

Chapter 4:

Wittgenstein's Suspicion and the

Insignificance of Philosophical Scepticism

'Tis evident, that so extravagant a Doubt as that which Scepticism may seem to recommend, by destroying every Thing, really affects nothing. -- David Hume, letter to John Coufts of May 8, 1745¹

I am sitting with a philosopher in the garden; he says again and again "I know that that's a tree", pointing to a tree that is near us. Someone else arrives and hears this, and I tell him: "This fellow isn't insane. We are only doing philosophy." -- Wittgenstein, On Certainty, 467

A. Accommodating the Sceptic: The Stine/Lewis View.

So, as I've argued, contextualism can be used to defeat the bold sceptic. But what are the results of applying contextualism to an evaluation of the not-so-bold sceptical stance? Stine and Lewis, who advocate contextualist forms of RA, both apply contextualism in a sceptic-friendly way. The sceptic to whom their application is friendly is the not-so-bold sceptic. We've already seen this more explicitly in the case of Lewis in our discussion of his proposed "rules of

¹From David Hume, An Enquiry Concerning Human Understanding / A Letter from a Gentleman to his Friend in Edinburgh, ed. Eric Steinberg (Indianapolis: Hackett, 1977); p. 117.

accommodation" (Chapter 3, section I). But Stine seems to be equally friendly to the not-so-bold sceptic. She writes:

In truth, in some sense skepticism is unanswerable. . . . The relevant alternative view does provide a kind of answer to the skeptic -- the only kind of answer which can be given. But the skeptic has an entering wedge, and rightly so. It is an essential characteristic of our concept of knowledge that tighter criteria are appropriate in different contexts. It is one thing in a street encounter, another in a classroom, another in a law court -- and who is to say it cannot be another in a philosophical discussion? . . . We can point out that some philosophers are very perverse in their standards (by some extreme standard, there is some reason to think there is an evil genius, after all) -- but we cannot legitimately go so far as to say that their perversity has stretched the concept of knowledge out of all recognition -- in fact they have played on an essential feature of the concept. On the other hand, a skeptical philosopher is wrong if he holds that others are wrong in any way -- i.e., are sloppy, speaking only loosely, or whatever -- when they say they know a great deal. (p. 254)

Thus, the sceptic is wrong if she wants to assert that our ordinary knowledge attributions are false. Thus, the falsity of bold scepticism. Likewise, according to Lewis, the "commonsensical epistemologist" is wrong if he asserts that the sceptic's claim that we know very little is false (see Lewis, p. 355). But, for the most part, everybody's happy. We can go on in our everyday lives truly claiming to know most of the things we think we know, and the sceptic can go on truly saying that we

know hardly anything at all. The "sense" in which contextualism provides an answer to scepticism is that it answers the bold sceptic; the "sense" in which scepticism is unanswerable is that the not-so-bold sceptic -- who claims in philosophical contexts that we don't know but who refrains from attacking the truth value of our ordinary claims to know -- is simply right. Some find this result to be very attractive -- so attractive, in fact, that it is thought to be a considerable advantage to the contextual theories that they allow for this result. The attraction, it seems, is supposed to consist largely in the fact that we can explain why sceptical arguments have such a strong pull on us while still holding on to the view that very many of our ordinary knowledge attributions are true.

But it can easily seem that this treatment is far too easy on the sceptic. Stine writes that "a sceptical philosopher is wrong if he holds that others are wrong in any way," but if the standards for knowledge do go up in philosophical discussions (or perhaps in philosophical discussions of what we can know) such that the sceptic can truly say we know nothing or almost nothing, then one might think that the "commonsensical epistemologist" would be saying something false not only

in saying that the sceptic is wrong, but also if he asserted, in such a context, that he knew a great deal. (Perhaps Moore was saying something false when he claimed to know this and that! He was, after all, giving philosophy lectures on the topic of scepticism.) Thus, the sceptic could perhaps truly say that the common sense philosopher is wrong.

This is all very tricky. Much depends (as we saw in section I of Chapter 3) on just what determines the standards for true knowledge assertions. If it is just the speaker's intentions, then perhaps the commonsense philosopher (intending relatively low standards²) and

²In the case of Moore, I get the feeling (and this feeling is shared by others I've talked to who have read Moore) that Moore was intending very high standards but thought he lived up to them with respect to certain of his beliefs about the external world. The matter is tricky, however, because Moore seems to cite as a reason for thinking that he knows certain things that it would be absurd to say otherwise. Consider the following well-known passage from Moore's "proof of an External World": "I certainly did at the moment know that which I expressed by the combination of certain gestures with saying the words, 'Here is one hand and here is another'. I knew that there was one hand in the place indicated by combining a certain gesture with my first utterance of 'here' and that there was another in the different place indicated by combining a certain gesture with my second utterance of 'here'. How absurd it would be to suggest that I did not know it, but only believed it, and that perhaps it was not the case!" (p. 146). One might think that Moore is being led to claim that he does know by the absurdity of suggesting in a normal context that he does not know that there is a hand present. Thus, one might think that it is relatively low standards that are governing Moore's use of 'know',

the sceptic (intending higher standards) are talking right past each other, each saying something true, their only mistake being their false belief that they are saying something at odds with the other's assertions. But if, as seems more plausible, more objective features of the situation (like what has already been said in the conversation and the setting: on the street vs. in a law court) set the standards, then it seems quite possible that if the standards allow for the sceptic to truly say we know hardly anything, then the commonsense philosopher, whatever standards he intends, is saying something false when he utters his knowledge claims.

Thus, despite Stine's assurances that the sceptic is wrong if he claims that others are wrong, I don't see how Stine has blocked the possibility that in a philosophical discussion, the sceptic simply wins: she is saying something true when she claims that we don't know and when she claims that her philosopher opponent is saying something false. Lewis seems to accept this; he writes that in such a context, "the commonsensical philosopher must concede defeat" (p. 355). The only sceptical move Stine seems to block is the sceptic's turning her results on ordinary, out-on-the-streets

after all.

knowledge claims. And many sceptics, I think, wouldn't be very interested in the truth values of these "vulgar" assertions.³

There are two ways in which an anti-sceptic might try to block more, and, in particular, to block the result that the standards go up so high in a philosophical debate involving a sceptic that the sceptic ends up saying something true when she denies that we know anything about the external world. One way is to make the conversational strength-governing mechanisms less sceptic-friendly. We've seen one way that could be done in section I of Chapter 3. In general, one may try to argue that it is quite difficult for the sceptic to establish and maintain the kind of contexts needed for her to truly deny that we know anything about the external world. One problem with this approach is that much of the motivation for contextualism comes from cases in which it seems that the standards for true

³Of course, as Lewis points out, the contextualist can object to the idea that the sceptic's standards are more legitimate than are the more ordinary standards (see Lewis, pp. 354-355). Thus, the sceptic might be blocked from writing off these "vulgar" knowledge claims as not being as legitimate as her own. All I am claiming on the sceptic's behalf is that, as far as the Stine/Lewis view goes, the sceptic may well win the day against his philosophical opponent, and that this may be all that many philosophical sceptics are really interested in.

knowledge attributions are raised fairly easily.⁴ To make the strength-governing mechanisms too resistant in general to standard-raising would be to sap the support of contextualism. The second -- and what seems to me more promising -- way of blocking the not-so-bold external world sceptic is to argue that, although in general, the standards of knowledge attributions can be fairly easily raised, there is something problematic about raising them so high as to make it false to attribute external world knowledge to anyone.

But we should consider the question of whether or not blocking the not-so-bold external world sceptic is at all an important or interesting task. Just how much of the interest of philosophical scepticism remains after it is limited to the not-so-bold variety is open to dispute. Stewart Cohen seems to think that most of the interest is thereby lost; he writes:

What is truly startling about skepticism, is the claim that all along, in our day to day lives, when we have claimed to know things, we have been

⁴Consider, for instance, the Bank Cases from Chapter 1. I do not mean to be suggesting that the standards can be raised by the mere mentioning of a sceptical hypothesis. Perhaps it is not that easy to raise the standards. In the Bank Cases, there are several factors which seem to favor the raising of the standards. But the contextualist should not make it so difficult to raise the standards that it becomes doubtful that they are raised in the examples similar to the Bank Cases.

wrong -- we have been expressing propositions that are literally false.

Contextualism, according to Cohen, deprives the sceptic of this startling, bold claim, because it shows that our everyday knowledge attributions -- properly interpreted -- are correct. The propositions we actually express and have been expressing all along are literally true. We do know relative to the standards that ordinarily govern those attributions.

Cohen goes on to write:

Now it is of considerably less interest (although not devoid of interest) that there is some other proposition involving stricter standards that is false. It is not the case that we know relative to the skeptical standards. This a fallibilist readily concedes.

Marilyn Adams, on the other hand, in the passage we looked at in Chapter 1, section D, seems to think that sceptics have traditionally been uninterested in our "low", everyday claims to know, and have all along been interested mainly in our lack of "high" knowledge. I do not want to come down on one side or the other on this point until after we investigate just what the sceptic's "high" knowledge is. Antecedently, it seems that there is plenty of room for the sceptic's not-so-bold

conclusion to be very interesting indeed. It will be my contention in this Chapter that this interest vanishes

⁵Cohen, "How to Be a Fallibilist," p. 117.

upon closer inspection.

One way the not-so-bold sceptic's contention can be uninteresting is for it to be meaningless. Accordingly, we shall begin by investigating Wittgenstein's suspicion that philosophical sceptics' denials of knowledge are just that. While I will not be able to find Wittgenstein's suspicion to be true, it will provide a good start towards illuminating the way in which I think the not-so-bold sceptic's denials are uninteresting.

B. Wittgenstein's suspicion. One of the main threads running through Wittgenstein's last notes (On Certainty), is his deep suspicion that claims to know are only meaningful in certain circumstances, and that "philosophical" uses of 'know' -- the philosophical sceptic's claims that we don't know this or that or the anti-sceptic's (particularly Moore's) claims in the face of sceptical doubts that he does indeed know this or that -- are meaningless or don't have a determinate meaning. Wittgenstein writes:

"I know that that's a tree." Why does it strike me as if I did not understand the sentence? though it is after all an extremely simple sentence of the most ordinary kind? It is as if I could not focus my mind on any meaning. Simply because I don't look for the focus where the meaning is. As soon as I think of an everyday use of the sentence instead of a philosophical one, its meaning becomes clear and ordinary.

Just as the words "I am here" have a meaning only in certain contexts, and not when I say them to someone who is sitting in front of me and sees me clearly, -- and not because they are superfluous, but because their meaning is not determined by the situation, yet stands in need of such determination. (OC, 347-348)

Early on in On Certainty, Wittgenstein gives the following illustration of a knowledge claim which is not made in suitable circumstances and is, therefore, nonsense:

I know that a sick man is lying here? Nonsense! I am sitting at his bedside, I am looking attentively into his face. -- So I don't know, then, that there is a sick man lying here? Neither the question nor the assertion makes any sense. (OC, 10)

Now, Wittgenstein admits that such sentences as "I know that that's a tree" and "I know that a sick man is lying here" can be used meaningfully in some easy-to-imagine circumstances. In fact, this serves as his explanation for why sentences like these may seem to be meaningful when uttered in unsuitable circumstances:

"I know that there's a sick man lying here", used in an unsuitable situation, seems not to be nonsense but rather seems matter-of-course, only because one can fairly easily imagine a situation to fit it. (OC, 10)

Then why don't I simply say with Moore "I know that I am in England"? Saying this is meaningful in particular circumstances, which I can imagine. But when I utter the sentence outside these circumstances, as an example to shew that I can know truths of this kind with certainty, then it at once strikes me as fishy.--Ought it to? (OC, 423)

Well, how shall we answer Wittgenstein's question? Should such knowledge claims strike us as fishy? And even if they are somehow fishy, what manner of fishiness is this? Should we accept the suspicion that such claims are meaningless or that their meaning is in some problematic way underdetermined?

Moore's "Defense of Common Sense" is, of course, in the background of Wittgenstein's reflections, but so is Norman Malcolm's 1949 attack on Moore, "Defending Common Sense" (henceforth abbreviated as DCS).⁶ In DCS, Malcolm responds to Moore's claim to know that "Here's a hand" as follows:

At this moment I am holding a pen, there is a desk before me, I am seated in a chair, and through the window I see a near-by tree. Let us imagine that there is another person in this room who has a clear view of me seated in this chair, before this desk, with this pen in my hand, and who has an unobstructed view of that near-by tree. Moore's assertion implies that it would be correct for me to say to that person "I know that I am holding a pen," "I know with certainty that I am sitting in a chair and before a desk," "It's perfectly certain that that [pointing at the tree] is a tree." I contend that I should misuse language if I were to make any of these statements. (DCS, p. 202)

⁶See pp. 171-172 of Malcolm's "Moore and Wittgenstein on the Sense of 'I Know'", where Malcolm relates the impact his DCS had on Wittgenstein.

Certainly there is something fishy about my behavior if, out of the blue, in the middle of an otherwise normal conversation, I say, "I know that I'm in Los Angeles." Malcolm seems to be imagining such an out of the blue situation (a situation in which the audience has no idea why the speaker is claiming to know the thing in question) when he imagines saying "I know that I am holding a pen" to the person who clearly sees that he is holding a pen. But we may wonder whether Moore's "philosophical" use of "I know" (in which we know why Moore is claiming to know -- in order to defend common sense from sceptical attacks) should be given the same treatment as such an out of the blue use. Malcolm, after contrasting his holding a pen example with some correct uses of "I know that that's a tree," lists three features of a situation which are present in the correct uses:

- (1) There is in each case a question at issue and a doubt to be removed.
- (2) In each case the person who asserts "I know that that's a tree" is able to give a reason for his assertion.
- (3) In each case there is an investigation which, if it were carried out, would settle the question at issue. (DCS, p. 203)

Malcolm claims that "all of these features are missing when Moore says in a philosophical context 'I know that that's a tree'" (DCS, p. 203). In fact, in DCS his case

against Moore's use of this sentence consists almost entirely of arguing that these features are missing when Moore says the sentence. So according to Malcolm, Moore's "philosophical" use of "I know" fails to be meaningful⁷ for much the same reason as does the out of the blue use. Throughout On Certainty, as Wittgenstein tries to locate what is wrong with Moore's claims to know, the lack of features very similar to the three listed above play a prominent role. And Wittgenstein, like Malcolm, compares Moore's claims with out of the blue utterances. Malcolm would agree with Wittgenstein's later reflection that saying "I know that that's a tree" is only meaningful "in particular circumstances," and according to Malcolm, the

⁷In DCS, Malcolm does not couch his charge against Moore's uses of knowledge claims in terms of their being meaningless; he instead says that Moore "misused" the expression "I know" and that "Moore's use of those expressions . . . is contrary to their ordinary and correct use" (p. 202). But Malcolm certainly meant more by his charge of "misuse" than that Moore had said something very queer given the circumstances, for Malcolm considers, as a response to his case against Moore, the claim that Moore's use of "I know" is "an odd use but not a misuse", and Malcolm argues that "this reply contains a mistake" (p. 215). Although Malcolm does not use the term "meaningless", he comes very close to saying that Moore's use of "I know" is meaningless when, after arguing that Moore's use is not properly connected with an investigation, he writes, "If you take away from that sentence its connection with actions of investigation you turn it into an empty utterance" (pp. 214-215).

circumstances Moore found himself in as he defended common sense in his famous lectures did not allow for a meaningful use of that sentence. So Malcolm and Wittgenstein (at least in some moods) question the meaningfulness of a) philosophical and b) out-of-the-blue uses of 'I know', largely because such uses lack some of the key features that accompany more "ordinary" knowledge claims.

C. Moore's "Ordinary Sense" Defense. Now I can work up a good deal of sympathy for the thought that the meaning of a use of either type is in some sense "not determined by the situation, yet stands in need of such determination," as I will explain more fully later. Like Wittgenstein, it often "strikes me as if I did not understand" Moore's claim, "I know that that's a tree." But many will go along with Moore's response to such a challenge. In a letter to Norman Malcolm,⁸ Moore describes a conversation he had had with Malcolm as follows: "Sitting in my garden two years ago, I pointed or nodded at the young walnut-tree & said 'I know that that is a tree.'" Moore continues:

⁸Reported in pp. 173-174 of "Moore and Wittgenstein on the Sense of 'I Know'".

You wanted then, and want now, to say that my use of that expression was a "misuse" & "incorrect"; but the only reason you give for saying so is that I used it under circumstances under which it would not ordinarily be used, e.g., under circumstances that there neither was at the moment nor had been just previously any doubt whether it was a tree or not. But that I used it under circumstances under which it would not ordinarily be used is no reason at all for saying I misused it or used it incorrectly, if, though this was so, I was using it in the sense in which it is ordinarily used -- was using it to make the assertion which it is ordinarily used to make.

Many will agree with Moore that the feeling one can get that one doesn't understand what's being said in the suspect uses of "I know" is an illusory feeling, which is explained by the fact that one can't understand why the speaker is claiming to know the thing in question. Moore writes:

If a person, under circumstances in which everybody would see quite clearly that a certain object was a tree, were to keep repeatedly pointing at it & saying "That's a tree" or "I know that's a tree," we might well say that that was a senseless thing for him to say; & even if he said it only once, under such circumstances, we might well say that it was a senseless thing for him to do -- meaning, in all these cases, that it was a sort of thing, which a sensible person wouldn't do, because, under those circumstances, it could serve no useful purpose to say those words. . . . But this is an entirely different thing from saying that the words in question don't, on that occasion, "make sense," if by this is meant something which would follow from the proposition that they were not being used in their ordinary sense.

The very end of the above passage is quite interesting: it shows how important it is to Moore that he was using the words in their ordinary sense. One would think that the charge that Moore's use of the words in question didn't make sense was stronger than anything that would follow from the proposition that they were not being used in their ordinary sense." The charge that his words don't make sense shouldn't just be the charge that they are not used in their ordinary sense, but that they are not used in any meaningful way. One could claim that Moore's use of the words made sense, though he did not use them in their ordinary sense but rather used them in a special, philosophical sense. Malcolm, in fact, does claim something very close to this in his later paper, "Moore and Wittgenstein on the Sense of 'I Know'" (henceforth abbreviated as MW):

It is evident that Moore was not giving any everyday employment to the words "I know. . . ." Nor was he using them in "the ordinary sense" -- there seems to be no such thing. But I certainly do not conclude that Moore was talking nonsense. He was giving a philosophical employment to the words "I know. . . ." He was saying something of deep philosophical interest. (MW, p. 185)

But Moore would have none of this. He might agree that he "was saying something of deep philosophical interest," but part of the reason why he would think that his use was of such great interest is that he was

using the words in their ordinary sense. Thus, Moore would find it essential that he not only be vindicated from the stronger charge (that Malcolm seems to make in DCS) that he was not using the words in any meaningful way, but also from the weaker charge (that Malcolm makes in MW) that he was not using the words in their ordinary sense. Moore would consider it a defeat to avoid the charge of senselessness only by giving up the claim that he was using the words in their ordinary sense.

The Moorean, then, will admit that there is something fishy about my saying out of the blue in an otherwise normal conversation, "I know that I'm in Los Angeles," but will insist that my statement would be both meaningful and true. And Moore's words -- "I know that that is a tree" -- it will be said, are in even better shape, because not only do they express something meaningful and true, but they are also, as Moore points out in his letter to Malcolm, being used

with a purpose -- the purpose of disproving a general proposition which many philosophers have made; so that I was not only using them in their usual sense, but also under circumstances where they might possibly serve a useful purpose, though not a purpose for which they would commonly be used.

Thus, while Wittgenstein suspects that the suspect uses of "I know" are in fact meaningless uses (or uses

in which the meaning isn't fully determined) but appear to be meaningful because we can easily imagine the same words being used meaningfully, Moore insists that the sentences have their ordinary sense in the suspect uses, that the suspect uses are perfectly meaningful, and may appear to be meaningless or senseless only because their purpose is unclear or unusual.

D. A Radical Response to Moore. There's something appealing to what Moore is insisting on. Malcolm and Wittgenstein admit its appeal. After quoting from Moore's letter, Malcolm writes:

These remarks of Moore's are lucid and strong. On first consideration they seem to be exactly right. (MW, p. 175)

And in On Certainty, Wittgenstein certainly flirts with accepting Moore's view. At one point, he writes:

So when Moore sat in front of a tree and said "I know that that's a tree", he was simply stating the truth about his state at the time. (OC, 532)

But Malcolm immediately goes on to attack Moore's claim to be using "I know" in its ordinary sense. And it seems to be Wittgenstein's dominant view in On Certainty that Moore was not saying something meaningful, ordinary, and obviously true. (It is interesting to note that the above passage from On Certainty is

immediately followed by the remark: "I do philosophy now like an old woman who is always mislaying something and having to look for it again: now her spectacles, now her keys."

So why shouldn't we just go along with Moore's response? Is there anything more to go on than Wittgenstein's feeling (that some of us share) that he can't "focus [his] mind on any meaning" when he tries to understand Moore's claims to know (OC, 347) or the "bafflement" and "mental paralysis" that Malcolm claims to have felt when he "could not get hold of the meaning" of Moore's sentence, "I know that that's a tree"? (MW, p. 185).

Malcolm provides a much more thoroughly worked-out Wittgensteinian response to Moore's defence, so let us look at what Malcolm has to say. He writes:

In response to Moore's assertion, in his letter, that he was using the sentence "I know that's a tree" in its ordinary sense, although not in ordinary circumstances, I want to protest: "How can that be? What conception do you have of 'the sense' of a sentence, if you think that the circumstances in which the sentence is spoken have nothing to do with its sense?" (MW, p. 179)

Malcolm suggests that Moore has "a particular philosophical picture of sense or meaning, namely, that the sense of a sentence is, as it were, attached to the sentence and is present whenever the sentence is spoken"

(MW, p. 178), and contrasts this picture with the post-Tractatus Wittgensteinian conviction that "the circumstances in which a sentence is used determine its sense" (MW, p. 178). Malcolm writes:

I confess that I do not have any exact feeling for alleged different "senses" of a sentence. But it is easy to illustrate different ways in which a sentence is used. I conceive of these different uses of a sentence as functions of the different circumstances in which the sentence is spoken or written. If someone wants to maintain, as did Moore, that when a sentence is employed in a number of different ways it nevertheless has the same sense, then I think that this contention is too nebulous to be capable of either proof or refutation; yet at the same time it can be shown to be a misleading philosophical image. (MW, p. 179)

Toward the goal of showing this philosophical image to be misleading, Malcolm lists several situations in which one can correctly assert a sentence of the form "I know that p," and in each case Malcolm says what the assertion "comes to" or "is like." Here are some examples of what such assertions "come to," according to Malcolm:

Here, "I know that not-p" is like saying, "I have the following evidence that not-p, namely. . . ." (p. 180)

"I know it's right" came to the same here as "I checked it." (p. 181)

In this case, "I know that p" is employed like "I don't need to be reminded that p." (pp. 182-183)

Here, "I know that p" is used like "I agree that p." (p. 183)

In this case, "I know that's a tree" is used like "My sight is still good enough to make out that tree." (p. 183)

In this example, "I know that p" comes to much the same as "you can rely on me that p." (p. 183)

Malcolm concludes that "the locution 'I know' does different jobs, achieves different things, in different circumstances" (p. 184). So far, so good. Who could deny that "I know" can do these different jobs? But Malcolm immediately goes on to write:

We could put it like this: Sometimes "I know . . ." can be paraphrased as "I have the following evidence . . ." sometimes as "I have checked it"; sometimes as "I can prove it"; sometimes as "I am certain . . ."; sometimes as "I agree . . ."; sometimes as "You can rely on me . . ."; and so on. (p. 184; emphasis added)

But can "I know" really be paraphrased in these different ways? Let us look at one of the situations Malcolm presents:

There is a committee meeting at which Mr. N was supposed to be present, but so far he has not arrived. One person says, "perhaps he forgot." Another replies, "I know he didn't forget, because he told me only ten minutes ago that he would be here."

In this example the expression "I know that not-p" was used to introduce the presentation of evidence that not-p, evidence that the speaker already had in his possession. Here, "I know that not-p" is like saying, "I have the following evidence that not-p, namely. . ." (MW, p. 184)

Well, saying "I know that not-p" here is something like saying "I have the following evidence that not-p,

namely. . ." But it is also importantly different. If it turns out that the Mr. N of Malcolm's example really did forget the meeting during the ten minutes, then we will say that the person who said "I know he didn't forget" said something false, whereas he would have said something true had he only said, "I have the following evidence that he didn't forget, namely. . ." I don't know if "I know he didn't forget" can be paraphrased as "I have the following evidence that he didn't forget, namely. . ."; that, of course, will depend on how loose one allows a paraphrase to get. But it seems that "I know he didn't forget" does not mean the same thing as "I have the following evidence that he didn't forget, namely. . ."

Would Malcolm agree or would he rather deny that "I have the following evidence that he didn't forget, namely. . ." would mean something different from "I know he didn't forget"? My feeling is that he would reject the question. Malcolm seems to think that there isn't anything more to the meaning of "I know" than what it does, what jobs it accomplishes, in a given situation. He would probably reply that the two expressions would do the same job, and that he doesn't know what else there is to inquire about. He writes:

At this point some philosophers will want to exclaim: "You can speak of different 'paraphrases' if you like. But these different paraphrases merely reflect different 'conversational implications.'" The sense, the meaning of 'I know...' remains the same throughout those different situations you described. Moore was right in declaring that he used the words 'I know that's a tree' in their ordinary sense."

I confess to finding this objection baffling. I have no perception of a constant, one and the same, sense. But of course this is a clumsy way for me to express myself -- as if I knew what to look for but could not find it. What I do perceive is that in one of the cases I described, the speaker who said "I know..." was, in and by saying that, agreeing with the other person; in another case he was not agreeing with but instead reassuring the other; in still another case he was doing neither of these things but instead informing the other that he had convincing evidence on the matter; and so on. All of this lies on the surface, clearly revealed in the different circumstances of the cases. There is nothing here for speculation. The view I want to reject appears to hold that these observable differences in the jobs that "I know..." does in the various cases do not disclose the essential meaning of "I know..." which is there behind the differing circumstances.

My response to this conception is to say that I do not understand what the "essential meaning" of "I know..." is, and that I do not see the necessity for believing in such a thing. How does one recognize it? How is it related to those different jobs that the locution "I know..." does in actual cases? (MW, pp. 184-185)

I have quoted this passage at length because, as we have seen, Malcolm claimed that Moore's picture, according to which there is a single ordinary sense of "I know," could be shown to be a misleading philosophical image, and the above is as close as Malcolm gets, at least in

MW, to showing this. But Malcolm seems to be not so much showing as simply saying that the image is misleading. There is not much more to this passage than Malcolm's profession of bafflement at the thought that there is such a sense, and his inability to perceive or understand such a sense. Of course, as Malcolm points out, we are dealing with two very different pictures of meaning here. So asking Malcolm to show his picture to be correct is perhaps asking too much.

Roughly, on Malcolm's picture, the meaning of a given claim is a function of what Austin calls the illocutionary force of the claim.⁹ Thus, the meaning of a claim to know is determined by the "job" it is being used to do -- to reassure, to introduce evidence, etc. This is contrasted with the view according to which the meaning or sense of the sentence is taken to be prior to the "job" the claim is being used to accomplish, and enters into the explanation of how the claim is able to accomplish that job.

Aside from Malcolm's profession of bafflement, the above passage also contains a challenge. Malcolm asks of the purported single sense of "I know": "How is it

⁹See Alan R. White, "Common Sense: Moore and Wittgenstein," pp. 325-326.

related to those different jobs that the locution 'I know . . .' does in actual cases?" But, of course, those who are inclined to think that there is a single sense of "I know" and who are inclined to appeal to "conversational implications," will not find any special difficulty in explaining how saying "I know that p" in its ordinary sense will have the conversational implications needed to do the various jobs that Malcolm cites. There is, of course, no widely accepted analysis of "I know that p," but it is widely accepted that the truth of p is necessary to the truth of "I know that p," and that it is also necessary that the speaker believe that p, and so part of what one is asserting in asserting "I know that p" is p and that one believes that p. And, interestingly enough, the assertion of p or of the speaker's belief that p is enough to explain how the speaker "does his job" in many of the cases Malcolm presents. For instance, it will be claimed that in the case we have looked at, simply asserting that Mr. N did not forget the meeting would be enough to introduce the speaker's evidence that Mr. N did not forget, because usually, for any p, after it has been said that perhaps p, one does not flat-out assert that not-p (as opposed to saying "perhaps not-p"), unless one

has some evidence that one is prepared to present. Thus, we explain how "He didn't forget" could do the job by appealing to a general fact about a conversational implication of asserting not-p after it has been suggested that perhaps p. And wouldn't "He didn't forget" do just as good a job of introducing the speaker's evidence as did "I know he didn't forget"? So, since "He didn't forget" could introduce the speaker's evidence and since part of what one asserts in saying "I know that he didn't forget" is "He didn't forget," it should not be surprising that "I know he didn't forget" is able to "do this job."

Of course, Malcolm would have none of this. (Though I don't see that he would be able to show that there's anything wrong here.) And perhaps part of his objection would be this. We haven't explained how "I know he didn't forget" does its job by appealing to its implication that "He didn't forget," which can do the job, because we haven't explained how "He didn't forget" could do this job. And now we reach a very important point. When Malcolm claims that it is misleading to think that there is some single sense over and above the uses to which a sentence is put on various occasions (see the above quotations from MW, pp. 178-179), he is

not making a point only about knowledge claims, but about all sentences. Similarly, in the following remarks from On Certainty, Wittgenstein is making a general point about language, and not just about knowledge claims:

A meaning of a word is a kind of employment of it. (OC, 61)

Isn't the question "Have these words a meaning?" similar to "Is that a tool?" asked as one produces, say, a hammer? I say "yes, it's a hammer". But what if the thing that any of us would take for a hammer were somewhere else a missile, for example, or a conductor's baton? Now make the application yourself. (OC, 351)

Just as "I know he didn't forget" can be used to do different jobs (eg., to introduce evidence, to reassure), so too can "He didn't forget" be used to do different jobs. And as long as Malcolm (and Wittgenstein) are relying on a general point about the relation of meaning to use, we should expect them to think it just as misleading to believe that there is some single sense "behind" the many uses to which "He didn't forget" can be put as it is to think that there is such a single sense behind the uses of "I know he didn't forget." Thus, just as with "I know he didn't forget," there will be no explanation needed of how "He didn't forget" does its various jobs, since there is no

single sense behind these various uses. It does these jobs, and that's all there is to it.

And as long as general views about the relation between meaning and use are behind the denials that sentences are not being used in an ordinary sense, we need not only worry about out of the blue and philosophical uses of "I know that that's a tree" and "I know that this is a hand," but also about such uses of "That's a tree" and "This is a hand." If these latter are being used in a non-ordinary way (to accomplish jobs for which they are not ordinarily used), we cannot claim that they nonetheless have their ordinary meaning, and if the purpose for which they are being used is unclear, we cannot claim that they nonetheless have a clear meaning. And it seems that Wittgenstein does worry about the meaningfulness of these latter. In fact, it is the simple "That's a tree" which Wittgenstein explicitly compares to an out-of-the-blue use of another sentence:

Someone says irrelevantly "That's a tree". He might say this sentence because he remembers having heard it in a similar situation; or he was suddenly struck by the tree's beauty and the sentence was an exclamation; or he was pronouncing the sentence to himself as a grammatical example; etc., etc. And now I ask him "How did you mean that?" and he replies "It was a piece of information directed at you". Shouldn't I be at liberty to assume that he doesn't know what he is

saying, if he is insane enough to want to give me this information?

In the middle of a conversation, someone says to me out of the blue: "I wish you luck." I am astonished; but later I realize that these words connect up with his thoughts about me. And now they do not strike me as meaningless anymore. (OC, 468-469)

I take it that "I wish you luck" did strike Wittgenstein as meaningless before he realized how these words connected up with the speaker's thoughts about him, and that Wittgenstein is hinting that the "irrelevant" utterance of "That's a tree" is also meaningless unless it can be connected up with, eg., the speaker's being struck by the tree's beauty -- ie., unless it is not irrelevant, but is rather used as an exclamation. A few sections earlier, Wittgenstein writes:

This is certainly true, that the information "That is a tree", when no one could doubt it, might be a kind of joke and as such have meaning. (OC, 463)

The implication here seems to be that in the circumstances in question (a tree is plainly in view), if the sentence "That is a tree" is not being used as a joke, then it does not have meaning.

Now Wittgenstein hints at many types of grounds for claiming that Moore's knowledge claims are meaningless, and we have only begun to mine On Certainty for the starts of such arguments. These are just hints, however, and Wittgenstein does not mount anything like a

sustained case for his suspicion. And many of his hints can be easily resisted by those who are attracted to the view that though Moore was using knowledge claims in unusual circumstances and for unusual purposes, he was nevertheless using them in their ordinary sense to say something obviously true. In many places in On Certainty, then, Wittgenstein will be written off as merely pointing out how odd the circumstances are for Moore's knowledge claims. Such oddness has been already conceded. And the main method of resistance (hinted at by Wittgenstein and developed by Malcolm) to Moore's claims of using "I know. . ." in its ordinary sense seems to be a general scepticism about their being any single sense behind the diverse uses of any sentence.

If this is how things stand, if so much weight is being put on this general scepticism about meaning behind uses, then it seems that very few people who aren't already confirmed Wittgensteinians will be tempted by Wittgenstein's suspicion. If, in his garden, Moore had only said "That's a tree" while pointing at a tree, and then had claimed to be using the sentence "That's a tree" in its ordinary sense to say something obviously true, I would find this claim to be not only appealing but absolutely compelling. No matter how in

the dark I was as to why Moore was saying "That's a tree," I would feel very confident of my understanding what he was saying. He would be saying what is ordinarily meant by the sentence. And since it would be a tree that he would be pointing at, he would be saying something obviously true.

E. Contextualism and Wittgenstein's suspicion. But I think that contextualism provides a case that philosophical uses of "I know that's a tree" are problematic in a way that a philosophical uses of "That's a tree" are not, a case which does not depend upon the general Wittgenstein/Malcolm scepticism about meaning behind use. This is best explained by means of examples.

Let us, then, again consider the Bank Cases of Chapter 1. Remember that in Bank Case A, the contextualist will say that I was saying something true when I said, "I know the bank will be open on Saturday." My having been at the bank two weeks before put me in a good enough epistemic position with respect to the bank's being open on Saturdays that I lived up to the standards for true knowledge attributions that were in place when I made my claim. In Case B, however, though

I was in an equally good epistemic position, that position was no longer good enough for me to truly claim to know, for the standards in place had gone up. What would put me in a good enough epistemic position for me to be able to truly say that I know in Case B? That may not be entirely clear, but presumably, if I go into the bank and ask a teller if the bank is still open on Saturdays, and he says that he is certain that it is, this would provide me with sufficient grounds for being able to truly claim to know. But suppose I report what the teller has told me to my wife and she is still not satisfied. "That's not enough," she says. "This check is very important. We still don't know."

My wife seems to have higher standards in mind than I do in this new case. How is the contextualist to evaluate the truth of my wife's claim, "We still don't know"? Well, again, a lot rides here on just what fixes the standards. One might deny that the standards have been raised enough to make her assertion true. But if, as Lewis has plausibly claimed, there are "rules of accommodation," according to which the "conversational score" (in this case, the standard of knowledge) tends to adjust in such a way as to make correct what someone says, one might think that my wife, in making her

statement, succeeds in raising the standards ("changing the score," in Lewis's terminology) and her statement is therefore true. Lewis suggests that the standards are raised more easily than they are lowered. Perhaps, then, if my wife insists, "You don't know," and I continue to say, "I do know," I am saying something false and she is saying something true.

It is, then, far from clear whether my wife succeeds in "changing the score" on me. But not only is it not clear whether or not my wife succeeds in raising the standards for true knowledge attributions, it is not even clear, in this case, what standards my wife is trying to put into place. And to the extent that I don't know what standards she is trying to install, I think the contextualist should say, I do not know what she means to be saying when she says, "We still don't know," because I don't know the intended standards of the 'know' in that claim.

Consider another pair of cases. Suppose two friends of mine are discussing a Dodgers baseball game. One friend was able to watch the game on television through the seventh inning, and knows that the Dodgers were winning by four runs at that point. The second person, however, says, "Still, the Dodgers may have lost. Their

pitcher tends to tire in the late innings." Now, suppose I have just heard the final score on the radio. I say, "I know the Dodgers won. I just heard it on the radio." In this case, I can be quite confident that the I live up to the standards for true knowledge attributions in place in this conversation that I am joining.

Suppose, in contrast, that a friend and I are together and we both hear the radio report that the Dodgers won. But he says, "I wish I knew whether or not the Dodgers won today." I say, "Didn't you hear the radio? The announcer just said that they won." He replies, "Still, for all that, they may have lost. We don't know that they won." Here, again, it is not clear what standards the speaker is trying to put in place (much less what standards actually are in place). This is illustrated by the fact that it is unclear what would satisfy my friend that the Dodgers had won. Would he be satisfied if we heard other reports on other radio stations? Or would we have to ask many people who had actually been to the game? Insofar as it is unclear what the intended standard is, the contextualist should say that it is unclear what the speaker is trying to say. And insofar as the context does not determine a

standard, the contextualist should say that, in an important way, the truth conditions of the knowledge sentence are not determined, and that, therefore, neither is the meaning of the sentence.

What, then, becomes of the claim that one is using 'know' in its ordinary sense? How should the contextualist respond if my wife or my friend in the above examples were to claim that the meanings of their sentences are perfectly clear because they are simply using 'know' in its ordinary sense? The contextualist should not claim that every time the standard moves, 'know' is being used in a different sense: There would then be too many senses of 'know'. It seems clear that it is best to say that 'know' has a single sense, but that it contains an indexical element. Roughly, "S knows that p" is true if and only if S has true belief that p and S is in a good enough epistemic position with respect to p to live up to the epistemic position standards in force, where these standards are determined by the context of the utterance of the knowledge attribution. Our charge is that we have no idea what standards are in place, or are even supposed to be in place, and that we therefore have no idea what the truth conditions are for these suspect knowledge claims. It

is no help in this situation to be told that 'know' is being used in its ordinary sense. It certainly isn't true that my wife or my friend intend the ordinary standards -- if there are any ordinary standards -- to be in place. On the contrary, they seem to be trying to put into place quite extraordinary standards.

The philosophical sceptic should perhaps be treated in the same way by the contextualist. If a sceptic claims that we do not know anything about the external world, we should not think we understand the meaning of what she is saying unless we understand what it would take to know something according to her standards. Since the standards that are in place constitute an important part of the truth conditions for knowledge attributions or denials, there is an important sense in which one does not know what is meant by 'know' in a given case if one does not understand the standards that govern its use -- if one does not know what it would take to "know" according to those standards.

Of course, the sceptic may claim that there are indeed fairly determinate standards behind her denial, but that they are just much higher than ordinary standards. We must investigate this claim. It would, of course, be unfair to demand that the sceptic specify

her standards in a precise way and then to conclude that there is something especially problematic about the meaning of her knowledge denial if she cannot do so. For it is difficult to specify in a precise way what the standards are for any knowledge attribution or denial. In fact, it is quite doubtful that there ever are precise standards for everyday assertions or denials of knowledge. Various contextual factors may make it clear that one does live up to the standards if one is in certain positions (eg. if one has heard the final score of the Dodger's game on the radio) but that one does not live up to the standards in other positions (eg. if one heard that the Dodger's were winning by four runs after the seventh inning). But it is very plausible to suppose that there are grey zones in which it is just indeterminate whether or not a person is in a strong enough epistemic position to live up to the standards in place. Demanding too much precision, then, would be a mistake. But in most normal contexts, though there will be some grey zones, it is clear enough how good a position one has to be in to live up to the standards in place to allow for a good deal of agreement in our assessments of whether or not someone "knows" in various imaginable circumstances. In some such circumstances it

will be clear that the subject "knows", and in other situations it will be clear that she doesn't. Thus, despite the existence of grey areas, we have a pretty good idea of how good a position one must be in to count as knowing. The problem with the philosophical sceptic's denials of knowledge, I think, is that we cannot even conceive of someone being in a good enough position to count as knowing according to her standards.

F. Knowing as God Would Know. This hooks up with a type of reason, present in On Certainty, for doubting the meaningfulness of philosophical uses of "know" which is quite different from the other Wittgensteinian reasons we have seen. At times Wittgenstein worries that when it is used outside its usual circumstances, the word "know" goes nonsensical because it takes on a "metaphysical emphasis" that it cannot tolerate (OC, 482). One seems to be claiming to know as God would know. Wittgenstein writes the following about a knowledge claim:

In its language-game it is not presumptuous. There, it has no higher position than, simply, the human language-game. For there it has its restricted application.
But as soon as I say this sentence outside its context, it appears in a false light. For then it is as if I wanted to insist that there are things that I know. God himself can't say anything to me about them. (OC, 554).

Elsewhere he writes:

Is God bound by our knowledge? Are a lot of our statements incapable of falsehood? For that is what we want to say. (OC, 436)

What are the sceptic's standards for knowledge? What type of epistemic position must one be in to satisfy her? She seems to have in mind a position so good that, in some very strong sense, one can't be wrong in it about the proposition in question. Here the sceptic can play on the fact that, as Austin puts it in "Other Minds":

'When you know you can't be wrong' is perfectly good sense. You are prohibited from saying 'I know it is so, but I may be wrong' (p. 66).

Now we may well wonder why this sceptic doesn't boldly claim that we always speak falsely when we ascribe knowledge to a subject, since it is not only in a philosophical setting that "I know it is so, but I may be wrong" is prohibited. If it is always a condition for knowledge that one can't be wrong, and if the sceptic thinks we fail to fulfill this requirement, then perhaps she should claim that we never truly ascribe knowledge. But, as we've seen, epistemic modal claims are as tricky as knowledge attributions. As in the case of knowledge statements, we seem to have varying standards for when we will assert epistemic modal

statements (eg. "I may be wrong", "I can't be wrong"). In fact, the standards seem to move up and down together in such a way that one can never say "I know it is so, but I may be wrong."¹⁰

As opposed to more ordinary standards for knowing or it being (epistemically) impossible that one be mistaken, the sceptic seems to have in mind standards for both such things that are quite a bit higher -- higher even than the standards in place in courts of law. This is shown by the fact that even in a court of law, some sceptical hypotheses can be brushed aside as too remote to be relevant challenges to a knowledge claim; as Barry Stroud writes,

Even when it matters a great deal, when it is literally a matter of life and death, as in a court of law, it is simply not true that the dream-possibility is always allowed to count as a relevant consideration for the claim to know some particular thing. If I testify on the witness stand that I spent the day with the defendant, that I went to the museum and then had dinner with him, and left him about midnight, my testimony under normal circumstances would not be affected in any way by my inability to answer if the prosecutor were then to ask, 'How do you know you didn't dream the whole thing?'. The question is outrageous; it has no tendency to undermine my knowledge.¹¹

¹⁰See Chapter 2, Part V and Chapter 3, especially section K.

¹¹The Significance of Philosophical Scepticism, p. 49. Of course, if there is some special reason to think, for example, that the witness often mistakenly

But the sceptic seems to have in mind some high or absolute way of knowing or of it's being impossible that one be wrong in which no sceptical hypotheses can simply be brushed aside as being too remote. What is the absolute sense in which a philosophical sceptic claims that it's possible that we're wrong?

G. Absoluteness, Sensitivity, and Sceptical Hypotheses.

The relevant alternatives theory can easily specify an absolute sense of "S knows (or does not know) that P" which one might think is intended by philosophical sceptics: A sense in which all alternatives to P are relevant. To know that P in this absolute way is to be able to discriminate the truth of P from any alternative to P.

I argued in Chapter 1 that the content of an attribution of knowledge cannot be given by specifying the range of alternatives that are relevant to that attribution. However, within the contextualist theory in which I've been working, we can specify an absolute way of knowing as follows. S knows that P in this

reports that things happened in reality that only happened in dreams, then this is another matter. Hence, Stroud's caveat concerning "normal circumstances."

absolute way only if S is able to have a true, sensitive¹² belief that not-H for any H such that the following conditional holds: If S does not know that not-H, then S does not know that P. So, for instance, since the conditional,

If I don't know that I'm not being deceived by an evil demon into falsely believing that I am holding a pen, then I don't know that I am holding a pen,

holds, I do not know that I'm holding a pen in the absolute way unless I am in a position to have a sensitive, true belief that I am not being deceived by an evil genius into falsely believing that I'm holding a pen. And since I am in no such position, I do not know in the absolute way that I am holding a pen.

Sceptical hypotheses specify ways in which a subject might go wrong about, or might fail to know, the proposition in question. So, for instance, in the Bank Cases of Chapter 1, the hypothesis that the bank's hours have changed since I was last there specifies a way in which I might be wrong in believing that the bank will be open on Saturday. Since this sceptical possibility is fairly remote -- I was at the bank quite recently -- there are contexts, like the one described in Bank Case

¹²See Chapter 3, section H.

A, in which I can truly claim to know that the bank will be open on Saturday even though I cannot rule out the hypothesis. In other contexts, like Bank Case B, I may have to be able to rule out this hypothesis, and other similarly remote hypotheses, in order to count as knowing that the bank will be open on Saturday. In the less demanding contexts, when I claim to know that the bank will be open on Saturday, I need to somehow depend upon, take for granted, or, as the sceptic might want to put it, take it on faith, that such ways in which I might be going wrong do not obtain. I claimed in Chapter 3 that in such cases, one is in at least as good a position to know that the sceptical hypothesis does not obtain as one is in to know what one claims to know. However, the fact that one cannot sensitively believe that the sceptical hypothesis does not obtain -- that one cannot "rule it out" -- gives rise to the feeling which is expressed by the phrases listed above: One is simply depending on it, taking it for granted, taking it on faith, that the hypothesis does not obtain.

Perhaps the absolute sense in which, according to the philosophical sceptic, we "know" nothing about the external world, is one which demands that a subject be in such a strong position with respect to the object of

putative knowledge, that she need not take any such thing on faith. Perhaps the sceptic's point is that we humans (as opposed to, perhaps, God), are never in such a position with respect to a fact about the external world.

H. Reid and the Necessity of "Living By Faith". Interestingly, Austin, a staunch anti-sceptic, writes, "It is naturally always possible ('humanly possible') that I may be mistaken" ("Other Minds", p. 66), but he does not explain his cryptic use of "humanly possible". This statement seems to have dark sceptical tendencies, especially when taken in conjunction with Austin's other statement from the very same paragraph of "Other Minds," that "'When you know you can't be wrong' is perfectly good sense. You are prohibited from saying 'I know it is so, but I may be wrong.'" What does Austin mean when he says that it is always "humanly possible" that one be wrong? Is there something about the human condition in virtue of which, in some sense, we are always liable to be mistaken?

The Eighteenth Century Scottish philosopher, Thomas Reid, viewed the human condition as follows. We humans are constituted in such a way that we form various

beliefs under certain conditions. For example, our nature is such that we all (or at least most of us) form the belief that $2+3=5$. And when we have various sensations, our nature is such that we often form beliefs that we are having sensations of the appropriate types. And some sensations produce in us beliefs that there are physical objects with certain properties which are situated in a certain way with respect to us; for example, a certain kind of tactile sensation will produce in most humans the belief that they are touching a hard object. Reid argued that if we are going to get anywhere, we must go along with our natural reactions of trust to at least some of these beliefs; we must trust our belief-forming faculties to be at least generally reliable. Furthermore, Reid viewed it as a contingent matter that we are set up in such a way that most of our naturally-formed beliefs are true.

Such an outlook seems to leave the door wide open to scepticism. The obvious sceptical response is to agree that if we are to get anywhere, we must go along with our natural reactions of trust, but to reply that, as far as absolute knowledge goes, we aren't going to get anywhere. Reid remarks in one place that we must all

"live by faith,"¹³ by which he seems to mean that we must simply trust that our nature is not such that we are constantly deceived. To this, the obvious sceptical response seems to be a simple "Amen." Reid seems to have no way of defending himself from the doubt that he might be constituted in such a way that he has radically and systematically false beliefs, that all of his faculties are unreliable. But things get worse. Reid writes that if one seeks, as Descartes did,¹⁴ to "remove this doubt" by showing that one's faculties are "true and worthy to be trusted", one will be unable to do so, because any such argument will be based on premises and

¹³ See Inquiry into the Human Mind upon the Principles of Common Sense (henceforth abbreviated as Inquiry), Dedication, p. 95b. All references to Reid's works will be to The Works of Thomas Reid, ed. by William Hamilton, Eighth Ed. (Edinburgh: James Thin, 1895), volume I. Page numbers will be followed by 'a' or 'b' to indicate the left or right column of the page respectively. Page numbers will be preceded by chapter number and section number in the case of the Inquiry, and by essay number and chapter number in the case of the Essays on the Intellectual Powers of the Human Mind (which will be abbreviated as Essays).

¹⁴ Here I am not interested in the accuracy of Reid's interpretation of Descartes. What exactly Descartes was up to in trying to remove the doubt that he had such a deceptive nature that he was liable to go wrong in any judgement and whether Descartes involved himself in circular reasoning in his attempt to remove this doubt are tricky issues which I address in my paper, "Epistemic Principles and Descartes's Quest for Scientia".

reasoning that are themselves called into question by the doubt. Reid asks, "For if our faculties be fallacious, why may not they deceive us in this reasoning as well as in others?", and writes:

Every kind of reasoning for the veracity of our faculties, amounts to no more than taking their own testimony for their veracity; and this we must do implicitly, until God give us new faculties to sit in judgement upon the old. (Essays, VI, 5, p. 447b)

But, of course, even if we were given new faculties by which we might try to validate the old ones, we would still be stuck with the same old problem: How could we know that these new faculties were reliable? Once we take seriously the doubt that, as Descartes puts it, "perhaps our nature is such that we go wrong even in the most evident matters"¹⁵ (which I will henceforth call 'Descartes's doubt'), it doesn't matter how many faculties we have at our disposal; the doubt will be irremovable.

This all seems to be playing right into the sceptic's hands. Isn't it a point in the sceptic's favor that the doubt in question could not be removed? But let us move

¹⁵Descartes, Philosophical Letters, p. 74. See also Meditation I, Descartes: Philosophical Writings (tr., ed. by Elizabeth Anscombe and P.T. Geach and henceforth abbreviated 'AG'), p. 64; and Meditation III, AG, p. 77.

beyond Reid¹⁶ and note that the reason this doubt is irremovable seems to have little to do with the human condition, but with facts that would hold for any believer we can clearly imagine: we have beliefs; there are processes by which we come to have these beliefs; if the processes are highly unreliable, the beliefs that result are quite epistemically defective; while we may be able to show some of our belief-forming processes to be reliable if we are allowed to use information from other belief sources in the validation process, if we tried to remove the general doubt that all of our sources of belief might be highly unreliable, we could

¹⁶Reid's own treatment of what he calls the "thorough and consistent sceptic," whom Reid construes as one who will not accept any belief until the belief-forming faculty from which it is derived is shown to be reliable, is to accept that since nothing can be proven or shown until some belief is accepted, thorough scepticism is invulnerable to attack. Reid writes that such a sceptic must "be left to enjoy his scepticism" (Essays, VI, 5, p. 447b), and admits, "To such a sceptic I have nothing to say" (Inquiry, V, 7, p. 130a). Against this sceptic Reid contents himself with the defense that neither can such a sceptic have anything to say to Reid. The consistent sceptic will end up not believing anything; in particular, he will not believe that Reid should not believe things (see Inquiry, I, 8, p. 104a and Essays, VII, 4, p. 485a). But this defense will not work against the "thorough" sceptic who does not resolve not to believe anything that comes from an unvalidated faculty, but who merely claims that since we cannot escape the possibility that our nature is such that we are constantly deceived, none of our beliefs are known in some high or absolute sense.

not do so in a non-circular¹⁷ fashion. What would it be like to be in such a position that one is not subject to this doubt, that one could rule out the "deceptive nature" hypothesis behind Descartes' doubt? I don't think we have any idea of what this would be like.

Let me not understate my point. It's not just that we don't know what it would be like not to be subject to this doubt in the same way that Nagel claims we don't know what it's like to be a bat. While we may not know what it's like ("from the inside") to, eg., navigate by sonar in the way a bat does, we do know quite a bit about how a bat does this. But how could any cognitive agent not be subject to Descartes's doubt? Mustn't there be some way or ways by which he comes to have beliefs? And mustn't these processes be reliable? And how could he show any of them to be reliable without relying on some belief that is the result of one of the very processes whose reliability is in question? Perhaps God is somehow immune to this doubt. But if He is, I have absolutely no idea how he manages to be. So it's not just that I don't know what it would be like "from the inside" not to be subject to this doubt, I

¹⁷The type of circularity in question is what William P. Alston has called "epistemic circularity." See Alston's "Epistemic Circularity."

have no idea of what a cognitive agent would be like that was not subject to this doubt. Does the sceptic, in claiming that we don't know, really mean only to be claiming that we don't live up to certain standards that we can't even imagine any cognitive agent living up to for any of his beliefs? Would this be an interesting claim? It would certainly be interesting if we had to live up to these standards in order to truly claim to know in any context. But the not-so-bold sceptic does not claim to threaten the correctness of our "lower" knowledge claims.

I recall an incident at a social gathering of philosophers and their families in which a young child entertained the guests by "doing philosophy." She asked her victim a simple question, and the victim answered by asserting the appropriate fact. She then asked, with just the right edge in her voice, "How do you know that?" The victim scrambled for evidence: "Because. . ." Then came the dreaded rejoinder: "But how do you know that?" After a couple of rounds, when it had become obvious that the child was going to be unrelenting, the victim wisely admitted defeat: "Oh, now I see, Socrates. I guess I didn't know that my name is John." This was all very cute in the hands of a five-year-old, but this

scepticism won by the simple and incessant demand that the victim prove his premises seems unworthy of much serious attention from professional philosophers.

Scepticism based on Descartes's doubt seems to be almost that trivial. If the opponent claims to know that p, one asks, "Isn't it possible that your nature is radically defective in such a way that you're liable to be wrong about anything, including what seems to you most evident and, in particular, about p?" If the opponent tries to prove that p or that his nature is not in that way defective, one simply subjects the premises of the opponent's "proof" to Descartes's doubt. Repeat as needed. Descartes's doubt, once it is allowed, cannot be removed. Usually, as I've pointed out, we allow ourselves to brush aside certain possibilities of error as being irrelevant or too remote -- to discount them despite the fact that one is not in a position to sensitively believe that they do not obtain. The sceptic, however, does not seem to approve of such brushings aside.

I. Descartes' Doubt and the Impossibility of "External Knowledge". Throughout his book, The Significance of Philosophical Scepticism, Barry Stroud tries to

distinguish between what he calls the sceptic's "external" and our everyday "internal" questions about and assertions and denials of knowledge. He, for example, accuses Moore of failing to take the sceptic's question about whether we know anything about the external world "in that increasingly elusive 'external' or 'philosophical' way" (p. 119), when Moore answers the question by saying, in an "internal" way, that he does indeed know this or that. Stroud never clearly states what the distinction is.¹⁸ His book is largely an attempt to elucidate this difficult distinction by means of examples of how it has popped up in various guises in the history of modern philosophy. Three of the marks of a "philosophical" or "external" use of 'know' seem to be a) that reasons for doubt cannot be brushed aside merely because they are too remote to be taken seriously, b) that the things one normally can take for granted are put into question and cannot be taken for granted, and c) that the sceptic's questions are general: she means to be questioning all

¹⁸In his review of Stroud's book, Richard Feldman complains, "I find myself on Moore's side here. The elusive philosophical sense of words that Stroud mentions eluded me clear through to the end of the book" (p. 306).

of our knowledge at once.¹⁹ Denials of knowledge based upon Descartes' doubt seem perfectly fitted to Stroud's characterization of "external" uses of "know". I must admit that I often find myself among those who share Wittgenstein's suspicion that we don't understand the philosophical sceptic's denials of knowledge. But on those occasions when I do think I have some grasp of what the sceptic is asking, it seems to me that if it's an "external" question about our knowledge that's at issue, then we can't simply brush aside Descartes's doubt. To "take it on faith" that we don't have a nature so deceptive that we are highly liable to go wrong about anything, even what seems to us most evident, seems far too vulgar and, if I have any feel for Stroud's distinction, far too "internal."

Thus, it seems impossible to conceive of a being who can "rule out" the "deceptive nature" hypothesis that's behind Descartes' doubt. And if the sceptic's standards for knowledge really do require that we be able to rule out all sceptical hypotheses that threaten our beliefs, including this one, then it seems impossible to conceive

¹⁹Feldman takes this third feature to constitute "the general idea" behind Stroud's notion of "external" uses of 'know'. See p. 306 of his review of Stroud.

of or imagine a being who knows according to the sceptic's standards.

J. The Contrast Strategy. The external world sceptic, if she allows that we know certain non-external-world propositions according to her standards, can say that we don't have to imagine what it would be like for a subject to know something according to her standards -- we ourselves know various things according to the standards. This external world sceptic can try to give us some idea of what status our external world beliefs lack by contrasting them with the beliefs that do live up to her standards. Her denial of knowledge to my external world beliefs can be the interesting claim that I do not know that there is a book in front of me in as wonderful a way as I do know that I am having such-and-such a sensation or as I know certain necessary truths self-evident to me, like $2+3=5$. I must admit that I may be wrong about there being a book in front of me in a way or sense in which I need not admit this of the other beliefs. Thus, limiting the scope of one's scepticism to allow certain beliefs to escape it may increase the interest of the scepticism that remains.

Thus, it is no accident, I think, that we find such staunch anti-sceptics as Reid and Austin²⁰ attacking the thesis that there is some special, absolute way in which one can't be wrong about one's own sensations, or more generally, one's own mental states.

In what way can't my beliefs about my own mental states be wrong? In the Third Meditation, Descartes distinguishes between what "the light of nature shows" him to be true and what "nature teaches" him or what he believes by a "natural impulse," which he later calls "blind impulse" (AG, pp. 79-80). His beliefs in "external objects" (p. 79) get classified as blind impulses of nature. And this classification has disastrous epistemic consequences for the beliefs in question, for Descartes writes that things revealed by the natural light are "absolutely beyond doubt," while he finds it all too easy to doubt his natural impulses. Descartes classifies various beliefs about his own mental states as things revealed by the natural light. But why the distinction? Why do external world beliefs invariably fall into the bad category? And what is the distinction here? Descartes writes that "there is a big

²⁰See "Other Minds", pp. 58-65. We will look at Reid's attack in more detail below.

difference" (p. 79) between being led to believe something by an impulse of nature and seeing something to be true by the natural light, but does not explain the difference. Clearly, though, Descartes thinks there is some sense in which beliefs formed by natural impulses can be wrong but in which he cannot be wrong about what is revealed to him by the natural light.

It is precisely such a distinction between natural impulses and seeing by the natural light that Reid is out to attack. Reid takes his predecessors -- Descartes, Malebranche, Locke, Berkeley, and Hume, or as Reid calls them, the "ideal philosophers" -- to be accepting of self-evident facts about one's own mental states and about necessary truths, and of what can be deductively proven on their basis, but as not being willing to accept anything else. (I am not here interested in the accuracy of Reid's treatment of his predecessors.) Reid's response is to claim that the "ideal philosopher's" belief in the existence of his sensations is on no better ground than are many of Reid's external world beliefs. Reid writes of Hume:

The author of the "Treatise of Human Nature" appears to me to be but a half-sceptic. He hath not followed his principles so far as they lead him; but, after having, with unparalleled intrepidity and success, combated vulgar prejudices, when he had but one blow to strike, his courage fails him, he fairly lays down his

arms, and yields himself captive to the most common of all vulgar prejudices -- I mean the belief of the existence of his own impressions and ideas. I affirm, that the belief of the existence of impressions and ideas, is as little supported by reason, as that of the existence of minds and bodies. No man ever did or could offer any reason for this belief. . . . And what is there in impressions and ideas so formidable, that this all-conquering philosophy, after triumphing over every other existence, should pay homage to them? (Inquiry, V, 7, pp. 129a, 130a)

Reid claims that beliefs in one's own sensations are on the same footing as some of one's external world beliefs: we cannot give any good reason for accepting beliefs of either kind. We are set up in such a way that upon having a sensation, we believe in its existence, and our constitution is also such that upon having certain sensations, we believe in the existence of material objects having certain qualities.²¹ In

²¹The beliefs in question are the results of what Reid calls "original perceptions." According to Reid's account of the perceptual process, human beings are, prior to any experience, set up in such a way that they will have certain "original perceptions" upon the occasions of having certain sensations. Thus, a particular tactile sensation produces in a human being both the concept of hardness and the belief in the existence of a hard material object. Reid also recognizes a class of "acquired perceptions," in which certain sensations indicate the presence of a material object with certain qualities because we have learned by experience that the sensation is usually accompanied by the presence of such an object. Thus, a certain kind of visual sensation causes one to believe that there is a hard material object present because one has learned by experience of the connection between this type of visual sensation and the presence of hard objects. In cases of original perception we do not have to experience a

either case, we must trust the process by which the belief is formed to be reliable. To "pay homage" to one belief while rejecting the other as unfounded is to be guilty of arbitrary partiality.

This seems right to me. It is often pointed out that there is a gap between the external world and our external world beliefs. If this gap has an inferential component, then we could evaluate the inference. But Reid claims that in some crucial cases, sensations produce our external world beliefs in such a way that nothing like an inference is going on.²² There is, to be sure, a process, a mechanism connecting the external world belief and its object, and we must trust this mechanism. But there is also this kind of gap between a

constant connection between the sensation-type and the presence of material objects of a certain type before the sensation can indicate the presence of such a material object. The belief is formed by what Reid calls "an original principle of the human constitution."

²²In fact, Reid's view is that in these cases, the sensation produces both the sensation belief and the external object belief, but that the process by which the external object belief is formed does not go through the sensation belief at all, though it does, like the process by which the sensation belief is formed, go through the sensation. If the process by which the external object belief is formed does not even go through the sensation belief, then it seems clear that there is nothing even remotely like an inference from the sensation belief to the external object belief involved in the process.

sensation and the subject's belief that he is having it. I take it as self-evident that the belief that one is having a sensation is distinct from the sensation itself. And wouldn't the belief in the sensation be seriously defective if it were not connected to the sensation by a reliable process? And how could we establish the reliability of the process without begging the question against the doubt that this process is unreliable?

Likewise, a belief in a simple necessary truth like $2+3=5$ seems to depend upon the reliability of the process by which it is formed for its status. My belief that I am not systematically mistaken about even what seems most evident to me -- even elementary arithmetic -- seems to be insensitive. I cannot "rule it out" that, as Descartes puts it, I go wrong "whenever I add two and three, or count the sides of a square, or do any simpler thing that might be imagined."²³ And so, insofar as I

²³First Meditation, AG, p. 64. Since Descartes feels he received his nature from God, he phrases his question in terms of God making him go wrong, but we find out in the next paragraph that Descartes's worry is a general worry about his having a deceptive nature, a worry that would apply regardless of what one thinks is the source of one's nature; in fact, ascribing my nature to some cause other than God only intensifies the worry. Descartes writes that whether one ascribes "my attaining my present condition to fate, or to chance, or to a continuous series of events, or to any other cause, delusion and error certainly seem to be imperfections,

continue to trust my judgments about elementary arithmetic, I must "take it on faith" that I am not so deceived.

What might mislead us here is a tendency to try to reduce epistemic modalities to other modalities, like logical or metaphysical modalities. Any such reduction which analyzes it is epistemically possible that p into something like it is metaphysically possible that something (where this 'something' can be filled in various ways) be the case and yet p be false, will have the result that any metaphysically necessary truth p will be epistemically necessary and its negation epistemically impossible. But a little reflection on the fact that both Goldbach's conjecture is true and Goldbach's conjecture is false are epistemically possible from my point of view right now should cure us of the desire to try such a reduction. More subtle reductions could be tried, but I am very sceptical about any effort to analyze the important notion of the ways things may be in terms of the also important but radically different notion of ways things could have been. Philosophers who have a bias in favor of

and so this ascription of less power to the source of my being will mean that I am more likely to be so imperfect that I always go wrong" (AG, p. 64).

knowledge of one's own mental states are attracted to analyzing it is epistemically possible (from a given subject S's position) that p into something like it is metaphysically possible that: S be in the mental state S is in now (having such-and-such sensations, etc.) and yet it be false that p. This type of reduction has the result that metaphysically necessary truths and truths about one's own mental states are epistemically necessary.

Reductions like the ones above have a built-in bias in favor of certain kinds of beliefs as candidates for epistemic necessities -- beliefs in metaphysically necessary truths and in one's own mental states. The notion of epistemic possibility, if I understand it correctly, does not have any such built-in bias. Furthermore, it seems to be a clear enough notion that it doesn't have to be reduced to something like metaphysical possibility in order for it to be an acceptable notion -- in fact, in its own right, it seems to me to be at least as clear as the notion of metaphysical possibility. The notion in question is most clearly and purely expressed by simple phrases of the form 'Maybe p' or 'P may be true' in cases in which one claims that p is epistemically possible from one's

own point of view. If so taken, the notion is clearly relative to points of view -- what "may be true" from one epistemic situation will not be possible from another.²⁴ The notion so understood is also extremely relevant to scepticism, because it is epistemically possible (from S's point of view) that not-p, so understood, intuitively seems to analytically entail that S does not know that p.²⁵ And, so understood, the notion does not seem to have the built-in biases mentioned above. That I could have been in the mental state I am in now without it being true that there is a book in front of me does not, by itself, analytically entail that I don't know that there is a book in front of me. Various non-analytic premises have to be added before the entailment will hold. If a philosopher wants to claim that it is exclusively beliefs in metaphysically necessary truths and in one's own mental states that one can't be wrong about, this philosopher will have to produce an argument for this claim.

²⁴See part I of Chapter 2.

²⁵Here I agree with Austin that "you are prohibited from saying 'I know it is so, but I may be wrong.'" That is, it seems to built into the language that, as Austin puts it, "When you know you can't be wrong." See Part III of Chapter 2.