

Chapter 2: Dream skepticism and externalism

Throughout ordinary life, I routinely take myself to know a great deal of mundane facts about my immediate external environment. At the time of this writing, I currently take myself to know that there are a laptop and a coffee cup before me, that I have two hands that are busy typing, and that it is raining outside today (in Vancouver, British Columbia). Indeed, the only occasion in which I would ever doubt that I can know such things is when I am actively engaging in philosophical reflection on skepticism. Otherwise, I would complacently go throughout my day, content that I know such facts.

Let's quickly remind ourselves of what we've learned about skepticism so far. In the last chapter, we saw a skeptical challenge which came in the form of Agrippa's Trilemma. This skeptical line of reasoning implored us to justify all our beliefs by finding a finite, non-circular argument to infer each of our beliefs. If such a task were indeed necessary for justification, then none of our beliefs could be justified or known. But we found that this task *isn't* necessary for justification or knowledge. That is because there can be such a thing as *foundational knowledge*. When a belief is an instance of foundational knowledge, it is genuinely justified, but its justification is not based on inference from other beliefs. Rather, its justification comes from some non-inferential basis (such as perception). In sum, a belief's foundational status does not thereby make it ineligible for knowledge.

This result provides a significant step of progress in our struggle with skepticism. But we are not out of the woods yet, for there is still another family of skeptical arguments that we have not yet considered. And unlike the last skeptical challenge, this one will not be so easily answered.

To see what we are up against, take any one of those mundane facts that I ordinarily consider myself to know. Take, for example, my claim to know that I have hands. (Indeed, I consider this belief to be a candidate for foundational knowledge, provided that it can be justified by perception without the aid of inference.) This would seem, in most circumstances, to be a perfectly safe claim to knowledge. How could I possibly fail to know that I have hands when I see them right here before me?

This new skeptic will aim to cast doubt on my knowledge by raising the possibility of hallucination and deception. They will invite me to attend to hypothetical scenarios in which I would be deceived into believing falsely. Consider, for example, the possibility that I am dreaming right now. Then I would be lying asleep in my bed, and the desk, the laptop, and the hands that I appear to see are really just figments of my dream. Or maybe I am the victim of an even greater scheme of deception. The skeptic is often fond of relating this example from science fiction. Imagine that my brain has been removed

from my body, placed into a vat of life-sustaining fluid, and hooked up to a supercomputer. The supercomputer is then sending electrical impulses into my sensory pathways in order to create a simulation of all of the experiences that I would normally have if I were still in my body. Within this scenario, it would appear to me (as the brain in the vat) as if I were looking at my hands before me. From my perspective, things would seem just as if I were experiencing the real world.

Upon reflection on these possibilities, we find that it is very difficult to rule them out. It is difficult to know that we are not actually in them. After all, for every experience that I have right now, my brain-in-a-vat counterpart will have a similar experience. If it appears to me that I have hands, then it will also appear to my BIV-counterpart as if it has hands. Between the two of us, everything will appear to be the same ‘from the inside.’ So it seems that I cannot rule out that *I* am a BIV right now, just by consulting my own first-personal experience. As far as my perceptions appear to me, I could either be experiencing the real world or the simulated world of my BIV-counterpart.

Once the skeptic has raised their hypothetical scenarios, they will claim that our knowledge is undermined. For if I really *knew* that I had hands, then I should be able to distinguish my own case from the case of deception. I should be able to rule out the possibility that I am a handless BIV. But since I (allegedly) cannot know that I’m not a BIV, it follows that I cannot really know that I have hands. Here is their reasoning put a bit more explicitly:

- 1 If I know that I have hands, then I know that I am not a handless BIV.
- 2 I do not know that I am not a handless BIV.
- 3 Therefore I do not know that I have hands.

Now, of course, there isn’t anything about this argument that depends on the belief being about hands; any belief about my external environment could have been used in its place. So the idea behind this argument gives rise to a much more general and radical conclusion: that I cannot know *anything* about the external world, no matter how plainly it lies before my face.

We have now revealed our primary skeptical foe. Unfortunately for us, this new argument does not rest on any elementary logical confusions or mistakes; hence it cannot be dismissed so quickly. It is a completely different animal from Agrippa’s trilemma. It thus presents a profound problem for non-skeptics, like myself.

What can we say for ourselves, in response to this argument? Since the argument is valid, we really only have three options: we either challenge the first premise, or we challenge the second premise, or we are forced to concede the conclusion. My plan for this chapter is to *grant* the first premise of the argument without much comment (save for a footnote¹). After all, the first premise does seem right to me. It seems right to say that *genuine* knowledge

¹Those who deny the first premise of this argument are called *fallibilists*. In their view, knowledge does not require infallibility; one can know something even if one hasn’t ruled out

requires us to rule out the possibility of deception. If, for all I know, my belief that I have hands could be the result of deception, then I do not really know that I have hands.

Having granted the first premise, our attention now turns entirely to the second. The key to overcoming this argument becomes a matter of explaining how we can *know* that we are not actually victims of radical skeptical scenarios. But how can we know such a thing? How *can* I know that I am not actually a brain in a vat right now? Eventually, I will offer my own answer to this question. But I will hold off on this answer until after the next section. Before I attempt to answer it directly, I first need to say a few words about other, highly influential attempts to solve this problem from the past. This historical digression will prove to be both necessary and instructive. It will illustrate the approaches to the problem that are both highly tempting, and yet ultimately doomed to failure. So they reveal to us the modes of thought that we mustn't be tempted by. They will also illustrate for us how the problem of "dream skepticism" intersects with the issue of foundationalism. How we confront the skeptical scenarios will vitally depend on what we take the foundational beliefs to be. So once again, we must confront the question, *what are the foundational beliefs?*

1 Descartes and the method of doubt

This far in our inquiry we have only established a few formal results. We have found that foundational knowledge is possible, but we've hardly said what it could be. We have also been made aware of the threat of radical deception. Now, when confronted with this predicament, it is natural to wonder what kinds of knowledge *wouldn't* be vulnerable to this kind of skeptical argument. Perhaps there is a (small) number of our beliefs that are safe and secure—or, for lack of a better word, *absolutely certain*. If there are any such beliefs, then perhaps it would be a good idea to find them, and then take them to be our foundational knowledge. And then once these certain foundations have been laid down, then perhaps we can work our way outwards from there.

The most famous attempt to carry out this programme comes from the 17th-century philosopher René Descartes in his *Meditations on First Philosophy*. In this work, Descartes makes it his avowed goal to seek out whatever can be known with certainty so that it can serve as the foundation for his overall worldview. His hope is that, by doing so, he can defeat the same skeptical doubts that we are confronted with here.

every possibility of being wrong. Although this view has been historically popular, I will not dwell on it for long. Suffice to say that it has some highly counterintuitive consequences. For instance, according to this view, we cannot always extend our knowledge by deductive inference. There will be cases where we do not know the conclusion of an argument, even though we know that the premises are true and the argument is valid. Case in point: I know that *if I have hands, then I am not a BIV*. The fallibilist accepts that I *know* that I have hands. And yet the fallibilist will not concede that I know the consequent, that I am not a BIV. Somehow my valid deductive argument failed to give me knowledge of the conclusion.

Within the context of Descartes's project, the skeptical hypotheses (e.g. dreaming, being a brain in a vat) do not serve merely as *problems* to be overcome. Rather, for Descartes, the skeptical hypotheses also serve the heuristic function of helping him to identify the beliefs that are fit to be foundations. This is essentially his famous "method of doubt." The idea here is that Descartes wants to rid his belief system of every falsehood and accept only certainties. To accomplish this, he first seeks to doubt everything that he formerly believed insofar as this is possible. And to ensure that he does this properly, he builds up the case for skepticism as best he can in order to weed out any belief that can possibly be doubted. At the end of his foray into skepticism, Descartes then vows to suspend judgment on everything that he can possibly doubt. If things go according to plan, then by the end of this phase, Descartes will have found the beliefs that cannot possibly be doubted. By pushing the limits of skepticism, he hopes to run up against the beliefs that are immune to skeptical attacks. And then finally, in the last phase of his plan, he will take these certain beliefs and make them his foundations. Having discovered his foundations, he will rebuild a structure of beliefs on top of it. Ultimately he hopes to recover his beliefs about the external world, of science, and of religion. But after destroying and rebuilding his belief system, he hopes that they will be grounded in certainty.

This is how the Cartesian method looks in theory, but how does it look in practice? At first, when Descartes seeks to doubt everything he possibly can, he finds that there are very few beliefs that pass the test of certainty. The most famous one is his belief that *he exists*. Descartes finds that even if he were dreaming, hallucinating or being radically deceived by an evil demon, he still could never doubt the fact that he existed. For as long as he is thinking, doubting, dreaming, and undergoing sensory experiences (even hallucinatory ones), then it follows that he must exist. In addition to this, Descartes also finds that beliefs about his own subjective mental states are immune to skeptical doubts. Let's switch back to my own case for the examples. When it comes to myself, I currently believe that *it appears to me that I have hands* and that *I think that I have hands*. Descartes would say that both of these beliefs pass his test for certainty. For even if I don't actually have hands—even if I am a handless brain in a vat—it would still be true that *it appears to me* that I have hands. As long as I'm attending to how things appear to me, I cannot be wrong about having this appearance. Since this appearance is part of my own subjective experience, I cannot be mistaken about it. The same goes for my belief that *I think that I have hands*. For even if I am wrong to think that I have hands, I cannot be wrong *about* the fact that *I am thinking this*. When it comes to my beliefs *about* my sensations and my thoughts, I cannot be wrong about them whenever I am attending to them, according to Descartes.

This idea of Descartes's has had a profound effect on subsequent thinking in epistemology. Many philosophers have accepted that our beliefs about our own private mental lives enjoy a special degree of certainty that isn't possible for any other sort of belief. For Descartes, this also means that these beliefs are uniquely suited to be foundational knowledge. In his view, we should begin

to construct our knowledge by starting with what we know about the states of our own minds. In order to know anything at all, we must begin by examining our own thoughts, our ideas, and our sensations. Once we have gathered our knowledge about our own inner states, then we may proceed to infer the rest of what we know from there. All of this amounts to a substantial thesis about the scope of our foundational knowledge. In short, it amounts to the claim that our foundational knowledge must be “internally accessible” to us. But what does this mean exactly? Well, for one thing, it means that the truth of our foundational beliefs must be discoverable by introspection or other a priori means. We must be able to attain this knowledge just by consulting our reason and the elements of our first-personal point of view. This also means, for the Cartesian, that our foundational knowledge must concern states that we cannot be deceived about. For them, this knowledge must be available even if we were brains-in-vats. Or, to put this idea differently, our foundational knowledge must be shared between ourselves and our BIV-counterparts.

When it comes to the historical Descartes, this only amounts to half of the story. The full story will also tell the tale of how he sought to prove, from these meagre foundations, the reliability of his senses and the existence of the external world. One would even expect this to be the most interesting part of Descartes’s work. After all, we are looking for a defence of our knowledge of the external world and Descartes claims to have provided one. But, alas, when we turn to this part of Descartes’s *Meditations*, the proof that we find is notoriously disappointing. Basically, Descartes inspects his idea of an omnibenevolent and omnipotent God, and upon reflection, concludes that he is too lowly of a creature to have come up with this idea on his own. He figures that the only possible explanation for him having this idea is that God actually exists and has imparted this idea to him. Once he is satisfied that God exists, Descartes then observes that an omnibenevolent and omnipotent God would never allow him to be radically deceived. And so Descartes thinks that he can trust his senses again, and he can regain his knowledge about external reality.

We do not need to dwell on Descartes’s theological solution to skepticism for long; suffice to say that his argument for God’s existence isn’t very compelling.² But Descartes’s real legacy to epistemology does not reside in his solution to skepticism. Rather, his lasting influence on the subject comes from how he approached the problem. Descartes took radical skepticism very seriously and he thought it was necessary to restructure his belief system to fend off this threat. To this end, he begins by retreating to the sanctuary of beliefs that are immune to skeptical refutation. He thought that these beliefs only include those that could be known by a subject who is victim to a radical skeptical scenario. So if we followed Descartes, we would be left with very few foundational beliefs.

²Here is the most obvious alternative explanation for how I came to the idea of God: I made it up. Or perhaps the idea was cumulatively made up by several people through a process of cultural evolution, and eventually passed on to me. Descartes considers these alternative proposals and he dismisses them as impossible. But his reasons for dismissing them rest on ideas from medieval philosophy and theology that are apt to strike the modern reader as baffling.

At most, our foundational knowledge would only include what we can gather from *a priori* reasoning and introspection on our private mental states. Then once we have recoiled into these ‘safe’ foundational beliefs, we face an immense task if we are to ever recover our knowledge of the external world. If we are to ever regain our knowledge of physics, chemistry, biology, politics, history, world affairs, etc., we must reason our way into these things by starting from our scant foundations.

It must be said that if *this* is how we acquire knowledge of the external world, then this knowledge is much more demanding than our common sense would have us believe. Take any one of the mundane beliefs that you think can be easily known. I think that I know that I have hands, and this seems to be something that I can know very easily. It seems that I can know this just by looking at them. But if Descartes is right, then this knowledge requires much more additional intellectual work on my part. I need to consult my private experiences and my *a priori* knowledge, and from these raw materials I must build my case for believing that I have hands. But how would this argument work? What can bridge the gap between my *a priori* and *perspectival* knowledge and the fact that I have hands? Must I somehow show *a priori* that my sense perception is reliable? Can I rule out the radical skeptical hypotheses by *a priori* means alone?

Besides Descartes’s attempt to thread this needle by appealing to theology, there has also been numerous secular attempts that have come from within a broadly Cartesian tradition. For example, both Bertrand Russell and Rudolf Carnap have developed elaborate systems where they define propositions about external reality in terms of *how things privately seem*.³ They also supply rules for inferring these propositions on the basis of sense experience using the methods of logic and probability theory. But despite the efforts of these two thinkers (and many other thinkers besides), there has always a suspicion that their procedures will not amount to genuine *knowledge*. For if, at best, all that I can say is that *I probably have hands* based on *how things (privately) seem to me*, then it doesn’t really seem as though I *know* that I have hands.

The need to explain our knowledge from within Cartesian constraints has invariably led to either disappointment or desperation. Some philosophers have become so desperate that they’ve succumbed to idealism. For them, it is *obvious* that I can know about the world of my experience: I can know that there is a coffee cup before me. But if, as Descartes suggests, the only kind of things that I can directly know about are my private, *subjective* experiences, then it becomes hard to square my knowledge of this coffee cup with the claim that it is an external, material object. So perhaps the reason why I can attain this knowledge so easily is because it *isn’t* really a mind-independent material object. Perhaps this coffee cup is really a product of my consciousness. Maybe this coffee cup is just an idea in my mind. This suggestion may be wild, but at least it has the virtue of explaining how I can have this knowledge—provided

³See Bertrand Russell’s *The Problems of Philosophy* for a beginner level exposition and his *Philosophy of Logical Atomism* for something more advanced. For Rudolf Carnap, the project is outlined in his *Aufbau*.

that we're understanding knowledge as Descartes did. I can know about this coffee cup *because* it is an idea in my mind.

I do not present this idea because I believe in it myself. Personally, I have always found idealism to be unattractive and ill-motivated. I present it only to illustrate the extreme measures that philosophers will take once they've succumbed to the Cartesian paradigm. In my view, this is a sign that something is wrong with the paradigm.⁴ It is time to now take a different approach to fending off skepticism.

2 The Moorean shift

Let's return to the main argument for skepticism. Thus far, I have willingly conceded the first premise, that *if I know that I have hands, then I know that I am not a handless brain in a vat*. Ultimately my intention is to deny the conclusion, that *I do not know that I have hands*. So as a matter of logical consistency, I am forced to deny the second premise. I must claim that I *do* know that I am not a BIV. But once again, merely claiming this is not enough to dispel the skeptical challenge. I must also explain why this is a rational thing to do.

To fully appreciate our situation, it is helpful to consider a famous "proof of the external world" from the English philosopher G.E. Moore.⁵ In one of his lectures, Moore claimed that he could *prove*, and thereby *know*, that external objects exist. In fact he thought that he could prove this in a very simple and flatfooted way. Moore lifted up one of his hands and said "here is a hand", he lifted up the other and said "here is another hand", and then concluded "therefore, external objects exist." And that's it. That's all it takes to refute skepticism, according to Moore.

Many first-time readers of Moore's "proof of the external world" are apt to find his proof very underwhelming. It strikes them as if Moore isn't taking the skeptic's challenge seriously, or that he fails to see what the problem is. Many readers accuse him of "begging the question."

But it's important to see that the real interest in Moore's proof doesn't reside solely within the flatfooted proof itself. The real interesting part of Moore is the higher-order discussion of *why* Moore's proof may or may not be an adequate response to skepticism. It forces us to take a step back and ask more fundamental questions about how we should approach the topic. We need to ask such questions as: who has the burden of proof here? Is it the skeptic or the advocate of common sense? Where should we start our reflections on skepticism? Should we start with some basic assumptions about what we know, or should we proceed with the Cartesian method of doubt? And what does it even mean to beg the question?

⁴Ever since Descartes—and especially in the 20th century—there have been numerous philosophers who have advocated for alternatives to Descartes's method of doubt.

⁵I assure you that the hilarity of an author named G.S. Moore advocating for G.E. Moore isn't lost on me.

Let's start with the last question and look at the allegation that Moore begs the question. To address this allegation, we need to remember a distinction that we had in §1.2. There we distinguished between two meanings of allegation that someone 'begs the question.' On one reading, when we accuse an argument to be 'question begging' we mean that the conclusion of the argument is buried within the premises. On another reading, when we accuse an assertion of 'begging the question', we mean that the assertion isn't appropriate within its conversational context because the interlocutor regards it as up for dispute.

Moore emphatically denies that his proof 'begs the question' in the first sense. His conclusion, that there are external objects, is in no way equivalent to his premises. His premises are "here is one hand" and "here is another"; both of these could very well be false while the conclusion remains true (say, if Moore lost his hands in an accident). Moore also stresses that his conclusion is a philosophical thesis, whereas his premises are items of common sense.

When we turn to the other sense of 'question begging', it may very well be true that Moore's premises 'beg the question' in that sense. It would be *illegitimate* for Moore to *assert* that he has hands in the context of a dispute with a committed skeptical philosopher, without further argument. But as we have already observed in §1.2, it doesn't thereby follow that Moore cannot *know* that his premises are true. Moore can *know* that he has hands even if he cannot rationally persuade a committed skeptic. Having knowledge does not require one to be capable of convincing every possible interlocutor of what one knows. You can know something even if you cannot convince others of its truth. So even if Moore 'begs the question' in this second sense, this fact alone does not prevent Moore from knowing both his premises and his conclusion, and so his proof is still relevant to the issue of skepticism.

Let us now go back to the skeptic's original argument. Remember, the skeptic has given us the following bit of reasoning; let's call it "the skeptic's argument":

- 1 If I know that I have hands, then I know that I am not a handless BIV.
- 2 I do not know that I am not a handless BIV.
- 3 Therefore I do not know that I have hands.

Moore, on the other hand, thinks that he has *proven* that the external world exists. So he must think that he *knows* that he is not a BIV, and he must think that he knows this on the basis of the premise that he has hands.⁶ So Moore offers us the following counterargument; let's call it the "the Moorean shift"⁷:

⁶Moore's original article does not address this specific skeptical hypothesis, but there's no harm in interpolating.

⁷It is now standard in philosophy to use the phrase "Moore shift" to counter your opponent's argument by reasoning from the denial of their conclusion to the denial of one of their premises. Sometimes this is a reasonable response to your opponent's argument, and sometimes it isn't. Here we will discuss whether it's a reasonable response to the skeptic's argument.

1 If I know that I have hands, then I know that I am not a handless BIV.

2* *I know that I have hands.*

3* Therefore *I know that I'm not a handless BIV.*

I hope that you see the differences between the two arguments. Both of them grant the skeptic's first premise, like I have done. But they differ in that their second premise and conclusion are contraposed. The skeptic reasons from the premise *I do not know that I am a handless BIV* to the conclusion *I do not know that I have hands*. Moore would instead prefer to reason in the other direction. He begins with *I know that I have hands* and then concludes that *I know I am not a BIV*.

As you can see, the difference between Moore and the skeptic is that they start from different places. Moore begins his thinking about skepticism with the premise that he knows that he has hands, whereas the skeptic begins by asserting that we do not know that we are not handless BIVs.⁸ The conclusion that we will end up with will depend on where we begin. This raises the question: which of the two alternative premises should we begin with?

Moore has his own reasons for preferring his starting point over the skeptic's. According to him, it's a matter of *sheer common sense* that he knows that he has hands. This, for him, is one of the most plainly obvious facts of all. The skeptic's premise, on the other hand, is a strange claim about our knowledge concerning a science fiction example. It concerns a remote possibility that we do not ordinarily concern ourselves with. So for Moore, we should be less sure about knowledge assessments when it comes to this strange example. So given the choice between the two, which should we prefer: our common sense claim to knowledge or a knowledge assessment about a strange and remote possibility? For Moore the choice is clear: we should side with common sense.

Moore's basic idea is that we should always start our reasoning with the premises that are most plausible. Any argument that reaches an implausible conclusion from implausible premises should be rejected; we should reject the premises on the basis of rejecting the conclusion. In more recent years, this broad methodology has been embraced by a philosophy called "phenomenal conservatism."⁹ Phenomenal conservatives claim that our starting point—our foundational beliefs—should consist of all of the propositions that *appear* or *seem* to us to be true, as long as there is no reason to doubt them. So if something seems to be true to me, then I have a defeasible justification for believing it. If, say, it seems to me that I know that I have hands, I am thus justified to believe this unless I also have appearances to the contrary. If the skeptic gives me their argument against this knowledge, then I am justified in rejecting their argument, provided that the falsity of their conclusion is more apparent to me than the truth of their premises.

⁸Perhaps it's not fair to say that they begin with that assumption. We will discuss their arguments for that premise shortly.

⁹E.g. Michael Huemer in *Skepticism and the Veil of Perception*.

Phenomenal conservatives may agree with Moore in their broad methodology, but they differ in some of the details. For one thing, Moore places a premium on common sense, whereas the phenomenal conservatives emphasize how things appear or seem. For them, the justified foundational beliefs are ultimately justified (albeit, defeasibly) on the basis of *how things seem to the subject*, from the subject's point of view. So, if from my point of view, it seems very strongly to me that, say, *morality consists of maximizing happiness and minimizing suffering*, and I have no other seemings that tell otherwise, then this belief is justified *for me*. And if it *seems* very strongly to you that *morality consists of following a certain preset list of rules* (e.g. don't lie, don't steal), and you have no other conflicting seemings, then that belief is justified *for you*. Phenomenal conservatives thus take the perspective of the subject to be crucial for the justification of beliefs. For them, it is *how things seem* to the subject, internally to their perspective, that ultimately determines the justification of their beliefs. This aspect of their view is nowadays known as *epistemic internalism*. The idea is that the justification of a subject's beliefs ultimately depends on how things stand from within their point of view—it depends on how things appear, from their perspective. We can also express this idea by appealing to brains in vats. Since *how things seem to you* will be the same whether you are normally situated or a BIV, it follows that you and your BIV-counterpart share all of the same justification for your shared beliefs, according to phenomenal conservatism. This is also what internalism says. The internalist thesis holds that the states that are eligible to confer justification *just are* whatever you share with your BIV-counterpart.

There is thus a sense in which the phenomenal conservatives also share sympathies with Descartes. Remember, Descartes held that our foundational knowledge must always be *about* the things that are *internal* to our perspective. The modern phenomenal conservatives do not exactly agree with this claim. They do not hold that our foundational beliefs must be *about* internal things. But they still emphasize the importance of internal things for justification. For them, the justification for our beliefs—whether they are *about* internal or external matters—must always depend on our internal states.

The reason I have brought up phenomenal conservatism is because it represents a contemporary internalist implementation of Moore's defence against skepticism. Also, because of its internalist character, it is highly attractive to those philosophers who maintain a Cartesian attachment to internal states as the locus of epistemological standing. But I will not be endorsing it myself. As I will explain in the rest of this chapter, I do not think that it is adequate to meet the skeptical challenge. I regard its attachment to internalism as a liability, not a virtue.

There are, however, elements of the Moorean response to skepticism that I do endorse. In particular, I think that the correct strategy for dealing with skepticism is to defend the Moorean shift argument against the skeptical argument. But to do this, I will not be appealing to either common sense or my appearances as the ultimate source of justification or knowledge. So as things stand right now, the question still remains before us: why think that I know

that I have hands? And how can I rationally use this premise to infer that I know that I'm not a BIV (using the Moorean shift argument)? In the next two sections I will take up each of these questions in turn.

3 A statement of my version of foundationalism

It is now finally time to lay out the version of foundationalism that I prefer. Remember, the task for a foundationalist is to explain what kinds of belief are suitable for foundational knowledge. My task, then, is to explain the conditions for when a belief can be known without inference.

My proposal, put very briefly, is simply this. A belief counts as foundational knowledge when it is the product of a channel that relates cognition to reality.¹⁰ Here is a picture to keep in mind to fix ideas. When a subject gathers their foundational knowledge, their mind will bear many links to the reality that surrounds them. They perceive the facts of their external environment, they gather further facts from the testimony of others, they store these facts to be remembered later, and (perhaps) they gain further insights through rational intuition.¹¹ That these links obtain, when they do, are objective facts about the relations the subject bears to their environment. When these links between the subject and the facts are secured, and the subject forms beliefs on their basis, then those beliefs count as foundational knowledge.

For now we can focus on the paradigmatic case of knowledge by contact with reality: perceptual knowledge. To be even more specific, we will focus on the kind of perception that is a species of *propositional attitude*. This is the kind of perception that we are talking about when we say that a subject perceives *that* something is the case. For example, I currently perceive *that* I have hands and I perceive *that* it is raining outside. Each of these is a *perceiving-that*; the complement of my perception is a proposition. Now, when it comes to perception in its propositional attitude form, it is widely recognized that perception is a so-called *factive state*. This means that the object that one perceives must always be a fact. In order for a subject to perceive *that p*, it must also be the case that *p*. For example, if I perceive that it is raining outside, then it must really be true that it is raining outside. If it weren't raining outside, then I couldn't possibly perceive *that it is raining outside*. If my perceptual faculties are faulty, then it could *seem* or *appear* to me *as if it is raining outside*. But that is not the same as *perceiving* that it is raining outside. In order for me to perceive

¹⁰To be clear, I will only claim this to be a *sufficient condition* for a belief to be foundational knowledge. I make no claims about necessary conditions. It is not my business here to give an *analysis* of knowledge.

Those who are familiar with the history of analytic epistemology might be reminded of the *causal theory of knowledge*, which is nowadays widely discredited. But I am not making any claim to have given an analysis of knowledge; I am only citing a sufficient condition for foundational knowledge.

¹¹The system I develop here is consistent with the existence of a priori rational intuition that places the subject in direct contact with the abstract reality that is represented by their intuition. But I'm not going to develop this into a positive thesis. I don't even know if I believe it myself.

that it is raining, it must really be raining. In short, perceiving-*that* is always a relation to a fact.

The central claim of this chapter is that when one perceives that p and thereby believes that p , then one possesses foundational knowledge that p . One can know that p on this basis, even when one cannot corroborate the truth of p by inference from other beliefs. Perception is one of my paradigmatic cases of foundational knowledge. When one perceives *that* something is true of their external environment, then one bears a direct cognitive relation to the facts of their external environment. This places one in a fairly robust position to know these facts. (Although I must admit, from a historical vantage point, this last claim will appear contentious. I thus devote the entirety of chapter 3 to defend it.) In fact, I think that perceptual knowledge is probably the most straightforward and secure kind of knowledge that is available to us. When we succeed in tuning into reality through perception, the knowledge we receive will have the highest epistemic credentials that are possible.

To put this view into perspective, note that there is actually an element of Descartes's foundationalism that I agree with. Within Descartes's method of doubt, he held that the only beliefs that are suitable for foundational knowledge are the ones that are immune to skeptical refutation. Foundational beliefs must be *infallible*, in the sense that, if you know them, then you can rule out every possibility in which they're false. In my view, the beliefs that are foundational knowledge *are* immune to the skeptical arguments of this kind. If I really perceive that I have hands (as opposed to merely *appearing to see that I have hands*), then I can thereby rule out the entire range of skeptical possibilities. If I really bear the perceptual connection to the states of my environment, then I can rule out the possibility that I'm sleeping, that I'm a BIV, and so on. (I haven't yet given a full argument for this, but I will begin in the next subsection.) The point is, I take our foundational knowledge to enjoy this aspect of Cartesian infallibility.

But with that said, there's another element of Cartesianism that I reject, and that is its commitment to *internalism*. Remember, for Descartes, the objects of foundational knowledge were supposed to be accessible by introspection from the subject's point of view. It's as if, for a subject to know that p as a matter of foundational knowledge, then it must *appear to them that p only if p obtains*. And I resoundingly reject this internalist constraint on foundational knowledge. As I said, when one genuinely perceives something, their perceptual state is borne out by the *relations* that they bear to their external environment. Perception is thus not an entirely internal state. It is an objective relation that depends on factors beyond the subject's perspective.

Another way to put this point is to contrast my state with my BIV counterpart. As a person situated in a normal environment, when I direct my visual attention towards my hands, it will *appear* to me that I have hands and I will *perceive* that I have hands. I can perceive that I have hands because they are really there to be perceived (and they are unobstructed, and my visual processes are working normally, etc.). But when my BIV counterpart directs its visual attention towards the portion of its simulated world that appears to be

its hands, it will not genuinely *perceive* that it has hands. It will *appear* to the BIV *as if* it has hands; but that is not the same as perceiving that it has hands. I therefore have a perceptual state that the BIV lacks. In my view, this means that I possess foundational knowledge that the BIV lacks. And this is despite the fact that, from within our own perspectives, things will appear to us in ways that are phenomenally indistinguishable. To summarize: in my view, our foundational knowledge depends on how we are related to the external facts. The upshot is that our foundational knowledge is not totally determined by how things appear to us. I will not necessarily share the same foundational knowledge with my BIV counterpart, even though things will seem the same to the two of us.

4 The skeptic responds

Given my version of foundationalism, my basic answer to the skeptical argument goes like this. I know that I am not a brain in a vat on the basis of the Moorean shift argument. That is, I know that I am not a BIV *because* I know that I have hands. The latter is an item of foundational knowledge; it is something I can know without inference. Indeed, I can know it without inference because it is based on the kind of state that affords foundational knowledge (see chapter 1, §2). To be specific, it is *based on perception*, which is a channel that relates my mind to the facts (see chapter 3). So as long as I perceive that I have hands, then I know that I'm not a BIV.

It is one thing to offer these assertions as a solution to skepticism, but that does nothing to assuage the concerns of those who harbour skeptical tendencies. The real uphill battle for my view is to defend it as an entirely adequate response to skepticism. In what follows, there will be a series of objections-and-replies between potential skeptics and myself. I hope that in showing how the view handles objections, I can display its theoretical virtue.

4.1 How do you know when you're perceiving?

According to my theory, I can know that I have hands provided that my belief is based on perception. So *if* I perceive that I have hands, *then* I can know that I have hands. Upon hearing this, a skeptic might be eager to ask me how I can ever know when the antecedent condition (the “if” part) obtains. Such a skeptic will want to know how I can ever *know* that I am perceiving that I have hands. They worry that if I cannot know this, then I cannot attain knowledge that I have hands.

If this skeptic is ambitious, they might even venture to argue that I *cannot* know that I am currently perceiving that I have hands. To argue this, they will invite me to consider how it currently *appears to me* as if I have hands. They will point out that I would have this appearance even if I were a brain in a vat. And so they will say that I cannot distinguish my current perceptual state from

mere appearance; hence I cannot know that I'm currently undergoing a genuine perception.

Regardless of whether we are dealing with the ambitious or moderate skeptic of this kind, it is important to see that their initial worry is mistaken. The skeptic worries that *if I cannot know that I'm currently perceiving, then I cannot know that I have hands*. But it is simply not a part of my view that one *must* have this higher-order knowledge *about* their perceptual state in order to know that they have hands via perception. My claim is simply that *if one perceives that they have hands, then they know that they have hands*. I never said that one must *also know that they are having this perception*. I say that they must *actually have* this perception. This harkens back to a point that I made ad nauseum throughout the first chapter: we mustn't confuse the *obtaining* of a fact with one's *knowing* of that fact. And in this case, I say that as long as it's a *fact* that I perceive that I have hands, then I know that I have hands.

I conclude that this skeptic poses no threat to my claim to know that I have hands. But what about the more ambitious version of skeptic? Must I also admit that I *cannot* know about my current perceptual state?

As far as I can tell, I don't need to admit this. I think that I can often have higher-order about my perceptual states, such as the knowledge that I perceive that I have hands. How can I know this? Well I don't know this simply based on how things appear. As the skeptic rightly observes, it will appear to my BIV counterpart as if it is perceiving its hands. And the BIV cannot know that it has this perceptual state, because it *doesn't* have this perceptual state. So my knowledge cannot be solely based on appearances. But presumably I can attain this knowledge based on other factors. When I know that I am perceiving that I have hands, I presume that this knowledge can be based on, firstly, my perceptual state, and secondly, my self-reflection upon this perceptual state. (These grounds for knowledge are not available to my BIV since it lacks the relevant perceptual state, and hence cannot reflect upon it.) I don't really have much detail to share as to how this happens; but for my purposes, that's fine. I also don't see any reason to suspect that I cannot attain this knowledge in this way. For all the skeptic has said, it's safe to say that I can have higher-order knowledge about my perceptual states.

4.2 How do you know that you know?

Besides worrying about how I can know about my perceptual state, there is another kind of skeptic that will worry about how I can know *that* I have this knowledge. This skeptic will see me claim that I can know that I have hands, despite the possibility of the envatted brain scenario. They will hear me say that I know this on the basis of perception. And I will also happily grant to them that whether or not I have this perception doesn't depend entirely on what's internal to my perspective; it also depends on how things are in my environment. It depends on whether the objects are there to be perceived. So I grant that my knowledge depends on factors that are external to my first-personal perspective. Upon hearing this, the skeptic will then claim that I can never *know that I know*

that I have hands. Because this knowledge depends on external states, I can never know when I have obtained it. This skeptic may even take things a step further and claim that this causes problems for my *first-order* knowledge that I have hands. They will say that if I cannot know that I know, then I cannot have the first-order knowledge that I have hands.

In their last step, the skeptic is appealing to an idea that epistemologists call “the KK principle.” The KK principle says that *if a subject knows that p, then they must be in a position to know that they know that p*. And of course, if this principle were true, then each time we have knowledge, we must be in a position to know an infinite number of things: we would have to know that we know, and know that we know that we know, and so on. One might worry whether this is really a plausible requirement for knowledge. Suffice to say that the KK principle is highly controversial. Many epistemologists take it to be false, and I agree with them. Knowledge about simple matters, like knowledge that one has hands, does not also require knowledge *about* one’s knowledge.

I am not going to present a rigorous refutation of the KK principle here.¹² I am content to state that my view rejects it. I will also say that, when the conversation is about skepticism, I am suspicious that the KK principle is often brought up for the wrong reasons. For in the context of a conversation with a skeptic, the basic Moorean line is to assert that *I know that I have hands*. But it’s also a rule of civil conversation to only ever make assertions when you know them to be true. If you assert something within a conversation and your interlocutor doubts what you say, they may ask you “how do you know that?”. And if you can’t explain to your interlocutor how you know something when they have expressed doubts, then it is *bad manners* to continue to assert it. So when the Moorean asserts to the skeptic that *I know that I have hands*, the skeptic has a conversational right to ask *how do you know that you know?*. And if the Moorean can’t provide an answer, then their assertion wasn’t appropriate in the context of that conversation. *But that doesn’t mean that their assertion is false*. Even if it’s against the norms of conversation to assert my knowledge to the skeptic, it doesn’t follow that I don’t know. I suspect that those skeptics who bring up KK to cast doubt on my knowledge are confusing *conversational politeness* with *truth*. If I am to appropriately assert that I know that *p* to a skeptic, then I must also explain *how I know that I know that p*. But it doesn’t follow that *my actual knowledge that p* requires *knowing that I know that p*.

With all of that said, I also emphatically deny that I cannot attain the relevant higher-order knowledge. The skeptic was worried that the *externalist* aspect of my theory disqualifies me from *knowing when I know*. But why would that follow? Why can’t I know that I know, if my knowledge depends on external factors? As far as I can tell, the only apparent reason why a skeptic would think this is because they doubt that I can attain knowledge of things that are external to me. But this is the very thing that I reject. (And any skeptic who would make this claim at this point would be begging the question.)

So then, how can I know when I have knowledge that I have hands? The story

¹²Timothy Williamson presents a refutation in *Knowledge and its Limits*.

that I would give will echo the last paragraph of the last subsection. First of all, I know that I have hands because I perceive them. I do not base this knowledge on mere appearances; instead, it is based on actual perception. Having attained this first-order knowledge, I can then self-reflect upon this knowledge and my perceptual state. Once again, I emphasize that I'm reflecting upon an actual state of knowledge. I'm not just reflecting on the appearance of having this knowledge. This is not something that a BIV could do, because the BIV would not have the relevant first-order knowledge to reflect upon. I presume that in many cases, I can thereby attain higher-order knowledge of my knowledge on the basis of this self-reflection. This story may be a bit hand-wavy, and I'm not sure how to state precisely when higher-order knowledge is available through self-reflection. But then again, that's not really my concern. The main point is: I don't *need* to have higher-order knowledge to know that I have hands, and yet, it's plausible that this knowledge is often available according to my picture.

4.3 Sensitivity to the truth

Each of the skeptical responses that we've seen so far have been largely independent of the skeptic's main argument. In one way or another, they have been attempts to undermine my Moorean strategy of baldly asserting that I know that I have hands. But this Moorean strategy could never work in the first place if the skeptic had decent arguments for their key premise, that we do not know we are not BIVs. So in this subsection and the next, we'll give the skeptic the opportunity to defend their premise. In a way, these are the most important arguments on behalf of skepticism. For if they can establish the premises of the main skeptical argument, the Moorean shift counter-strategy would be irrational.

So how could the skeptic argue that I do not know that I'm not a handless BIV? Historically there have been a couple of arguments for this conclusion. The one that we'll consider first has to do with the notion of *sensitivity to the truth*. Basically, the skeptic argues that *if* I were a BIV, then I wouldn't know any better. Since everything would still appear to me as if I were a normal non-envatted human, I would continue to believe as if I were going through my normal life. In particular, I would continue to believe that I'm not a BIV. So if my current belief that I'm not a BIV *were false*, then I would continue to believe it all the same, none the wiser. But that concession seems to undermine my current claim to know that I'm not a BIV. For it seems that if I were to really know it, then I couldn't be so easily tricked into believing it if it were false.

The skeptic is now making use of a principle that epistemologists call "sensitivity." This principle was popularized by the philosopher Robert Nozick, who argued that it is an essential ingredient to knowledge. The simplified version of the principle states that *if a subject S knows that p, then if p were false, S would not believe p*. Sometimes this is glossed by saying that knowledge requires *sensitivity to the truth*. If you genuinely know something, then you must be sensitive to when it is true and when it is false, which means that you wouldn't continue to hold onto the belief if it were false.

In the hands of the skeptic, this principle can lend itself to a straightforward justification of skepticism. That is because the simplified version of the sensitivity requirement is not met for the belief that *I'm not a BIV*. That is, if I were a BIV, then I would still believe that I wasn't a BIV. So according to the simplified sensitivity principle, my belief isn't sensitive to the truth, and hence it isn't known. But then the skeptic can run their main argument afterwards. When we put all of the pieces together, their argument would go like this.

- 4 If I know that I'm not a BIV, then if I were a BIV, then I wouldn't believe that I'm not a BIV. (This is an instance of the simplified sensitivity principle.)
- 5 But if I were a BIV, then I would still believe that I'm not a BIV.
- 6 So I do not know that I'm not a BIV.
- 7 If I know that I have hands, then I know that I'm not a BIV.
- 8 Therefore I do not know that I have hands.

And of course, this argument is completely general. If it were sound, then it would show that I do not know any of my beliefs about the external world.

(Historically, the main proponents of the sensitivity principle have often sought to deny premise (7), that *if I know that I have hands then I know that I'm not a BIV*. But as I remarked in the introduction of this chapter, I am not going to pursue denying this premise. I do not believe that's a promising way to go.)

So how can we respond to this new argument for skepticism based on the sensitivity principle? The first thing to say in response to this argument is that it is now widely accepted that the simplified version of the sensitivity principle is false. The principle's first advocate, Robert Nozick, has given us this counterexample. Suppose that there is a grandmother, "granny", who cares deeply about her grandchildren but doesn't see them often. Because she is caring, she wouldn't be able to bear the news if any of her grandchildren were to fall ill. If she were to hear that her grandchildren were unwell, this news would cause her to be distressed. One day, her grandson visits her in person and she sees that he is well. She forms a belief, on that basis, that he is well. But as it turns out, granny would still have believed this even if he were unwell. Because in that case, the kid's parents would never have let him visit, and they would have lied to granny, telling her that her grandson is fine, in order to prevent her from being distressed.

It's plausible to think that granny knows that her grandson is well. After all, she sees him with her own eyes. But her belief also fails to meet the condition of the simplified sensitivity principle. If the grandson were unwell, then she would still believe that he was well, because she would never see him and she'd be deceived by the parents. So granny has this knowledge, even though she is not 'sensitive to the truth' in the simplified sense. As a result of this example, it is plausible to think that there's something wrong with the simplified version of

the sensitivity principle. Moreover, it is easy to see where it has gone wrong. In the case where granny knows that her grandson is well, she has a sound *basis* for her knowledge—she sees him in person and sees that he is well. But in the case where her grandson is unwell, she no longer has this basis available to her. In the case where she believes falsely, she believes this because the parents prevent her from seeing her grandson and have lied to her. These observations show that the simplified version of the sensitivity principle was too simple. If there is any sense in which sensitivity is important for knowledge, it must also take into account *the basis* on which the subject forms their beliefs.

When epistemologists nowadays formulate the sensitivity principle, they recognize the need to incorporate the subject's basis for believing. Once the appropriate adjustments are made, we get a new principle, which we can call the "basis sensitivity principle." The basic idea here is that a subject can attain knowledge only if their basis for believing is sensitive to the truth. To put this more explicitly, we get:

Basis sensitivity principle. If a subject S knows that p on basis B , then if p were false, then S wouldn't believe that p on basis B .

Unlike the simplified sensitivity principle, this principle doesn't have any obvious counterexamples, like the granny case. Granny can know that her grandson is well by seeing him in person, and it's not true that *this particular* basis for believing is insensitive to the truth. Perhaps if the grandson were unwell and this fact was concealed from her, granny would still believe, but she would no longer believe *on the same basis*.

Within the contemporary epistemological scene, there is a great deal of controversy over whether the revised basis sensitivity principle is actually true. Some philosophers champion this principle as correct and other philosophers reject it. I am not going to bother to weigh in on this debate here. Suffice to say that I think it is highly plausible. But my response to skepticism will not rest on accepting the principle. Rather, I'm going to freely offer the principle to the skeptic, and then argue that my own version of foundationalism isn't threatened by their resulting argument.

Let's see how the skeptic's argument would go once the needed adjustments are made. Their entire argument would then become something like this.

- 4* If I know that I'm not a BIV on basis B , then if I were a BIV, then I wouldn't believe that I'm not a BIV on basis B . (This is the basis sensitivity principle.)
- 5* But if I were a BIV, then I would still believe that I'm not a BIV on basis B .
- 6 So I do not know that I'm not a BIV.
- 7 If I know that I have hands, then I know that I'm not a BIV.
- 8 Therefore I do not know that I have hands.

Once the argument is put this way, it becomes clear that it matters a great deal as to what is the actual basis for our knowledge. Perhaps given some candidate bases, the argument from (4* & 5*) to (6) is sound, and given others, it is unsound.

Here is my own view on the matter. I believe that I am not a BIV ultimately on the basis of my perceptual knowledge. I perceive that I have flesh-and-bone hands, that I have a body, that I'm sitting in a chair, that I am breathing air, and so on. All of this perception-based knowledge allows me to conclude that I am not a brain in a vat.¹³ Moreover, this belief of mine is sensitive to the truth *with respect to that basis*. In any situation in which I really perceive these things, it will also be true that I'm not a BIV. For if I were a BIV, then I could not have the same perceptual knowledge to base my belief on. For as we have already observed, my BIV-counterpart does not share my perceptual states. (They share the same *appearances* and *seemings* as I do, but those aren't the same thing as perceptual states and they aren't what my knowledge is based on.) So it might very well be true that my BIV-counterpart believes that it isn't a BIV. But it doesn't follow that the basis for my belief is insensitive, for my BIV doesn't share that basis.

The main result is that my view allows me to reject premise (5*). I can thus accept the basis sensitivity principle, along with premise (4*), and it will not give rise to any obstacle for my knowledge that I'm not a BIV. I take this to be a tremendous virtue of my view. I can accept the basis sensitivity principle without descending into skepticism.

The same cannot be said for other versions of foundationalism. Recall the popular internalist theories that I have spoken of earlier. Many of the most important rivals to my view are internalist theories. They will seek to understand knowledge and justification as ultimately based on factors that are internal to the subject's perspective. In particular, these views will hold that the justification for our beliefs will depend entirely on factors that would be available to us even if we were BIVs. They will then typically claim that our beliefs qualify as knowledge if they are justified and true (and perhaps meet a few other conditions). Take, for example, the phenomenal conservative position. This view holds that the justification for our beliefs is always ultimately rooted in how things appear or how things seem to us. These states are supposed to be entirely accessible by introspection and shared with our BIV-counterparts.

Suppose that we identify the basis for our beliefs (in the sense of 'basis' that's relevant for assessing knowledge and justification) with the states that are favoured by internalists. Whatever internalist bases we choose, my belief *that I'm not a BIV* will be based on the same states whether I am a normal non-envatted human or the victim of the brain-in-a-vat scenario. Moreover, I would still retain this belief even if it were false and I were a BIV. So if internalism is right about the bases for our beliefs, then the basis for believing *I'm not a BIV* is not sensitive. We can also think about this point in terms of

¹³The inference that I make may be no different from Moore's; "Here is a hand; here is another; therefore external objects exist."

phenomenal conservatism. The phenomenal conservative holds that our basis for belief is always rooted in our subjective appearances and seemings. But then if I believe that I'm not a BIV because of how things appear to me 'from the inside', then even if I were a BIV I would still believe for the same reason—it would still appear as if I wasn't a BIV. Thus the appearances and seemings are not sensitive guides to the truth.

Consequently, the internalist views in general, and phenomenal conservatism in particular, are more vulnerable to the kind of skeptical argument that stems from the basis sensitivity principle. If we grant the skeptic their two main premises, (4*) and (7), then internalism will lead straight into skepticism about the external world. In my opinion, this is a fairly significant reason to be dissatisfied with the internalist views. It is a reason to prefer my kind of account instead.

Of course, the internalist won't see it that way. They will come up with their own ways to keep hold of their position. Many of them will simply grant the skeptic's argument and then adopt a more pragmatic attitude towards epistemology. They will admit that "maybe we can't achieve *true* knowledge of the external world", but then insist, "that's fine; we should instead aim for beliefs that are probable and justified." Another tactic that the internalist might employ is to use another Moore-shift inference against premises (4*) or (7). They will insist that it is *obvious* that we have knowledge, and so the basis sensitivity principle must be false, or the requirement of ruling out every skeptical scenario must be unnecessary for knowledge. Perhaps this latter internalist is right to hold on to our knowledge as an obvious fixed point and then reconsider the assumptions that got them into their skeptical predicament. But if that's the way that they're going to go, then I would urge them to reconsider another one of their assumptions: namely, their internalism. After all, is it really so obvious that we share the same bases for knowledge as our BIV-counterpart? Is this really more obvious than premises (4*) and (7)? Some of them will dig in their heels and say yes; but I'm not going to agree with them on this point.¹⁴

As I've said, I take this all to be some reason to prefer my kind of view over the internalist views. My view neatly avoids the basis-sensitivity skeptical argument precisely because it rejects the internalist restrictions. Instead, my own view is *externalist* in nature. Now what exactly this means is not easy to make precise. But for one thing, it means that the basis for our knowledge is not simply how things appear to us. It means that we do not share the same basis for knowledge and justification as the victims of radical schemes of deception. Very roughly, it means that our knowledge and justification are partly dependent on our environment; in order to attain knowledge and justified beliefs, we must be situated in an environment that is conducive to our knowledge. We must be

¹⁴Suppose that for Moorean reasons, we are compelled to give up one of the following three propositions: (i) we share the same knowledge-basis with our BIV-counterparts, (ii) if I know that I have hands then I know that I'm not a BIV, and (iii) if I know that p on basis B , then if p were false, then I wouldn't believe that p on basis B . And suppose that I'm right that the weakest link *seems* to be the internalist claim (i). That would be highly ironic. For it would mean that the phenomenal conservative methodology would undermine itself!

related to the objects of our environment in such a way so that we can actually perceive them. This, for one, means that the objects must *actually be there* in our field of vision, unobstructed from our view, so that our perception can work properly. (See the next chapter for more about perception). When these environmental factors are in place, then the basis for our knowledge—e.g. our perceptual states—will be sensitive to the truth.

4.4 The ability to distinguish the truth

Having now said all of this, it is still possible that sensitivity, with the precise construal that philosophers have given it, does not speak to the skeptic's underlying concerns. I still have one more attempt to interpret their concerns to offer. Once again, this will be another argument for their premise that *I cannot know that I'm not a BIV*. But this time, the idea is centred around the idea of being *able to distinguish truth from falsity*. The skeptic could be thinking that, since things would appear to be normal to my BIV-counterpart, it follows that my BIV-counterpart cannot distinguish their situation from normal life. But if envatted life and normal life cannot be distinguished, then it follows that *I* cannot rule out the possibility of being a brain in a vat. Then the skeptic's argument would be off and running again. Here is how the skeptic might reason, put into a fully reconstructed argument.

- 9 If I were a brain in a vat, then I wouldn't be able to distinguish my case from the normal, non-envatted case.
- 10 Thus even if I'm in a normal, non-envatted case, I am never able to distinguish my case from being a brain in a vat.
- 11 If I cannot distinguish my case from being a brain in a vat, then I cannot know that I'm not a brain in a vat.
- 12 Thus I cannot know that I'm not a BIV.
- 13 If I know that I have hands, then I know that I'm not a BIV.
- 14 I do not know that I have hands.

And just like that we have another full-blooded argument for skepticism. This one is not based on sensitivity, but rather the ability to distinguish one's case from the sceptical scenario.

Before we go any further, we might want to ask what it means to have the ability to distinguish one case from another. Presumably it's an epistemic ability: it's an ability to gain a certain kind of knowledge. In fact, I think it's plausible to interpret *S can distinguish case A from case B* as *S can know (or is in a position to know) that A obtains and B doesn't obtain*. But if I'm right about this interpretation, then (11) is pretty much true by definition and (12) follows directly from (10). We should also accept (9) as true, since a BIV can

probably never come to know that it is BIV. So there is really only part of this argument to consider: does (10) logically follow from (9)?

To answer this question, we have to consider whether the *ability to distinguish A from B* is symmetric. If one can distinguish A from B, does it thereby follow that one can distinguish B from A? This is the move that the skeptic tacitly relies on when they infer (10) from (9). The skeptic argues that I could not distinguish (as a BIV) the envatted situation from the normal non-envatted case. They then infer that I cannot distinguish (as a non-envatted human) my normal case from being a brain in a vat. But is this inference warranted?

We must admit that, at first glance, this is a fairly intuitive inference. Oftentimes if I can distinguish A from B, then I can also distinguish B from A. I can distinguish when I drink IPAs from when I drink saisons; I can also distinguish when I drink saisons from when I drink IPAs. But despite the fact that this symmetry is common, it is not a universal feature of the logic of distinguishability. To see that distinguishability is not generally symmetric, consider the case of being slightly tipsy and the case of being drunk. If I've only had a couple of beers and I'm only slightly tipsy, then I'm often able to distinguish my state from being drunk. But as I have more beers, I will eventually become drunk; and once I am, I will not generally be able to distinguish my state from being only slightly tipsy. The point is, it's not generally true that being able to distinguish A from B implies being able to distinguish B from A. This inference won't be licensed by the logic of distinguishability.

If this is right, then the inference from (9) to (10) in the skeptic's argument is invalid. Premise (9) can be true while (10) is false. And in fact, I think we have excellent reason to think that this is one of the times where distinguishability fails to be symmetric. If I am a normal, non-envatted human, then I enjoy many epistemic abilities that I wouldn't have if I were a brain in a vat. I can perceive my flesh-and-bone hands, I can perceive my body, I can perceive that I'm sitting in a chair, and so on. All of this allows me to distinguish my own case from being the victim of the BIV scenario; that is, I can *know* that I'm not a BIV. But if I were a BIV, then I wouldn't enjoy any of these perceptual states. I would thus have very little to go off of to know my own fate.

5 Externalism in general

Throughout this chapter, we have seen the skeptic try many subtle ways to make the same general point. So you think you know that p , where p is some fact about the external world? Well the skeptic would like you to reconsider. They describe a scenario q where *if q were the case, then things would appear as if p and yet p would be false*. (Maybe you are dreaming; maybe you are a brain in a vat.) They implore you to rule out the possibility that q . If you can't rule out q , then they say that you don't really know that p .

We have also seen, throughout this chapter, many different strategies for combating this kind of skepticism. Perhaps the most enticing strategy is the highly ambitious one that was innovated by Descartes. For this strategy, we

attempt to reclaim our knowledge by rebuilding it upon the foundations that are immune to skeptical doubt, which for the Cartesians, means our *a priori* knowledge and the appearances. In contrast to this approach, we also saw Moore's advice to begin our inquiry into knowledge and skepticism by starting with what we know as a matter of common sense, which for him, included some beliefs about particular external objects (e.g. that I have hands). Lastly, we saw the phenomenal conservatives, who take our foundationally justified beliefs to be the ones that are based on the most coherent set of our strongest seemings and appearances.

In my view, no anti-skeptical strategy will completely succeed as long as it subscribes to the internalist requirement that we must share the same knowledge-basis as our BIV-counterparts. The Cartesians tried to run their programme with this in mind, but I've never seen them succeed in recovering what was recognizably *knowledge*. The phenomenal conservatives overtly claim to justify their beliefs based on *how things seem* and *how things appear*, but I worry that they won't yield knowledge, since these seemings and appearances aren't sensitive to the truth. The only strategy that doesn't explicitly abide by the internalist constraint is Moore's stubborn insistence that he knows that he has hands. However, when interrogated, Moore would defend this assertion as an item of common sense. But I do not think that a sole appeal to common sense, without anything else to offer, is sufficient to respond to the skeptic's worries. That is because the skeptic has arguments to offer, and so we need to develop some philosophical theory to engage with them on the theoretical level.

If I'm right, then the key to solving the skeptical problem is to appreciate how the ordinary person with working perception is in a very different epistemological situation than their brain-in-a-vat counterpart. The former stands in a direct cognitive relationship to the facts of their environment, whereas the latter is cognitively closed off from their surroundings. And according to my account, this makes a world of difference when it comes to their ability to attain knowledge. It is true that things may appear the same to the both of them, but there's more to knowledge than how things appear. Knowledge is a matter of being in contact with the world in the right way.

5.1 Internal defeat

Before I close this chapter, I want to flag two sources of discomfort that are common against this externalist solution to skepticism. At this stage, we are no longer considering the skeptic's complaints. We are now moving on to the complaints of those who find the externalist approach to epistemology unsatisfactory for reasons that are intrinsic to it.

One aspect of my view that is bound to bother some people stems from my claim that *believing that p on the basis of perceiving that p is sufficient to know that p*. A critic of mine could say that this doesn't pay sufficient attention to the rest of the subject's evidence or beliefs. They will say that *knowing that p* requires the subject to be *rational* in believing that *p*, and being rational in believing something requires it to cohere with the rest of one's beliefs and

evidence. If the subject has other evidence or beliefs that undermine p , then they cannot know that p , even if they perceive it.

Let's give my critic an example to see what they're working with.¹⁵ Suppose that I visit an art gallery and I view a red painting. I perceive that it is red and thereby form the belief that the painting is red. (I thus meet my proposed criteria for knowing that it is red). But as I'm viewing the painting, a deceitful gallery attendant tells me that it isn't really red; it just appears to be red because of a clever trick done with the lighting. The attendant doesn't strike me as a deceitful person, but nonetheless I ignore their testimony and continue to believe that the painting is red. It's a good thing that I did, because the painting really was red and my perception really was working properly; the attendant was lying to me.

Many epistemologists take examples like this to count against the externalist conception of knowledge.¹⁶ According to them, I don't really know that the painting is red in this case, despite the fact that I perceived it to be red and my perception is working properly. I don't really know it because my belief is irrational. I'm dogmatically holding onto a belief that I have evidence against (albeit misleading evidence). And if these epistemologists are right about this, then there's a problem for my view: perceiving isn't sufficient for knowing.

My answer to this conundrum will be shocking to any reader who harbours internalist sympathies. I maintain that I actually *do* know that the painting is red, despite the misleading testimony. Not only that, but I also maintain that there's a sense of "rational" in which I *am* believing rationally, and that this is the sense of "rational" that's important for knowledge.

Of course, in order to defend myself in this way, we are going to have to get clearer on what we mean by *rationality*. In the context of this discussion, there are two things that we could mean. When we say that a belief is *rational*, we could mean that the belief is supported by the subject's other beliefs, *whatever those happen to be*. (We sometimes speak of subjects *rationalizing* their beliefs by corroborating them with other beliefs of theirs.) In this sense, a belief would be *irrational* if it doesn't cohere with the subject's other beliefs—either it is inconsistent with their other beliefs or it is otherwise made implausible by their other beliefs. This sense of rational is internalist because it depends on how things stand within a belief structure, regardless of how things stand in the world. However, according to the other sense of *rationality*, we say that a belief is *rational* if it is what the subject *should believe* (has reason to believe), *given how things are objectively*. That is, given the subject's environment, and given their perceptual states, reliability of their interlocutors, pre-existing knowledge, etc., what should they believe in their situation? Whatever is the right thing to

¹⁵There are many examples like this throughout the epistemological literature, but this one is closest to Maria Lasonen-Aarnio's from "Unreasonable Knowledge", and I am agreeing with her verdict.

¹⁶It isn't so clear how this example could work to count against externalism *per se*—the mere claim that we do not share the same knowledge basis or justification as our BIV counterparts. But this example does directly challenge my particular version of externalism which claims that perception is sufficient for knowledge.

believe—objectively considered—is *rational* in the second sense.¹⁷

In order to see the difference between these senses of rationality, simply consider someone who, for whatever reason, has previously come to accept some patently bad beliefs. At the time of my writing this, there's a regrettably large (and loud and annoying) group of people who believe that the COVID-19 pandemic is a hoax, and they have their batch of conspiracy theories and junk science to back up this belief. Well, given that they have this belief, corroborated with further beliefs of theirs, they can then rationalize the further belief *that it is safe to congregate in large numbers*. This belief could exhibit *internal rationality* (in the first sense of rationality) within their overall belief system.¹⁸ But it is *not* something that they *should* believe, objectively speaking. The scientific evidence and facts of the world tell against the various beliefs that they use to back up their belief. The point is, just because something is *consistent with* (or even supported by) one's previous beliefs, it doesn't follow that one *should* believe it. Perhaps one's other beliefs are bad, too.

If this is right, then internal consistency within one's belief system (the first sense of rationality) does not provide a positive *objective* reason for why one *should* believe something (the second sense of rationality).¹⁹ These two senses of rationality come apart, and so they should be distinguished. But having now distinguished them, we can ask, which one is relevant for knowledge? If a belief is to count as knowledge, then in what sense must it be rational? My proposal is that knowledge implies rationality in the *second* sense, and not the first. That is, if I know that *p*, then *p* is something that I should believe, given my evidential situation, objectively considered. But my belief is not required to be rational in the second sense of cohering with everything else that I believe. Coherence is only required with the beliefs of mine that I *actually know*—it isn't required with every other irrational or false belief that I have. To see this, consider the philosophy student who has become convinced, in his philosophy class, of skepticism about the external world. He becomes convinced that he cannot know that he has hands, and that he ought to suspend judgment on the matter. But like the rest of us, he cannot shake the belief that he has hands in his life outside the classroom. Does it follow that he cannot know that he has hands, because of the internal incoherence between his beliefs? I submit that it does not follow. As long as he perceives that he has hands, he can know that he has hands—despite whatever false undermining beliefs he might buy into.

Okay, now let's go back to my belief about the painting in the art gallery. If I am to know that the painting is red, then this belief must be *objectively rational*;

¹⁷You might wonder whether a rational belief, in the second sense, must always be true. I'm not sure about that; I take no stand on the matter. Perhaps it is objectively rational to believe that you won't win the lottery, even if this turns out to be false.

¹⁸Some readers will respond that it's unlikely that these people have internally coherent belief systems. If we dissect their overall beliefs systems, they will be riddled with incoherence and contradiction. Perhaps this is true; but the point still stands. Just pick another example of a belief system that is internally coherent but objectively irrational. Clearly it's possible for there to be such belief systems.

¹⁹This is the conclusion of a very long, technical, but profound paper by Niko Kolodny, called "Why be rational?"

it must be something that I *should* believe, given my perceptual state and the reliability of my informant. So is my belief rational in that sense? As it turns out, the answer is *yes*. Objectively speaking, I *should* trust my perceptual state and disregard the testimony of the gallery attendant. After all, my perceptual states are actually accurate, and the attendant is trying to deceive me. Put it this way: if you had a god's-eye-view of the situation, what advice would you whisper into my ear about what I should believe? You would tell me to ignore the attendant and that my perception is working just fine. You would say that, objectively speaking, I should trust my senses and continue to believe that the painting is the colour I perceive it to be. From the god's-eye-view, my belief is *rational*.

5.2 Knowledge and doxastic decisions, redux

When an internalist hears my response to the previous objection, they are bound to become very uncomfortable with my view. I have just appealed to a god's-eye-view of the subject to determine what is 'objectively' rational for them to believe. But the internalist will be quick to point out that *we*, as inquiring subjects, do not enjoy this mystical god's-eye-view of our situations. And since we do not possess this god's-eye-view, we do not have a perfect ability to discern what is *objectively rational* in all situations.

This is a consequence of my view that I'm perfectly willing to accept. We don't *always* have a perfect ability to discern what's the objectively rational thing to believe. Sometimes what we ought to believe will depend on factors that aren't apparent to us, such as the reliability of our informants. That is because *what we ought to believe* is objective. It can depend on factors that extend beyond our first-personal perspective. This isn't to say that we're *never* in a position to know our epistemic obligation. Again, we're often in a position to know what we can know. But there will also be situations where we will lack the higher-order knowledge as to what is the most rational thing to believe at the time.

Yet, despite my willingness to admit this, the internalist will still worry that there is something missing in my view. They will worry that my view will become *practically useless*, since it claims that our epistemic obligations will depend on factors that aren't entirely apparent to us. Such an internalist will say that a proper epistemology should offer some sort of *guidance* as to what we should believe in any situation. For example, it should be able to offer the person in the art gallery some practical advice as to what they should believe, *given how things seem from their perspective*. In other words, a useful theory of epistemology should provide a practical guide to belief, such that it is always possible *in principle* to know that one is following it faithfully.

To respond to this source of discomfort, I would like to remind the reader of a distinction that I made in the first chapter between three tasks for epistemology. There I called them the three tiers of epistemological inquiry. First, there is the task of applied epistemology, which concerns the practical questions about what we should believe in real life situations. Secondly there is the tier of normative

epistemology, which concerns fundamental questions about belief systems in general—such as, *can we know anything at all?* and *what are our most basic epistemic obligations?* And thirdly there is the tier of *meta-epistemology*, which concerns the nature of knowledge and epistemic obligation, and the sorts of states that they depend on. I will continue to maintain, as I said in the last chapter, that each of these tiers of inquiry ought to be kept separate at first. We shouldn't discredit an answer to one type of question because it doesn't lend itself to any quick answers to the others. If there is any connection between these types of epistemological questions, then these connections need to be *discovered*, not presupposed or forced.

Throughout this chapter my primary concern has been with the second tier normative questions. Particularly, we've been concerned with the question of whether we can know anything at all about the external world. I answered yes and I gave my reasons. I would also note that this ties into the other normative question about our basic epistemic obligations. It is plausible to think that if we're in a position to know something, then we ought to believe it. So if perceptual knowledge is possible, then *we should believe what we in fact perceive*.

It has not yet been my concern to say anything about applied epistemology. As I hinted at in the last chapter, I do think there's a connection, but it is not as straightforward as the internalist hopes for it to be. What exactly this connection is will have to wait for chapter 5.²⁰ There are still a couple more skeptical problems that we have to deal with at the more fundamental level, and we should probably think about those before we turn to more practical matters. Suffice to say that these internalist scruples are not at all a good reason to reject my view. One shouldn't dislike a theory because it doesn't offer quick answers to questions that it was not designed to answer.

Finally, what of the tier-III meta-epistemological questions? The ultimate concern of my project has never been with these questions. However, within this chapter, we found an important and unexpected connection between meta-epistemology and normative epistemology. In order to answer the question *can I know anything at all about the external world?*, we found that it's crucial to think about what kinds of bases there are for knowledge. We found that if knowledge depends on *how things appear to the subject*, then our claims to knowledge of the external world are precarious and doubtful. But if, as I've insisted, knowledge depends on whether the subject is tuned into their environment—for instance, by perceiving the facts—then skepticism can be defeated and knowledge of the external world is possible.

²⁰But if you're feeling impatient, you could check out my paper "Doxastic Decisions and Epistemic Internalism."