than the rivals but is a very good explanation. Defenders of modest foundationalism who appeal to best explanationism are likely to defend this stronger claim as well.
25. This is not to say that they justifiably believe (i.e., have a well-founded belief) that (CS) is the better explanation; it is just to claim that their evidence does support that proposition.
26. For a defense of this kind of contextualism, see Stewart Cohen, "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13 (1999): 57–89.
27. A version of this theory is defended by David Lewis in "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (1996): 549–67.
28. This is the view David Lewis defends in "Elusive Knowledge."
29. The material in this section is drawn from Richard Feldman, "Skeptical Problems, Contextualist Solutions," *Philosophical Studies* 103 (2001): 61–85.
30. For discussion of this point, and additional objections to the relevant alternatives theory, see Jonathan Vogel, "The New Relevant Alternatives Theory," *Philosophical Perspectives* 13 (1999): 155–80.

CHAPTER EIGHT

Epistemology and Science

This chapter examines *The Naturalistic View*. *The Naturalistic View* is not a single thesis about the conditions for knowledge and justification. Instead, it encompasses a general view about the proper role of science in epistemology. It holds that science should play a much more significant role in epistemology than advocates of *The Standard View* have traditionally given it. We will examine two issues arising out of *The Naturalistic View*.[1] The first issue concerns the implications of some research results that seem to show that people systematically reason badly and, perhaps, know less than *The Standard View* suggests. The second issue arises because epistemologists usually defend and discuss *The Standard View* without paying attention to scientific results. Naturalists think that this is a methodological mistake.

I. EVIDENCE OF HUMAN IRRATIONALITY

On the basis of a large body of research into the ways people form beliefs, some philosophers conclude that people are systematically irrational.[2] The charge is that people have a deeply rooted tendency to make a variety of logical blunders, errors concerning probability, mistakes involving causation, and so on. Although results such as these are not likely to support skeptical conclusions as broad and general as those discussed in Chapters 6 and 7, they do call into question the extent of our knowledge. If we make as many errors as the critics say we do, then perhaps we do not have knowledge even when we manage to get things right. If we make as many mistakes as some critics charge, then it is unlikely that we know as much as *The Standard View* says we do.