Chapter 11
Skepticism About the External World

Abstract One way to challenge our scientific knowledge is to challenge all of our knowledge of the world around us. This chapter explores the classic philosophical problem of external world skepticism. This philosophical problem challenges whether we can know anything at all about the world outside of our own minds. In the process of examining and responding to arguments for external world skepticism important insights about the nature of scientific knowledge are revealed. One of the foremost of these insights is that knowledge in general does not require evidence that makes the believed proposition absolutely certain—beyond all *possible* doubt. Instead, significant, yet fallible, evidence is all that is required for knowledge, both scientific and mundane. Another insight is that the explanationist account of evidential support developed in the previous chapter helps to show that, despite initial appearances, external world skepticism is not a significant threat to our knowledge after all. We do have good evidence for believing the world is roughly the way we commonsensically take it to be.

We have discussed many features of knowledge as well as what is required for us to come to possess knowledge, both scientific and mundane. Throughout this discussion we have been working under the assumption that we do in fact know a great many things. Despite the fact that this assumption is exceedingly plausible, our knowledge of scientific claims faces many challenges. In this part of the book we will examine some of the major threats to our scientific knowledge. We will begin by exploring threats which challenge our scientific knowledge by way of challenging some broader category under which our scientific knowledge falls. Afterward, we will examine threats arising from the debate between realists and anti-realists which specifically target our scientific knowledge.

This chapter begins our consideration of the more sweeping challenges to our scientific knowledge. The first such challenge we will think through is one that threatens all of our knowledge of the world around us. It is particularly worrisome when we realize that included in this body of knowledge is not just scientific knowledge, but all sorts of things we typically take ourselves to know: there are trees, other people exist, we have bodies, and so on. This challenge comes from a classic philosophical problem: external world skepticism. Although the threat of external world skepticism is one that can be overcome, consideration of this

K. McCain, *The Nature of Scientific Knowledge*, Springer Undergraduate Texts in Philosophy, DOI 10.1007/978-3-319-33405-9_11

philosophical problem is worthwhile because it illuminates important features of the nature of scientific knowledge and knowledge more generally.

11.1 Challenge I: Lack of Certainty

One common way skeptics challenge our scientific knowledge is by pointing out that we cannot be *certain* that scientific claims are true. The sense of "certain" that is relevant here is not merely a psychological state of being completely sure or exceedingly confident that something is true. Some people are completely sure that humans never landed on the moon. They believe the moon landing is a hoax, and they believe this without a shadow of a doubt. Nevertheless, despite their confidence the claim that humans never landed on the moon is definitely not certain for these people! The relevant sense of "certain" here is the sense in which you are certain that p just in case your evidence for p guarantees p's truth. In other words, to be certain in this sense requires that your evidence for a particular claim is so strong that it is impossible for you to have that evidence when the claim is false.

Of course, it is at least possible that despite the evidence we have in support of any particular scientific claim we could be mistaken. It is often awareness of this tentativeness that leads some to the misguided objection that evolutionary theory is "merely a theory".[1] The thought being that the evidence in support of evolution does not make it absolutely certain that it is true, so it is merely a theory and not something we know to be true. This is sometimes a bit disconcerting because it is partly correct. We do not have sufficient evidence to be absolutely certain of the truth of evolution, or any other scientific claim for that matter. Our evidence is always such that it is possible for us to have that evidence and yet the scientific claim in question turn out to be false. Recognition of this fact is one of the reasons why it is commonly accepted that science is tentative—since we lack certainty for scientific claims, the scientific claims we accept as true are always revisable in the light of new evidence.

The fact that our evidence does not make the truth of scientific claims certain for us should not be surprising. After all, as we noted in Chap. 5, the same is true of any claim about the world around us—in a scientific setting or not. For example, it is possible for you to have the same evidence you do now and yet not be reading this book. After all, you could be having a particularly vivid hallucination. So, your evidence for thinking that you are currently reading this book, while outstandingly strong evidence, is not so strong that it makes it absolutely certain for you that you are reading this book. This fact about our situation with respect to propositions about

[1]For more on this sort of objection to evolutionary theory and why it is misguided see McCain and Weslake (2013). For extensive discussion of various misconceptions concerning evolutionary theory and analysis of why some resist evolution in spite of the evidence see Kampourakis (2014).

the external world leads us to the first argument for external world skepticism we will examine:

CERTAINTY

C1) If you know some proposition about the external world, then your evidence makes the truth of that proposition certain for you.
C2) Your evidence does not make the truth of any proposition about the external world certain for you.
C3) Therefore, you do not know any proposition about the external world.

A couple points about CERTAINTY should be noted immediately. First, its second premise is clearly true. For any proposition about the external world your evidence in support of that proposition could be misleading. As noted above, it is at least possible that you could have all of the same evidence you do now, but you are not in fact reading this book. It is possible that you are really being deceived by a malicious demon of the sort philosopher René Descartes (1641/1988) described. It is possible you are really simply having an especially vivid dream. It is possible that unbeknownst to you someone has injected you with a hallucinogenic drug. Admittedly, these possibilities are far-fetched. Nevertheless, they are genuine possibilities—these scenarios are not impossible. Since they are possible, there is a very slight chance that your evidence for thinking you are reading a book is misleading, and you are not reading a book at all. But, if your evidence could be misleading with regard to whether you are reading a book, you cannot be *certain* that you are doing so. Second, C2 clearly applies to our scientific knowledge. Most (perhaps all) scientific claims are to some degree based upon observations we have made of the world around us. If we cannot be certain that any of those observations are accurate, we cannot be certain that the scientific claims based on those observations are true.

At this point you may be getting slightly worried. CERTAINTY concludes that we cannot know anything about the external world, which entails that not only do we lack scientific knowledge, we cannot know mundane things like that we have bodies, that other people exist, and so on. CERTAINTY has a very threatening conclusion. To make matters worse we have seen that one of its two premises is clearly true! Fortunately, it is easy to adequately respond to CERTAINTY because its other premise is clearly false. Knowledge simply does not require certainty. This is something we noted in Chap. 5, and it is something that nearly all epistemologists accept (Cohen 1988). This is why the dominant view of the evidence required for knowledge is fallibilism. Fallibilism is the idea that knowledge does not require evidence which guarantees truth (evidence that makes the proposition in question certain). According to fallibilism, it is possible for you to know that you are reading this book even though there is a remote chance that your evidence is misleading. Accepting fallibilism allows us to reasonably deny C1, and so, given fallibilism, CERTAINTY fails to commit us to denying that we have knowledge of the external world.

Although CERTAINTY does not pose a significant threat to our scientific knowledge or our mundane knowledge of the world around us, consideration of this argument helps illuminate important facts. First, we have to acknowledge that if knowledge really does require evidence which makes propositions certain, then we have very little (if any) knowledge. Perhaps we can have evidence which is this strong for thinking that we exist when we are currently thinking. Perhaps we can even have this sort of evidence for believing things like we are in pain when we are currently experiencing an intense pain. Regrettably, that is pretty much the full extent of the propositions about whose truth we can be absolutely certain. We definitely cannot have this sort of evidence for propositions about the external world. Thus, although we may wish for certainty when it comes to the things we believe, this is something that we simply cannot have for nearly any of our beliefs.

Second, consideration of CERTAINTY reveals why we should be fallibilists about knowledge. As we noted in Chap. 5, it is difficult to say precisely how much evidence is required for knowledge, but we can clearly recognize that it is something less than what is required for certainty. Additionally, as we noted in that chapter, a plausible answer to the question of how much evidence is required for knowledge is the criminal standard. Recall, the idea with the criminal standard is that in order to know that *p* your evidence must make the truth of *p* beyond a reasonable doubt for you. This means knowledge requires a significant amount of evidence, but nowhere near what is needed for certainty. By examining CERTAINTY we have not only taken an important first step in defending our scientific knowledge, we have furthered our understanding of the nature of knowledge in general. Unfortunately, there are more threatening skeptical challenges to our knowledge of the external world.

11.2 Challenge II: Underdetermination

Consider our commonsense view of the world, what we might call the "Real World Hypothesis" (RWH).[2] The RWH will include various claims, perhaps things like: "The earth is more than 3 minutes old", "Fire is hot", "Water is wet", and so on. The exact claims making up the RWH are not important for our purposes here. Rather, all that matters is that the claims of the RWH have two significant features. First, the truth of these claims entails an external world which is independent of our minds. That is, the truth of these claims requires that there is a world which is a certain way regardless of what we believe or take ourselves to know about that world. Second, if the claims of the RWH are known to be true, they constitute empirical knowledge. That is to say, we cannot come to know all of the claims of the RHW without investigating the world around us—not all of them are things we can know a priori (by reasoning alone). One of the classic ways of attacking our

[2]This name for our commonsense view of the world was put forward by Vogel (1990).

knowledge of the external world is to present skeptical alternatives to the RWH and argue that our evidence fails to determine which of these hypotheses is correct.

This is a broader threat than the underdetermination challenge that often arises in the context of debates between realists and anti-realists about science, which we will discuss in Chap. 14. This broader underdetermination challenge is particularly threatening for two reasons. One reason is that the skeptical argument which claims our evidence underdetermines whether the RWH or some rival skeptical alternative is true threatens all of our knowledge of the external world, not just our scientific knowledge. The other reason this challenge is particularly threatening is that it seems at least initially plausible that such rivals to the RWH do exist. For example, the skeptical alternative which claims that all of our experiences are caused by a manipulating demon who wants us to erroneously believe the RWH is consistent with all of our sensory experiences. The skeptical alternative that we are really brains in vats being deceived by the workings of a supercomputer is consistent with all of our sensory experiences too. A host of other skeptical alternatives fit with our sensory experiences as well. So, we seem to have genuine rivals to the RWH that at least initially seem to fit our evidence just as well as the RWH does. This is worrisome.

The skeptical argument from underdetermination can be formulated more precisely as follows:

UNDERDETERMINATION

U1) If you know the RWH, then you have evidence that favors the RWH over all rival skeptical alternatives.
U2) You do not have evidence that favors the RWH over all rival skeptical alternatives.
U3) Therefore, you do not know the RWH.

Three points about UNDERDETERMINATION are worth keeping in mind. First, knowing the RWH should be understood as knowing the claims of the RWH are true. This does not mean that one has ever thought that there is such a hypothesis or that she thinks of the claims of the RWH as hypotheses. Instead, the idea is simply that when one "knows the RWH" she knows (most of) the claims which make up the RWH such as "The earth is more than 3 minutes old", "Fire is hot", "Water is wet", and so on to be true. Second, to say that the evidence favors the RWH, over its rivals simply means that the evidence supports the RWH more than it does the rival skeptical alternatives to our commonsense picture of the world. Third, it should be noted that UNDERDETERMINATION cannot be dismissed simply by accepting fallibilism like CERTAINTY can. UNDERDETERMINATION claims that we do not even meet the fallibilist standard of evidence for knowledge—it holds that our evidence does not even make the RWH beyond a reasonable doubt for us. This makes UNDERDETERMINATION a much more worrisome threat to our knowledge of the external world. Fortunately, we have already equipped ourselves with the tools to respond to this skeptical threat.

11.3 The Explanationist Response

The first step to responding to UNDERDETERMINATION is to recognize that premise U1 is true. If your evidence does not support the RWH over its skeptical rivals, then you are not justified in believing the claims of the RWH. And, of course, if you are not justified in believing these claims, you do not know that the RWH is true. Since we recognize that U1 is true, we can focus all of our attention on showing that U2 is false.

The explanatory account of evidential support that we explored in Chap. 10 provides a very promising way of attacking premise U2. Recall, according to the explanationist account of evidential support discussed in that chapter, *Explanationism*, the evidence supports believing a particular proposition when that proposition is part of the best explanation of the evidence, or when that proposition is an explanatory consequence of the best explanation of the evidence. A similar approach can be used to respond to UNDERDETERMINATION. This *Explanationist Response* involves arguing that the truth of the RWH is a better explanation of relevant features of our sensory experiences than any of the available competing explanations (the various rival skeptical alternatives). Since the RWH is the best explanation of relevant features of our sensory experiences, the evidence we gain from those sensory experiences supports it.[3] As a result, the combination of explanatory considerations and our sensory experiences provides us with evidence that favors the RWH over all skeptical rivals. Thus, the Explanationist Response shows that premise U2 is false.[4]

Of course, before we can be confident that the Explanationist Response satisfactorily defends our knowledge from the threat of UNDERDETERMINATION we must first get clear on a couple things. We have to be clear about what exactly the relevant features of our sensory experiences are, and we have to get clear about why

[3] Putting things more precisely, inferring to the best explanation of our sensory experiences supports the RWH because the RWH is not only the best available explanation of the relevant features of those experiences, it is also a very good explanation in its own right. Though it can be difficult to spell out what exactly is required for an explanation to be good enough to be legitimately inferred, it is plausible that the RWH satisfies this requirement because the RWH accounts for all of the relevant features of our sensory experiences in a highly unified manner. See the previous chapter as well as Lipton (2004) for an explanation of why it is necessary to limit legitimate inferences to the best explanation to only those where the inferred explanation is of sufficient quality.

[4] There are other ways of responding to UNDERDETERMINATION. However, we will only focus on the Explanationist Response for a couple reasons. The Explanationist Response fits very nicely with the explanatory account which we developed in the previous chapter. Additionally, this response has the most promise of providing a satisfactory response to the threat of external world skepticism. Readers interested in other responses to the threat of external world skepticism are encouraged to see Greco (2000) for a response that is particularly appealing to those who like externalist views of epistemic justification, Huemer (2001) for the sort of response to this argument provided by the internalist account of epistemic justification known as "phenomenal conservatism", as well as Moore (1939), Pryor (2004), and Willenken (2011) for what has come to be known as the "Moorean" response to external world skepticism.

we should think the RWH really does explain these features better than its skeptical rivals. We will briefly consider the first point in the next section and the second point in the section after that.

11.3.1 Relevant Features of Sensory Experience

There are a number of features of our sensory experiences that are relevant to evaluating the explanatory merits of the RWH and its skeptical rivals. First of all, our sensory experiences come to us in an involuntary, spontaneous fashion (BonJour 1999; BonJour and Sosa 2003; Vogel 2008). Assuming that you do not have problems with your vision, when you open your eyes in normal lighting conditions you spontaneously have various visual experiences. These spontaneous visual experiences will also be involuntary. After all, while your eyes are open, the lighting is good, and you are looking at a tree you cannot make it so that you no longer have the experience of seeing a tree by just willing to no longer see it.

Additionally, our sensory experiences exhibit a high degree of regularity and coherence with one another. For example, our sensory experiences of one kind, say vision, fit with one another coherently, and they also fit coherently with our other kinds of sensory experiences. When you have a visual experience of one car crashing into another car nearby, your auditory sensations of a loud sound correspond to your visual sensations of the cars colliding.[5] Finally, there is coherence between our sensory experiences and our volitional activities. For instance, when you decide to grab (what you take to be) a glass of water your visual experiences usually change in a regular way which the hypothesis that you have limbs and are using one of them to voluntarily reach for a glass explains very well.

The RWH explains each of these facts about our sensory experiences. Given the RWH, our sensory experiences come to us in an involuntary and spontaneous fashion because they are caused by the interaction of mind-independent objects in the external world and our sense organs, and this interaction is not directly under our voluntary control. The reason our sensory experiences fit together coherently with one another is that the various kinds of sensory experiences are caused by the same mind-independent objects in the external world affecting our sense organs. In order to see how the RWH works it will be helpful to consider how it provides explanations of our sensory experiences in a particular example.

Consider the sensory experiences you are having right now. You are having visual sensations as of a roughly rectangular object as well as tactile sensations as of an object with edges some of which are longer than others (assuming that

[5]This is not to assume that we have different sense organs responsible for different kinds of sensory experiences. Such an assumption would surely stack the deck in favor of the RWH. Instead, here it is simply important to pay attention to the fact that we have sensory experiences of different kinds (visual, auditory, tactile, etc.), and these sensory experiences fit together coherently.

you are holding a hardcopy of this book in your hand rather than reading on an e-reader or your computer). You do not have to put forth effort to have these sensory experiences—they come to you spontaneously and involuntarily. The RWH explains these sensory experiences by positing a three-dimensional, rectangular object (the book), as well as your sense organs, and an appropriate light source. It is these same external objects which account for both kinds of sensations you are having.

The RWH also explains why our sensory experiences are coherent over time. Consider again your sensory experiences while reading this book. Your visual experiences from a few moments ago cohere with those you are having now. They are not exactly the same because you are reading different words, but they are very similar, as one would expect, if the same external object caused your sensations now as before. Likewise, your tactile sensations now are similar to those you had just a short while ago. You have a very similar tactile sensation of the weight of the book throughout this period (assuming that you have not set the book down, of course). According to the RWH, you have similar tactile sensations throughout this period of time because there is an actual three-dimensional object in your hand the whole time.[6]

Clearly, the RWH explains the relevant features of our sensory experiences quite well. Unfortunately, the mere fact that the RWH provides a good explanation of these features is not enough for the Explanationist Response to successfully defeat UNDERDETERMINATION. In order to provide a reasonable defense of our knowledge of the external world the RWH needs to not only provide a good explanation of these features, it needs to be the *best* available explanation of them. In other words, the RWH must be a better explanation of these features of our sensory experiences than its skeptical rivals. We will now see that it is reasonable to think the RWH is in fact the best explanation of these features of our sensory experiences.

11.3.2 The Superiority of the RWH

The RWH offers a very good explanation of why our sensory experiences have the features that they do. By comparison the classic skeptical alternatives put forward as rivals of the RWH are clearly inadequate. These skeptical alternatives claim that our sensory experiences are caused by a deceptive demon or by our brains being directly stimulated by a supercomputer without providing a story as to how the demon or the supercomputer produces our sensory experiences. Classic skeptical alternatives leave us without much of an explanation at all for why our sensory experiences have the features that they do. They simply tell us that whatever features our sensory experiences do have are the result of the actions of the demon or the supercomputer.

[6]For further discussion of how the RWH explains various features of our sensory experiences see McCain (2014).

Additionally, the RWH allows us to make accurate predictions concerning what our future sensory experiences will be like. For example, the RWH predicts that in a few seconds you will have similar visual sensations to those you are having now because you will be reading the same book. That is, the same external object will be interacting with your sense organs in a very similar way to how it is right now. Classic skeptical alternatives cannot make such predictions though. These skeptical alternatives simply claim that a demon or a supercomputer or whatever causes your sensory experiences; this does not allow for predictions at all (let alone accurate ones). Instead, these skeptical alternatives can only account for your sensory experiences after the fact by claiming that your sensory experiences a few seconds from now are also caused by the demon or the supercomputer. These classic skeptical alternatives cannot match the explanatory power of the RWH. Thus, the RWH is clearly a better explanation of the relevant features of our sensory experiences than the classic skeptical alternatives.[7]

Since the classic skeptical alternatives are clearly inferior to the RWH, it seems the only rivals to the RWH that might be thought to pose a genuine threat to the RWH's superiority are what Jonathan Vogel (1990, p. 660) calls "improved skeptical hypotheses" (ISHs).[8] ISHs are improved versions of the classic skeptical alternatives. These versions of classic skeptical alternatives to the RWH satisfy two constraints. First, ISHs "should invoke items corresponding to the elements" of the RWH (Vogel 1990, p. 660). Second, they "should also posit, as holding of these items, a pattern of properties, relations, and explanatory generalizations mirroring" the RWH's (Vogel 1990, p. 660). Essentially, to generate an ISH we simply "extract the explanatory skeleton or core from the RWH—that there are some entities bearing some properties that are related in ways exactly analogous to those specified by the RWH—and then to add that the entities and their properties are somehow different from the ones mentioned in the RWH" (Vogel 1990, p. 661). ISHs seem to provide some explanation of how the objects of our sensory experiences are related to one another and why they behave as they do because ISHs are isomorphic to the RWH. We will focus on ISHs when evaluating whether the RWH is the best available explanation of the relevant features of our sensory experiences because they are the only skeptical rivals that seem to have any hope of matching the RWH's explanatory power.

Although ISHs perhaps seem to be a match for the RWH when it comes to explanatory power, the RWH is a better explanation of the features of our sensory experiences than these skeptical rivals. Despite the fact that the RWH and ISHs are isomorphic and capable of explaining the same facts, the RWH offers a *simpler* explanation of those facts. All theories, the RWH included, must posit some fundamental explanatory regularities, which cannot be reduced to some other

[7] See McCain (2012) for further elaboration of why these sorts of skeptical alternatives are inferior to the RWH.

[8] See Vogel (1990, 2005) for further reason to think that ISHs are the only skeptical rivals to the RWH that have the potential to match the RWH's explanatory power.

regularity or set of regularities. In any theory some facts have to be taken as brute or else we will be caught in an infinite regress of regularities that can be reduced to other regularities. As we will see, however, the RWH can make use of necessary truths in its explanations in ways that ISHs can only mimic by positing additional contingent fundamental regularities. That is to say, we will see that ISHs can only truly match the explanatory power of the RWH by positing fundamental regularities beyond those required by the RWH. This is very important because, all other things being equal, an explanation which posits fewer fundamental regularities is better than an explanation that posits more. Thus, we will see that the RWH is superior to ISHs because the RWH is simpler in this important way.[9] Since it is superior to its skeptical rivals, the RWH is the best available explanation of the relevant features of our sensory experiences.

In order to see how it is that the RWH is simpler than the ISHs it will be helpful to consider a particular ISH. One commonly presented skeptical rival to the RWH is the "Brain in a Vat" hypothesis. According to this skeptical hypothesis, it may be that we are brains floating in vats with all of our sensory experiences simply resulting from the stimulation of parts of our brains by electrodes connected to a supercomputer. Of course, as it is, this barebones skeptical alternative is clearly inferior to the RWH—it does not offer any explanation of why our sensory experiences have the features they do beyond claiming that the supercomputer causes them to have those features. However, it is possible to construct an improved version of this skeptical hypothesis that at least has a chance of matching the RWH's explanatory power. Let us call the ISH version of the Brain in a Vat hypothesis "BIV".[10] In order to be an ISH, BIV must "invoke items corresponding to the elements" of the RWH, and it must "posit, as holding of these items, a pattern of properties, relations, and explanatory generalizations mirroring" the RWH's (Vogel 1990, p. 660). For example, the RWH posits a number of external objects and causal relations between them and your sense organs to explain your current sensory experiences while you are reading this book. BIV explains those same sensory experiences by positing sequences of code in the supercomputer's operating program corresponding to the various external objects of the RWH and relations holding between these sequences of code which mirror the causal relations of the objects posited by the RWH. With

[9]There are other explanatory virtues we might use to compare the RWH and ISHs. Our focus is only on explanatory simplicity because it is an explanatory virtue that is clearly relevant to determining which theory is best (there is controversy concerning the importance of many other purported explanatory virtues). Additionally, comparing the RWH and ISHs in this manner will be sufficient for us to see that the RWH has features which make it superior to these skeptical rivals. For discussion of other explanatory virtues, see Beebe (2009), Lycan (1988), Quine and Ullian (1978), Thagard (1978), and Vogel (1990). For further discussion of the superiority of the RWH over ISHs and comparisons between the two with respect to additional explanatory virtues see McCain (2014).

[10]The conclusions drawn from the comparison of the RWH and BIV are equally applicable to comparisons of the RWH and other ISHs.

our understanding of roughly how BIV explains various sensory experiences in hand we are ready to compare it to the RWH.

The RWH is simpler than BIV because it can make use of necessary truths in explanations in a way that BIV can only mimic by positing additional contingent fundamental regularities.[11] Since both the RWH and BIV are committed to the necessary truths (these truths hold regardless of whether the RWH or some ISH is true) and BIV has to posit additional contingent fundamental regularities that the RWH does not, the RWH offers the simplest explanation of the features of our sensory experiences.

In an effort to illustrate how the RWH can make use of necessary truths in explanations in a way that BIV cannot it will be helpful to consider an example. Consider a case where you have sensory experiences as of a baseball field in normal conditions. In particular, consider the sensory experiences of first base, third base, and home plate. These three bases visually appear to be arranged in a triangular pattern.[12] According to the RWH, three external world objects which occupy three physical locations in a triangular pattern cause your sensory experiences of these three bases. Given the RWH, a particular necessary truth can help explain why when you have sensations of moving at the same speed your sensation of walking lasts longer when you have the sensation of walking from first base to third base to home plate than when you have the sensation of walking directly from third base to home plate. The relevant necessary truth here is the triangle inequality theorem. The triangle inequality theory states that the sum of the lengths of any two sides of a triangle is greater than the length of the remaining side.[13] According to the RWH, part of the explanation of why your sensations of walking differ in duration is that, according to the triangle inequality theorem, the distance from first base to third base to home plate in this case *must* be greater than the distance from third base directly to home plate.

BIV cannot make use of the triangle inequality theorem in its corresponding explanation of your sensations. The counterparts of the three external objects (the bases) in the RWH that BIV posits do not have genuine locations because they are sequences of code in the supercomputer's software. Since these "BIV-bases" do not have genuine locations, the triangle inequality theorem does not entail that the distance from BIV-first base to BIV-third base to BIV-home plate is greater than the distance from BIV-third base directly to BIV-home plate. Consequently, in order for BIV to explain the differing duration of your sensations when you seem to take different routes between what appear to you to be bases it must posit some contingent regularity which governs the relations between BIV-first

[11] See BonJour (1999), BonJour and Sosa (2003), and Vogel (1990, 2005, 2008) for further support of this claim.

[12] This example is adapted from Vogel (2008, pp. 547–548).

[13] The triangle inequality theorem is normally regarded as axiomatic for Euclidean geometry as well as many non-Euclidean geometries. In fact, in any geometry with a well-defined distance function the triangle inequality theorem holds. That is, the triangle inequality theorem is a theorem in any geometry that is a metric space.

base, BIV-third base, and BIV-home plate in a way that is similar to the triangle inequality theorem. For example, BIV will need to posit something like a regularity that when the sequence of code which produces sensations of walking from base to base is ran that code is paired with specific codes which produce various sensations of duration. This pairing will need to be such that the sequence of code which produces sensations of walking from first base to third base to home plate is paired with code which produces sensations of duration X. The sequence of code which produces sensations of walking directly from third base to home plate will need to be paired with code which produces sensations of duration less than X. Of course, this additional contingent regularity of the code pairings in BIV does not have a counterpart in the RWH. Thus, BIV will have to posit contingent fundamental regularities that the RWH does not in order to explain your sensory experiences in this situation. There are numerous other situations where the triangle inequality theorem (or other necessary truths) does explanatory work in the RWH's explanations that it cannot do in BIV's explanations. Since the set of necessary truths are fundamental regularities of both the RWH and BIV, and BIV has to posit more contingent fundamental regularities than the RWH, the RWH is simpler than BIV in an important way.

In light of the fact that they have the same explanatory power while the RWH is simpler, the RWH is a better explanation of the features of our sensory experiences than BIV. Further, all other ISHs face this same problem. Thus, the RWH is the best explanation of the relevant features of our sensory experiences. Since it is the best explanation of the relevant features of our sensory experiences, our evidence does favor the RWH over the skeptical alternatives. Hence, U2 of UNDERDETERMINATION is false. Thus, although UNDETERMINATION poses more of a challenge to our knowledge of the external world than CERTAINTY, it is a challenge that we can overcome.

11.4 Conclusion

In this chapter we have seen that one way our scientific knowledge can be challenged is by threatening all of our knowledge of the external world. We have also seen that both of the major arguments for external world skepticism are unsound. Consideration of these arguments has revealed some important points about the nature of knowledge though. First, knowledge does not require certainty. Since fallibilism is true, we only need evidence that makes a proposition beyond a reasonable doubt for us in order to know it. Second, the challenge posed by underdetermination of theory by evidence can be expanded into a seemingly powerful attack on all of our knowledge of the external world. Third, our knowledge can be successfully defended from the threat of underdetermination by recognizing that the RWH better explains our experiences than any of its skeptical rivals. Although we have seen that our knowledge, both scientific and mundane, withstands the threat of external world skepticism, other challenges to our knowledge await in the remaining chapters of this part of the book.

References

Beebe, J. (2009). The abductivist reply to skepticism. *Philosophy and Phenomenological Research, 79*, 605–636.

BonJour, L. (1999). Foundationalism and the external world. *Philosophical Perspectives, 13*, 229–249.

BonJour, L., & Sosa, E. (2003). *Epistemic justification: Internalism vs. externalism, foundations vs. virtues*. Malden: Blackwell Publishing.

Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives, 2*, 91–123.

Descartes, R. (1641/1988). *Meditations on first philosophy*. In J. Cottingham, R. Stoothoff, & D. Murdoch (Trans.), *Descartes: Selected philosophical writings* (pp. 73–123). Cambridge, UK: Cambridge University Press.

Greco, J. (2000). *Putting skeptics in their place: The nature of skeptical arguments and their role in philosophical inquiry*. Cambridge, UK: Cambridge University Press.

Huemer, M. (2001). *Skepticism and the veil of perception*. Lanham: Rowman & Littlefield.

Kampourakis, K. (2014). *Understanding evolution*. Cambridge, UK: Cambridge University Press.

Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). New York: Routledge.

Lycan, W. G. (1988). *Judgement and justification*. Cambridge, UK: Cambridge University Press.

McCain, K. (2012). A predictivist argument against scepticism. *Analysis*, *72*, 660–665.

McCain, K. (2014). *Evidentialism and epistemic justification*. New York: Routledge.

McCain, K., & Weslake, B. (2013). Evolutionary theory and the epistemology of science. In K. Kampourakis (Ed.), *The philosophy of biology: A companion for educators* (pp. 101–119). Dordrecht: Springer.

Moore, G. E. (1939). Proof of an external world. *Proceedings of the British Academy*, *25*, 273–300.

Pryor, J. (2004). What's wrong with Moore's argument? *Philosophical Issues, 15*, 349–378.

Quine, W. V. O., & Ullian, J. S. (1978). *The web of belief* (2nd ed.). New York: Random House.

Thagard, P. (1978). The best explanation: Criteria for theory choice. *Journal of Philosophy, 75*, 76–92.

Vogel, J. (1990). Cartesian skepticism and inference to the best explanation. *Journal of Philosophy, 87*, 658–666.

Vogel, J. (2005). The refutation of skepticism. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 72–84). Malden: Blackwell.

Vogel, J. (2008). Internalist responses to skepticism. In J. Greco (Ed.), *The Oxford handbook of skepticism* (pp. 533–556). Oxford: Oxford University Press.

Willenken, T. (2011). Moorean responses to skepticism: A defense. *Philosophical Studies, 154*, 1–25.