

22. Knowledge-first epistemology isn't limited to the defense of $E = K$. Williamson (2000) argues knowledge is the norm of assertion. We discuss norms of assertion in Chapter 5, Section 5.3.
23. For further discussion, see Schroeder (2008) and Comesaña and McGrath (2014).
24. The safety theorist might reply that these beliefs are safe because by the time they are formed the belief couldn't easily be false. At the last minute, in Comesaña's example, the subject decides not to disguise himself as Michael, and we might think, given that fact, the subject wouldn't easily go wrong in taking Judy's word about the route to the party. However, this sort of reply is a double-edged sword. Consider Henry in the barn case. Given the fact that he is looking at the (lone) barn, could he easily go wrong? If the safety theorist wants to appeal to safety to explain why Henry doesn't know, it seems he can't reply to Comesaña by this sort of "by the time the belief is formed" gambit.
25. For a recent argument that one can know something while believing it unreasonably, see Lasonen-Aarnio (2010).

Skepticism About Knowledge

Matthew McGrath



The skeptic claims we lack knowledge, or justified belief, or both, across some important domain. In the extreme case, the domain is universal—that is, we lack any knowledge at all, or any justified belief at all. Such skepticism is sometimes called *global skepticism*. For instance, the regress argument from Chapter 1 is an argument that no one is ever justified in believing anything. But there are local skepticisms as well, such as skepticism about knowledge of other minds, about the future, about the external world (the world outside your own mind).¹ These skepticisms are certainly radical, even if they do not question all our knowledge. Out in the street or in the coffeehouse, all of us claim we know things about other minds, about the future, and certainly about the external world. Think of what we say: "I know the waiter recognized me from last week" (implying knowledge of another mind) or "I know Beth will get to the restaurant before us" (implying knowledge of the future), and of course both of these imply knowledge of an external world, the world outside your own mind. Could all of these statements, and the beliefs they express, be *wrong*? It seems ludicrous to think so. What makes the issue of skepticism interesting, though, is not the plausibility of skeptical claims about lack of knowledge but the compelling arguments in favor of those claims.

In this chapter, we will focus on what is probably the most widely discussed sort of skeptical argument since Descartes: arguments that we lack knowledge of the external world, based on our inability to rule out "skeptical possibilities." We'll call these *skeptical possibility arguments*. Such arguments can be given for skepticism about justified belief, but they are particularly powerful for skepticism about knowledge. There is something about the concept of knowledge that makes this sort of argument compelling. After illustrating how these arguments go in more detail, we will examine a number of traditional responses to them.

4.1 DESCARTES AND SKEPTICAL POSSIBILITY ARGUMENTS

To get the flavor of skeptical possibility arguments, let us begin with the master, René Descartes, in his *Meditations on First Philosophy* (first published 1641). Descartes begins Meditation I by resolving to rid himself of all the opinions he had adopted. Having observed so many of his previous opinions to be in error, he set about “commencing anew the work of building from the foundation” in order to “establish a firm and abiding superstructure in the sciences.” He tells us that reason has convinced him that he ought to withhold belief from what is not entirely certain no less than from what is manifestly false. His strategy is to try his best to doubt as many of his former beliefs as he can, striking at the foundations of these beliefs.

You might wonder about the wisdom of Descartes’ search for certainty, or whether his method of doubt would help establish a firm basis for the sciences. However, you need not agree with him on these matters in order to find what comes next very compelling.

We now will quote one of the most powerful paragraphs in all of Western philosophy. Descartes has just suggested that perhaps because the senses sometimes deceive (e.g., about distant objects or minute objects) we can never trust them, even in the best cases (e.g., about whether Descartes is now before a fire holding a piece of paper). He proceeds to mock this suggestion by comparing such doubts to those of a madman. And yet he continues . . .

Though this be true, I must nevertheless here consider that I am a man, and that, consequently, I am in the habit of sleeping, and representing to myself in dreams those same things, or even sometimes others less probable, which the insane think are presented to them in their waking moments. How often have I dreamt that I was in these familiar circumstances, that I was dressed, and occupied this place by the fire, when I was lying undressed in bed? At the present moment, however, I certainly look upon this paper with eyes wide awake; the head which I now move is not asleep; I extend this hand consciously and with express purpose, and I perceive it; the occurrences in sleep are not so distinct as all this. But I cannot forget that, at other times I have been deceived in sleep by similar illusions; and, attentively considering those cases, I perceive so clearly that there exist no certain marks by which the state of waking can ever be distinguished from sleep, that I feel greatly astonished; and in amazement I almost persuade myself that I am now dreaming.

Let us look at the structure of Descartes’ argument a bit more carefully. Descartes raises a *skeptical possibility*, a possibility in which all his experiences systematically mislead him about the world around him. The possibility he raises is that he is now dreaming. If he were now dreaming, his beliefs about his surroundings would be badly mistaken. In order to be certain about his surroundings based on his sensory experiences, it seems, he must rule out the possibility that he is now dreaming. And yet there are “no certain marks by which the state of waking can ever be distinguished from sleep.” So, how can he be certain he is sitting before the fire, with a piece of

paper in his hand? He can’t. He has found a reason to doubt even his present perceptual beliefs.

Descartes seems to intend his argument to concern certainty. But we can recast it in terms of knowledge. The argument goes as follows:

The Dream Argument

1. You don’t know you’re not now dreaming.
2. Unless you know you’re not now dreaming, you cannot know anything about the world around you based on your current sensory experiences.
3. Therefore, you cannot know anything about the world around you based on your current sensory experiences.

The argument for (1) might proceed much as Descartes’ does. How can you know you’re not dreaming unless you have a test that you could apply to distinguish waking from sleeping? But you don’t have any such test.

If you’re like the author of this chapter, you’ll feel that clearly you *do* know much about things around you and that you *do* know you’re not dreaming. But you’ll worry about exactly *how* you know these things. Most epistemologists, starting with Descartes himself, see the problem this way. The final goal is not to convince people they don’t know much; rather, it is to explain how it is possible that we do know, given these powerful arguments to the contrary.

You might hope for an easy victory over the skeptic. Couldn’t we question the argument for step (1) once we quit certainty for knowledge? Why can’t you know you’re not now dreaming? You know dreams rarely are so orderly, coherent, what we call “realistic.” Some are, but comparatively few. So, isn’t it very likely that you’re not dreaming? If we don’t require absolute certainty for knowledge, isn’t this enough for knowledge, assuming your belief is true and you are not right by the sort of luck found in Gettier cases (see Chapter 3, Section 3.2)?

There are several argumentative strategies the skeptic might use in response. We’ll consider two. First, the skeptic might point out that dreaming could affect memories, too, so that what appears to be a genuine memory in a dream might well not be. Given this, how do you know that what you call your “knowledge” of what dreams are typically like isn’t just a dream-memory and not knowledge at all? A second strategy, perhaps more compelling, is to argue that unless you know now that you weren’t dreaming in the past when you acquired beliefs about what dreams were like, you don’t know what dreams are typically like and so don’t have information that would allow you to know you’re not now dreaming. How do you know you weren’t dreaming on these past occasions? In fact, we might push the dreaming hypothesis further: How do you know your entire life hasn’t been spent dreaming? If you don’t know you haven’t spent your life this way, how can you know anything based on sensory experience either in the past or present?

Descartes himself didn't stop with the dreaming argument. He raised the skeptical stakes higher, with his evil genius possibility. The argument, transformed into an argument about knowledge, is this:

The Evil Genius Argument

1. You do not know you haven't spent your life being deceived by an evil genius.
2. Unless you know you haven't spent your life being deceived by an evil genius, you cannot know anything about the world around you (or even about mathematics or logic).
3. So, you cannot know anything about the world around you (or even about mathematics or logic).

This is a very powerful skeptical argument. Indeed, early in Meditation II, Descartes suggests that only knowledge of his current thoughts—of the fact that he thinks, that he seems to see things, etc.—can elude it.²

Let's sketch the general plan of a skeptical possibility argument. The person giving the argument—call him or her the "skeptic"—targets some ordinary belief or class of beliefs that we ordinarily think we know. Call this belief or class of beliefs *O* (short for "ordinary"). The skeptic then formulates a skeptical possibility (SK)—that is, a possibility in which everything is the same as far as you can tell from the inside (you have the same experiences, apparent memories, beliefs, etc.) but in which the target beliefs are false. The skeptic then argues as follows:

Premise 1: You do not know that not-SK.

Premise 2: Unless you know that not-SK fails, you do not know *O*.

Conclusion: You don't know *O*.

4.2 RESPONDING TO SKEPTICAL POSSIBILITY ARGUMENTS: REJECTING PREMISE 2

One way to respond to such arguments is to insist that you don't need to know not-SK in order to know *O*. To respond in this way is to reject Premise 2. It might be asked: Why exactly must you know you're not dreaming or not a brain in a vat (BIV) in order to know that there is a desk over there, a tree out the window, and so forth? If accounts of knowledge such as those discussed in Chapter 3 are on the right track, all you need to have knowledge is to have a justified true belief that isn't right by luck. Call the no-luck condition, whatever it is, *X*. Couldn't you have such JTB + *X* for the proposition that there is a desk over there without knowing much of anything about skeptical scenarios?

It can't be this simple, though. For one thing, some skeptical possibilities are incompatible with some ordinary propositions that we believe. Consider the skeptical possibility that you are a (handless) BIV. This skeptical possibility can't obtain if you have hands, because having hands entails that it's false

that you are a BIV. Moreover, you can see the entailment. How, then, can you know that you have hands but not know that you are not a BIV, when you can plainly see that if you have hands, you're not a BIV?

The point just made seems to depend on a principle like that of "deductive closure" from Chapter 3. According to the deductive closure principle, if you know *P*, then if you competently deduce *Q* from *P*, believing *Q* as a result of this deduction, then you know *Q*. Applied to the case at hand, deductive closure would tell us that if you know that you have hands and you deduce from this that you aren't a BIV believing you aren't a BIV as a result, you end up knowing you aren't a BIV.

However, as we noted in Chapter 3, deductive closure is controversial when the relevant deduction is question-begging. It was fine to appeal to it in Chapter 3, where we were concerned only with ordinary cases in which we expand our knowledge through deduction (e.g., "my 1983 Toyota Celica has a timing chain problem, so *some do*"). As we noted, if you deduce *Q* from *P* where this deduction is question-begging or circular, it's not clear that you could thereby come to know *Q*. And the deduction from having hands to not being a BIV seems question-begging or circular. We need to formulate a closure principle designed to be neutral about these issues. Let us pause to try to get straight on just what the needed closure principle looks like.

4.2.1 Formulating a Closure Principle for Knowledge

Instead of supposing one has deduced something from something else one knows and asking about the epistemic status of the resulting belief, let's ask about the implications of knowing a proposition with respect to knowing further things that it entails. In some cases, you won't be able to *come to know* something entailed because you *already* know it. Suppose you haven't measured your little brother's height in a few years. He's nine years old. You measure it and so come to know that he is a certain height, say four feet five inches tall (135 cm). But there is no question of your coming to know by deduction from this piece of information that he is less than seven feet tall. You *already* knew, prior to finding out about his exact height, that he isn't anywhere near seven feet tall, based on looking at him.

In other cases, it's not merely that you happen already to know *Q* based on grounds independent of *P*; it's that you *need* to already know *Q* in order to know *P* in the first place. For instance, to use an example of Crispin Wright's (2002, 333), suppose, when watching people mark an X on a sheet of paper, that you can know on the basis of seeing them do this that they are voting only if you already know that an election is taking place. We might well think that this applies to the skeptical case. In order to know I have hands, I have to already know—so it is plausible to think—that I am not a BIV. That's why I can't come to know I'm not a BIV by deducing it from my knowledge that I have hands.

Notice, though, that in both sorts of cases mentioned (where you happen already to know the conclusion and where you had to already know it in

order to know the premise), although you can't come to know the entailed proposition Q from the original proposition P, it is nevertheless true that if you know P then you also do know Q. Even if you already knew Q before knowing P (as in the little brother's height case), it's still true that if you know P, you know Q. Similarly, if you had to already know Q in order to know P, again it's still true that if you know P, you know Q (as in the voting case). Given that these seem to be the only ways you might know P, know that P entails Q, but not be able to come to knowledge of Q through deduction from P, and given that in all of them you do know Q, we can propose a closure principle:

If you know P, and you know that P entails Q, then you know Q.

One further tweak is needed before we continue. Closure, as we've formulated it, seems not to allow for failing to "put two and two together." But I might know that *bazaar* is a word that has a *z* in the middle. And I might know that if *bazaar* is such a word, then some English word has *z* in the middle. But if I don't put these together, I might not believe and so not know that some English word has *z* in the middle. We don't want our closure principle to declare otherwise. Something similar might well be true of skeptical hypotheses. Perhaps we have strong grounds to believe them false, but we might never have assembled those grounds and "put two and two together." To avoid this implication, we'll use this formulation:

(Closure): If you know P, and you know that P entails Q, then you are in a position to know Q.

We call this *closure*, as distinguished from the principle of Chapter 3 we called *deductive closure*. The difference between the principles is irrelevant except when we are considering question-begging or circular deductions.

We now have a closure principle suitable for use in our discussion of skepticism. The principle still has real bite to it. We can appeal to it to defend Premise 2 in skeptical possibilities in which the skeptical possibility (SK) is incompatible with the target ordinary proposition (O). So, for instance, consider this skeptical possibility argument:

- i. You are not in a position to know you're not a BIV.
- ii. In order to know you have hands, you must be in a position to know that you're not a BIV.
- iii. So, you don't know you have hands.

The closure principle is used to defend (ii), thus blocking the sort of response to skeptical possibility arguments that we have been considering. For, if you know you have hands, then one of three situations is the case: (A) you can deduce and come to know that you're not a BIV; (B) you already

knew you aren't a BIV; or (C) you were already in a position to know you're not a BIV from grounds independent of any deduction from *I have hands* but hadn't put two and two together. Whichever of the three situations obtains, if you do know you have hands, you are also in a position to know you are not a BIV. And this is what (ii) says.

What about the original thought behind rejecting Premise 2 of skeptical possibilities arguments—the thought that JTB + X is enough for knowledge of ordinary propositions such as *I have hands* regardless of whether one has any idea of what to think about skeptical possibilities? Appealing as this thought is, we can see that closure makes it doubtful. For, if closure is true for knowledge, and K is JTB + X, then JTB + X itself obeys closure as well: If you have JTB + X for P and you have JTB + X for the proposition that P entails Q, then you're in a position to have JTB + X for Q. But then we can see that we can't happily claim we have JTB + X for ordinary propositions while not worrying much about whether we have JTB + X for claims that we are not in skeptical scenarios. Closure disturbs the consoling thought that "philosophical" worries about skeptical scenarios have nothing to do with whether we know ordinary facts. This is one reason it is such an interesting principle.

One small matter deserves attention before turning to ways one might reject closure. You might notice that argument (i)–(iii) differs from the skeptical possibility arguments discussed above, and from the general schema for them, by including "in a position to know" instead of "know" in several places. That substitution was needed in order to make use of the closure principle to justify (ii). This substitution requires the skeptic to defend the claim—(i) here—that you aren't in a position to know the relevant skeptical possibility fails to obtain. A question arises here: Is it harder for the skeptic to defend this claim than to defend the weaker claim that you simply don't know that possibility fails to obtain? If so, appealing to closure would come at a cost to the skeptic.

Fortunately for the skeptic, substituting "position to know" for "know" in the first premise doesn't make that premise any harder to defend. The guiding idea behind the skeptic's claim that you don't know that you're not dreaming, that you're not a BIV, and so forth was always that you don't have the materials to know, and never merely that you don't know. A person might fail to know something simply because he hasn't put two and two together. For example, suppose it's Sunday. You know the library is closed on Sundays, and you know that today is Sunday. But you fail to put two and two together and so, failing to realize it's closed today, you head off to the library. In this case, you don't know it's closed today, but you *would* know this if only you would put these other pieces of information together in your mind. You're in a position to know it's closed today. However, according to the skeptic, given what you have to go on, you can put all your evidence together all you like and you still won't end up knowing you're not a BIV. For the skeptic, you're not in a position to know such things. The switch from "knows" to "position to know" if anything *better* expresses what the skeptic is trying to say.

Recall the skeptical argument:

- i. You are not in a position to know you're not a BIV.
- ii. In order to know you have hands, you must be in a position to know that you're not a BIV.
- iii. So, you don't know you have hands.

From this discussion, we may conclude that in order to block the second premise of arguments like (i) – (iii), one would need to reject closure. We need to see whether there is a plausible way to do this.

4.2.2 Nozick on Skepticism: A Way to Reject Closure?

Robert Nozick doesn't simply declare closure false when applied to these cases. Rather, independently of concerns about skepticism, he defends an account of knowledge under which it turns out that closure fails exactly in these sorts of cases. On his account, which we discussed in Chapter 3 (Section 3.4.2), in order to know *P*, your true belief in *P* must be sensitive—that is, if *P* were false, you wouldn't believe *P*. (We can ignore methods for the purposes of this discussion.)

Working with this account of knowledge, Nozick gives a clever response to skeptical possibility arguments. Such arguments are appealing because the first premises are true: You don't know, and you're not in a position to know, that you're not a BIV, that you're not dreaming, and so forth. Were you in one of these skeptical scenarios, you would still think you weren't. That's why you don't know. (He thinks, moreover, this captures our sense of why it seems we don't know these things.) So, for Nozick, the skeptic is right about that much. However, the skeptic is wrong to conclude from this that you lack knowledge of ordinary matters, such as that you have hands. Just because your belief that you're not a BIV isn't sensitive, it doesn't follow that your beliefs in ordinary propositions that entail this proposition are not sensitive. Sensitivity isn't closed under known entailment. You can know *P*, know *P* entails *Q* but not know *Q*, and yet not be in a position to know *Q*. The closure principle for knowledge, on Nozick's account of knowledge, has many true instances, but it fails in precisely these sorts of cases. Premise (ii) above and all Premise 2's in skeptical possibilities arguments are therefore false. So Nozick argues.

How could one sensitively believe *I have hands* without sensitively believing *I am not a BIV*? Let's say that the actual world is "normal" just if it is roughly what you think it is—that is, a world in which you perceive tables, chairs, and so forth in all the normal ways. If the actual world is normal, then the BIV world is very remote from actuality, unlike, say, worlds in which Mitt Romney won the 2012 U.S. election. Now, a statement of the form "had *A* not been the case, *C* wouldn't be the case" is true in a world *W* if and only if in the closest worlds to *W* in which *A* is false, *C* is also false. Assume the actual world is a normal one. Then the closest worlds in which you don't have hands are ones in which you lost them in an accident, say. In such worlds you can plainly see

and feel that you don't have hands and so you wouldn't believe you do have hands. Thus, if the actual world is normal, your belief that you have hands is sensitive. However, whatever the character of the actual world, your belief that you're not a BIV is insensitive: If you were a BIV, you'd still think you weren't a BIV. So, if the actual world is normal, your belief that you have hands is sensitive, whereas your belief that you aren't a BIV is insensitive. Putting this together with Nozick's account of knowledge, it follows that if the actual world is normal, you know you have hands but don't know you're not a BIV.

Is Nozick begging the question against the skeptic by appealing to the fact that our world is a normal one? No. His point is that it is not a necessary condition of having knowledge of ordinary matters that you know or are in a position to know that you're not in skeptical scenarios, contrary to what Premise 2 in skeptical possibility arguments asserts. To show one thing, *A*, isn't a necessary condition of another, *B*, it's enough to give a possibility in which *B* obtains but *A* doesn't. This is what he does. His claim is not that ours is a normal world (although he of course believes this to be true), but that normal worlds are possible, and that if ours is a normal one, then ordinary beliefs such as your belief that you have hands would be sensitive, even though the belief that you aren't a BIV wouldn't be sensitive. Given his theory of knowledge, this would show that it is possible for you to know that you have hands but not know or be in a position to know that you are not a BIV. Thus, Premise (ii) is false in the above skeptical possibility argument, and generally whatever the skeptic chooses for Premise 2 will be false. Thus, skeptical possibility arguments fail—all of them.

As we saw in Chapter 3 (Section 3.4.2), there are apparent counterexamples to the sensitivity requirement on knowledge that have nothing to do with closure. So anyone following Nozick will need to find something plausible to say in response to these counterexamples.

But even putting such counterexamples aside, it seems impossible to limit counterexamples to closure to skeptical cases. In Chapter 3, we noted that on Nozick's account, you could know *P* but not be in a position to know through deduction that *you do not falsely believe P*. This was a violation of the deductive closure of knowledge. The same example can be used to show that Nozick's account violates the closure principle more appropriate to theorizing about skepticism, the principle we are calling "closure." You do know that *P* is true (for some *P*, such as *there is a chair over there*), and you do know that this entails that you don't falsely believe *P*. But you do not sensitively believe the latter. There are also other counterexamples that link the ordinary to mildly skeptical propositions. Consider this one: I know there is a maple tree standing in my front yard. I'm at my office now, several miles from home, but I still know this. Had it been false, I would have known about it, it seems, because in the closest worlds in which it is false I would have arranged to have the tree cut down or I would have witnessed the fallen tree in the morning after a storm. Now consider the possibility that the maple tree hasn't been cut down moments ago by rampaging hooligans. This proposition is entailed by *there is a maple tree standing in my front yard*. But my belief in it is not sensitive.

Had the maple just been cut down moments ago, I wouldn't have known about it (yet). So while I know there is a maple tree standing in my front yard, Nozick must say that I don't know that rampaging hooligans haven't cut it down moments ago. You can invent many other similar examples (e.g., you know your car is parked in such and such lot, but you aren't in a position to know it hasn't been stolen and driven away).

Stemming the tide of counterexamples to closure is one problem for Nozick's account of knowledge. Another is the puzzle of why, if the instances of the closure principle involving skeptical scenarios are false, we feel a pressure to retract our claim that we know ordinary matters such that we have hands after we concede we don't know we're not BIVs. Why does conceding you don't know you aren't a mere BIV make you feel you really can't just continue to maintain that you know that you have hands, feet, and so forth? The truth of the closure principle would explain why we feel this pressure: We feel it because we sense that knowing something (P) that one can see entails something else (Q) isn't possible unless one is in a position to know Q.³

We have seen there are serious problems with the attempt to block skeptical possibilities arguments at Premise 2, that is, the premise that unless one knows (or is in a position to know) \sim SK, then one does not know O. At least when the ordinary propositions (O) entail the falsity of the skeptical hypothesis SK, it seems Premise 2 is in good shape. In this author's view, a more promising line of response is to grant Premise 2 to the skeptic but to reject Premise 1—that is, to insist that we *do* know that the various skeptical hypotheses are false. We *do* know that we're not BIVs. We *do* know that we're not dreaming and so forth.

4.3 REJECTING PREMISE 1: YOU KNOW YOU'RE NOT A BIV!

Fine, you might say, but *how* do we know these things? Notice what this question seems to be asking for: a reason, a justification. We might hope to reply dismissively by claiming that the issue of justification is irrelevant, because justification is not necessary for knowledge. So long as you have true belief that is not true by sheer luck, then you know. In Chapter 3, we discussed various proposals that might justify this sort of reply, in particular proposals featuring reliability conditions of one sort or other on knowledge (though not the sensitivity proposal!). The idea, in general, would be to eliminate the justification condition on knowledge and think instead of knowledge as requiring only true reliable belief, and then claim that we can know the skeptical possibilities are false without having any reason or justification.

Yet it has seemed to many philosophers, including this author, that this is not enough. If all we can justifiably say in reply to the skeptic is that maybe we're hooked up nonaccidentally to the truth in this way, then all we can justifiably say is "maybe we know." But this seems problematic in at least two ways. First, it's an awfully weak answer. I think I'm justified in thinking that I do know, not only that maybe I do. Second, it's an awfully odd answer to

the skeptic. It is the sort of answer that we might think relies on a false theory of knowledge. Knowledge and justification are not so cleanly separable as this response would have it. To see this, imagine a team of detectives working on a murder case. They have some good evidence, let's suppose, that the suspect did it. They might say, "We have some reason to believe he did it, but we don't know yet that he did it." Later, after much stronger evidence comes in, they might say, "Now we don't merely have reason to believe he did it; we know it." It seems the detectives are treating degrees of justification as a kind of scale with knowledge on top or at least at the upper end. This way of thinking about knowledge cannot simply be dismissed by saying that knowledge is one thing, justification is something very different. They seem importantly related, and the skeptic is drawing on this relationship.

Well, then, how do we know we're not living in a skeptical scenario? Here are three traditional proposals for explaining how we pull this off:

1. We know such things a priori; that is, our knowledge is not based on experience.⁴ Call this the *a priori proposal*.
2. We know them a posteriori, but not based on knowledge of the world we've gained through experience, only based on knowledge of what our experience is like together with our apparent memories. Call this the *straightforward a posteriori proposal*. It is straightforward because it doesn't seem to beg any questions against the skeptic, unlike the next proposal.
3. We know them a posteriori, based on knowledge of the world we've gained through perception. Call this the *bold a posteriori proposal*.

The remainder of the chapter will consider each of these three proposals in turn. We postpone to the following chapter a fourth relatively new proposal concerning the semantics of "know."

4.4 THE A PRIORI PROPOSAL

The eighteenth-century Scottish philosopher Thomas Reid argues that we are justified in believing and presumably know certain "first principles" concerning contingent matters of fact. The justification is a priori because it is based on the self-evidentness of the principles rather than experience.

Some of the first principles Reid lists concern a person's own mental life, but most on his list concern the world outside one's mind. Consider these examples (Reid 1785, Essay 6, Chapter 5):

First Principle Concerning Memory

That those things did really happen which I distinctly remember.

Concerning Perception

That those things do really exist which we distinctly perceive by our senses, and are what we perceive them to be.

Concerning Reasoning

That the natural faculties, by which we distinguish truth from error, are not fallacious.

Concerning Other Minds

That there is life and intelligence in our fellow-men with whom we converse . . . That certain features of the countenance, sounds of the voice, and gestures of the body, indicate certain thoughts and dispositions of mind.

None of these propositions is true as a matter of necessity. There are possible worlds in which what we distinctly remember (i.e., distinctly seem to remember) didn't really happen; worlds in which what we distinctly perceive (what we experience) isn't reflective of what there is; and so forth. Reid claims, nevertheless, that we are justified in believing that memory and perception are reliable.⁵ So, Reid claims our beliefs in first principles are justified; and if they turn out to be true nonaccidentally, they constitute knowledge, at least according to the sorts of views of knowledge discussed in Chapter 3. The knowledge would be *a priori* because it is not based on experience, but rather on the self-evidentness of the principles.⁶ Using this knowledge, we could clearly rule out a great many skeptical hypotheses. If your senses are reliable, you are not a BIV, for instance. If your reason and memory are reliable, a Cartesian demon is not constantly deceiving you in your reasoning about mathematics and logic.

Reid's claims stand in stark contrast to the traditional view of the *a priori* going back at least to Kant. On the traditional view, only necessary truths can be known *a priori* and contingent truths can only be known with the help of experience (understood broadly to include introspection of one's own mind). One reason for thinking contingent truths cannot be known *a priori* is that when we put aside experience, all that remains as a source of knowledge is pure reason—and pure reason by itself can't tell you that you are in one possible world rather than another. It can only tell you about what is true in all possible worlds or none. You have to, so to speak, "look to see" what is *special* the case in the possible world that is actual.⁷

Although Reid maintains there is no way to prove a first principle via "direct or apodictical proof," he does offer some suggestive remarks about how we might, as he says, "reason even about them." He gives us some marks by which we might decide what is a first principle and what isn't. Such marks would be quite useful, for if we could determine that a proposition *P* is a first principle, then since first principles are true, we could conclude that *P* is true. Some of the marks Reid lists could only be verified *a posteriori*, such as the "consent of ages and nations" or the appearance of certain beliefs early in our lives before education, as well as the agreement in the testimonies our faculties give us. You can't know *a priori* what the consent of ages and nations is, nor when beliefs appear in normal development; and you must use knowledge of what your experiences are over time to know about the agreement of our faculties' testimonies. However, two of the marks he mentions are things we can arguably know *a priori*, and we will focus on these.

First, Reid suggests that first principles might "admit of proof by *reductio ad absurdum*." To prove something by *reductio ad absurdum*, we assume the proposition is false and derive an absurdity from it; we then conclude, from the fact that the assumption led to absurdity that the proposition must be true. Here is an example from elementary logic. How can we prove the following is true: "If *P* & *Q*, then *P*"? Like this: Assume it is false. Then *P* & *Q* is true and so *P* is true; but *P* is also false. But that is absurd. Therefore, the statement "If *P* & *Q*, then *P*" is true. This strategy seems appropriate for proving *necessary* first principles, such as those of logic, but the ones important for us at the moment (e.g., the first principle concerning perception) are *contingent*. If we assume them to be false, how can we derive an absurdity?

The second consideration concerns consistency. Reid writes:

. . . [it] is a good argument *ad hominem*, if it can be shewn that a first principle which a man rejects, stands upon the same footing with others which he admits: for, when this is the case, he must be guilty of an inconsistency who holds the one and rejects the other.

Thus, the faculties of consciousness, of memory, of external sense, and of reason, are all equally the gifts of nature. No good reason can be assigned for receiving the testimony of one of them, which is not of equal force with regard to the others. The greatest sceptics admit the testimony of one of them, and allow that what it testifies is to be held as a first principle. If, therefore, they reject the immediate testimony of sense or of memory, they are guilty of an inconsistency." (Reid 1785, Essay 6, Chapter 4)

It is a striking argument, but the skeptic has a number of replies available.

On the one hand, the skeptic might heartily agree that consciousness (which for Reid amounts to introspective beliefs about one's mental life) is in the same boat as sense experience: We can devise Cartesian-style skeptical arguments that attack it, too. For instance, perhaps the Cartesian demon can make it seem that you have certain experiences, beliefs, desires, and so forth when you don't. If this is possible (an interesting topic in itself), then the skeptic will argue as follows: You don't know there isn't a demon deceiving you about these things, and therefore your introspective beliefs aren't knowledge.

On the other hand, the skeptic might reply that Descartes was right that skeptical possibilities must include stipulations about aspects of our mental lives—for example, that things at least seem certain ways to us—and these aspects therefore can't be targeted by skeptical possibility arguments. If so, then we have a good basis for treating consciousness differently than perception ("external sense") or memory, contrary to what Reid claims.⁸

Finally, as Reid acknowledges, he proposes consistency only as a basis for an *ad hominem* argument (i.e., only to point out irrationalities on the part of certain philosophers). If we tried to turn his remarks into a defense of the principles about sense perception, consciousness, memory, and so forth, by suggesting that they concern natural faculties and that natural faculties are all reliable, we would face a suite of difficult questions: What is naturalness? Can we know *a priori* what is natural? Even if we can know this *a priori*,

what basis could we have a priori for thinking that naturalness conduces to reliability rather than unreliability?

The a priori strategy is by no means ruled out, but its proponents have much work to do.⁹

4.5 THE STRAIGHTFORWARD A POSTERIORI PROPOSAL

We next turn to the *straightforward a posteriori* strategy, which appeals to facts about what our experiences and apparent memories are like. This strategy attempts to do just what the skeptic thinks we must do: to argue, based on the character of how things seem from the inside, to the conclusion that we're not in the various skeptical scenarios. The main worry about this strategy is that this is just too meager a set of data from which to draw the desired conclusion.

How might the argument from our experiences and apparent memories to the conclusion that we're not BIVs go? There seems no way to deduce from the fact that you have certain experiences and apparent memories that you are not a BIV. To deduce is to draw a conclusion that *must* be true if the premises one uses in the deduction are true. But the truth of premises about your experience and apparent memory does not guarantee you are not a BIV, for BIVs, too, can enjoy the very same experiences and memories you have (this is how the BIV hypothesis is specified.) We might do better to appeal to *inference to the best explanation*.¹⁰ Inference to the best explanation is a mode of inductive inference whereby one infers the truth of the best explanation of the data. The best explanation isn't strictly deducible from the data. There might well be rival explanations that are compatible with the data (but not with the best explanation). But still these rival explanations may well not explain the data as well. Indeed, when, in the Sherlock Holmes stories, Holmes "deduces" things, he doesn't use what philosophers call deduction; he uses inference to the best explanation. Of course, it's not only detectives who use this mode of argument. We use it when we try to figure out what's wrong with our cars or why the soup we cooked for dinner turned out so bland.¹¹ The claim we want to consider, then, is that your experience and apparent memories are better explained by the hypothesis that you have a body and interact perceptually and agentially with a real world in much the way you think we do (call this the *Real World Hypothesis* [RWH]) than by any skeptical hypothesis, and therefore that you can know the skeptical hypothesis to be false.

In Chapter 1 (Section 1.9), we considered a principle about inference to the best explanation:

(JEI) Justification by Explanatory Inference

If hypothesis H purports to explain S's evidence E, and there is no incompatible hypothesis H' that provides a better or equally good explanation of E, then S is justified in believing H on the basis of E.

Numerous questions arise concerning JEI. How should we understand "no incompatible hypothesis"? Should we take this to subsume all hypotheses, whether anyone has ever formulated them or not? If so, it could be difficult to

know in many cases whether we are justified in believing H on the basis of E; perhaps there is some incompatible hypothesis out there that we haven't thought of that better explains E. If, instead, we take "no incompatible hypothesis" to subsume only those incompatible hypotheses we know about, only those available to us, we face a different problem. Couldn't H be the "best of a bad lot," as Bas Van Fraassen puts it (1989), so that it provides the best explanation compared with the available rival hypotheses, but still not a particularly good one, with the consequence that we aren't justified in believing H? Suppose there are ten possible suspects for a burglary. One of these ten had to do it, because they were the only ones with opportunity. We don't have any good explanation, for any of them, though, why they committed the burglary. There presumably is some excellent explanation out there, but it isn't available to us ready to hand. However, we do know that one of the ten suspects, Scott, once shoplifted. The proposition *Scott committed the burglary* therefore offers a slightly better but still weak explanation of burglary. Are we therefore justified in believing that poor Scott did it? That seems far too strong.

There are further difficulties. What if H is in fact the best explanation of E, but the subject S has no idea this is so? The best explanation for the observation of rainbows is one appealing to refraction and reflection of light within water molecules at a certain angle to the observer. But human beings haven't always been justified in believing this hypothesis! It took major developments in science to discover it. We could revise JEI by adding a qualification to the effect that the subject S know or justifiably believe that H better explains E than the competing hypotheses. But this is asking a lot, in the skeptical case. Must an ordinary person, to be justified in believing RWH, let alone ordinary propositions like *this is a tree*, know that these provide the best explanations of the evidence? Is an ordinary person even in a position to construct such explanations?

Aside from these reservations, we have to ask whether RWH is a better explanation than any of the competing hypotheses, including skeptical hypotheses. If we could help ourselves to our accumulated scientific knowledge, the answer would be easy. But to do this is to beg the question against the skeptic; it is to appeal to knowledge that the skeptic calls into question. What makes the straightforward a posteriori approach straightforward is that it plays fair with the skeptic and thus disallows such question-begging argumentation. Suppose, though, we limit ourselves to the character of our experiences and apparent memories and whatever else we know a priori. Within these constraints, does RWH provide a better explanation of our experiences and memories than skeptical hypotheses? Many philosophers, going back at least to John Locke, have argued in the affirmative.¹² Reid, too, although he says that first principles are beyond evidence, speaks of memory and "external sense" as confirming one another. Here are some of the regularities these philosophers cite:

- Regularities within any sense-experience modality (e.g., if you seem to see someone pick up a rock and make a throwing motion, you can expect to seem to see it flying through the air toward the window)

- The cross-modal regularities of sense-experiences (e.g., if you seem to see a rock heading in the direction of glass, you can expect to seem to hear a sound of the glass breaking)
- Regularities in the connection between decision, effort, and sense-experience (e.g., if you decide to throw the rock, you can expect to feel as if your muscles contract and then release with the rock seeming to leave your hand, and then as an apparent result to seem to see it flying through the air, hear it strike the window, etc.)
- Regularities between apparent memories, effort, and current experience (e.g., if you have an apparent memory of having placed your keys on the mantel, you can expect that, upon seeming to walk over to the mantel and seeming to look at it, you will seem to see keys on it)

Such regularities seem well explained by RWH. This is because they seem to be just what we would expect if there were real objects we perceived and which had at least roughly the character we believed them to have. By contrast, consider the hypothesis that your experiences and memories are simply “random” (i.e., that they have no explanation at all). How likely would it be that they take this coherent form?

But this “random” hypothesis is not the skeptic’s; the skeptic appeals to lifelong dreams, evil geniuses, and super-neuroscientists. We can distinguish versions of skeptical hypotheses, weaker and stronger versions. Strong versions build into the hypothesis assumptions about the character of your experience and memories. The weakest versions don’t. An example of a strong version is the hypothesis that you are a BIV with these very experiences and memories. An example of a weaker version is that you are a BIV stimulated by a neuroscientist of some sort or other. Where not otherwise noted, in this book we have in mind strong versions of skeptical hypotheses.

With this strong/weak distinction in mind, let’s ask: Could we make sense of the coherence of your experiences and memories assuming stronger or weaker variants of the BIV possibility? Start with the weak hypothesis that you are simply a BIV. Being a BIV seems compatible with just about any run of experience, coherent or incoherent. Why think your experiences would be coherent if you were simply a BIV? Suppose we strengthen the hypothesis by adding to it that the scientists give you orderly experiences of a kind that make it seem that you perceive tables, trees, people, and so forth. Still, there are many sorts of coherent regularities involving tables, trees, and people that are a priori possible. Why does your experience exhibit these particular coherent regularities rather than others? Suppose we strengthen the hypothesis further so that it builds in that you have the very experiences and memories you have. Then we have no difficulty seeing that if that hypothesis were true, your experiences and memories would exhibit the sorts of regularities they do. The hypothesis clearly predicts the regularities. But are they explained by the hypothesis? It seems they aren’t.

The skeptic can point out that the evil neuroscientists themselves have plans and execute these plans. Could this provide the materials for an equally

good explanation of the particular sorts of coherence exhibited in our experience and memory? The skeptic’s general strategy might be to create a mapping between the supposed real-world objects and a realm of “substitute” entities that preserves properties and relations (Vogel 2005), or to use a term from mathematics to create an *isomorphism*. Thus, just as there are rocks and windows, with their particular properties and modes of interaction, the skeptic would posit as part of her hypothesis ideas in the mind of the deceiver with corresponding properties and relations; or perhaps instead of ideas she might choose files in a supercomputer, or what have you. The idea is to piggyback off the RWH but in a way that would make our experiences and memories very much inaccurate.

Is the RWH better than such a complex skeptical hypothesis? You might think the latter is ad hoc, complicated, forced, and so forth, and so a worse explanation. But supposing all this is true, are such factors epistemically relevant? Do they make the hypothesis less credible for us? Or are these features of explanation merely ones we find pragmatically useful—we can work better with theories with these “virtues,” we can better understand them, and so forth?

The explanationist response to skepticism is by no means hopeless, but it is fair to say that making the case that the response succeeds is not a straightforward matter. And we need more than just a reason to think the evidence tips in favor of RWH over skeptical hypotheses. We are looking for a reason good enough to provide us with knowledge of the external world, and it is not clear that the explanationist response gives us this.

4.6 THE BOLD A POSTERIORI PROPOSAL

Finally, consider the bold a posteriori proposal. On this view, you can use information about the world gained through present and past sensory experience to know skeptical hypotheses are false. Here is one way the argument that we’re not BIVs might go:

1. I have hands.
2. So, I’m not a (handless) brain in a vat.

(You might find this reasoning humorous: Is it a joke? It’s interesting to ask why it seems humorous.) Can I know (2) through this reasoning from (1)? This example borrows from early-twentieth-century philosopher G. E. Moore’s famous “proof of an external world.” In his 1939 lecture to the British Academy, Moore told the audience:

It seems to me that, so far from its being true, as Kant declares to be his opinion, that there is only one possible proof of the existence of things outside of us, namely the one which he has given, I can now give a large number of different proofs, each of which is a perfectly rigorous proof. I can prove now, for instance, that two human hands exist. How? By holding up my two hands, and saying, as I make a certain gesture with the right hand, “Here is one hand,” and adding, as I make a certain gesture with the left, “and here is another.” (1959, 145–146)

Moore addressed his proof to idealists—that is, philosophers who think there are nothing but minds and ideas (and so no physical mind-independent world). It would take us too far afoot to discuss its merits against idealists. But his argument, his “proof,” inspires the reasoning (1)–(2) above.

Such Moorean reasoning (reasoning via Moore-style proofs) clearly begs the question if given in response to the skeptic. The skeptic argues that you don’t know (1) because you don’t know (2). If, following Moore, you respond by arguing from (1) to (2), you’ll be using a premise the skeptic has just claimed you can’t know unless you already know (2)! In arguing from (1) to (2), you’re simply ignoring his arguments, not engaging with them. To the extent that someone found the skeptic’s arguments compelling, that person could not be convinced out of them by your reasoning. Your reasoning from (1) to (2) is thus dialectically useless.

However, think about arguing against a determined flat-earthier. The flat-earthier claims the earth is flat and explains to you how there is a massive conspiracy afoot that extends to satellite photography. Suppose you reply to him by citing satellite photography, without explaining how it is that this photography in fact isn’t a hoax. You simply appeal to the photography. Clearly you beg the question against this flat-earthier. To the extent that someone found the flat-earthier’s arguments compelling, that person couldn’t be convinced by your merely appealing to the satellite photography. Still, such photography can help a person come to know the earth isn’t flat.¹³ Might it be similar with Moorean reasoning? When we discussed closure in Section 4.2.1, we put aside questions about whether deductions that beg the question can be ways of expanding our knowledge. We wanted closure to be neutral on that matter. Here we return to those questions. Might Moorean reasoning, despite its question-begging character, still be a way of giving us knowledge that we’re not in skeptical scenarios? It might be speculated that the historical Moore had something like this in mind. He surely knew his “proof” wouldn’t satisfy idealists, because it begged the question, but perhaps he thought it was a perfectly good proof nevertheless, one through which a person could know that idealism is false.

So, in evaluating the Moorean reply to the skeptic (i.e., the bold a posteriori reply), we must distinguish dialectical issues from epistemic ones.¹⁴ The key question is whether Moorean reasoning could give us knowledge that we are not in skeptical scenarios, or whether, instead, we’d have to already know we’re not in skeptical scenarios in order to know the premises of that reasoning. If one already has to know or be in a position to know the truth of the conclusion of a piece of Moorean reasoning like (1)–(2), then such reasoning couldn’t enable us to know we’re not in a skeptical scenario.

Let’s try to get clearer on just how the Moorean reasoning is supposed to give us knowledge we’re not in skeptical scenarios. The Moorean claims that reasoning from (1) to (2) is a way of coming to know that one isn’t a BIV. So far, so good. But how is it that the reasoning is supposed to pull off this feat? Let’s ask the Moorean: Does the fact that one has the experience of having hands by itself provide sufficient grounds for knowing the conclusion that

one isn’t a BIV, without going through the Moorean reasoning? If the Moorean answers “yes,” it seems the Moorean is sliding back to the straightforward a posteriori proposal. We would have to explain how it is that facts about experience could provide strong reasons to think skeptical hypotheses are false. The Moorean proposal is supposed to be distinct from this. It is supposed to be bold in a way that the straightforward proposal isn’t. So, the Moorean’s answer should instead be: No, the experience by itself isn’t strong evidence that one isn’t a BIV, but nonetheless if one begins with the experience and goes through Moorean reasoning, one gains and relies on strong evidence that one isn’t a BIV and so comes to know that one isn’t a BIV.

So, the truly novel idea in the Moorean proposal is that what the experience can’t do by itself can be done when one goes through the Moorean reasoning. You might ask: How does this work? How, by going through the Moorean reasoning, does one gain more evidence for believing one isn’t a BIV—more empirical evidence, not a priori evidence (for recall this is the bold a posteriori proposal)?

This might seem like magic. It might seem too easy a way of acquiring knowledge that one isn’t a BIV if one didn’t have it before. It’s as if one begins with weak evidence (the experience) and somehow bootstraps one’s way up to strong evidence simply by reasoning, without any further information coming in.

Here is one way the Moorean might attempt to explain the power of Moorean reasoning. Your experience by itself is sufficient to give you justification that you have hands and, when all goes right (you do have hands, you’re not in a Gettier case), knowledge that you have hands. You don’t need already to know that you aren’t in a skeptical scenario. Then, once you have the knowledge that you have hands, you have a new reason, beyond the mere experience, to believe that you are not a BIV. This new reason is that you have hands. This is very strong support indeed for the conclusion that you are not a BIV. That you have hands entails that you aren’t a BIV, whereas that you have an experience doesn’t.

What the Moorean might suggest, then, is a two-part proposal. First, experience gives us grounds for knowledge of ordinary propositions about the world, regardless of whether we are already justified or know that we are not in skeptical scenarios. Second, once we do have this knowledge, we have evidence that we know entails that we are not BIVs. We can then deduce and come to know—if we didn’t already know it—that we are not BIVs. In Chapter 6, we will examine in detail whether this two-part proposal is too good to be true.

So, we have reached some tentative conclusions about what the bold a posteriori proposal must look like and what problems it faces. We have found that if it is to distinguish itself from the straightforward a posteriori proposal, it must attribute some epistemic power to Moorean reasoning: Experience alone isn’t a sufficient basis for knowing one isn’t a BIV, but the Moorean reasoning somehow gives one a sufficient basis. The key question about it is how Moorean reasoning could have this epistemic power.

4.7 CONCLUSION

We have examined a number of traditional responses to skepticism and found none clearly correct. This does not mean we ought to be skeptics. As the historical Moore is also famous for remarking, it would not be rational to be as confident of the premises of any skeptical argument as we are that we know a lot about the world around us. In response to a skeptical argument of Bertrand Russell's based on four assumptions, Moore writes:

I cannot help answering: It seems to me *more* certain that I *do* know that this is a pencil . . . than that any single one of these four assumptions is true, let alone all four . . . Nay, more: I do not think it is *rational* to be as certain of any one of these four propositions, as of the proposition that I do know that this is a pencil. (Moore 1959, 226)

Moore is exactly right here, in the view of this author. We should not give up our belief in our knowledge of the external world in the face of skeptical arguments. Similarly, when we read about the liar paradox or Zeno's paradoxes of motion, the right reaction to take when we can't see exactly what goes wrong is not to stop believing in truth or motion; rather, it is to hold on to those beliefs and think there must be something wrong in the arguments. For example, consider the following argument from Zeno:

For any two places, A and B, which are not right next to one another, to move from A to B, a thing must first go halfway from A to B. But to go halfway from A to B, one must first go halfway from A to the halfway point between A and B—that is, one must go a quarter of the way from A to B. Similarly, to go a quarter of the way from A to B, one must go an eighth of the way. This series is infinite. Each distance corresponding to an element in the series takes some finite amount of time to cover in one's motion. An infinite sum of finite amounts of time is itself an infinite amount of time. Thus, in order to move from any place A to another separated place B, it takes an infinite amount of time!

It takes the concepts of calculus to see where and why precisely it goes wrong. If you aren't a calculus pro, still you shouldn't be taken in by the argument. You *know* there is motion. The rational thing to conclude, even if you don't know where the argument goes wrong, is that it does go wrong somewhere. This is exactly what Moore thinks we should think in response to skeptical arguments like the ones we have considered in this chapter.

Of course, such a response is not all that we want. We want to know what goes wrong in the skeptical arguments and more important how and why it goes wrong. In this chapter, we focused on attempts to explain why Premise 1 of the skeptical possibility arguments goes wrong. (Recall that Premise 1 is the step that denies that we know that the skeptical possibility fails to obtain.) We discussed in some detail three strategies for

explaining how we know we're not in skeptical scenarios: the a priori strategy, the straightforward a posteriori strategy, and the bold a posteriori strategy. But, speaking for myself, I can say that, like Moore, I am much more confident, and I think rationally so, that I do know about the external world than I am that any particular one of these strategies succeeds. (Easy to say, you might think! But consider your own case. What would *you* say?)

This chapter has discussed the problem of skepticism about knowledge from a traditional perspective. It has not brought to bear any of the innovations of recent years. In the next chapter, we will discuss the antiskeptical potential of one such innovation, viz. contextualism about "knows."

QUESTIONS

1. Play the part of the skeptic. What would you, as skeptic, say in response to the following argument that you know you are not dreaming?

Dreams typically are disjointed and don't make use of memories of events the day before. Waking life is typically not like this. So, it's very likely that right now, I'm not dreaming, because my experience right now is coherent and I'm recalling events the day before. Maybe I can't be absolutely certain on this basis that I'm not dreaming, but I can be very justified—enough to know I'm not dreaming.

2. Nozick argues that if our world is normal, then we know we have hands but don't know we're not BIVs. This is because the BIV world is so very far removed from actuality. But some skeptical hypotheses aren't so far removed, it would seem. Consider the possibility that you're having a very realistic dream right now that you're reading a book. Is this at all "abnormal"? If it isn't abnormal, does Nozick have to concede that you don't know you're reading a book unless you know you're not merely having a realistic dream that you're reading a book? Or is there a way Nozick could apply his general strategy to the dreaming possibility as he does to the BIV one? If so, explain how he might do this.
3. Is it inconsistent, as Thomas Reid claimed, to think introspection delivers us knowledge of our minds while thinking that skeptical arguments show that perception cannot deliver us knowledge of the world outside our minds? Try to devise a skeptical possibility argument against introspection. Can you do it?
4. In his "Proof of an External World," Moore (1959, 146) gives three conditions he takes to be individually necessary and jointly sufficient for a proof: (i) the conclusion must be different from the premises; (ii) the conclusion must follow from the premises; and (iii) the premises must be known to be true. Do you think these conditions are indeed individually necessary for a proof? Do you think they are jointly sufficient? Explain.
5. Moore claimed that his "proof of an external world," which you recall took the form of holding up his hands and arguing "here is one hand, and here is another, therefore there are external objects" is exactly analogous to a perfectly good proof that there are three misprints in a manuscript that one might offer by pointing to each of the three and remarking "here is a misprint, here's another, and here's a third" (Moore 1959, 147). Citing the three misprints settles the

question of whether there are any misprints in the manuscript; it proves that there are some. Does the proof of an external world, by the same measure, prove that there are external objects? Why or why not? What is the relevant difference between the two "proofs" if there is any?

6. One might argue that it doesn't matter to you whether you're a BIV; all that counts is your psychological state (i.e., your experiences, feelings, and apparent memories). That's what makes life worth living. Since, by stipulation, you're in the same psychological state if you're a BIV, who cares? Consider the following response. You *do* care about things other than your psychological state; you care about having real friends, not just "virtual" ones. You care about being in good physical shape, not just "seeming" to be in good physical shape from the inside. And so on. If the Real World Hypothesis is true, you're achieving some things you care about that you wouldn't be achieving if the BIV hypothesis were true. Thus, it does matter whether you're a BIV. Do you agree with this response? Why or why not?
7. For the sake of argument, assume that it does matter whether you're a BIV. It's better not to be a BIV than a BIV, other things equal at least. Even still, does it matter whether you *know* you're not a BIV? Try to give the best case you can for answering yes, and then try the same for answering no. Which is the better case? Why?
8. We focused in this chapter on skeptical possibility arguments for skepticism. There is another sort of skeptical argument called an "underdetermination" argument. The argument goes like this. We do not know that the Real World Hypothesis (RWH) is true unless we have evidence that favors it over its rivals, such as the BIV hypothesis. But our evidence does not favor RWH over BIV. Our evidence "underdetermines" which of these two hypotheses is correct. Therefore, we do not know the RWH hypothesis is true. Where is the weakest point in this argument?

FURTHER READING

- Blumenfeld, David, and Jean Beer Blumenfeld (1978). "Can I Know I Am Not Dreaming?" In Michael Hooker (ed.), *Descartes: Critical and Interpretive Essays*. Baltimore: Johns Hopkins University Press.
- Descartes, René [1641] (1996). *Meditations on First Philosophy*. In J. Cottingham and B. Williams, eds., *Meditations on First Philosophy and Selections from the Objections and Replies*. Cambridge: Cambridge University Press.
- Logue, Heather (2011). "The Skeptic and the Naïve Realist." *Philosophical Issues* 21(1): 268–288.
- Moore, G. E. (1959). *Philosophical Papers*. London: George Allen and Unwin. This volume includes the key epistemological papers by Moore: "Proof of an External World," pp. 127–150; "Four Forms of Skepticism," pp. 196–225; and "Certainty," pp. 226–251. Key excerpts are reprinted in Kim, Sosa, Fantl, and McGrath (eds.), *Epistemology: An Anthology*. Oxford: Blackwell.
- Nozick, Robert (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Reid, Thomas [1785] (1983). *Essays on the Intellectual Powers of Man*. In R. E. Beanblossom and K. Lehrer, *Reid: Inquiry and Essays*. Indianapolis: Hackett Publishing.
- Stroud, Barry (1984). *The Significance of Philosophical Skepticism*. Oxford: Oxford University Press.

NOTES

1. It's important to recognize that skepticism about knowledge concerning some domain is not the same as the denial of the existence of that domain. To be a skeptic about our knowledge of other minds is to think that we cannot know that there are minds other than our own. This is consistent with not denying their existence.
2. This is the point of the cogito, his famous "I think, therefore I am." Even if you are in a skeptical scenario, and so you are being deceived about many matters, you are still thinking and indeed you *are* (i.e., you *exist*). Thus, Descartes thinks he finds his Archimedean point, one from which he can "move the earth" and build his foundation for firm knowledge.
3. To admit that you don't know you aren't a BIV while insisting that you do know you have hands, feet, and a body is to affirm what Keith DeRose calls the "abominable conjunction" (1995, 27–29): "Yes, I know that I have hands, feet, and body; but no, I don't know that I'm not a brain in a vat." This isn't merely an odd thing to say; it seems absurd.
4. A priori knowledge, going back at least to Kant, is traditionally defined negatively as knowledge that is "independent" of experience.

You may naturally ask, "What is meant by 'independent' and what is meant by 'experience'?" "Independent" should be understood to concern epistemic independence, not causal independence. Perhaps you couldn't know this proposition (call it T)—if one thing *x* is taller than another *y*, then *y* isn't taller than *x*—without having had some experiences to give you concept of *taller than*. So, your knowledge that T is true isn't causally independent of experience. But the source of your knowledge—the epistemic basis—is not experience. Given that you have the experience necessary to think about whether T is true, it isn't experience that helps make this knowledge.

How broadly we understand "experience" is to some extent a matter of choice. Do you know a priori that you're thinking right now? If we understand "experience" to include introspection of thoughts, then *yes*. If we understand it to refer exclusively to perceptual experience, then we might conclude that it is a priori (assuming that we don't literally perceive our thoughts). Here we understand "experience" to include not only perceptual experience but also introspection.

For a more thorough introduction to the a priori/a posteriori distinction, see Bruce Russell (2013).

5. The justification is *prima facie* only. It can be defeated if you acquire special evidence that your senses, in a particular situation, or perhaps even in general, are unreliable. See Chapter 1 (Section 1.8).
6. There is debate among interpreters of Reid just how his first principles are to be understood. William Alston (1985) and Keith Lehrer (1989) maintain that they are general truths, and we, too, adopt this interpretation. However, Van Cleve (1999) argues that they are rather principles of evidence. To give an example, Van Cleve understands the case concerning memory as follows: If you distinctly remember that *p*, then it is a first principle for you that *p*. As Van Cleve puts it, his interpretation posits indefinitely many first principles of memory, one for each of the testimonies of memory. He then understands "It is a first principle for you that *p*" as meaning that you are immediately justified in believing that *p* (in the

sense of “immediately justified” discussed in Chapter 2 [Section 2.1]). If Van Cleve’s interpretation is correct, Reidian first principles do not provide an *a priori* way to know that skeptical hypotheses are false. I thank Marina Folescu and Patrick Rysiew for advice on the interpretation of Reid.

7. Saul Kripke (1980) argued that contingent *a priori* knowledge is possible, indeed actual. He claimed that we have *a priori* knowledge of such matters as “the standard meter bar in Paris is one meter long.” The thought is that the definite description “the length of the standard meter bar in Paris” fixes the reference of “one meter.” But is this a case of *a priori* knowledge of a contingent truth? Or is the most we can say this: Merely by knowing what you mean, you can know that the sentence above expresses something true? But knowledge of what you mean seems to depend on introspection of your intentions and so not to be *a priori*. See note 4.
8. Can the skeptic give a similar response for reason? Can she say that we have a good basis for treating reason differently than perception?
9. See Chapter 11 (Section 11.4.2) for a ray of hope for the *a priori* approach. The ray of hope is not of the sort Reid had in mind, however, but rather one concerning probability. Essentially, the thought is that building so much about one’s actual experiences and memories into the skeptical hypothesis reduces its prior probability severely. This prior probability, a probability not grounded in empirical information, is an *a priori* probability.
10. The discussion to follow overlaps with Chapter 1 (Section 1.9).
11. For more on how exactly to spell out what inference to the best explanation amounts to, see Peter Lipton (2004).
12. Their ranks include Bertrand Russell, C. D. Broad, and A. J. Ayer, as well as the contemporary philosophers Laurence Bonjour and Jonathan Vogel. The eighteenth-century philosopher George Berkeley famously appealed to something like inference to the best explanation in favor of an idealist version of RWH.
13. In fact, for a naïve flat-earthier, one simply going on how things look from Earth or from the fact that we don’t “fall off” Earth, satellite photography is a way to learn the earth is round.
14. Relatedly, we should distinguish questions about whether one could resolve one’s doubts by going through a piece of reasoning from the epistemic question of whether the reasoning could be a way of coming to expand one’s knowledge if one lacked the relevant doubts. In some cases, one cannot resolve one’s doubts with a piece of reasoning even though that piece of reasoning if you lacked those doubts could give you knowledge. Suppose I doubt whether there were 250 years of relative peace during the Ming Dynasty. However, my reasons for doubt are not the usual ones. I doubt this, suppose, because I, quite irrationally, doubt current historical methods are at all reliable. I cannot then resolve this doubt by reasoning in accord with an argument appealing to premises justified by those very historical methods. However, this isn’t to say that someone reading a book on the Ming Dynasty couldn’t expand his or her knowledge by reasoning from those very premises to the conclusion that there were 250 years of relative peace during the Ming Dynasty. An argument’s being incapable of resolving one’s doubts about the conclusion seems to be an intrapersonal counterpart of an argument’s begging the question against an opponent. See Pryor (2004), Section 5, for a clear and detailed account of these matters.

PART II

Justification and Knowledge: Special Topics