# 8
# *An Abductive Response to the Skeptic*

We should be careful to get out of an experience only the wisdom that is in it—and stop there; lest we be like the cat that sits down on a hot stove-lid. She will never sit down on a hot stove-lid again—and that is well; but also she will never sit down on a cold one anymore.

—Mark Twain, *Following the Equator*

## 8.1 The Skeptical Challenge

Explanationists faced the challenge of showing that reasoning abductively in a way that conflicts with Bayesian reasoning could *ever* be rational. In the previous chapters, I have addressed that challenge at some length by showing arguments that allegedly demonstrate the irrationality of abduction to be seriously flawed and, on the positive side, by showing that abduction, although perhaps not universally recommendable, is sometimes an ecologically rational mode of reasoning.

There are still some open issues, however. We demonstrated the ecological rationality of abduction in computer simulations. What about abduction in the real world? In section 3.5, we saw that the participants in the experiments reported in Douven and Schupbach (2015a) tended to be more accurate as they gave more weight, in updating, to explanatory considerations. That is *some* evidence for the real-world value of abduction. We also discussed some successes of abductive reasoning in section 2.1. Showing that abduction is a reliable mode of inference in a broad range of contexts, in that judgments of explanatory superiority reliably point in the direction of the truth in those contexts and hence that giving weight to such judgments often tends to make us more accurate, might just be a matter of going beyond anecdotal evidence of the kind examined in section 2.1 and launch a systematic empirical investigation. For instance, it would probably not be very complicated to see whether the extent to which medical doctors rely on explanatory considerations in their work correlates with their success rates, a project that could show the real-world significance of the findings reported in sections 6.3 and 7.4.

There is an additional hurdle, mentioned in section 6.2, which is especially pressing because it concerns the question of whether it is rational to reason abductively in contexts in which philosophers have hoped that abduction could make a decisive dialectical difference. I refer to contexts in which abduction is supposed to break a deadlock due to underdetermination by the evidence, as explained in section 2.3. To revisit the hurdle in these contexts, consider that while scientific realists like to claim that Thomson was led to the discovery of the electron via abductive reasoning—the existence of the electron best explained the phenomena he had observed in his cathode tube experiments—antirealists deny realists the right to speak of a discovery here. And antirealists will not be impressed by our argument for the ecological rationality of abduction, because as they would rightly point out, that argument did not concern the kind of abduction relevant to the realism debate. Indeed, it did not concern any debate in which abduction is supposed to justify existence claims going beyond what at least one party to the debate holds that we have direct epistemic access to.

In this chapter, I aim to show how we can defend abduction in a context in which such a defense is arguably hardest to accomplish, to wit, in the debate about Cartesian skepticism, according to which we must remain agnostic about the existence of an external world. After all, in that debate one faces an opponent who is unwilling to grant one almost anything, unlike, for instance, in the debate with the scientific antirealist, in which it is typically possible to find some common ground that one might then hope to leverage in an attempt to convert one's opponent.

More specifically, Cartesian skeptics (henceforth simply "skeptics") are typically opposed to admitting as evidence anything that goes beyond the purely phenomenal, and equally typically, they disown the use of any inferential principles, including any form of abduction, that might enable one to move from premises about the phenomenal alone to a conclusion about the external world. Not all nonskeptics may find this troublesome. Although the aforementioned facts may stand in the way of a *resolution* of the skepticism debate, which I understand to require that the debate be settled in a way that is acceptable to both parties involved, those facts need not bar an *argument* against skepticism. Supposing a less restrictive conception of evidence than the skeptic's or stronger rules of inference than she is willing to countenance may enable one to argue on empirical grounds for the external world hypothesis. Such an argument may leave the skeptic cold and therefore may not bring any nearer a resolution of the debate, but if we nonskeptics can convince or reassure ourselves of the tenability of our own position, then what further good is there to be gained by convincing *others* of it? As Psillos observed (see section 6.2), in philosophy the goal is not always to convince one's opponent and can instead be to justify one's position to one's own satisfaction. Philosophy is not politics.[1]

The chief aim of this chapter is to show that a resolution of the debate becomes possible once it is recognized that the skeptic embraces overly restrictive and ill-motivated positions vis-à-vis both evidence and inference, and that more reasonable alternatives are available and that embracing those alternatives will crucially, and inter alia, allow us to commit to abductive reasoning. Naturally, for this aim it will be imperative to make clear that these alternatives should appear more reasonable from the skeptic's point of view. A secondary aim is to show that the main aim is likely to be more important than may appear to those who believe that all the nonskeptic should care about is an argument against skepticism; as I argue, whether or not nonskeptics *need* no more than such an argument, it is quite unclear that they are in a better position to give one than to give a resolution of the skepticism debate.

Focusing on the latter aim first, I begin in section 8.2 by considering two well-known arguments against skepticism that rather explicitly do not seek to resolve the debate, and by arguing that one of them assumes a conception of evidence and the other a conception of inference that the skeptic can reasonably reject, and that should seem problematic even from the nonskeptic's perspective. In section 8.3, I argue that the skeptic's views on evidence and inference are problematic as well, however, and I propose alternative views that differ from both the skeptic's and those supposed by the aforementioned

1. For other examples of this line of thinking see Pryor (2000) and Byrne (2004). In Pryor's terms, it is enough for nonskeptics to pursue the "modest anti-skeptical project," which requires that we "establish to our satisfaction that we can justifiably believe and know such things as that there is a hand" (Pryor, 2000, p. 517).

arguments but that should by anyone's lights appear more reasonable than either. Finally, section 8.4 presents a convergence result showing that given those more reasonable views, it is at least in principle possible to resolve the skepticism debate.

## 8.2 Mooreans and Russellians against the Skeptic

To see that an antiskeptical argument is not automatically pointless if it presupposes a more liberal conception of evidence than the skeptic's, recall G. E. Moore's (1962) famous "proof of an external world." A common initial reaction to this argument is that in it Moore is begging the question by assuming that he has perceptual evidence of his having two hands. But note that one may admit this yet insist, as some argue Moore himself does, that by making the said assumption he is no more begging the question against the skeptic than the skeptic is begging the question against her opponent by assuming that our evidence is limited to the purely phenomenal. Suppose that this can be successfully maintained. Then the least Moore's argument would seem to yield is a kind of standoff between the skeptic and the nonskeptic. And given that to many the skeptic so clearly seems to have a winning hand—undeservedly, it is typically added—it would be a remarkable achievement to show that the skeptic in any case does not win the debate (even if he does not lose it).[2]

Some nonskeptics might want to go further and claim that considerations of question-begging should not detain us here. All that matters, they might say, is that the premises and presuppositions of Moore's argument are acceptable to the nonskeptic; whether Moore is begging any skeptical questions or whether, if he does, this is offset by the skeptic's begging nonskeptical questions in turn, is neither here nor there. To these nonskeptics' eyes, Moore has offered us a successful argument against skepticism, and it is none the worse for failing to persuade the skeptic.

2. Epistemic conservatives could even argue that inasmuch as the one position is as good (or bad) as the other, dialectically speaking, the sensible thing for us to do is not to abandon what we believe anyway—namely, that we do have lots of interesting knowledge, including knowledge of an external world—in favor of a position that no one has ever seriously held (as opposed to temporarily entertained for philosophical purposes). But epistemic conservatism is controversial; see, e.g., Christensen (1994) and Vahid (2004).

Much the same could be said about the Russellian response to skepticism, which argues for the external world hypothesis on the basis of this hypothesis's explanatory power.[3] The best explanation of our perceptual experiences and of the complex patterns and regularities manifest in those experiences is that the *apparent* causes of our experiences—mostly external objects—typically are their *real* causes. Therefore, we are justified in believing that there is an external world. The response assumes, patently, that the external-world hypothesis *is* the best explanation of our experiences. Some authors seem to think of this assumption as immediately compelling. Others, most notably Laurence BonJour and Jonathan Vogel, have argued for it extensively—and quite ingeniously, to my mind.[4] While this is not to say that the assumption is entirely undisputed, I am willing to grant it here.[5] Note that this is far from enough to settle the skepticism debate. The Russellian response also relies on some qualitative version of abduction, and why should skeptics want to sanction abduction in *any* of its forms? Even so, the Russellian response is not necessarily a failure. It, too, might achieve a standoff between the skeptic and her opponent, and for reasons previously mentioned, that would be of considerable significance.

Some might even think that whether the Russellian response achieves a standoff is immaterial; what is important is simply that the argument is acceptable from a nonskeptical perspective, however many skeptical issues it may beg and however few nonskeptical issues the skeptic may have to beg in order to defend herself against it (if she has to beg any at all).

But it is not at all obvious that the Moorean and the Russellian responses to the skeptic *are* successful in either of the senses just explained, that is, whether they yield standoffs with the skeptic or at least arguments that are acceptable from a nonskeptic's viewpoint. In fact, the following considerations give some grounds for thinking not merely that the skeptic has more right to her position regarding evidence than the Moorean has to hers, as well as more right to reject certain forms of abduction, or indeed any form, than the Russellian has to endorse abduction generally—which seems to preclude that

3. I call it "Russellian" because the first response of this type (or at least the first I know of) is to be found in Russell (1912, ch. 2). More elaborate and refined versions of it have been propounded since, however; see, for instance, Harman (1973, chs. 8, 11), Goldman (1988, ch. 9), Vogel (1990, 2005), and chapter 5, written by BonJour, of BonJour and Sosa (2003).

4. See Vogel (1990, 2005) and BonJour and Sosa (2003, ch. 5).

5. Sosa (1997, pp. 413–414), for one, disputes it.

the foregoing arguments yield standoffs—but also that the Moorean and Russellian positions, and concomitantly the Moorean and Russellian responses to skepticism, should, upon reflection, be unappealing to the nonskeptic herself.

Whereas the skeptic may grant that we have an instinctive trust in our senses, she will deny that our instincts are worthy of trust on this point, and she can seemingly do so without having to commit herself to any implausibly strong assumptions. In particular, her denial need not involve anything as strong as the principle that we ought to have good reasons for trusting our senses, or perhaps even know that they are trustworthy, before we can trust them. To defend her claim, the skeptic could make do with what appears to be a rather minimal and unproblematic assumption, namely, that it is perfectly all right for us to trust our senses, *provided* that we have no reasons for *doubting* their trustworthiness. After all, she will have no difficulty pointing to such reasons. For example, we probably all have had conflicting experiences of the sort in which what at one moment appeared to be a bent stick the next moment appeared to be a straight one, or in which what at first appeared to be a sheep later appeared to be a rock. Such experiences fall far short of *establishing* that our senses are untrustworthy: trustworthiness does not require infallibility. Yet even the Moorean will have to concede that they do give us *some* reason to doubt our senses' trustworthiness and concomitantly the admissability of sentences such as "I have two hands" as possible evidence reports.

Note, that the skeptic can thus grant a lot to nonskeptics. For instance, nothing seems to prevent her from concurring with Crispin Wright (1991) that barring evidence against it, we have an unearned warrant to believe the hypothesis that sensory perception is reliable. Also, she can treat that hypothesis "liberally" in Jim Pryor's (2004, p. 354) sense, meaning that the hypothesis can be among the conditions that justify one defeasibly in believing a given proposition $A$ even if one lacks antecedent justification for the hypothesis, although, importantly, evidence against the hypothesis would undermine one's defeasible justification for $A$. So, it seems that a skeptic who argues along the lines suggested in the text is not threatened by antiskeptical responses—such as Wright's and Pryor's but also Alex Byrne's (2004)—which proceed on the assumption that the skeptic must be assuming that if we are to know anything via perception then we must antecedently know, or at least be antecedently justified in believing, that perception is reliable. As may be noted, I am assuming that Wright and Pryor would count the reasons for doubting the trustworthiness of our senses as (possibly inconclusive) evidence against the hypothesis that our senses are trustworthy. Although they do not discuss in any detail the matter of possible counterevidence, they would seem hard-pressed to give a generally acceptable account of evidence on which reasons of the kind mentioned in the text do *not* count as evidence against the said hypothesis. (On the popular Bayesian conception of evidence, which we also assume in the following sections, it would seem near to obvious that those reasons count as evidence against that hypothesis.)

Again, basically the same that was just said about the Moorean response can be said with respect to the Russellian response. It has happened more than once that I have assumed something because it furnished the best explanation for things that I knew to be the case only to find out later that the assumption was false. I am not an exception in this respect, I am sure. In fact, similar failings of abduction are known from the history of science. In section 2.1, we discussed that although from the viewpoint of nineteenth-century astronomy the small deviation from Newtonian predictions of Mercury's orbit around the sun was best explained by assuming the existence of a yet-to-be-discovered planet, that planet never was discovered, and the deviation is now generally thought to be explained by the General Theory of Relativity as due specifically to the curvature of space-time. Could Russell or anyone else be said to have as much right to endorse a mode of inference whose reliability has come to appear at least somewhat questionable as the skeptic has a right to disavow that rule? In order to back up a negative answer to this question, all the skeptic would have to assume is a principle according to which it is all right to reason abductively, as long as no reasons for doubting abduction have arisen.

At this juncture, some might want to reply that there is an inconsistency in the skeptic's position because she can raise her concerns about the reliability of sensory perception and abduction only by presuming sensory perception to be reliable. ("How else could she claim with any confidence that our senses as well as abduction at least sometimes fail us?") This would be wrong, however. As to sensory perception, note that one can register to have had conflicting perceptions without being able to sort the veridical ones (if such there be) from the nonveridical ones; that there were conflicts between them is enough to aver that at a minimum, they cannot all have been veridical. As to abduction, suppose that sensory perception is *not* reliable. Then that itself is a failure of abduction. For if, as the Russellian assumes, our perceptual experiences are best explained by assuming that their apparent causes are their real causes, then abduction justifies us in believing that sensory perception is reliable (that is precisely what the reliability of our senses amounts to—that typically the apparent causes of our perceptual experiences are their real causes). So, the reliability of abduction is called into question either way, whether it is assumed that sensory perception is reliable or whether it is assumed that sensory perception is not reliable.

A more sensible reply would point to the fact that, as we have briefly noted, occasional failures are compatible with reliability (just not with infallibility). Thus, to start with abduction, it might be that, its registered failures notwithstanding, this rule is reliable. Some externalists might hold that if abduction is in fact reliable, then inferring to the best explanation is *ipso facto* a legitimate procedure even in the face of doubts with regard to the rule's reliability. I am not sure whether this hard-nosed kind of externalism has any actual sympathizers. If it does have them, then I would expect this to be due to the fact that some believe it to be our only hope for resisting skepticism. As it follows from this chapter that this is not so, I set aside the designated externalist position without a hearing. That is not necessarily to embrace internalism; for all I need to assume, it may be epistemically unassailable to infer to the best explanation even lacking any reason whatever to believe that abduction is reliable. Nor is it necessarily to say that once failures of the rule have come to our attention we can nevermore utilize it, or at least not legitimately so. What I do think is that on any plausible epistemology, be it internalist or externalist, once such failures have come to our attention, we face the task of showing, or at least making plausible, that these failures are mere accidents and not symptoms of a systematic defect before we can utilize abduction or utilize it again. But how to do this?

It seems that unless we somehow *define* truth in terms of (best) explanation—an option that we have already rejected—the only way to show that abduction is reliable, if it is, is by means of an empirical argument. At first glance, such an argument may not seem too hard to come by. Are we not often enough in a position to verify a conclusion that we had previously reached on the basis of explanatory considerations? Newton was *right* about the Dutch having beaten the British. Le Verrier was *right* about what caused the perturbations in Uranus's orbit. I and doubtless many others were *right* about the head and the headless body found in Jackson belonging together. These conclusions had all been reached abductively. Most of the pertinent evidence may in fact be of a more humdrum type. Suppose that you conclude that your neighbors are back from vacation, as that would best explain the presence of their car in the driveway. If soon afterward you see them busy in the garden, it is very tempting to classify this as a success of abduction, lending some credibility to the hypothesis that abduction is a reliable mode of inference. Still further support for that hypothesis could easily accrue from a more systematic examination of applications of abduction.

There is a problem, however, that we already noted in the introduction to this chapter. As we can now put the problem, it is that an inductive approach as just sketched relies on a Moorean conception of evidence. And so far it is still an open question whether the nonskeptic can embrace that conception with as much right as the skeptic can reject it.

Some may think that an empirical argument for abduction need not rely on a Moorean conception of evidence. Indeed, it seems that we could restrict our evidence to sense data (or impressions or appearances or something else) and use abduction to arrive at predictions about future sense data. For instance, given that it appears to me that my neighbors' car is in their driveway, I predict that it will soon appear to me that I am seeing my neighbors. If time and again such predictions are borne out, then would that not be enough to warrant the conclusion that abduction is reliable? The problem is that this would at most warrant the conclusion that abduction is reliable when it is applied to sense data statements for obtaining justification in further sense data statements and not that abduction is reliable when it is applied to sense data statements for obtaining justification in external world statements. And eventually we want the latter.

To return, then, to the question of whether the nonskeptic has as much right to rely on a Moorean conception of evidence as the skeptic has to reject it, I expect that with respect to our senses' reliability, too, all but perhaps the most radical externalists will agree that once doubts about that reliability have arisen, we face the task of showing, or at least making plausible, that our senses are in fact reliable before we can rely on them or rely on them again. Naturalized epistemologists have long been urging that on this matter we turn to psychologists for assistance, and it seems that so far no one has been able to come up with a more promising idea. But in trying to answer the question of whether sensory perception is reliable, psychologists make claims of the following sort (and now I am exaggerating but not essentially so): "I put up my hands ten times before the subject's eyes and in all cases it seemed to the person that there were hands before her." That research practice is perfectly warranted from their viewpoint because they are not engaged in a debate with skeptics. From our present philosophical viewpoint, however, the practice is *not* warranted: by judging that we put up our hands before the subject we are unarguably relying on our senses, which is to violate the previously stated principle that in the face of doubts about their reliability we should suspend such reliance. That is to say that by admitting the said kind of claims as evidence we are begging the question, not just against the skeptic but against all except perhaps advocates of the variety of externalism that I have set aside for our purposes.

The foregoing considerations exhibit a seemingly telling asymmetry between the skeptic and the Moorean and also an asymmetry between the skeptic and the Russellian. They show that the skeptic can raise doubts about the Moorean conception of evidence as well as about the Russellian conception of inference, that her opponents will have to acknowledge as such and that, for all they have hitherto shown, cannot be removed in a non-question-begging way; and for all that we have been told so far, the Moorean and Russellian are not in a similar position to discredit the skeptic's positions regarding evidence and inference. That seems enough to conclude that the skeptic has a greater right to her positions vis-à-vis evidence and inference than the Moorean and Russellian have to theirs—on *any* reasonable understanding of what it is for a party to have greater right to a position than another party has to a contrary position. As a result, it seems hard to maintain that the Moorean and Russellian responses to skepticism succeed in realizing a standoff with the skeptic. Moreover, given that, as we have said, anyone except (perhaps) the most extreme externalists may be expected to concur that once doubts about the reliability of our senses or of abduction have arisen they must first be removed before we can rely (or rely again) on these faculties, and given that so far no one has been able to think of a way for removing such doubts that does not assume the Moorean or Russellian conceptions of evidence and inference, respectively, it seems that these conceptions of evidence and inference, and thereby the Moorean and Russellian responses to skepticism, should appear highly problematic even from the perspective of the nonskeptic.

For all we presently know, what goes for the Moorean and Russellian conceptions of evidence and inference goes for every other nonskeptical conception of evidence or inference that anyone might ever seriously consider. As a result, it is far from clear whether there could be *any* conceptions of evidence and inference that in the final analysis the nonskeptic, but not the skeptic would be comfortable to adopt. It thus seems that nonskeptics cannot be quite so confident as some of them appear to be that there exist any successful arguments against skepticism that do not at the same time resolve the debate.

What *is* clear is that if both skeptics and nonskeptics are left with the skeptical conceptions of evidence and inference as their only options, then we must all end up being skeptics. For it has been long known that starting from these conceptions, one just has too slender a basis for the construction of *any* successful argument for the external world. As I further argue, however, other conceptions of evidence and inference *are* available—conceptions that should seem preferable to all hands, including the skeptic.

## 8.3 Evidence, Inference, and Expert Functions

If the registered fallibility of our senses should be of epistemic significance and should discomfit the Moorean position vis-à-vis evidence, one still wonders how it could motivate the skeptic's drastic retreat to a sense data conception of evidence. In fact, the Moorean and the skeptical responses to recorded failures of our senses both seem disproportionate. Once such failures have come to our attention, we cannot, contra the Moorean, simply ignore them and go on trusting our senses undeterred. But by distrusting our senses altogether or at least by effectively treating them as though they are to be distrusted, skeptics seem to err in the opposite direction; they seem to overreact. Speaking figuratively, paraphrasing this chapter's epigraph, it seems that the skeptic, unlike the Moorean, is willing to get out of the experienced failures of our senses the wisdom that is in them. Unfortunately, she does not stop there; rather she is like the cat who once sat on a hot stove-lid and now will never again sit even on a cold one.

By way of analogy, consider how ordinary people react when they realize that their memories have betrayed them, that something they seemed to remember with great clarity was actually false, such as when they were certain that they had attended some party but then later discovered that they simply could not have been there. I think that no reasonable person ever concluded on the basis of such failures that she would better remain agnostic about everything she seems to recollect from memory. Nor, however, is it likely that after she has been confronted with its fallibility, such a person would go on trusting her memory in exactly the same way or to exactly the same extent as she did before that event. What such failures *are* likely to occasion in her, I believe, is a sense of caution vis-à-vis her memories, a kind of metacognitive awareness that her recollections are not fully dependable. I, for one, do trust my memory, but due to experiences of the aforementioned kind I do not trust it 100 percent but only about 95 percent. That does not make me an agnostic about my recollections, although it does mean that for me they no longer have the status of infallibility that I once thought could be accorded to them.

The skeptic may be expected to respond with something like the following: "Evidence ought to be absolutely infallible: anything accepted as evidence should be guaranteed to be true. This requires a sense data conception of evidence, for our senses may betray us but we cannot go wrong in judging how things appear to us. So I'm not overreacting: cats that want to be sure that they will never sit on a hot stove-lid again do well not to sit on *any* stove-lids." To the best of my knowledge, however, no skeptic has ever succeeded in producing a convincing argument for this absolute infallibility requirement. It might be asserted that the requirement is plausible on intuitive grounds alone, so that no argument for it is needed. Yet the intuition (if it is one) seems hardly decisive on its own.

For consider that another intuition that we have about evidence is that evidence ought to be informative and useful to at least some extent. Why invest the huge amounts of time, effort, and money in gathering evidence that we do invest, when we get nothing useful in return? This means trouble for the skeptic if Hans Reichenbach is right that the requirement that evidence be informative (however minimally) is at odds with the requirement that evidence be infallible. As most other logical empiricists, Reichenbach had long thought that the only sentences that we could not wrongly judge to be true, and therefore the only sentences that could serve to report evidence, were of the sort of "I have the impression of there being a table in front of me." But in *Experience and Prediction* (Reichenbach, 1938) he came to reconsider this view, and in chapter 3 of that book he argues convincingly that it is *false* that we cannot be mistaken about such sentences and that we can make our "basic statements"—the sentences that we accept as evidence reports—totally secure only by making them devoid of any content and thereby depriving them entirely of any use they might otherwise have. At bottom, the problem as Reichenbach describes it is that in making the judgment that I have the impression of there being a table in front of me, I am relying on my memory in classifying the impression as being that of a table. I am thereby implicitly comparing it with previous impressions I have had. But my memory may be confused, and the type of impression that I am now having might yesterday have made me judge that I was having the impression of there being a chair, or whatever, in front of me. As my judgment that my memory is not right now confused in this way is fallible, so is my judgment that I am currently having the impression of a table. To circumvent this problem, Reichenbach points out, one would have to stop classifying one's impressions in ways which rely on one's memory, but in so doing one's evidence would become utterly uninformative.[6]

Skeptics could retort that whereas the preceding may show that the requirement of infallibility cannot be maintained unimpaired, we can still insist that evidence should be as secure as is compatible with the requirement that it should not be entirely useless. She could then claim that because, for instance, "It appears to me that there is a table in front of me" is not entirely empty of content but is more secure than "There is a table in front of me," the attenuated requirement still favors the skeptical position regarding evidence over the Moorean one.

However, conceding that a trade-off will have to be made between certitude and informativeness almost automatically invites the question whether we should not be willing to sacrifice a little more certitude for some additional informativeness. The intuitively plausible answer might seem to be *yes*, but then the problem is that we will somehow have to balance the different values involved, and that it is unclear how to do this in a principled way.

This much does seem clear: if we were fully convinced that our senses are reliable, then it would be eminently reasonable for us to admit "I have two hands" and similar sentences as possible evidence reports, just as the Moorean does, because thereby we would seem to be making an excellent deal in which a little loss of certitude is traded for a significant increase in informativeness. If by contrast we were fully convinced that our senses are *un*reliable, then it would be eminently reasonable for us to restrict the possible evidence sentences to sense data statements or some such, just as the skeptic does, because thereby we would seem to get a lot of certitude at the cost of no significant loss of informativeness.

As intimated previously, however, we must be honest enough in our response to the skeptic to admit that there are reasons to doubt the reliability

6. For a more recent and more general argument for the unreliability of our introspective judgments, see Schwitzgebel (2008).

of our senses. Acknowledging as much seems incompatible with being fully convinced that our senses are reliable—at least as long as we have not been able to neutralize those reasons. On the other hand, whereas the registered failures of our senses may be sufficient to shake our instinctive trust in our senses, they hardly suffice to *establish* that our senses are unreliable. As a result, at the present stage of inquiry it would be as unreasonable to be fully convinced of the reliability of our senses as it would be to be fully convinced of their unreliability. What to do in this situation?

Someone might conceivably be able to devise a method that would allow us to choose in a principled manner between, for instance, "I have two hands" and "It appears to me that I have two hands" as possible evidence reports, depending, perhaps among other things, on our degree of conviction that our senses are reliable. Although it is, before the actual specification of such a method, hard to say anything definitive about it, I suspect that it could never be wholly satisfactory because however we should opt according to the method, it would always seem that in going along with the skeptic we would be according the registered failures of our senses too much weight, whereas in going along with the Moorean we would be according them too little. Be this as it may, a satisfactory solution to our problem does seem to be available. Here, I propose an account of evidence that allows us to assume a sort of middle position between the skeptic's and Moore's, and thereby to get out of the registered fallibility of our senses only the wisdom that is in it and stop there.

The proposal makes use of the notion of uncertain evidence. Let $E$ be some sentence that according to the Moorean can report evidence, such as the sentence "The sun is out." Then $e$ is the sentence "It appears that $E$," that is, in the present case the sentence "It appears that the sun is out." Now suppose that it comes to appear to you that the sun is out. Then, presumably, if you were a Moorean, you would take $E$ to report your evidence. If instead you were a skeptic, you presumably would take $e$ to report your evidence. But how should you express the evidence that you are receiving if you think—as, we said, would be reasonable to do—that the Moorean is not being quite sensitive enough to the fallibility of our senses but the skeptic is being oversensitive in this respect? The following would seem a natural suggestion in that case: "I am uncertain about what exactly the evidence is I am receiving. Were I a Moorean, I would say that it is that the sun is out, but were I a skeptic, I would say that it is that it *appears* that the sun is out." Abbreviate this "uncertain evidence report" (as we may call it), or any other to the same effect, by "$E/e$." The important question now is how such a report could be useful and, more specifically, how one could learn anything from it.

The proposed answer to this question draws on the machinery of Bayesian epistemology. As discussed, Bayesians hold that we ought to accommodate new evidence via the rule of conditionalization, which dictates that a person's new probability for a given sentence $A$ after she has come to learn another sentence $B$ (and nothing else) be equal to her earlier probability for $A$ conditional on $B$. The rule is silent about how to proceed with the kind of uncertain evidence considered here. However, two relatively recent additions to Bayesian epistemology help us on this point.[7]

The first is the theory of expert functions as developed by, among others, Haim Gaifman (1986) and van Fraassen (1989). In van Fraassen's presentation of the theory, the central definitions are the following:

**Definition 8.1** *Let* Pr *be person $P$'s probability function. Then $q$ is said to be an* expert function *for $P$ concerning family $\mathcal{F}$ of sentences iff, for all $A \in \mathcal{F}$,* $\Pr\big(A \mid q(A) = x\big) = x$.

Notice that it follows from this and the standard ratio definition of conditional probability that if $q$ is an expert function for $P$ concerning $\mathcal{F}$, then for all $A, B \in \mathcal{F}$: $\Pr(A \mid B) = q(A \mid B)$.

**Definition 8.2** *Let persons $P_1, \dots, P_n$ have probability functions* $\Pr_1, \dots, \Pr_n$. *Then $P_1, \dots, P_n$ constitute a* panel of experts *for person $P$ on $\mathcal{F}$ iff there are functions $q$ and $w$ such that (1) $q$ is an expert function for $P$ in the sense of*

7. Because considerations similar to those that led us to doubt both the Moorean and the skeptical conceptions of evidence led Jeffrey to the generalized rule of conditionalization that now goes by his name (see section 3.2), one might think that Jeffrey's rule could be of help in solving the problem at hand. Perhaps it can; see the first appendix in Douven (2005). However, the theory of expert functions subsequently appealed to allows us to state a rule for updating on uncertain evidence that *is* mathematically very similar to Jeffrey's rule but has a bit more structure; and the extra structure enables us to relate an update on uncertain evidence directly and in a seemingly natural way to the probability that is most directly related to the uncertainty of the evidence, namely, the probability that sensory perception is reliable. Another reason for preferring the theory of expert functions to Jeffrey's rule has to do with unity of presentation. As shown further on, the theory of expert functions can be utilized to model one's accounting for explanatory considerations in learning from evidence when one is uncertain about how much weight to give to such considerations; I do not see how Jeffrey's rule can be used for that.

*definition 8.1, and (2) w is a weight function such that* $q(A) = \sum_{i=1}^{n} w(P_i) \Pr_i(A)$ *for all* $A \in \mathcal{F}$.

To say that $w$ is a weight function is to say, first, that $w(P_i) \geqslant 0$ for $i = 1, \ldots, n$, and, second, that $\sum_{i=1}^{n} w(P_i) = 1$. In principle, this function could stand for anything, although it is most naturally interpreted as measuring person $P$'s confidence that the respective expert is right about the sentences in $\mathcal{F}$; for instance, if $w(P_i) = w(P_j)/2$, then she trusts in expert $P_i$ only half as much as she trusts in expert $P_j$.

A consequence of definition 8.2 is that whereas $P$'s probability for $A$ conditional on $B$ (for arbitrary $A, B \in \mathcal{F}$) will be some average of the various $\Pr_i(A \mid B)$'s, it will in general not be the average as weighted by the weight function $w$ but one weighted by some other weight function $w_B$, which may differ from $w$, as $B$ may provide evidence for or against the expertise of the various members. (Naturally, the functions *need* not differ: $B$ may be irrelevant to the question of the members' expertise.)

The second addition to Bayesian epistemology that we need consists of so-called counterfactual probabilities, which were proposed by philosophers of science in the 1980s in the context of what is now commonly known as "the problem of old evidence." Roughly, the problem is that as a simple consequence of the probability calculus, no data that are already known at the time that a theory is conceived can confirm that theory, which seems intuitively wrong. To do justice to our intuitions, various authors have proposed that confirmation by known data is to be measured relative not to our actual probabilities but to the counterfactual probabilities that we would have had if the data had been unknown to us.[8] To see what role counterfactual probabilities play in our account, suppose that it appears to you that the sun is out. Even though you are unsure, we still assume, how to express the evidence you are receiving, it seems not difficult at all to determine your probability for a given $A$ conditional on your present evidence $E/e$ on the counterfactual supposition that you endorse a Moorean conception of evidence nor to determine your probability for $A$ conditional on $E/e$ on the counterfactual supposition that you share the skeptic's standpoint regarding evidence. Let us use "$\mathfrak{m}(A \mid E/e)$" to designate the former, and "$\mathfrak{s}(A \mid E/e)$" to designate the latter. Because you share neither the Moorean nor the skeptical standpoint

8. See, e.g., Horwich (1982) and Howson (1984); also Sprenger and Hartmann (2019, ch. 5).

regarding evidence, these are not actual probabilities but counterfactual ones. In a sense, this is like having at hand expert opinions from two experts who hold radically different views about which sentences are suitable for reporting evidence. The obvious idea now is to apply the theory of expert functions to these counterfactual probabilities and make your conditional probability $\Pr(A \mid E/e)$ for $A$ given $E/e$ a weighted average of $\mathfrak{m}(A \mid E/e)$ and $\mathfrak{s}(A \mid E/e)$.

The next question is how the weights are to be determined. A plausible suggestion is that the weights should reflect your confidence in the reliability of sensory perception. Specifically, it seems that to the extent that you believe sensory perception to be reliable, you will agree with the Moorean on what can be accepted as evidence and concomitantly think it correct to set $\Pr(A \mid E/e) = \mathfrak{m}(A \mid E/e)$. Equally, to the extent that you believe sensory perception to be unreliable, you will agree with the skeptic that evidence is restricted to the phenomenal and concomitantly believe it to be correct to set $\Pr(A \mid E/e) = \mathfrak{s}(A \mid E/e)$. So, if we call the hypothesis that sensory perception is reliable "RSP," the natural thought is that the weights should be determined by your probability for RSP. If we want to follow strictly the theory of expert functions, we cannot simply let the weights be Pr(RSP) and Pr(¬RSP): according to this theory we should be sensitive to the fact that new evidence may alter the confidence that you have in the "experts." This requirement is satisfied by letting $\Pr(\mathrm{RSP} \mid E/e)$ and $\Pr(\neg\mathrm{RSP} \mid E/e)$ be the weights. Thus, on the current proposal conditional probabilities on uncertain evidence of the kind that we are considering are to be determined by the following formula:

$$\text{(8.1)} \quad \Pr(A \mid E/e) = \Pr(\mathrm{RSP} \mid E/e)\,\mathfrak{m}(A \mid E/e) + \Pr(\neg\mathrm{RSP} \mid E/e)\,\mathfrak{s}(A \mid E/e).$$

Once it is clear how to determine probabilities conditional on uncertain evidence, we can conditionalize on such evidence.

Naturally, (8.1) cannot be fully general, for it cannot provide a way for determining one's probability for RSP conditional on some piece of uncertain evidence. This may seem to be a problem, given that we need such conditional probabilities in order to apply (8.1). But the problem is only apparent. We proposed to formulate uncertain evidence reports in the way that we did because we wanted to have a type of evidence that does not presuppose an answer, one way or the other, to the question of the correctness of RSP. It seems that precisely because of this deliberately built-in neutrality as regards

RSP, pieces of uncertain evidence cannot be telling either for or against that hypothesis. That is to say, it seems that we can set Pr(RSP | $E/e$) equal to Pr(RSP) for all uncertain evidence $E/e$ (and thus also set Pr(¬RSP | $E/e$) equal to Pr(¬RSP)). I should add that even if this easy solution were unavailable, (8.1)'s not being entirely general could not be held against it, at least not if we accept the theory of expert functions. For we cannot defer to experts for *all* our opinions; eventually we will have to provide answers ourselves to such questions as how we should weigh the various experts relative to one another and whether it would not be better to consult an additional expert. But this could not be a valid objection to the theory: that we are not always in a position to consult experts does not make expert advice less valuable when it can be had.

Because the preceding account of learning from uncertain evidence may seem fairly abstract, I briefly illustrate it with the aid of an example which also prepares for the use that we subsequently make of (8.1).

**Example 8.1** Call the hypothesis that abduction is a reliable mode of inference "RABD." I now show how uncertain evidence may affect one's probability for RABD, provided that one accepts (8.1). To that end, let Pr be a given person's probability function, where this person accepts (8.1) and where $\mathfrak{m}$ and $\mathfrak{s}$ are the probability functions of her imaginary "Moorean" and "skeptical" experts, respectively; just for purposes of illustration, say that Pr(RSP) = Pr(RABD) = .5, $\mathfrak{m}$(RABD) = .7, and $\mathfrak{s}$(RABD) = .3. Furthermore, let $E/e$ be the following evidence report: "Were I a Moorean, I would express the evidence that I am presently receiving as 'My neighbors are home'; were I a skeptic instead, I would express it as 'It appears to me that my neighbors are home'." Now suppose that our person had earlier deemed that the presence of the neighbors' car in their driveway, or at least the seeming presence of the seeming neighbors' seeming car in the seeming driveway, was best explained by the supposition that they were home. Then, we may assume, at least from her imaginary Moorean expert's viewpoint, $E/e$ counts as a success of abduction and thus as supporting RABD: application of the rule has led, or would have led, the person to conclude that her neighbors are home, and now, from the said viewpoint at least, the evidence vindicates that conclusion. Say that $\mathfrak{m}$(RABD | $E/e$) = .8. From the viewpoint of the imaginary skeptical expert, on the other hand, $E/e$ will at least not be negatively relevant to RABD: that her neighbors *appear* to be home may not confirm what the person had earlier judged to be the best explanation of the seeming presence of their car in the driveway—namely, that the neighbors *are* home—but the evidence certainly does not show that conclusion to be false nor, we may assume, does it render the conclusion less likely. For simplicity, say that, from the skeptical expert's viewpoint, the evidence is irrelevant to RABD, that is, $\mathfrak{s}(\text{RABD} \mid E/e) = \mathfrak{s}(\text{RABD})$. Finally, for the reason given in the preceding paragraph, we may set $\Pr(\text{RSP} \mid E/e)$ equal to $\Pr(\text{RSP})$. Thus, if our person comes to accept $E/e$ and conditionalizes on it, then her probability for RABD will go up from .5 to .55, for $\Pr(\text{RABD} \mid E/e) = \Pr(\text{RSP} \mid E/e)\, \mathfrak{m}(\text{RABD} \mid E/e) + \Pr(\neg\text{RSP} \mid E/e)\, \mathfrak{s}(\text{RABD} \mid E/e) = .5 \times .8 + .5 \times .3 = .55$. ♦

Next let us turn to the question of the confirmation-theoretic status of abduction. Just as conceding to the skeptic that the Moorean conception of evidence cannot be maintained is not automatically to concede that only sense data can qualify as evidence, so conceding to her that the Russellian position vis-à-vis abduction cannot be maintained is not to concede that explanatory factors can play no role whatever in our inferential practices. With respect to the role of such factors, a middle position would seem to be desirable as well. To endorse abduction categorically would, in light of its registered failures, seem reckless. Yet to reject it wholesale would seem an ill-motivated move, for in light of much that has been discussed in previous chapters it is hard to escape the impression that explanatory considerations can guide us toward the truth frequently enough. It appears that the tools that we used in the preceding also allow us to state the desired middle position with respect to the question of which confirmation-theoretic role we are to grant to abduction.

In chapter 4, we saw that, contrary to what most Bayesians hold, updating probabilities by dint of conditionalization is compatible with assigning bonus points for explanatory reasons. The trick, we explained, is to account for explanatory relations between sentences when one determines one's conditional probabilities; by doing so, one can coherently (i.e., non-Dutch-bookably) give weight to explanatory considerations in updating probabilities. In chapter 4, this point was illustrated by means of EXPL, but in the meantime we have encountered other probabilistic versions of abduction as well. We also saw that even if abduction came with Dutch-bookability, it still would not follow that abductive reasoning makes us irrational. That would be so only if there could be no compensating advantages. And as was amply shown, there *can* be such advantages.

Here, I will not make any assumptions about *how* weight is given to explanatory considerations. Indeed, if you believe that there is a meaningful way to bring explanatory considerations to bear on the determination of priors—contrary to what was suggested in section 1.2.2—then you could assume that *that* is how explanatory considerations are taken into account in the following.

That we *can* rationally, even coherently, give weight to explanatory considerations does not mean that we *must* or *will* give weight to such considerations. Realistically speaking, whether we will want to assign a role to explanatory considerations will depend on how confident we are that such considerations guide us toward the truth. Someone convinced that they do point us to the truth will presumably want to assign bonus points on the basis of explanatory considerations (how big a bonus she will want to assign will depend, among other factors perhaps, on how strong in her view the connection between truth and explanation is). On the other hand, someone who thinks that although they may perhaps be of pragmatic concern, explanatory considerations are unrelated to truth, will in determining probabilities not want to take them into account at all. But what is one to do if one is convinced only to some intermediate degree that explanatory considerations and truth are related in a way that would make giving the former a confirmation-theoretic role a prudent epistemic policy?

Suppose that you want to determine your probability for a hypothesis $H$ conditional on evidence $E$, where you believe that $H$ explains $E$ better than any rival hypothesis. Now try to imagine what probabilities you would assign to $H$ given $E$ were you to be fully convinced that explanatory considerations reliably guide us toward the truth, respectively were you to be fully convinced that such considerations are at best of pragmatic interest. It is not too unrealistic to suppose that you can come up with the requisite counterfactual probabilities; designate them by "$\mathfrak{R}(H \mid E)$" and "$\mathfrak{S}(H \mid E)$," respectively. We can again think of these as constituting a panel of experts in the sense of definition 8.2, consisting of a "Russellian" expert who in determining probabilities gives full weight to explanatory factors, and a "skeptical" expert who in doing the same disregards them completely. Unsurprisingly, then, the proposal is that you make your actual probability for $H$ conditional on $E$ a weighted average of the counterfactual conditional probabilities $\mathfrak{R}(H \mid E)$ and $\mathfrak{S}(H \mid E)$, where for obvious reasons this time the weights are your probabilities for the reliability of abduction and its negation, respectively, both conditional on $E$ (for the reason explained in the comment on definition 8.2). Thus, on the current proposal your conditional probability for $H$ given $E$ is given by

$$(8.2)\quad \Pr(H \mid E) = \Pr(\text{RABD} \mid E)\,\mathfrak{R}(H \mid E) + \Pr(\neg\text{RABD} \mid E)\,\mathfrak{S}(H \mid E).$$

For basically the same reason that (8.1) could not be entirely general, (8.2) cannot be entirely general either. In the case of (8.1), it seemed plausible to think that $\Pr(\text{RSP} \mid E/e) = \Pr(\text{RSP})$ for all $E/e$. A parallel assumption for RABD would seem unwarranted, though: (8.2) is not restricted to evidence that is formulated in a special way so as to guarantee neutrality regarding RABD's truth. Sometimes it will be clear, or can be readily assumed, that there is no explanatory connection between RABD and the evidence; this will be so for our applications of (8.2) in the following. But if not, then in determining the probability of RABD conditional on the evidence you are left to your own devices. For reasons already pointed to, that is no defect of (8.2); not all our opinions can come from experts anyway.

By way of further comment on (8.2), note that the question of the status of this rule is orthogonal to the question of whether we admit only sense data statements as evidence, or also uncertain evidence reports, or even statements such as "I have two hands." Whether one thinks that only sense data statements are admissible or whether one also wants to admit "I have two hands" and its ilk, (8.2) will permit one to give weight to any possible explanatory connections between one's evidence and the hypothesis or hypotheses one is interested in, where the exact weight reflects one's confidence that explanatory power is a mark of truth; in the former case, the variable $E$ will merely have a narrower range than in the latter.[9]

9. But how is one to proceed in case there is an explanatory connection between a hypothesis and some uncertain piece of evidence, which would seem to require the involvement not only of an imaginary Russellian expert and an imaginary expert skeptical about the confirmation-theoretic significance of explanatory considerations, but also of a Moorean expert and one skeptical about the reliability of sensory perception? Neither (8.1) nor (8.2) provides for the possibility of a consultation of all these experts together. However, in Douven (2005, p. 276) it is shown that (8.1) and (8.2) can be plausibly combined in a single rule, which does provide for the said possibility. Yet for our present purposes such a (more complicated) combined rule can be dispensed with, as we are interested only in explanatory connections between a certain hypothesis—RSP, as it happens—and certain sense data statements. To account for these connections, (8.2) suffices.

To end this section, let me also briefly illustrate by means of an example how (8.2) works in practice.

**Example 8.2** Consider a person who accepts (8.2). Say that this person has probability function Pr with the associated expert functions $\mathfrak{R}$ and $\mathfrak{S}$, and say that $\Pr(\text{RSP}) = \Pr(\text{RABD}) = .5$, $\mathfrak{R}(\text{RSP}) = .7$, and $\mathfrak{S}(\text{RSP}) = .3$. Let $E$ report some pattern in the person's sense data; for instance, let it be the sentence "I am currently having sensorimotor and tactile experiences as if I am typing on a keyboard, auditory experiences of soft clicks as if keys are being struck, and visual experiences as if letters pop up on my computer screen."[10] This is a sentence that both the Moorean and the skeptic may be supposed to admit as a possible evidence report. Further, the person takes $E$ to be best explained by the hypothesis that her senses reliably inform her about her surroundings. Indeed, what better (simpler, more straightforward) explanation could there be for the occurrence of the pattern than that things actually are as they appear to be, that is, that the person is really typing on a keyboard?[11] That means nothing for the person's imaginary expert skeptical about explanatory considerations; so, $\mathfrak{S}(\text{RSP} \mid E) = \mathfrak{S}(\text{RSP})$. But things look quite different from the perspective of her imaginary Russellian expert, who is fully convinced of the confirmation-theoretic significance of explanatory power. Given that RSP best explains $E$, the expert's probability for RSP conditional on $E$ exceeds her unconditional probability for that hypothesis, specifically $\mathfrak{R}(\text{RSP}) = .8$. Finally, the person deems $E$ irrelevant to RABD, that is, $\Pr(\text{RABD} \mid E) = \Pr(\text{RABD})$. Then, similar to what we had in example 8.1, we find that $\Pr(\text{RSP} \mid E) = \Pr(\text{RABD} \mid E)\,\mathfrak{R}(\text{RSP} \mid E) + \Pr(\neg\text{RABD} \mid E)\,\mathfrak{S}(\text{RSP} \mid E) = .55$. Hence, if the person comes to learn $E$ and subsequently conditionalizes on it, her probability for RSP will increase from .5 to .55. ♦

10. Many similar examples can be distilled from Vogel's 1990 paper as well as from BonJour's rather extensive catalog of characteristics of sense data that seem best explained by RSP (in BonJour & Sosa, 2003, pp. 88–91).

11. Are these experiences not already explained by assuming that her senses are *now* correctly informing her about her surroundings? They are, but we are after the *best* explanation, and potential explanations are to be assessed, among other things, on the basis of their scope or generality. On this basis, the hypothesis that her senses *typically* inform her correctly about her surroundings is to be preferred to the much less general hypothesis that her senses *now* inform her correctly about her surroundings.

## 8.4 How to Resolve the Skepticism Debate?

I will now show that if we accept the middle positions regarding evidence and abduction presented previously, then, given some very weak further assumptions that no party to the skepticism debate should find disputable, a resolution of this debate becomes possible, at least in principle.

The hypotheses in which we are interested are RSP, which states that our senses are reliable, and RABD, which states that abductive reasoning is reliable. The main thing to be shown is that someone whose probability function satisfies certain (weak) constraints and who accepts both (8.1) and (8.2) may come to assign a high probability to both RSP and RABD. Example 8.1 shows how, assuming (8.1), RABD may be confirmed by uncertain evidence of a particular type, and example 8.2 shows how, assuming (8.2), RSP may be confirmed by evidence reporting certain patterns in a person's sense experiences. What may not be immediately obvious from those examples and what the following will make clear is that, given the use of both (8.1) and (8.2) and supposing that there will be "enough" evidence of the kinds encountered in the examples for us to arrive at probabilistic conclusions about RSP and RABD that are strong enough to allow for the refutation of skepticism, the initial probabilistic assumptions about evidence and hypotheses (and their interrelations) can be exceedingly weak.

The person that we consider has an initial probability function Pr with associated imaginary expert functions $\mathfrak{m}$, $\mathfrak{s}$, $\mathfrak{R}$, and $\mathfrak{S}$, as introduced in the previous section. Our first assumption is that the person herself as well as all her imaginary experts assign positive initial probabilities to both RSP and RABD; in other words, for all $p \in \{\Pr, \mathfrak{m}, \mathfrak{s}, \mathfrak{R}, \mathfrak{S}\}$ and $H \in \{\text{RSP}, \text{RABD}\}$:

$$(8.3) \qquad p(H) > 0.$$

Note that this is compatible with the person's finding initially vastly more probable that our senses are not reliable than that they are, and also with her finding initially vastly more probable that abductive reasoning is not reliable than that it is; she may in fact be nearly certain of the unreliability of both (but not entirely certain). I should also note that neither (8.3) nor any further assumption to be made is incompatible with the person's (and her imaginary experts') assigning a very high prior probability to, for instance, "We are brains in a vat," and a very low one to the external-world hypothesis.[12]

The second assumption concerns the various pieces of evidence our person comes to know in the course of her investigations.

(8.4) Insofar as the evidence that becomes available to our person is relevant to the hypotheses at issue, it can be represented as a sequence $\sigma$ each of whose elements is either a piece of uncertain evidence that is relevant to RABD—using our earlier notation, we designate these elements by "$E_i/e_i$" ($i \in \mathbb{N}$)—or a sense data statement that is relevant to RSP—these elements are denoted by "$E'_i$" ($i \in \mathbb{N}$).

Informally, the idea is that the pieces of uncertain evidence all qualify as successes of abduction at least from the perspective of a Moorean; so among them could be the uncertain evidence report called "$E/e$" in example 8.1. The $E'_i$, on the other hand, all report patterns in our sense experiences that are best explained by supposing that our senses reliably inform us about our surroundings; among them could be, for instance, the evidence sentence called "$E$" in example 8.2. If $i < j$, then $E_i/e_i$ comes before $E_j/e_j$ in $\sigma$, and similarly for $E'_i$ and $E'_j$. Apart from these, no assumptions need to be made about the ordering of the elements of $\sigma$.

The remaining assumptions make more precise how the elements of $\sigma$ bear on RSP and RABD relative to Pr and the various expert functions. To state these assumptions, I adopt the notational convention that if among the first $m + n$ elements of $\sigma$ there are $m$ evidence reports of the sort that we indicate by "$E'_i$" and $n$ of the sort that we indicate by "$E_i/e_i$," then, with $p \in \{\mathrm{Pr}, \mathfrak{m}, \mathfrak{s}, \mathfrak{R}, \mathfrak{S}\}$, $p_n^m$ is $p$ updated, consecutively, on the first $m + n$ elements of $\sigma$, in the order in which they appear in the sequence. So if, for example, $\sigma$ should begin

$$\left\langle E_1/e_1, E_2/e_2, E'_1, E_3/e_3, E_4/e_4, E'_2, E'_3, E'_4, \ldots \right\rangle,$$

then $p_2^0$ is the result of updating $p$ consecutively on the first two elements of $\sigma$, and $p_4^3$ the result of updating it consecutively on the first seven elements. To give a slightly more concrete example (supposing still the same beginning

12. However, as can easily be verified, given assumptions (8.3), (8.4), and (8.5), she could not coherently assign probability 0 to the external-world hypothesis (nor could her imaginary experts).

of $\sigma$), $\mathfrak{S}^3_4(\text{RSP} \mid E'_4)$ is the probability assigned by the person's imaginary expert skeptical about explanation to RSP conditional on $E'_4$ after the person has come to know the first seven elements of $\sigma$.

The third assumption is that no element of $\sigma$ is from the perspective of any of the imaginary experts negatively relevant to either RSP or RABD, that is to say, for $p \in \{\mathfrak{m}, \mathfrak{s}, \mathfrak{R}, \mathfrak{S}\}$, $H \in \{\text{RSP}, \text{RABD}\}$, and all $E'_m, E_n/e_n \in \sigma$

$$\text{(8.5)} \qquad \begin{aligned} p^{m-1}_n(H \mid E'_m) &\geqslant p^{m-1}_n(H); \\ p^m_{n-1}(H \mid E_n/e_n) &\geqslant p^m_{n-1}(H). \end{aligned}$$

Given that, informally, we think of the $E'_i$ as reporting experiences that are best explained by RSP, it would seem plausible to assume that at least $\mathfrak{R}^{m-1}_n(\text{RSP} \mid E'_m) > \mathfrak{R}^{m-1}_n(\text{RSP})$ for all $E'_m$. Similarly, because we think of the $E_i/e_i$ as successes of abduction from the standpoint of the Moorean, it would seem plausible to assume that $\mathfrak{m}^m_{n-1}(\text{RABD} \mid E_n/e_n) > \mathfrak{m}^m_{n-1}(\text{RABD})$ for all $E_n/e_n$. Nevertheless, for the result we are aiming at, the weaker (8.5) is enough.

Next, we have that for all $E'_m, E_n/e_n$

$$\text{(8.6)} \qquad \begin{aligned} \mathfrak{R}^{m-1}_n(\text{RSP} \mid E'_m) &\geqslant \mathfrak{S}^{m-1}_n(\text{RSP} \mid E'_m); \\ \mathfrak{m}^m_{n-1}(\text{RABD} \mid E_n/e_n) &\geqslant \mathfrak{s}^m_{n-1}(\text{RABD} \mid E_n/e_n). \end{aligned}$$

This will be immediately obvious from our description of the elements of $\sigma$ in relation to RSP and RABD, as perceived by the various experts; that description in effect seems to warrant $>$'s instead of the $\geqslant$'s, but again the weaker assumption suffices.

A further assumption is that for all $E'_m, E_n/e_n$,

$$\text{(8.7)} \qquad \begin{aligned} \Pr^m_{n-1}(\text{RSP} \mid E_n/e_n) &\geqslant \Pr^m_{n-1}(\text{RSP}); \\ \Pr^{m-1}_n(\text{RABD} \mid E'_m) &\geqslant \Pr^{m-1}_n(\text{RABD}). \end{aligned}$$

It is reasonable to suppose, as previously stated, that pieces of uncertain evidence are probabilistically irrelevant to RSP; the $\geqslant$ in the first part of (8.7) is merely to make the assumption as weak as possible. As to the second part, note that it is perfectly conceivable that evidence that is at least from a Russellian perspective explanatorily relevant to RSP, such as the $E'_i$'s in $\sigma$ are supposed to be, is from *any* perspective completely irrelevant to RABD; in any event there is no reason to suppose that such evidence, or at least some such pieces of evidence, should be negatively relevant to RABD. Take our example, "I am currently having sensorimotor and tactile experiences as if I am typing on a keyboard, auditory experiences of soft clicks as if keys are being struck, and visual experiences as if letters pop up on my computer screen." One would seem hard-pressed to give an argument to the effect that this sentence, or any sentence like it, is *bound* to be negatively relevant to RABD, or at least that by supposing otherwise we would be begging skeptical issues.

The last two assumptions are

$$\lim_{m,n \to \infty} \mathfrak{R}_n^{m-1}(\mathrm{RSP} \mid E'_m) = 1 \tag{8.8}$$

and

$$\lim_{m,n \to \infty} \mathfrak{m}_{n-1}^{m}(\mathrm{RABD} \mid E_n/e_n) = 1. \tag{8.9}$$

That is, we assume that, as the person comes to know more and more elements of $\sigma$, her imaginary Russellian expert's probability for RSP converges to 1 and her imaginary Moorean expert's probability for RABD converges to 1 also. Because these assumptions are clearly more substantive than the previous ones, I want to emphasize that the skeptic should have no difficulty granting them, for they are still exceedingly weak. To see just *how* weak, note that, for instance, (8.9) assumes that *in the infinite limit* (and not necessarily sooner) after our person has come to know infinitely many pieces of evidence *all* of which (and not merely some, or most, of which) report *from a Moorean perspective* (and not necessarily from anyone else's perspective) successes of abduction and *none* of which report from *any* (and not just some) perspective failures of abduction, the imaginary *Moorean* expert will be certain of RABD conditional on any further such piece of evidence (for all we are supposing, the imaginary *skeptical* expert's corresponding probability may well be no higher than her initial probability for RABD). What could one be denying, by making this assumption, that the skeptic might hold dear? Similar remarks apply to (8.8).[13]

The link between these assumptions and the skepticism issue is established by the following theorem, which is proved in appendix D:

13. I would expect the Humean skeptic to find both (8.8) and (8.9) disputable. She presumably would think that any Bayesian ought to choose an initial probability distribution guaranteeing that no claim not entailed by any finite conjunction of evidence sentences can have its probability altered in the course of her (the Bayesian's) coming to know more evidence. A fortiori, it would from a Humean perspective not be reasonable for the Russellian to become

**Theorem 8.1** *If (A1)–(8.9) hold, then*

$$\lim_{m,n\to\infty} \Pr_n^m(\mathrm{RSP}) = \lim_{m,n\to\infty} \Pr_n^m(\mathrm{RABD}) = 1.$$

From an epistemological point of view, what happens in the infinite limit may seem to be of academic interest at best. Note, though, that by the definition of a mathematical limit, it follows from theorem 8.1 that

**Corollary 8.1** *If (8.3)–(8.9) hold, then for any c such that* $0 \leqslant c < 1$

(1) *there are* $E'_k$, $E_l/e_l \in \sigma$ *such that for all* $E_n/e_n \in \sigma$ *with* $n \geqslant l$ *and all* $E'_m \in \sigma$ *with* $m \geqslant k$ *it holds that* $\Pr_n^m(\mathrm{RSP}) > c$, *and*
(2) *there are* $E'_{k'}$, $E_{l'}/e_{l'} \in \sigma$ *such that for all* $E_{n'}/e_{n'} \in \sigma$ *with* $n' \geqslant l'$ *and all* $E'_{m'} \in \sigma$ *with* $m' \geqslant k'$ *it holds that* $\Pr_{n'}^{m'}(\mathrm{RABD}) > c$.

So, given some threshold value such that probabilities exceeding it qualify as high, it follows from theorem 8.1 that if (8.3)–(8.9) hold, then after the person has come to know some *finite* initial segment of $\sigma$ it will be, and remain, highly probable to her that sensory perception is reliable and also highly probable that abduction is a reliable mode of inference.

How does this result bear on skepticism? Suppose, for now, that RSP does become highly probable to me and that that justifies me in believing RSP. Then, assuming that I am justified in believing how things appear to me, I am justified in believing that (1) my senses reliably inform me about my surroundings, and hence things generally are as they appear to be in my perceptual experiences, and (2) most or even all of my perceptual experiences appear to be of an external world. Together beliefs (1) and (2) justify my belief that there exists an external world. The same conclusion is warranted on the supposition that I become justified in RABD, still assuming the explanatory superiority of the external-world hypothesis. Thus, supposing the conditions of theorem 8.1 to be satisfied, we can become justified in believing that there is an external world, along two different routes, in effect.

---

increasingly confident in RSP as she comes to know more and more elements of $\sigma$, nor would it be reasonable for the Moorean to become increasingly confident in RABD under the same circumstances. So, by assuming that the external-world skeptic is willing to grant (8.8) and (8.9), and thus to grant that learning from experience is possible at least for the Russellian and the Moorean, I am effectively assuming that the external-world skeptic is not also a Humean skeptic. Again, I take this assumption to be justified in light of Schurz's work on meta-induction (Schurz, 2019).

I assumed that high probability would be enough for justification. According to some authors, that is indeed the case.[14] But most nowadays take Henry Kyburg's (1961) so-called lottery paradox to show that the connection between probability and justification cannot be so straightforward. Many of these authors believe that high probability is only *nearly* sufficient for justification and that some defeater has to be added to the condition of high probability to keep the lottery paradox at bay (e.g., Korb, 1992; Pollock, 1995; Ryan, 1996; Douven, 2002b; Olin, 2003). There is no agreement among them about what this defeater should come to, but they do generally agree that it should apply as selectively as possible to the type of propositions that engender the lottery paradox (such as, most typically, propositions to the effect that a given ticket in a particular lottery is a loser). As neither RSP nor RABD is of this type, these authors would probably grant that our result shows that we can become justified in believing these hypotheses.[15]

As mentioned in section 2.2.2, the connection between probability and justification may actually be still more complicated than philosophers have so far believed. But even in the worst case, in which we cannot infer anything about the justificatory status of a statement however probable it is (barring, I suppose, its being certain), the nonskeptic would still seem to have gained the upper hand if RSP and RABD had been shown to be highly probable. What is left of the pull of skepticism if we are in a position to claim that, in all likelihood, there is an external world?

I end with a number of comments on the previous result. First, I have tried to be as minimalistic as possible in choosing the inferential machinery to be used in the debate against the skeptic. Specifically, I assumed Bayes's rule, the theory of expert functions, and counterfactual probabilities. Might

14. See, among others, Kyburg (1961), Klein (1985), and Foley (1992).

15. Depending on one's account of knowledge, this may suffice to show that we can come to know that in all likelihood there is an external world (supposing that there *is*). I will not argue for this claim here, however, as it should be enough to persuade the skeptic to show that, starting from probabilistic assumptions she should find unobjectionable, we can become *justified in believing* that there exists an external world (perhaps it is already enough to show that this can become highly probable to us). And according to Wright (1991), it is justified belief in, rather than knowledge of, an external world that is at stake in the skepticism debate. Specifically, he maintains that "[k]nowledge is not really the proper central concern of epistemologico-sceptical enquiry. . . . We can live with the concession that we do not, strictly, *know* some of the things we believe ourselves to know, provided we can retain the thought that we are fully justified in accepting them" (p. 88). In the same vein, see Russell (1912, ch. 2), but see also Nagel (2014, ch. 2) for a dissenting opinion.

the skeptic still object that I have thereby been loading the dice against her? I do not see how. Of course, the *Humean* skeptic might complain about the use of Bayes's rule in particular, but the foregoing addressed the *Cartesian* skeptic, who, we may assume, is perfectly comfortable using Bayes's rule to make inferences about future sensory experiences.[16] Similarly for the theory of expert functions and for the concept of counterfactual probability: I can think of no reason why either should be objectionable from the skeptic's viewpoint or why adopting one or both of them would, in itself, tip the scale in favor of the nonskeptic. These elements were needed for formalizing views of evidence and inference that are different from those held by skeptics and that the skeptic *might* want to reject. But I have done my best to argue that the new proposals are more reasonable than the skeptic's views on these matters and should appear as such also to the skeptic.

Second, it cannot be stressed enough that we do not have a resolution of the skepticism debate here. We have no more than *stipulated* the existence of the sequence $\sigma$ with its specific connections to the person's probability function. One may wonder how sensible it is to suppose that if we were to do the empirical research required for a resolution of the skepticism debate, the data would turn out as favorably for the nonskeptic as is required for our result to apply. It may be said that while perhaps it is sensible to suppose that by far most of our experiences are best explained by assuming that our senses are reliable and also that a systematic investigation of the success rate of abduction would reveal that this mode of reasoning mostly serves us well, $\sigma$ contains only nonnegative evidence for our hypotheses RSP and RABD and that this is too much to assume. First, however, note that the preceding is compatible with there being finitely many pieces of evidence *negatively* relevant to either RSP or RABD (or both), as long as it is assumed that after the last such piece of evidence has become known to her, our person and her imaginary experts still assign some positive probability to both RSP and RABD. Second, I strongly suspect that in particular (8.5) can be relaxed so as to permit occasional negative evidence to pop up even in the long run while still allowing us to derive the requisite convergence result. Third, and most importantly, it should not be forgotten that the claim that I am making in this section is very much an in-principle claim: it is in principle possible to resolve the skepticism debate empirically, that is, to settle it in a way that should be acceptable to all parties.

16. As for the Humean skeptic, I would once more refer her to Schurz (2019).

To the extent that at the beginning of this chapter such a resolution seemed *im*possible in principle, the result obtained here should still count as a step forward.

Third, while the convergence result may very well *describe* how we nonskeptics actually come to rely on our senses as well as on abductive reasoning and thereby come to be nonskeptics, the result is in the first instance meant to have *normative* implications. In sections 8.2 and 8.3, reasons were given for preferring the middle positions regarding evidence and inference to the traditional "extreme" positions regarding these matters (both the skeptical and the nonskeptical positions). Further, supposing these middle positions and a number of very weak assumptions which, I have argued in this section, should seem reasonable to all parties involved, and supposing "favorable" data, a convergence result could be stated to the effect that accumulating evidence will justify us in believing that our senses are reliable as well as in believing that abduction is a reliable mode of inference. And that, as was also argued, is enough to justify us in believing in the existence of the external world. Hence, if the data turn out favorably, then it follows from the result that we *ought* to be non-skeptics regarding the external world.

Fourth, Wright (1985) and others have argued against Moore's proof of an external world that although Moore does have justification for "I have two hands," this justification does not transmit to "There exists an external world," even if we assume the entailment of the latter sentence by the former to be known. The reason for this transmission failure is that, to these authors' eyes, Moore must be antecedently justified in the conclusion of this argument if he is to have justification for "I have two hands." If they are right, then even if justification is closed under known entailment, that still does not mean one's justification for certain premises justifies one in believing a conclusion that is known to follow from these premises. Whether they are right is still a matter of controversy.[17] Be this as it may, Wright's and others' appraisal of Moore's argument is irrelevant to the two arguments that, I claimed, justify us in believing that there is an external world if assumptions (8.3)–(8.9) hold. Consider that the empirical argument that in that case would justify us in believing RSP and RABD merely assumes that we assign a positive probability, *which may otherwise be as small as one likes*, to the external-world hypothesis. It would certainly be wrong to say that the external world hypothesis has

17. See, for instance, Pryor (2004) for some incisive criticisms.

been *assumed* in an argument if, for all the argument must suppose, that hypothesis's probability may be vanishingly small.

Finally, and most relevantly to our overarching concern, it has been seen how, from minimal assumptions, we can defend abduction in a way that makes it useful in the kind of context in which philosophers had *hoped* it would be useful. The skepticism debate may well provide the most challenging of these contexts, given that (as said) there is really not much one can start from without begging any issues. In particular, while the skeptic might be happy to admit abduction to make inferences about sensory impressions, she might object to the use of abduction in order to infer the existence of objects and processes causing, but not reducible to, such impressions. That would seem to make the prospects of an empirical defense of abduction bleak. But taking on board conceptions of evidence and inference that, I have argued, should make sense to all parties, it was seen how we can more or less simultaneously build confidence in the reliability of our senses and in applications of abductive reasoning that let us infer the existence of real objects and processes.

This same construction can be reapplied in the scientific realism debate, which provides an only slightly less challenging context than that of the skepticism debate. Scientific antirealists are typically willing to grant the existence of observables—anything that can be seen with the naked eye—but regard any claim postulating unobservables as a case of epistemic overreach. For instance, van Fraassen (1985) has raised skeptical concerns about realist readings of micrographs, in particular about taking such micrographs as evidence for the existence of bacteria, viruses, molecules, and other such things. According to him, there is no non-question-begging evidence showing that our current microscopes are veridical. But few realists have been swayed by his arguments. The following excerpt from an interview with Ronald Giere expresses a widely shared realist sentiment:

> An overall kind of nonrealism or antirealism I find very hard to even take seriously. I know people do it. But look at electron microscope pictures of chromosomes, of DNA—we have pretty nice pictures of DNA! The overall idea that there is an observational realm and that we should not go beyond that I think is just not tenable. (Giere in Callebaut, 1993, p. 171)

However, even though most people with some background in science will be sympathetic to Giere's view, it would seem prudent to be at least initially somewhat more cautious in our interpretations of micrographs. In the case of perception, we all have first-hand evidence of its fallibility, evidence which, we said, should make us wary of going along with Moore's proof of the external world even if it should not make us side with the skeptic. Although not as generally known, however, there is also evidence showing that modern microscopes are not completely trustworthy and that, for example, procedures used for preparing specimens can produce various sorts of artifacts, some of which can go unrecognized for quite some time, perhaps even forever (Rastogi et al., 2013). So here, too, a kind of intermediate position between van Fraassen's and Giere's would seem the more reasonable one, at least as a starting point.

Then, in order to test empirically whether abduction is reliable also when we infer, or become more confident in, the existence of certain unobservable kinds of entities postulated by best-explaining theories, we can imagine having at hand an expert deeming microscopes at least somewhat trustworthy, next to an expert who instead shares van Fraassen's view on microscopes. And also parallel to what we had before, we have our pair of experts holding different views on the reliability of abduction as applied to unobservables. The rest then unfolds as it did previously. Suppose that using some type of microscope, we "discover" a virus that had earlier been postulated on explanatory grounds. That would constitute evidence for the reliability of abduction from the perspective of the expert who believes microscopes to be at least somewhat trustworthy. Similarly, Hacking (1981, pp. 144–145) points out that when we study blood cells both by means of a low-resolution electron microscope and by means of a fluorescent microscope—types of microscopes that exploit different physical principles—the visual outputs we get from these microscopes tend to be very similar. As he rightly remarks, that would seem best explained by assuming that these types of microscopes are veridical and thus would constitute evidence for that assumption at least from the perspective of the expert who has some confidence in the reliability of abduction.[18]

18. For further details, see Douven (2005).