

## Chapter 2

### The Lure of Radical Skepticism

Philosophical skeptics hold that one cannot know anything about the external world.<sup>1</sup> The phrase “the external world” (a.k.a. “the objective world”) refers to everything that exists outside of one’s own mind. So, according to the skeptic, you can not know how many fingers you have, if any. You cannot know whether the book you seem to be reading from really exists, or if it is just a convincing illusion. You also cannot know anything about the minds of other people—you cannot know whether there are any other conscious beings in the world, nor, if there are, what kinds of thoughts they might be having. You can only know, at most, what is going on in your own mind. Every person (assuming there are more than one person) is in the same situation—that is, each person knows only of his own mind.

As an aside, notice that the skeptic does not say you only know what is going on *in your head*. He says you only know what is going on in your *mind*. “Heads,” just like fingers and books and brains, are objects in the alleged physical world of whose existence we can never be certain. You think you have a head only because you think you can see and feel it; but what you really, directly experience is only a mental representation of a head (just as, in the case of a finger, you experience a mental representation of the finger), which could be an illusion. For all you know, you might be only a disembodied spirit, subject to a massive hallucination of existing in the physical world.

Notice also that skepticism is not the same thing as idealism. The skeptic need not *deny* the existence of an external world, as the idealist does. Rather, the skeptic maintains that we *do not know* there is an external world. To use an analogy, the skeptic is like the agnostic, whereas the idealist is like the atheist. This puts the skeptic in a stronger position than the idealist, for this reason: the skeptic does not need to prove that our beliefs about physical objects are false or even probably false; the skeptic only needs to create some reasonable doubt about those beliefs.

Even so, the skeptic's position seems extravagant. Is there a serious doubt, for example, about whether rocks exist? People who harbor such doubts for too long are liable to end up in mental institutions. Why, then, does this position merit serious discussion? I will return to this question later (section 3.4). For now, I will simply say that skepticism deserves serious philosophical attention because of the arguments that have been developed on its behalf. In philosophy, it is not good enough merely to find your opponents' views absurd and dismiss them. Even if you are right (one might say, especially if you are right), the important thing is to understand *Why*. In the present case, this means understanding the lines of thought that lead to skepticism, and knowing precisely where they go wrong. This turns out to be a much more difficult task than it sounds.

I shall present four skeptical arguments below. Each of these arguments purports to show, at a minimum, that there is no good reason for thinking external objects exist. In fact, the first two arguments try to show even more: that there is no good reason for believing anything whatsoever. My aim in this chapter, playing devil's advocate, is to present the skeptic's case in its strongest form. After that, we must endeavor to develop a theory of perceptual knowledge that avoids the skeptic's attacks. My positive theory, I believe, will be the stronger for having confronted its opposition forthrightly.

## 2.1 The Infinite Regress Argument

The first argument goes as follows. In order for me to *know* something to be true, I must have an adequate reason for believing it. This is one of the things that distinguishes knowledge from mere belief. To take an example from Richard Fumerton, suppose I announce that the world is going to come to an end in the year 2100.<sup>2</sup> You ask me, “How do you know that?” What you are asking for is a *reason* (specifically, some *evidence*) for believing that the world will come to an end in the year 2100. Now suppose I say, “Oh, it’s just a whimsical hunch I have. I don’t really have any reason for thinking that’s true.” In that case, you could conclude that, although I may *believe* that the world is coming to an end in the year 2100, I certainly do not *know* that it is. Beliefs like that—beliefs held for no reason—are typically referred to as “arbitrary assumptions.”

In addition, in order for my reason to be adequate, it too must be something that I know to be the case. Again, suppose I announce that the world will end in the year 2100. This time, when asked why I believe this, I say, “I believe the world will come to an end in the year 2100, because the Plutonians are going to launch a lethal nuclear strike against us in that year.” When asked how I know about the Plutonians’ plans, however, I reply that I don’t know any such thing; it was just a whimsical hunch. Once again, you would conclude that I do not know that the world is going to end in the year 2100, since the reason I gave for this hypothesis was inadequate.

These requirements on knowledge create the threat of an infinite regress. For suppose I claim to know some proposition, A. You ask me my reason for believing it. It turns out that my reason for believing A is another proposition, B. You ask me my reason for believing B, which turns out to be a further proposition, C. You ask

me my reason for believing C .... It is clear that this cannot go on forever. I cannot actually have an infinitely long chain of reasons standing behind my original assertion A. Nor is it permissible for me to rely on circular reasoning. For instance, suppose that, when asked my reason for believing A, I say I believe A because I believe A. In that case, I have not really given any (legitimate) reason for believing A. Nor are matters improved if I say B is my reason for believing A, and A is my reason for believing B. Nor, again, will matters be better if I simply expand the circle to include more beliefs; no exercise in circular reasoning will help me gain knowledge, however many steps the circle contains.

There is one remaining possibility, then. Every chain of reasoning must have a beginning point. In other words, all of my beliefs must rest, sooner or later, on propositions that I believe *for no reason*. Now, how can I know whether these starting beliefs are true? From what we have said above, it is clear that I cannot. By definition, I have no argument or evidence for my starting beliefs—if I did, then they would not be “starting beliefs.” But without an assurance of the truth of my starting assumptions, the derivation, however rigorous, of other propositions from them is worthless. This is an obvious point, but it bears stressing. After all, any proposition whatsoever can be derived from *some* premises or other. The mere fact that I can derive my belief that A from some assumptions does nothing whatever to establish the truth of A; if it did, one could also establish the negation of A by *its* derivability from certain, other assumptions.

In general, conclusions are only as good as the premises they are based on. So if I do not know my starting premises to be true, then even more surely I do not know my conclusions to be true. Therefore, I cannot know anything.

This argument can be summarized as follows:

1. In order to know something, I must have a good reason for believing it.
2. Any chain of reasons must have one of the following structures:
  - (a) it is an infinite series, or
  - (b) it is circular, or
  - (c) it begins with a belief for which there are no further reasons. But,
3. I cannot have an infinitely long chain of reasoning for any of my beliefs.
4. Circular reasoning cannot produce knowledge.
5. Nor can I gain knowledge by structure 2c, for
  - (a) I would not know my starting beliefs to be true (from 1), and
  - (b) I cannot gain knowledge by deriving it from assumptions that I do not know to be true.
6. *Therefore*, I cannot know anything.<sup>3</sup>

What is wrong with the above argument? The majority of philosophers and others who hear the argument say that it is premise (1) which is mistaken. They say that there are certain self-evident, or *foundational* propositions. A foundational proposition, by definition, is one that we can know to be true *without* having a reason for it, and the people who believe in such things are called “foundationalists.” According to most foundationalists, propositions such as “ $2 = 2$ ” and “I am now conscious” would be good examples of foundational propositions. I do not have to give an argument, or engage in a process of reasoning, to know that I am conscious, or that the number 2 is equal to itself. I merely think about these propositions and thereupon find their truth immediately obvious.

The skeptic, of course, will deny the existence of self-evident propositions. But why? Obviously, the skeptic cannot say, “It is self-evident that self-evident propositions do not exist.” To be

consistent, he will have to produce an *argument* against the idea of self-evident propositions. What argument can he give?

Well, the skeptic can argue that the foundationalist has no way of distinguishing *self-evident* propositions from merely *arbitrary* propositions. A self-evident proposition, we have said, is one that we need have no reason for in order to be fully justified (or rational) in accepting it. An arbitrary proposition, on the other hand, is a proposition that we have no reason for and would be wholly unjustified in believing. For instance, suppose I suddenly decide, completely out of the blue, that I think there is a twelve-headed purple dragon living on Venus. This would be an arbitrary belief. The foundationalist must explain what *differentiates* a foundational proposition like “ $2 = 2$ ” from an arbitrary proposition like “There is a twelve-headed purple dragon on Venus.” That is, he must identify some feature of the foundational proposition that the arbitrary proposition lacks, and that explains why the foundational proposition is justified. Let “F” denote this feature.

Assume, then, that I have a belief, A, which is a legitimate foundational belief. And assume that I have another belief, B, which is merely arbitrary. By hypothesis, A has F, while B lacks F. Now, either I am *aware* of feature F, or I am not. But if I were completely unaware of feature F, then how could its presence serve to make it rational for me to accept A? If the presence of F is to explain why I am rational (or justified) in accepting A but not rational in accepting B, it must be something that I am aware of (in the one case, but not in the other). Otherwise, A and B will be, *from my point of view*, equally good (or equally arbitrary) assumptions. In that case, given the information available to me, it would be equally reasonable for me to accept one as to accept the other.

So the foundationalist position will have to be that it is reasonable for me to accept A, because I am aware that A has feature F. But

then A is not a foundational proposition after all, because I *do* have a reason for accepting A—namely, that A has F. Thus, foundationalism is reduced to absurdity: from the supposition that A is legitimately foundational, we can derive the conclusion that A is not foundational after all. Therefore, the very idea of a foundational proposition is self-contradictory. Therefore, it appears, the skeptic’s argument stands.

## 2.2 The Problem of the Criterion<sup>4</sup>

I have on my desk an epistemologically interesting toy called “the Magic Eight Ball.” It is a plastic ball painted like an eight ball, and it is meant to be used as follows. You ask the eight ball a yes/no question. Then you tum it over and see an answer float up to a window in the bottom. Answers include the likes of “Yes, definitely,” “Very doubtful,” and “Cannot predict now.”

Now, imagine there were a community in which use of the eight ball was an accepted method of arriving at conclusions.<sup>7</sup> Suppose you meet one of these eight-ball reasoners, and you ask him why he believes that the eight ball is a reliable informant. He swiftly takes out his Magic Eight Ball, says, “Are you reliable?” and turns it over. At this point, if the answer “No” floats up to the window, then the eight-ball reasoner is in trouble. But suppose a definite “Yes” answer appears, and the eight-ball reasoner triumphantly declares that the reliability of the eight ball has been established. Would this be legitimate?<sup>5</sup>

Evidently not. You would no doubt object, rightly, that there is a problem of circularity here. If we already knew the eight ball was reliable, then we would be justified in accepting the answers it produces. But if the eight ball is unreliable, then we should not trust its answers. And if we don’t know whether it is reliable, then we likewise should not trust its answers *until* its reliability has

been established. The method of eight-ball reasoning presupposes that we know the eight ball to be reliable, in the sense that it would not be reasonable to use the method unless we already knew (or at least had reason to believe) it to be a reliable method. Therefore, we certainly cannot use eight-ball reasoning to establish that the eight ball is reliable.

Now consider an analogous case. Suppose some skeptic comes along and asks you why you believe the senses to be reliable. Why do you think that, when you seem to see, hear, or feel things, this is a reliable indicator of the way things really are, in the external world? How would you respond?

Here is one thing you might try. You go to an eye doctor to have your eyes examined. He gives you a series of tests, and at the end he assures you that your eyesight is perfect. Then you go to another doctor to take some hearing tests. He assures you that you have excellent hearing. (You might have difficulty finding doctors to test your taste, smell, and sense of touch, but let's pass over that difficulty.) You then explain to the skeptic that the reliability of your senses has been established. Would this be legitimate?

Apparently not. You would be engaging in just the same sort of circular reasoning that the benighted eight-ball reasoner used, for you can only collect the results of your tests by using your senses. You may seem to hear the doctors tell you that your hearing and eyesight are normal, but how do you know they are really saying that? Indeed, if you were in doubt as to the reliability of your senses in general, you could not even be sure that the doctors really existed, let alone that they were reliable informants.

So you will have to use some other method to verify the reliability of your senses—you will have to rely on some cognitive faculty other than the senses. But—here is the problem—whatever method you try to use to verify that your senses are reliable, the skeptic can always ask why you believe *that* method to be reliable.



For instance, suppose you wanted to prove the reliability of the senses through the exercise of pure reason (though I have no idea how such a proof would go).<sup>6</sup> In that case, the skeptic could ask why you think reason itself is reliable. You could not use reason to establish the reliability of reason, nor could you use the five senses to establish the reliability of reason, again on pain of circularity. So you will need to find yet a third belief-forming method. At which point the skeptic will question the reliability of this third method as well. At some point, and probably sooner rather than later, you will have to either resort to circular reasoning or else give up on answering the skeptic's question.

But this means that ultimately you cannot establish the reliability of your cognitive faculties. And all of your beliefs are formed through one or another of your cognitive faculties, whether it be through the five senses, or reason, or memory, or introspection. Since you cannot know whether any of your belief-forming methods is reliable, it seems, you cannot know whether any of your beliefs is true. In short, you are in the same position as the eight-ball reasoner. The eight-ball-generated beliefs were all unjustified since the eight-ball reasoner could not (noncircularly) establish that the eight ball was reliable. Similarly, all of your beliefs are unjustified since you cannot (noncircularly) establish that your belief-forming methods in general are reliable.

This argument can be summarized as follows:

1. All my beliefs are formed by some method.
2. I am justified in accepting a belief formed by method M only if I *first* know that M is reliable.
3. I do not have an infinite series of belief-forming methods.
4. Thus, all my beliefs must rest on beliefs formed by methods whose reliability has not first been established. (from 1 and 3)
5. Therefore, none of my beliefs are justified. (from 2 and 4)

This argument is similar to the argument of section 2.1. Again we have a threat of infinite regress or circularity, though this time it would be a series of belief-forming methods, rather than a series of beliefs. Once we rule out both the infinite regress and the circularity possibilities, the only remaining possibility is that I have belief-forming methods whose reliability is not established by any method. This is analogous to the beliefs of section 2.1 that are not supported by any reasons. I pointed out that even the most impeccable derivation of a conclusion from such assumptions would do nothing to establish the truth of the conclusion, given that we have no reason to accept the starting assumptions. Similarly, even the most scrupulous exercise of method M, whatever that may be, will do nothing whatsoever to establish the truth of any conclusion, given that we have no reason to think M itself is reliable. As a result, it seems, we have no way of knowing anything whatever.

### **2.3 How Can You Get outside Your Head?**

Most of the things we think we know, including everything we think we know about the physical world, we learn through sensory perception, which includes sight, hearing, taste, touch, and smell. Of course, this does not mean that everything we know about the physical world is something we actually observe. A lot of what we know of the physical world is the result of scientific theorizing or inference, but those theories and inferences are ultimately *based on* observations. For instance, we know of the existence of atoms through inferences from the observed results of experiments. Similarly, I know that the Battle of Hastings took place in 1066, not because I personally observed it, but because I read that in a history book—but I knew what the book said only because I could

*see the book.* In that sense, my belief was acquired through the exercise of my senses.

If you think about it, then, you will probably realize that everything you think you know about the external world is dependent on your senses. So in order to determine how much we really know about the physical world, we must first ask what the senses really tell us about the physical world. This question can be separated into two sub-issues: First, what is it that the senses make us *directly aware* of? Second, what can be *inferred* from what we are thus directly aware of?

The skeptical argument we are about to consider seeks to establish, first, that the senses do not make us directly aware of the physical world; and second, that no conclusions about the physical world can be inferred from what we are directly aware of either. It will follow that we can have no knowledge of the physical world.

To put that another way: the skeptic will seek to show, first, that direct realism is false. Second, he will try to show that indirect realism is false as well. It will then be clear that we have no knowledge of the physical world, since we do not know about it directly, and we do not know about it indirectly either.

The first part of the argument—the falsity of direct realism—is supposed to be established by the sort of argument we began with in chapter I. There are, actually, quite a few arguments against direct realism, which we shall consider more fully in chapter 6, but for now, let us stick to the argument from double vision. As you recall, we considered a case in which, though there is only one physical finger in front of you, you seem to see two fingerlike things. This was supposed to show that what you are immediately aware of is mental images, rather than the physical finger. The reasoning can be summarized as follows:

1. As your focus shifts to the background, the finger like thing you are seeing splits in two.
2. No physical object splits in two at this time.
3. Therefore, the thing you are seeing is not a physical object.

If the thing you are seeing is not a physical object, the next natural candidate is that it is a mental image (what else could it plausibly be?). These mental images are traditionally called “sense data,” so that is what I will call them here.

(Aside: They have also been called “impressions,” “ideas,” and “perceptions” by various people at various times. I think those terms are misleading, so I prefer the technical term “sense data.” Despite my use above, “mental image” is also misleading, because it suggests that we’re talking specifically about visual sense data—in ordinary language, there’s no such thing as an ‘image’ of taste or smell—but in fact, we’re talking about mental items that allegedly exist whenever we exercise any of the five senses.)

Now, given that all you ever directly perceive is your own sense data, can you infer anything about the external world? At first glance, this doesn’t seem too difficult. Your sense data must come from somewhere, and you know that you didn’t create them, since you have no direct control over your sense data. (If you did, you could just decide to stop hearing that horrible music your neighbor is playing.) So they must have been caused by external objects. Suppose you are having a sense datum of a tree. The simplest explanation of why you’re having this experience—and normally the correct one—is that there is a tree in front of you, which is causing your experience. Granted, it is possible for a person to hallucinate a tree; however, that is not the normal situation, and there is no special reason for thinking you are hallucinating now.

The great skeptic David Hume neatly exposed the problem with this line of thought:

It is a question of fact, whether the perceptions of the senses be produced by external objects resembling them. How shall this question be determined? By experience surely, as all other questions of a like nature. But here experience is, and must be entirely silent. The mind has never anything present to it but the perceptions, and cannot possibly reach any experience of their connexion with objects. The supposition of such a connexion is, therefore, without any foundation in reasoning.

Hume believed, plausibly enough, that the only way of knowing that A causes B (where A and B are any two types of events) is by having some experience of A and B—specifically, you must observe A being followed by B on a number of occasions. For instance, suppose there's a light switch on the wall in front of me. In order to find out what the switch does, I'll have to try it out. I flip it a few times, notice the light go on and off, and conclude that flipping the switch causes the light to go on or off. If I never observed the flipping of the switch, I would not have been able to know this. Now, it is true that, when I enter a room I've never been in before, I can often predict that the switch on the wall will turn on the lights. However, this is because I am relying on *past* experience with light switches. If I had never had any experience with any light switches, I would have no idea what it would do.

Now, Hume says, for the reasons given above, that we never actually see physical objects, only our representations of them. Therefore, we have certainly not observed the presence of physical objects being followed by the occurrence of sense data. Therefore, we cannot claim to know that physical objects cause sense data. In particular, we could not claim to know that physical trees cause tree-like sense data, because we have never actually had any direct experience of a physical tree; all we have seen is the tree-representing sense data. It is as if I saw the lights go on and off

periodically, but I never saw the light switch (suppose the switch was located in another room of which I was unaware). In that case, I would never know what was causing the lights to go on or off.

This argument can be summarized as follows:

1. In order to have knowledge of the physical world, we must be able to know that our sense data are caused by physical objects.
2. In order to know that A causes B, one must have experience of A and B.
3. We have no experience of physical objects.
4. Therefore, we do not know that physical objects cause our sense data. (from 2,3)
5. Therefore, we have no knowledge of the physical world. (from 1,4)

[Sections 2.4 and 2.5 are not included here]

## Notes

---

<sup>1</sup> More precisely, that we cannot know any *contingent* truths about the external world; most skeptics do not question our knowledge of necessary truths. Henceforth, my talk of “knowledge of the external world” is to be understood as referring to knowledge of contingent truths about the external world.

<sup>2</sup> This illustration, including the Plutonian hypothesis in the following paragraph, is from Fumerton, *Metaphysical and Epistemological Problems*, 39 (I have altered the date of the hypothesized end of the world).

<sup>3</sup> Versions of the infinite regress argument for skepticism can be found in Sextus Empiricus, 72-75, and I. T. Oakley.

<sup>4</sup> I call the argument of this section “the problem of the criterion” because it is closely related to the ancient skeptical argument in which the skeptic says that one needs a criterion for distinguishing accurate perceptions from illusions, before one can reasonably rely on one’s senses. The skeptic goes on to argue that

---

there is no noncircular way of establishing that a given criterion is correct. See Sextus Empiricus, 145-46, and Chisholm, “The Problem of the Criterion.”

<sup>5</sup> The Magic Eight Ball example is from Fumerton, *Metaepistemology*, 50-51.

<sup>6</sup> Descartes attempted such a proof in the *Meditations*, but no one other than Descartes seems to have found his argument convincing. Davidson (“A Coherence Theory of Truth and Knowledge”) also gives an argument of this kind, but I think it no better than Descartes’s (see Foley and Fumerton for a successful criticism of Davidson).