Knowledge of Perception

Daniel Stoljar

> I'm not asking whether you know you are not a zombie. Of course you do. I'm asking how you know it. The answer to that question is not so obvious. Indeed it is hard to see how you can know it. Wittgenstein . . . didn't think he saw anything that allowed you to infer he saw it. The problem is more serious. There is nothing you are aware of, external or internal, that tells you that, unlike a zombie, you are aware of it. Or, indeed, of anything at all. (Dretske 2003, 1)

1

Suppose I am watching my son do somersaults in the living room; and suppose (having nothing better to do) I ask myself whether I am watching my son do somersaults. Would I then come to know that I am?

Of course the answer is yes. I am not asking you to suppose that my living room is populated with the sort of things epistemologists like to discuss: fake sons, gurus, mad scientists, etc. The imagined situation is not that at all. Rather it is just a normal day in which I am watching my son do somersaults. If you like, imagine also that I fully understand the notions involved—I, watching, son, living room, etc.—that I have asked myself seriously (and not just in a half-hearted way) whether I am watching my son, and that I am not so busy as to have my mind crowded with other things. If it is distracting, forget also that watching is some sort of action, and so that to know that I am watching is

to know that I am acting; concentrate just on seeing since watching entails seeing.[1] Do I know that I see my son? In the imagined situation, of course the answer is yes.

But there is an apparently simple argument for the opposite conclusion. The first premise of the argument is E1:

Evidence Argument (Version 1)

E1. If I know that I am watching my son, there is an answer to the question "How do I know that I am watching my son?"

The claim here is not that I must *know* the answer to the question; I do not have to know how I know. The claim is rather that there must *be* an answer. What motivates this claim? Well, not everyone who is watching his or her son knows it. Perhaps somewhere in Brazil a new world monkey—a pygmy marmoset, say—is watching (and so seeing) his son do somersaults; but maybe he does not know that he is. Or perhaps on another day I am so preoccupied with one thing or another that I do not know that I watching my son even though I am. However, it is not a miracle that I know that I am watching my son while my pygmy marmoset friend does not know that he is watching his; likewise it is not a miracle that I know I am watching in some circumstances but not in others. Hence, there must be an answer to the question "How do I know?"

The second premise of the argument is E2:

[1] I will myself move back and forth between watching and seeing in the discussion to follow; but nothing substantial will turn on this.

E2. If there is an answer to the question "How do I know that I am watching my son?" that answer must be a proposition that (a) I am aware of and (b) "tells me" that I am watching my son.

Very often a question of the form "How to you know that p?" is a request for evidence. If I tell you that the Dean is going to close down the physics department, and you ask me how I know, in many instances you would be satisfied if I responded with something like, "Well, she told Professor X in the strictest confidence that she was going to, and he immediately told me." These further claims are things I am aware of and are things that support or confirm or speak in favor of the thing that I know. To the extent that this case is typical, what you are looking for when you ask me how I know something is a proposition or fact that in a certain sense I am aware of and which supports what I claim to know; that is, you are looking for my evidence.

The third premise of the argument is E3:

E3. There is no proposition that meets conditions (a) and (b).

This premise is motivated by a plausible idea about the nature of perception. When I watch, and so see, my son, what facts to I come to be aware of? Well, if we set aside skepticism about the external world, the facts that I come to know are facts such as the following: that my son is doing somersaults, that my son is in the living room, that my son exists, and so on. These are the facts that in the first instance I come to know on the basis of perception. However, these facts seem to bear no particular relation to the fact that I am watching my son. For one thing they do not entail or necessitate that I am watching my son. My son might perfectly well be doing somersaults without my

watching him—if I closed my eyes, say. Moreover, they do not confirm or support or tell me that I am watching my son. Imagine someone advancing an argument whose premise is that X's son is doing somersaults to the conclusion that X is watching his son do somersaults. That looks like a spectacularly terrible argument; the premise not only fails to entail the conclusion, it fails to give it any reasonable support.

E1 and E2 tell us that if I know that I am watching my son there must be an answer to a certain question and also that the answer must meet certain conditions. E3 says that no answer meets the conditions. Putting these together, I do not know that I am watching my son. But, since—obviously—I do know, something is wrong. How is it possible that I know I am watching my son?

The argument I have just given—the evidence argument, as I will call it—is Dretske's or at any rate is Dretske's as I understand him (cf. Dretske 2003, 2010; the argument is also mentioned by other writers or at any rate by other writers as I understand them—I will set them aside here.[2]) In the passage quoted above, Dretske puts the point in terms of (philosophical) zombies, i.e., creatures by definition physically the same as us but lacking conscious experience; his question is "How do you know you are not a zombie?" But as he himself notes, zombies are not really to the point and raise distracting questions. For one thing, a lot of philosophers think that zombies are strictly speaking

[2] It seems to me that M. G. F. Martin is suggesting an argument along these lines, when he writes: "why should the evidence that the subject has about a how the world is have any bearing on what beliefs a particular person has? The kind of evidence that grounds our beliefs about the world would seem in general not to be relevant to the question of what believes we ourselves have" (1998, 110). That Martin is discussing belief rather than perception lends support to the idea mentioned in comment 5 below, namely, that the evidence argument is a very general argument and is not limited to the case of perception. Dretske's argument has also received extensive discussion in Byrne 2005 and Byrne, this volume.

impossible, or, at any rate, that we have no reason to suppose them possible. The point comes out better if you adopt one of Dretske's more mundane examples (which I have) and frame matters in terms of that.

My aim is to present a response to the evidence argument and to explore the conception of introspection that emerges from that response. I will begin, in section 2, with some preliminary comments before presenting the response in sections 3 to 7. In sections 8 and 9, I will contrast my proposal with an alternative.

2

Comment 1: You might think that the evidence argument is generated by a careful choice of example. Watching my son involves seeing him, and this involves standing in a relation to him, something I could hardly do unless he exists. But then you might think that the evidence argument is generated by the possibility that I do not know that he exists, and as such is simply the problem of the external world all over again. However, while it is true I cannot see my son unless he exists, this is not what generates the problem. Many views of perception postulate perceptual facts that do not require the existence of my son, and one could perfectly well discuss the basic issue in terms of these facts rather than the fact that I am seeing my son. Talk of seeing simply makes the problem easier to state.

Which theories of perception involve facts that do not require the existence of my son? Consider the intentionalist view. At least on a cartoonish understanding of that view, I see my son only if he plays the right sort of explanatory role in my being in a certain representational state, the state of visually representing (the proposition) that my son

exists. But we could if we like focus on the question "How do I know that I am in the state of visually representing that my son exists?" That question does not presuppose that he exists. Likewise, consider the disjunctivist view. At least on a cartoonish understanding of that view, I may be in the situation of either seeing my son or it seeming to me *as if* I am seeing my son. This disjunctive fact does not require the existence of my son. But we could if we like focus on the question "How do I know that either I am seeing my son or it seems to me that I am seeing my son?" So far as I can see, adopting either of these theories of perception does not affect the issue I want to focus on, and so I will concentrate here on the simpler idea that I am seeing my son.

Comment 2: It sometimes happens that a person knows something but does not know (because they cannot remember) how. Suppose I know that the Battle of Marston Moor was fought in 1644 but I do not remember that Mrs. Beamish (my fourth grade teacher) told me this, and in fact I have no idea how I know. If so, I may not be (currently) aware of any proposition that speaks in favor of what I know. Does this show that the condition on evidence set out in E2 is mistaken? Well, as stated it does. However, I take this to be a somewhat superficial problem for a proponent of the evidence argument; for example, E2 could be adjusted to accommodate this by replacing "is" with "is or was." I will therefore ignore this complication in what follows.

Comment 3: It might be that if either a perceptual model or the inferential model of introspection were true, E3 would be false. Nevertheless I will set both aside here. As

regards the perceptual model, Dretske himself assumes that Shoemaker has refuted it.[3] I am not sure about that, since it seems to be doubtful that the perceptual model is clear enough to be refuted. But I do agree with both Dretske and Shoemaker that we should do without it if we can, and that will be my policy. As regards the inferential model, as this is usually developed I infer that I am watching my son from premises about my behavior. This idea may be stipulated to be irrelevant in the case at hand. Even if there are occasions on which I come to know that I see my son by inferring this from my behavior, I did not in the case I began with; or at any rate so we may imagine. In recent work, Alex Byrne has developed and defended the inferential model in a nonstandard version.[4] For Byrne, I infer that I see my son from certain premises that Byrne calls 'v-propositions', where a 'v-proposition' encodes information about how my son looks as he is doing somersaults. Byrne's proposal is interesting and controversial; however, as I understand matters, one of its main motivations is the idea that it solves the problem presented by the evidence argument. If, as I think, that problem may be solved in a different way, this part of the motivation for Byrne's view evaporates.

Comment 4: I said that support for E3 derives from a fact about the nature of perception, viz., that when I see my son, in the normal case I come to know certain facts about my son, e.g., that he exists, that he is doing somersaults, etc. It might be thought that support for this premise derives also or instead from another source, viz., the so-called

[3] "I can't improve on Shoemaker" (Dretske 2003, 11 n. 12). For Shoemaker's criticism of the perceptual model, see Shoemaker 1996, and for other criticisms, see Siewert, this volume. For defenses of the view, see Gertler, this volume, and Horgan, this volume.
[4] See Byrne, this volume; see also Byrne 2005.

diaphanousness of perception.[5] This claim is somewhat notorious in philosophy of mind. Some people take it to mean that you are never aware of the fact *that* you see but only the fact (or thing) *which* you see. In my view, this claim is too incredible to be true. For surely we are sometimes aware *that* we see; for example, I know (and so am aware) that I am seeing my son. Others (more plausibly) take the diaphanousness claim to be primarily a claim about attention, viz., that when I try to attend to *my seeing* of my son, I wind up attending to my son. If this claim is true, it raises a number of questions—in which of the many senses of 'attention' is it true, for example, and what is it about attention in the relevant sense that makes it true? But while these questions are important and interesting, they also seem remote from the issues we are discussing and so I will set them aside.

Comment 5: One might think that the evidence argument is limited in scope to perception since that is our lead example. If so, the issue we are discussing might appear narrow. However, a similar argument can be mounted in the case of many other mental states. To illustrate, take imagination. Suppose I imagine that my son is doing somersaults in the living room. In the normal case I know perfectly well that I am imagining this. How do I know? Well, imagining my son doing somersaults might make me aware of many propositions: e.g., that it is possible that my son is in the living room, that it is possible that my son is doing somersaults, etc. But none of these propositions, even if true, make it likely that I am imagining him. Or take intention. Suppose I intend to join my son doing somersaults in the living room in ten minutes. In the normal case I know perfectly well

[5] It has often been suggested to me (e.g., by David Chalmers in conversation) that Dretske 2003 is relying on the diaphanousness of perception. I am not sure that this is so, but even if it is, the evidence argument seems to me to be stronger if we decouple it from any implausible version of diaphanousness. For my own treatment of the diaphanousness issues, see Stoljar 2004.

that I intend this. How do I know? Well intending to do somersaults might make me aware of various propositions: e.g., that I will in ten minutes do somersaults in the living room, that there will in ten minutes be two people doing somersaults in the living room, etc. But none of these propositions, even if true, make it likely that I am *currently* intending to do somersaults. The question of whether the proposal I will make about the evidence argument can be extended from perception to other cases is interesting but I will not pursue it here.

Comment 6: The question "How do you know?" is a close cousin of the question "Why do you believe?" The first presupposes that the addressee knows the thing in question, while the second presupposes only that he or she believes it. But both questions are normally requests for evidence. If you ask me why I believe that the Dean is going to close down the physics department you would very often be asking me for some propositions that confirm or support what it is that I believe. In consequence, the evidence argument could be formulated in terms of the question "Why do I believe that I see my son?" rather than the question "How do I know that I see my son?" So far as I can see, the argument stated in these terms would have the same persuasiveness of the original and so I will set it aside as well.

Comment 7: What about an uncritical appeal to the notion of introspection? That is, why not try to close down the issue by saying that the answer to the question "How do I know that I am seeing my son?" is "by introspection" and leave it at that? Dretske thinks that introspection is by definition that (distinctive) way in which I come to know (e.g.) that I am watching my son. But, as he says, an appeal to introspection in this sense gives us no insight into what if anything is wrong with the argument. It may be true that I come to

know that I am seeing my son by introspection, but this tells us only that the evidence argument is unsound. It does not tell us *why* it is unsound, and it is this latter question that will be our focus.

3

Why then is the evidence argument unsound? My answer will proceed as follows. First, I will consider a particular proposal about how to respond, which I will call "the foot-stamping view." Then I will consider a different proposal, which I will call "the Austin-inspired view." Neither the foot-stamping view nor the Austin-inspired view is correct on its own. However, if they are combined together in the right way, we have (I think) a plausible proposal about what is wrong with the evidence argument.

3.1. The Foot-Stamping View

According to the foot-stamping view, the third premise (E3) of the evidence argument is false. In the case we are imagining, E3 says that there is no proposition (a) of which I am aware and (b) which supports the proposition that I am watching my son. Against this premise, the foot-stamper insists that there *is* such a proposition, viz., the very proposition that I am watching my son. How so? Well, in the first place, I *know* that this proposition is true and in this sense I am aware of it. After all, nobody is denying that I know that I am watching my son; "Of course you do," says Dretske in a related context. Second, it is a proposition that tells me that I am watching my son, because sometimes propositions tell me about, or speak in favor of, themselves. (Both the "telling" metaphor and various accounts of epistemic support permit this possibility.)

What if anything is wrong with the foot-stamping response? Dretske offers one objection when he remarks, "To insist that we know it despite having no identifiable way we know it is not very helpful. We can't do epistemology by stamping our feet" (2003, 9). However, while this remark is what inspired the name of the view, it is in fact hard to see exactly what objection Dretske intends by it. For the foot-stamper (as least as I understand him) is not insisting that I know that I am watching my son and yet I have no identifiable way that I know. On the contrary, the foot-stamping response agrees that there *is* an identifiable way (at least if this means: an identifiable answer to the question "How do I know?"). That way is the (easily identifiable) proposition that I am watching my son. It is puzzling, therefore, that Dretske can say otherwise.

We will be in a better position to respond to this puzzle at a later stage in our discussion. For the moment, I want instead to suggest that a proponent of the evidence argument might respond to the foot-stamping response in a way that is different from Dretske, viz., by finessing it rather than confronting it directly. Right at the start of the argument we agreed that the question "How do I know that I am watching my son?" has a nontrivial answer. If so, the argument can be adjusted to avoid the foot-stamping response. The reformulation would go as follows:

> E1. If I know that I am watching my son do somersaults, there is a nontrivial answer to the question "How do I know that I am watching my son?"
>
> E2. If there is a nontrivial answer to the question "How do I know that am watching my son?" that answer must be a proposition that (a) I am aware of; (b)

"tells me" that I am watching my son; and (c) is distinct[6] from the proposition that I am watching my son.

E3. There is no proposition that meets conditions (a), (b), and (c).

E4. Therefore, I do not know that I am watching my son.

On the face of it, this version of the argument—which I will call version 2 of the evidence argument, to distinguish it from the version already in play—has whatever plausibility the original had, and yet foot-stamping is no response to it. Hence, if the foot-stamping response is the only response available to the evidence argument, it looks as if we so far have no response.

3.2. The Austin-Inspired View

To this point we have been mostly uncritical about the question "How do I know that I am watching my son?" assuming only that it is often a request for evidence. But there is a famous discussion, J. L. Austin's, that bears on what this question may mean. Here is Austin:

> Suppose I have said 'There is a bittern at the bottom of the garden', and you ask 'How do you know?' my reply may take very different forms:
>
> a. I was brought up in the fens
> b. I heard it
> c. The keeper reported it

[6] As Mark Jago pointed out to me, there are a number of questions about how to interpret 'distinct' here. I acknowledge this, but I will leave this notion intuitive in the discussion to come.

d. By its booming

e. From the booming noise

f. Because it is booming.

> We may say, roughly, that the first three are answers to the questions 'How do you come to know?', 'How are you in a position to know', or 'How do you know?' understood in different ways: while the other three are answers to 'How can you tell?' understood in different ways. (Austin 1961, 79)

There are various interpretative questions about what Austin says here. For one thing, the answers (a–f) seem to concern, not the question "How do you know that there is a bittern at the bottom of the garden?" but rather the question, "How do you know that it (that) is a bittern?" For another thing, it is not quite clear how precisely to characterize the distinction at issue. Elsewhere Austin says that (a–c) provide my "credentials" while (d–f) provide my "facts" (i.e., the facts that "prove" (as he says) it is bittern). Finally, it is natural to see a connection between Austin's use of the idea of "telling" and Dretske's: on the one hand, there is something I am aware of that tells me it is a bittern, viz., its booming; but, on the other hand, there is (apparently) nothing I am aware of that tells me I am watching my son. If so, answers (d–f) are responses to a request for evidence.

All of these interpretative points may be developed in different ways. For our purposes it is sufficient to draw from Austin (without any claim to scholarly accuracy) a distinction between (what I will call) an *evidence-seeking version* of the question "How do I know?" in which we ask for evidence for what I know—this is the version we have been operating with all along—and an *explanation-seeking version*, in which we ask for an explanation for the fact that I know. So, for example, if you ask me how I know that

the Dean is going to close down the physics department, you *might* be asking me for evidence for the proposition that I know, viz., that the Dean is going to close down the physics department, or you might also or instead be asking me for an explanation for a certain fact, i.e., the fact that I know that the Dean is going to close down the physics department.[7]

Armed with this distinction, a straightforward objection to the argument comes into view, viz., that it does not go through if what is at issue is an explanation-seeking question:

> E1. If I know that I am watching my son do somersaults, there is a nontrivial answer to the explanation-seeking question "How do I know that I am watching my son?"
>
> E2. If there is a nontrivial answer to the explanation-seeking question "How do I know that am watching my son?" that answer must be a proposition that (a) I am aware of; (b) "tells me" that I am watching my son; and (c) is distinct from the proposition that I am watching my son.
>
> E3. There is no proposition that meets conditions (a), (b), and (c).

[7] Objection: The question "How do you know the Dean will close down the physics department?" is certainly a request for evidence. But it is also a request for explanation, at least of a sort. Hence, there is no distinction of the sort that you want to draw from Austin. Response: It is true that requests for evidence might be thought of as requests for explanation too; at least it does not stretch the concept of explanation to suppose so. But there are nevertheless requests for explanation that are not also requests for evidence, at least not in the sense of 'evidence' that is at issue in the evidence argument. To mark this, we could if we liked work with a distinction between evidence-seeking and explanation-but-not-evidence-seeing questions rather than the simpler distinction I have adopted in the text. However, since doing so would affect anything I have to say, I will continue with the simpler distinction.

E4. Therefore, I do not know that I am watching my son.

The problem with the argument in this version—version 3 of the evidence argument, as we can call it—is that E2 is false: it is not true that if there is an answer to the explanation-seeking question "How do I know that I am watching my son?" then I must be aware of that answer. Of course, I do need to be aware of the answer if I am to *give* the answer to someone else; but for it to *be* an answer I do not need to be aware of it.

To illustrate, consider again the explanation-seeking question, "How do I know it is a bittern?" What might answer this question, and so explain the fact that I know that it is a bittern? Well, one answer is provided by Austin, namely, that I was brought up in the fens, or rather, to fill out the story a little bit, that I was brought up in the fens and people who are brought up in the fens know a bittern when they see one. These facts, assuming for the moment (implausibly!) that they are facts, would indeed explain that I know that it is a bittern. However, nowhere here is it required that I am aware of the facts in question. For one thing, I do not need to know that I was brought up in the fens, even if I was; I may be completely mistaken about my origins and insist that I was brought up in the tropics. Nor do I need to know that people who are brought up in the fens know a bittern when they see one, even if this is so; I may insist that there is some sort of ambient gas in the fens that blocks people from knowing bitterns when they see them, and that the only way to be able to do this is by growing up in the tropics.

What is true for the question about bitterns is true also for the question about seeing my son—at least if what is at issue is a request for explanation. What explains the fact that I know that I am watching my son? Here is an initially plausible suggestion. I

know that I am watching my son because (a) I am watching my son; (b) I am suitably situated—here, "suitably situated" summarizes the facts of the case I began with: viz., that I understand the concepts involved, I have asked myself whether I am watching my son, and I am not distracted by other things; and (c) that suitably situated people who are watching their sons know that they are. Taken together these claims explain the fact that I know that I am watching my son. On the other hand, it is obviously not a requirement on (a–c) answering the explanation-seeking question that I am aware of (a–c). If so, E2 is false, and the argument can be rejected.

3.3. Foot-Stamping Mingled with Austin

The Austin-inspired response is clearly not going to do on its own. By itself, the response reveals only that any successful version of the evidence argument will employ *one* legitimate version of the "How do you know?" question, i.e., the evidence-seeking version. That there is *another* legitimate version—i.e., the explanation-seeking version—is (one might think) beside the point. However, at this stage we may fruitfully combine the Austin-inspired response with the foot-stamping response considered earlier. For if we combine these two ideas we may present a proponent of the evidence argument with a destructive dilemma, as follows. If you want to advance the evidence argument, you must operate either with the explanation-seeking version of the "How do you know?" question or with the evidence-seeking version. If you operate with the explanation-seeking version—that is, if version 3 of the argument is in play—then the second premise (E2) is false, for answers to the relevant requests for explanation do not need to meet the conditions laid down in that premise; in particular one does not need to be aware in any sense of these answers for them to be answers. If, on the other hand, you operate with the

evidence-seeking version, then either the first premise (E1) or the third premise (E3) is false, depending on whether you require a nontrivial answer to the question; that is, depending on whether version 1 or version 2 is in play. If version 1 is in play, then (E3) is false, for there is a proposition of which I am aware and which speaks in favor of what I know, viz., the proposition that I am watching my son. If version 2 is in play, then (E1) is false, for it is not the case that there must be a nontrivial answer to the evidence-seeking question "How do I know that I am watching my son?" Of course, there must be a nontrivial answer to the explanation-seeking question, but that is a different matter.

4

How plausible is the Austin-inspired foot-stamping response just sketched? In my view it provides the right *form* of a solution to the problem posed by the evidence argument. However, before it is adopted we need to bring out and deal with two major issues. The first arises when we look more closely at the explanation that I have just offered of the fact that I know that I am watching my son. The second returns us to Dretske's objection to the foot-stamping view. In the next few sections, I am going to address both issues. It is convenient to start with the first issue, for thinking it through will provide materials for dealing with the second.

To bring out this first issue, let us focus again on the explanation-seeking question "How do I know that I am watching my son?" The answer we have given to this question so far is this: (a) I am watching my son; (b) I am suitably situated; and (c) suitably situated people who are watching their sons know that they are. Now, what is the status of the claim (c)? The first thing to say is that it seems not to be just an accidental

generalization but is rather a claim with a certain sort of modal force. The main reason for this is that it "supports counterfactuals," as it is usually put. For example, I have no idea if (e.g.) Barack Obama is watching his son (or daughter in his case) doing somersaults right now. However, I do know that *were* he to do so, and *if* he were suitably situated, then he *would* know that he is.

If we agree that the claim has modal force, the question arises as to what the strength, or nature, of that modal force is. An initially plausible proposal is that the claim is *naturally* necessary; that is, it is or involves an empirical counterfactual-supporting generalization—a psychological law, though presumably of a rough and ready kind, and one that tolerates exceptions, an "other things being equal" law. On this view the claim that suitably situated people who are watching their sons know that they are is analogous to the (rough and ready, other things being equal) claim that people who are brought up in the fens know a bittern if they see one.

However, while this assumption is certainly plausible, it creates some difficulties for the response to the evidence argument that we have put forward. To illustrate this, consider a possible world W. W is very much like the world I imagined at the outset in which it is normal day and I am watching my son. The difference between W and the actual world is that, at W, the psychological law we just described does not obtain. Now, is it true at W that I know that I am watching my son? On the assumption that the law in question does not obtain at W, the answer would seem to be "no." After all, the thing that explained why I know that I am watching my son in the original case is missing in W and nothing has taken its place. Hence (barring a miracle) I do not know that I am watching my son in W.

However, it is extremely implausible to say that I do not know that I am watching my son at W, for two reasons. In the first place, what would stop me doing so? After all, I am watching my son, I have asked myself whether I am, I understand the relevant concepts, and—finally—I am rational, or at least as rational as I am in the actual world. But why then would I not come to know that I am watching my son? What else do I have to do? In the second place, regardless of whether I know that I am watching my son in W, it is very plausible that I ought to believe this; that is, I ought from a rational point of view, to believe at W that I am watching my son. But if it is the case that I ought to believe that I am watching my son, then we have a further reason for supposing that I know this at W. For it seems plausible to say that if I am rational, then other things being equal I will believe what I ought to believe, and the belief in question will amount to knowledge. Hence, if it is true in W that I ought to know that I am watching my son, then, since I am rational there, I will know this.[8]

The problem for the suggestion that (c) is merely naturally necessary, therefore, is that it entails that at W I do not know that I am watching my son. But, since it seems obvious that I do know this at W, it cannot be that (c) is merely naturally necessary. This is not to deny that (c) *is* naturally necessary of course; it remains plausible that it is. But it is to deny that it is *merely* naturally necessary, i.e., it must be necessary in some further way as well. In the face of this, a very natural suggestion is that (c) is not merely naturally necessary but is in addition metaphysically necessary; that is, it is true in all possible worlds (in the widest sense of "all possible worlds") that suitably situated people

[8] That (c) is not merely nomologically necessary I take to be a main lesson of Shoemaker's discussion in his Brown lectures; see Shoemaker 1996. Shoemaker's "self-blindness argument" is (as I understand it) precisely an argument that one should deny that at W, I do not know that I see my son.

who see their sons know that they are. On this view the claim that suitably situated people know that they are watching their sons if they are is analogous to the claim that water is H_2O.

However, just as it is peculiar to say that (c) is merely naturally necessary, it is likewise peculiar to say that (c) is metaphysically necessary. The problem this time is not the implication that at W I do not know that I am watching my son. If it is metaphysically necessary that suitably situated people who are watching their sons know that they are, then, since I am suitably situated at W, and am watching my son there, then I will indeed know that I am. The problem rather is that the idea that (c) is metaphysically necessary rules out lots of *other* cases that seem on the face of it to be genuinely possible.

To illustrate, consider a possible world W*. W* is very much like the world I imagined at the outset with the exception that at W* I do not believe, and so do not know, that I am watching my son. If you like, imagine that at W* God has plucked out the fact that I believe that I am watching my son (and so has plucked out the fact that I know that I am) and yet has left everything else intact. Is W* a metaphysically possible world? It is hard to see why not. The fact that I know (or believe) that I am seeing my son is distinct from the fact that I see my son. (Of course, the reverse is not the case because of the factivity of knowledge; but that is not the direction we are interested in.) Hence, it does not seem impossible (at least in the sense of metaphysical possibility) that God could have brought the second fact into existence and not the first. Hence, W* is possible. On the other hand, if W* is possible, then (c) cannot be metaphysically necessary.[9]

[9] That (c) is not metaphysically necessary I take to be a main lesson of Armstrong's discussion of introspection in, most famously, *A Materialist Theory of the Mind*; see

5

At this point, therefore, we seem to be confronted with a dilemma. The claim (c)—"suitably situated people who are watching their sons know that they are"—is clearly a modal claim in some sense for it supports counterfactuals. But the necessity in question is apparently neither merely natural nor metaphysical. It is not merely natural, because to say that is to say that at W I do not know that I am watching my son. And it is not metaphysical, because to say that is to exclude outright W*. How to respond?

The way out of this dilemma, I think, is to agree that claim (c) is neither merely naturally necessary nor metaphysically necessary, and to propose instead that (c)—or at least something very close to (c)—is *rationally* necessary or, equivalently, that it is a requirement of rationality. To say that it is rationally necessary is to say that it is on a par with such principles as "if you believe p and believe if p then q, then you will believe q" and "if you intend to F and know that you can only F if you G, then you will intend to G." These claims (or, better, instances of these claims, or even better, qualified instances of these claims) are not metaphysically necessary. It is not metaphysically necessary that if I believe p and believe if p then q, then I will believe q. But nor is it true that it is merely nomologically necessary. It is not a mere fact of empirical psychology that I tend to believe q when I believe p and if p then q. Rather such principles are rationally necessary, that is, they are principles that one is required by rationality to follow. One may not follow them, of course, and one may follow them mostly but not in every case;

Armstrong 1968. Armstrong's "distinct existences argument" is (as I understand it) precisely an argument that W* is possible.

but this is just to say that one may not be in these respects rational, and may not be in these respects rational on every occasion.[10]

So what I would like to propose is that (c)—or at least something like (c)—is a rational necessity or is a requirement of rationality. Of course, this idea is just that: a proposal; it is hard to *prove* that it is true. But one argument in favor of its truth is that it allows us to avoid the dilemma about the modal status of (c) that we have been considering. On the one hand, the thesis that (c) is rationally necessary does not imply that at W I do not know that I am seeing my son. The reason is that at W I am as rational as I am at the (imagined) actual world with which we began. If we assume that I abide by or follow the relevant principle of rationality at the actual world, I will likewise abide by or follow it at W. But then it will be true at W, no less than at the actual world, that I will know that I am seeing my son. By the same token, the thesis does not imply that W* is impossible. That it is a rational necessity that suitably situated people who are seeing their sons know that they are does not entail that it is a metaphysical necessity. In particular, it is quite consistent with the idea that this is a rational necessity, and that I follow it, that there is a case in which I am suitably situated but have been prevented (by God, for example) from believing that I am seeing my son. To be sure, that will be case in which I have been prevented by God from following (at least on this occasion) a requirement of rationality, but since it is not metaphysically impossible to do that, W* is not metaphysically impossible.

---

[10] Why rational necessity as opposed to some other sort? One possible reason is that, following Fine 2004, it is reasonable to suppose that there are three fundamental kinds of necessity: natural, metaphysical, and normative. If the connection we are looking for is neither natural nor metaphysical, and if Fine is right that there are three fundamental varieties of necessity, it must be normative.

At the end of section 3 we offered an answer to the explanation-seeking question "How do I know that I am seeing my son?" The answer was this. I know that I am seeing my son because (a) I am seeing my son; (b) I am suitably situated; and (c) suitably situated people who are seeing their sons know that they are. This answer ran into difficulty on the issue of what the modal status of (c) is. A better answer, which solves this problem, is this. I know that I am seeing my son because (a) I am seeing my son; (b) I am suitably situated; (d) there is a rational requirement that suitably situated people who are seeing their sons know that they are doing so; and finally (e) I am rational at least to the extent that I abide by this requirement. It is the combination of (d) and (e) that sustains counterfactuals. Even if Barack Obama is not watching his daughter doing somersaults right now, it is true that were he to do so and if he were suitably placed, we would know that he is watching his daughter. Why so? The answer is that it is a requirement of rationality that suitably situated people who are watching their daughters know that they are, and that Barack Obama just as much as me abides by this requirement.

6

I have argued so far that if we view (c) or something like it as a rational necessity rather than either a metaphysical necessity or a mere natural necessity, we may avoid a major problem for the Austin-inspired foot-stamping view. However, this suggestion faces a number of objections. In this section I will try to deal with some of the main ones.

First, one might object that to say that there is a rationally necessary connection between the fact I am seeing my son and the fact that I believe or know that I am is

simply to say that the first is evidence for the second. If so, we seem to have taken back the Austinian distinction between explanation and evidence that was one of the planks of our discussion.

However, in general there is no reason to assume that if p rationally necessitates q then p is evidence for q, and this is especially true if we are using the notion of evidence that is at issue in the evidence argument, i.e., the notion according to which a proposition is evidence only if I am aware of it and it supports what I know.[11] Suppose it is a requirement of rationality that if I believe that p and believe that if p then q, then I believe q. And suppose now I believe both p and if p then q but on the basis of no evidence at all; imagine if you like, that both have implanted in my mind on the whim of some bored archangel. In that case it remains a requirement of rationality that I believe q. But it does not follow that I have evidence for q; on the contrary, I have absolutely no evidence for q just as I have no evidence for p or if p then q. Hence, one can say that there is a rational connection between various states without the first being evidence for the first.[12]

Second, one might object that to say that (c) is required by rationality is implausible because it is too demanding. Might there not be possible defeaters, i.e., factors that make it the case that I do not know that I am seeing my son, despite the facts that I am doing so, I am suitably placed and am rational? For example, imagine I am an eliminativist about vision and so think that nobody ever sees his son at all.[13] Or imagine

[11] I will consider a different notion of evidence in sections 8 and 9 below.
[12] The argument given in this paragraph is a version of Broome's argument against bootstrapping; see Broome 2005, 2007. Broom does not apply the argument to the case of self-knowledge.
[13] For some discussion of the eliminativist case, see Siewert 1998, and Silins, this volume.

that I am a victim of reverse Anton's Syndrome, and so am convinced (falsely) that I am blind. Or imagine that someone I trust (or some collective of people I trust) has just informed me that, whatever I am inclined to think, I am not seeing my son. In such cases, I may not know or believe that I am seeing my son even if I am rational and am suitably placed.

However, while it is true that we should allow for defeaters of this kind, it is a mistake to suppose that this is a reason to abandon the idea that (c) or something like it is rationally necessary. Instead it is a reason to develop a more discerning interpretation of the precise content of the requirement in question; indeed, it is for this reason that I have spoken so far of "(c) or something like it." There are a number of different ways in which one might try to develop this more discerning interpretation. For example, when considering the suggestion that (c) is naturally necessary we noted that on such a view, the law in question would presumably be an "other things being equal" law, a law that tolerates exceptions. Perhaps what is true for the laws of empirical psychology is true also for the laws of rational psychology? On this view, the requirement is not that *without exception* suitably situated people who are watching their sons know that they are, but is rather that that *other things being equal* suitably situated people are in this position. Another possibility is to weaken the requirement in a different way and say that it is not that suitably situated people who are watching their sons are *always* such that they know that they are, but is rather that suitably situated people who see their sons are *mostly* such that they know that they are. These two ideas—the other things being equal idea and the mostly idea—are different, and it would take further discussion to decide which of them is right, or indeed to decide whether a suggestion distinct from both might not be

preferable. I will not try to resolve the matter here. The main point is only that it does not seem impossible to accommodate cases of defeat by being careful to formulate (c), or the intention behind it, in the right way.[14]

Third, one might offer a simpler reason to reject the idea that it is a rational requirement that suitably situated people who are watching their sons know that they are, viz., that it is not plausible that rationality really cares about watching sons in particular.

However, while it is implausible to think that rationality cares about watching sons in particular, it is not implausible to see the requirement we have been working with as a specific instance of a more general requirement. This is not to deny the difficulty in formulating the general requirement. For example, suppose one suggested that it is true of *all* mental states that, mostly, if one is in them, and one is suitably situated, one will know that one is in them. This will obviously not do. Some mental states—e.g., those postulated by psychology or linguistics, or by Freudian psychoanalysis—do not exhibit this property. Or suppose instead one suggested that it is true of all *conscious* mental states that, mostly, if one is in them, and one is suitably situated, one will know that one is in them. This is more plausible but is nevertheless subject to counterexamples. Consider cases in which one is completely absorbed in looking at a particular object; for example, suppose I am seeing my son doing somersaults and am totally absorbed in his

[14] One important consequence of weakening the rational requirement in the second way mentioned in the text is that if this is done, the position I am describing will not face the anti-luminosity argument developed by Williamson (2000). If seeing my son were a luminous condition in Williamson's sense, then in every case in which I satisfy that condition I would be in a position to know that I am in it. If this is true simply in most cases, however, seeing my son is not luminous in Williamson's sense, even in a rational individual. For some discussion of this argument, see Silins, this volume, and Greenough, this volume.

doing somersaults. That state—the conjunctive state of seeing X do somersaults and being totally absorbed in X's doing somersaults—is conscious in the main senses that dominate philosophy of mind. It is *phenomenally* conscious because there is something it is like to be in it, and it is *access* conscious because it controls rational thought and action. On the other hand, it is not true that if I am in it, I will know that I am in it on the condition that I am suitably placed and am rational. The reason is that if I ask myself whether I am in this state, I will inevitably shift some of my attention from my son's doing somersaults to the question of whether I am in the state. But in that case, I will no longer be in the conjunctive state of seeing my son and being totally absorbed in doing so! So the state of seeing my son do somersaults and being totally absorbed in his doing somersaults does not have the property that the state of seeing my son does.

Instead of saying that in a rational individual the property I have attributed to seeing my son is true of all mental states, or true of all conscious mental states, I think the better thing to say is that in a rational individual this property attaches to *many* states that the subject is in, without trying to specify the kind of states in question. So on this view it is a rational necessity that many mental states are such that that mostly when one is in them, and one is suitably placed, one will know that one is in them. This is certainly a bit of a mouthful. But one can capture the general idea more pithily by saying that it is a requirement of rationality that individuals have a capacity for expertise about their own first-person present-tense psychological states. Expertise does not mean omniscience. Expertise tolerates a certain level of ignorance and error. Someone can be an expert on Russian literature without knowing everything about Russian literature. And it need not be that there is some fact about Russian literature that every expert knows. Rather

expertise means knowing a lot of facts about Russian literature, and perhaps knowing fairly deep organizational facts as well. Likewise, I am suggesting, a suitably situated rational subject can be an expert on their first-person present-tense psychological facts consistently with failing to know (and failing to be in a position to know) any few of these facts, and consistently with having some false beliefs about them.

Finally, one might think it implausible that rational people should exhibit this sort of self-expertise in the first place. But as against this, I think there are examples that support it.[15] Suppose George goes on Mastermind with the stated topic of Russian literature. The compère asks various questions: "Did Tolstoy write *Anna Karenina*?" "Did Chekhov die young?" "Did Dostoyevsky get sent to Siberia?" and so on. George replies as best he can, but after a while a startling fact emerges: not only is George not an expert on Russian literature, he is an almost complete ignoramus. What explains George's situation? Well it might be that he does not understand the questions, or that he is irrational, or that he is distracted. But in this case it need not be anything like this. Another possibility is that he simply has not had enough exposure to Russian literature and has gone on Mastermind only because of some false belief that he has—e.g., about what the rules of Mastermind are. An explanation along these lines makes sense of what George is doing, even though of course in another way we might wonder what on earth he was thinking.

[15] Examples such as these provide support for the idea that the self-expertise is a requirement of rationality, but they do not by themselves provide an explanation of why this is so. However, the question of why the self-expertise rule is a rule of rationality is not something I will attempt to answer in this chapter.

All right so far; but suppose now George goes on Mastermind with the stated topic of his own first-person present-tense psychological states. (A somewhat strange topic! But recall that just as the topic of Russian literature is in some sense a system of propositions, so too is the topic of George's first-person present-tense psychological states.) The compère again asks various questions: "Do you see your son?" "Do you have a tingle in your elbow?" "Do you believe that there will be a third world war?" and so on. George replies as best he can, but once again a startling fact emerges: not only is George not an expert on his first-person present-tense psychological states, he is an almost complete ignoramus. What explains George's situation this time? Well, this time there is no possibility of saying that George has not had enough exposure to his first-person present-tense psychological states. Either it makes no sense to talk of exposure in this case at all or if it does George has been exposed as much as it is possible to be. But if the lack of exposure possibility is ruled out, our set of options for explaining George's situation is somewhat narrow. It remains possible of course that George does not understand the questions put to him, or that he is distracted. But assuming that we can rule out these hypotheses, it would seem that the only option left is to suppose that he is irrational. To put it another way: on the assumption that George is suitably placed, then either he has considerable expertise on the topic of his first-person present-tense psychological states or he is irrational. But he does not have such expertise. So he must be irrational.

7

I said earlier that our Austin-inspired response to the evidence argument had the right form but that there were two issues that needed to be confronted before it was adopted.

The first, which we have just been discussing, is that the explanation for the fact that I know should itself be a rational explanation. The second, as we noted above, returns us to Dretske's objection to the foot-stamping view, the objection that we "can't do epistemology by stamping our feet." We have so far made no real attempt to answer that objection. However, it might be insisted that we must make an attempt to answer it. For in developing our own response to the evidence argument, we have relied on the foot-stamping view. If you rely on something, you cannot very well note a major objection to it, and say next to nothing about how that objection might be met.

How then might Dretske's objection be met? Well, what *is* that objection? As we saw earlier, if the objection is that it is a mistake to assert that one knows without offering an identifiable way that one knows, it does not seem very forceful. For a proponent of the foot-stamping view (and so a proponent of the Austinian foot-stamping view) *has* an identifiable way, namely, the very proposition at issue. There is, however, a different way to take Dretske's objection. On this different way of understanding it, the objection does not rely on the highly questionable claim that one cannot ever offer the very proposition known as evidence for what one knows. Rather it relies on the very plausible claim that normally one cannot do this; the objection then challenges the foot-stamper to explain why it can be done in this case. To put it differently, when Dretske says "we can't do epistemology by stamping our feet" what I think he means (or at least is best taken to mean) is this: we do not normally do epistemology by stamping our feet; why then can we do it in this case?

Dretske's objection as understood here is closely related to an objection to the foot-stamping view that comes up in a dialectical context. When you normally ask "How

do you know that p?" you are fishing around for a proposition distinct from p. Now a proponent of the foot-stamping view is committed to the idea that at least in certain contexts this is mistaken; at least sometimes it is appropriate to repeat the proposition p in answer to a question of that form. How can this be so if the question "How do you know that p?" presupposes that an answer to it is distinct from p? What a proponent of the foot-stamping view must hold here is that this presupposition is a pragmatic presupposition, something that is cancelled in the present (highly eccentric) case. However, even if that is right, it would seem that there is a challenge here nevertheless, viz., *why* is the presupposition cancelled in this case? Unless that further question is answered, we have as yet no good answer to the opponent of the foot-stamping view. It is this challenge that I think is the key idea behind Dretske's objection.

If that is the key idea behind Dretske's objection, however, then the proposal I have advanced meets it. For on this proposal, there *is* a reasonable account of why the presupposition is cancelled. Suppose I know that I am seeing my son, and suppose that the explanation for this fact is that it is rational necessity that people who are seeing their sons know that they are, that I abide by this necessity, and that I am suitably situated and that I am seeing my son. Suppose now I am asked "How do you know that you are seeing your son?" As we have seen, this question can be understood in various ways. If it is understood as a request for explanation for the fact that I know, the answer is the one I just gave. If it is understood as a request for evidence of a nontrivial sort, then it is a question whose presupposition I reject. There is (or need be) no evidence of a nontrivial sort that supports what I know. On the other hand, it is in this case reasonable of me to reject the presupposition. Mostly, if one fails to have evidence of a nontrivial sort for

what one claims to know, this would be an epistemic embarrassment, for it would suggest intuitively that one has no rational basis for one's claim to knowledge. But not in this case: for in this case I have arrived at my knowledge by abiding by rational principles. Finally, suppose the question is understood as a request for evidence of either a trivial or a nontrivial sort, i.e., as a question that permits a trivial answer. In that case, I can simply repeat the proposition that I know, and so provide that trivial answer. Once again, in the circumstances this is a reasonable thing for me to do. Mostly, if one can only provide evidence of a trivial sort for what one claims to know, this would be an epistemic embarrassment, for it again would suggest that one has no rational basis for one's claim to knowledge. But not in this case: again in this case, I have arrived at my knowledge by abiding by rational principles.

8

The Austin-inspired foot-stamping response I have offered to the evidence argument alleges in effect that the argument exhibits a confusion of evidential and explanatory concerns. This response was of course always on the cards. For the evidence argument is framed in terms of "how do you know" questions, and such questions notoriously oscillate between these two readings. What I have tried to bring out in the course of our discussion, however, is that in drawing the evidence/explanation distinction, it is crucial to acknowledge that there is a significant admixture of rationality in the explanatory side of the ledger. Unless one does that, one will so far as I can see have no response to the problem about natural and metaphysical necessity that I reviewed in sections 4 and 5. Nor will one have a response to Dretske's objection to the foot-stamping view, which I mentioned in section 3 and again in section 7. I have also said a little bit about what the

underlying requirement of rationality is, and in particular that it involves the idea that rational people exhibit a certain capacity for self-expertise.

The picture of introspection that emerges from the response to the evidence argument might be summarized in the following terms. Suppose we say that I know that I am watching my son by *introspection*. What do we mean, i.e., what is the 'by'-phrase doing? Answer: at least in central cases, it is indicating that there is a distinctive answer to the question "how do I know?" What is distinctive about the answer is this:

1. If we fix it so that the 'how' question has its evidence-seeking reading, the answer will be trivial if there is one at all (and there will only be a trivial answer if I know merely by introspection).

2. If we fix it so that the 'how' question has its explanation-seeking reading, the answer will involve the idea that it is requirement of rationality that people exhibit a capacity for self-expertise, and that I am rational in at least this respect, and that I know what I do through an exercise of this capacity.

This kind of picture leaves a lot of questions unanswered: (1) it does not say what the underlying scientific facts are; (2) it leaves open how extensive knowledge by introspection might be; (3) it does not (so far) take into consideration cases in which we arrive at introspective knowledge by inference; (4) it does not (so far) take into consideration cases in which I believe mistakenly that I am watching my son; (5) it does not speak to the sense in which the notion of introspection may be introduced as a mode of attention; (6) it does not connect up this notion introspection with consciousness in any of its many varieties; and finally (7) as we have seen it does not go very far into the analysis of what the requirement is exactly.

All of these questions deserve detailed discussion, but we do not have the space for that here. Instead I will close the chapter by briefly contrasting the position I have outlined with what is an apparently different response to the evidence argument. This alternative response focuses, not on what is meant by the 'how' question, but rather on what is meant by 'evidence'. In particular, the response says, there are two slightly different notions of evidence in play. The first notion is the one that we have been using all along. According to this notion, something is evidence just in case (a) I am aware of it, and (b) it supports what I know—we may call this evidence 1. But, says a proponent of this response, there is also a different notion of evidence that is sometimes in play in philosophical discussions. According to this notion, a proposition is evidence just in case the truth of the proposition makes it appropriate for me to believe what I know—we may call this evidence 2.[16]

Armed with this distinction, a different response to the evidence argument comes into view, viz., that it does not go through even if the question "how do I know" is given its evidence-seeking interpretation. For it might be that the evidence sought is in evidence 2 rather than evidence 1:

[16] In introducing the concept of evidence 2, I am following Pryor 2005, though his main focus is on justification rather than evidence. Likewise, the concept of evidence 1 is very closely related to an idea about evidence (and justification) which Pryor (2005, 189) calls "the Premise Principle" and states as follows: "The only things that can justify a belief that P are other states that assertively represent propositions, and those propositions have to be ones that could be used as premises in an argument for P. They have to stand in some kind of inferential relation to P: they have to imply it or inductively support it or something like that."

E1. If I know that I am watching my son do somersaults, there is a nontrivial answer to the evidence 2-seeking question "How do I know that I am watching my son?"

E2. If there is a nontrivial answer to the evidence 2-seeking question "How do I know that am watching my son?" that answer must be a proposition that (a) I am aware of; (b) "tells me" that I am watching my son; and (c) is distinct from the proposition that I am watching my son.

E3. There is no proposition that meets conditions (a), (b), and (c).

E4. Therefore, I do not know that I am watching my son.

The problem with the argument in this version—version 4 of the evidence argument, as we may call it—is that E2 is false; the conditions set out in E2 will not apply. In particular, that I am watching my son is on this view a proposition that makes it appropriate to believe what I know, but it does not follow from this that I am aware of that proposition or of its truth.

9

How plausible is this 'two notions of evidence' response, and how does it compare with the response I have given? Clearly it is only a legitimate response if evidence 2 is a legitimate notion. Is it? Well, a major problem for evidence 2 is that it is very unconstrained. For example, there is a sense in which the total physical state of the universe at a certain time will make it appropriate to believe that I am watching my son at that time (if this is appropriate for me), but the total state of the universe is not my

evidence on any reasonable conception of evidence. In response, what proponents of evidence 2 will say[17] is that the fact that makes it appropriate for me to believe I am watching my son must be available to me. In other words, on this refinement of the notion, evidence 2 is a proposition (a) whose truth makes it appropriate for me to believe what I know and (b) which is available to me. So for example, if I am watching my son, and this fact is available to me, then it is appropriate for me to believe that I am watching my son.

At this point, however, the key question is this: What could this 'availability' come to? So far it is only a label for a relation I must bear to a fact if that fact makes it appropriate to believe that that fact obtains and is therefore my evidence. What is it? One answer explains availability directly in terms of awareness, and in turn explains this either in terms of perceptual or inferential model of introspection. However, first, this is an option that we set aside long ago (see section 2), and second, it makes it a bit of a mystery why the detour through evidence 2 was required.

A second answer explains availability in terms of rationality. In fact this is exactly what Pryor—whose discussion of justification we have been adjusting to our own concerns—does when he explains the notion of availability. First, he connects availability and the idea of a ground:

> A justifier is available to you at a given time—it will be something you can "take as" a reason—if it's something that could then ground a belief of yours. (Pryor 1995, 195)

[17] Here again I am following Pryor (2005, 195).

He then argues that the relevant notion of grounding can be explained in terms of following a rule or norm of rationality:

> What does it take for your belief to be grounded on some fact or condition C that you're in? A natural thought is that your belief counts as so grounded iff it's formed (or sustained) in a way that's guided by the epistemic norm "When in C, believe P." If that's right, then the best way to understand the grounding relation is by inquiring into what it takes to be guided by such a norm. (ibid., 195)

I do not want to assess this proposal in detail here. Rather I want simply to note that, if 'availability' is understood this way, then it looks as if the 'two notions of evidence' response is a variation on the response I have given, for a norm as Pryor intends it is just the notion of a requirement of rationality. In turn, if that is right, we are in the presence here, not of two responses to the evidence argument, but of one. On the other hand, if what I have argued is right, that response is successful.[18]

[18] I gave (in some cases rather different) versions of this chapter to audiences at ANU, Melbourne, Oxford, and Sydney. I am very grateful to the extremely good suggestions I received from the audiences at those universities. I am also indebted to specific suggestions from Alex Byrne, David Chalmers, Mark Jago, John Maier, Tori McGeer, and Declan Smithies. I am very grateful also to two anonymous referees for Oxford University Press who gave extremely helpful comments on the chapter that (I think) improved it considerably.

References

Armstrong, D. M. 1968. *A Materialist Theory of the Mind*. London: Routledge and Kegan Paul.

Austin, J. L. 1961. "Other Minds." Pp. 76–116 in *Philosophical Papers*. Oxford: Clarendon Press.

Broome, J. 2005. "Does Rationality Give us Reasons?" *Philosophical Issues* 15: 321–337.

———. 2007. "Requirements." In *Homage a Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*, edited by T. Ronnow-Rasmussen et al.

Byrne, A. 2005. "Introspection." *Philosophical Topics* 33: 79–104.

Dretske, F. 2003. "How Do You Know You Are Not a Zombie?" Pp. 1–14 in *Privileged Access: Philosophical Accounts of Self-Knowledge*, edited by B. Gertler. Burlington: Ashgate Publishing.

Fine, K. 2004. "The Varieties of Necessity." Pp. 254–282 in *Conceivability and Possibility*, edited by T. Gender and J. Hawthorne. Oxford: Oxford University Press.

Martin, M. G. F. 1998. "An Eye Directed Outward." Pp. 99–121 in *Knowing Our Own Minds*, edited by C. Wright, Barry C. Smith, and Cynthia MacDonald. Oxford: Clarendon Press.

Pryor, J. 2005. "There Is Immediate Justification." Pp. 181–201 in *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa. Oxford: Blackwell Publishing.

Shoemaker, S. 1996. *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.

Siewert, C. 1998. "Self-Knowledge and Rationality: Shoemaker on Self-Blindness." Pp. 131–146 in *Privileged Access: Philosophical Accounts of Self-Knowledge*, edited by B. Gertler. Burlington: Ashgate Publishing.

Stoljar, D. 2004. "The Argument from Diaphanousness." Pp. 341–390 in *New Essays in the Philosophy of Language and Mind,* supplemental volume of *The Canadian Journal of Philosophy*, edited by M. Ezcurdia, R. Stainton, and C. Viger. Calgary: University of Calgary Press.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.