

Moorean Facts and Belief Revision, or Can the Skeptic Win?

Thomas Kelly
Princeton University

Forthcoming in John Hawthorne (ed.), *Philosophical Perspectives*, vol.19: Epistemology.

1. Introduction

A **Moorean fact**, in the words of the late David Lewis, is ‘one of those things that we know better than we know the premises of any philosophical argument to the contrary’.

Lewis opens his seminal paper ‘Elusive Knowledge’ with the following declaration:

We know a lot...We have all sorts of everyday knowledge, and we have it in abundance. To doubt that would be absurd...It is a Moorean fact that we know a lot. It is one of those things that we know better than we know the premises of any philosophical argument to the contrary (1999: 418).

I take it that these remarks are put forth in a methodological spirit. ‘Elusive Knowledge’ stands as Lewis’ fullest statement of his own positive views on knowledge and skepticism. It is here that he lays out the details of his own contextualist epistemology. In beginning the paper in this way, I take Lewis to be endorsing a thought along the following lines: Before we commence with the Hard Epistemology, here is something that we can lay down at the outset: We know a lot. The proposition that *We know a lot* can serve as a fixed point in our inquiry, and we should make adjustments elsewhere as needed in order to hold on to this fundamental commitment. Simply put, we should build the rest of our theory around this fixed point.¹

In proceeding in this way, Lewis is far from alone among prominent contemporary philosophers. Compare David Armstrong on the proposition that *Things move*:

It is a very fundamental part of the Moorean corpus that there is motion. Things move. Perhaps we have still not, after two and a half thousand years, got to the full bottom of Zeno's brilliant arguments against the existence of motion...But certainly Zeno should not persuade us that things do not move. Neither should anybody else (1999: 79).

Elsewhere, Armstrong suggests that since the belief that *inductive inference is rational* similarly enjoys the status of 'Moorean knowledge', Hume's formidable argument to the contrary is powerless to show otherwise. Indeed, Armstrong insists that the fundamental rationality of inductive inference can legitimately be taken as a datum for philosophy and used as a premise for theorizing in metaphysics (1983: 53-54).

According to the conception advanced by Lewis and Armstrong, belief in a Moorean fact has a kind of epistemic standing which renders it peculiarly resistant to being rationally undermined. Indeed, it is occasionally suggested that belief in a Moorean fact cannot be rationally undermined at all. According to a more common conception—and the one to be explored here—belief in a Moorean fact is invulnerable to being undermined *by means of philosophical argument*. This leaves open the possibility that such a belief could be rationally undermined by particular courses of experience or by the right sorts of scientific discoveries. This invulnerability to being undermined by philosophical argument then, is a reflection not only of the epistemic standing that is claimed for such beliefs but also of the perceived impotence of philosophy itself.

This line of thought is made explicit by Kit Fine:

In this age of post-Moorean modesty, many of us are inclined to doubt that philosophy is in possession of arguments that might genuinely serve to undermine what we ordinarily believe. It may perhaps be conceded that the arguments of the skeptic appear to be utterly compelling; but the Mooreans among us will hold that the very plausibility of our ordinary beliefs is reason enough for supposing that there *must* be something wrong in the skeptic's arguments, even if we are unable to say what it is. In so far then, as, the pretensions of philosophy to provide a world view rest upon its claim to be in possession of the epistemological high ground, those pretensions had better be given up. (2001: 2, emphasis his).

The strategy of attempting to defuse skeptical challenges by appeal to the alleged sanctity of Moorean facts is not without its detractors.² Nevertheless, the popularity of the strategy runs high.³ Moreover, it is not simply the sheer number of prominent philosophers who follow Moore in relevant respects which is noteworthy but also the philosophical orientations of those who do. After all, Lewis, Armstrong, and Fine hardly resemble common sense philosophers of yore: they are not, for example, purveyors of paradigm case arguments or even practitioners of an essentially conservative Strawsonian descriptive metaphysics. Rather, all three stand squarely in the venerable tradition of speculative metaphysics. The fact that philosophers of their general orientation nevertheless pay homage to Moore does as much as anything to bolster the credibility of Fine's claim that ours is an age of 'post-Moorean modesty'.⁴

I believe that the topic of Moorean facts raises deep questions of both philosophical method and first order epistemology. How should we respond to arguments that challenge beliefs of which we are extremely confident? To what extent can such arguments—or rather, those who put them forth—hope to get some rational grip or traction on us? To what extent, if at all, does one's starting point constrain the kinds of revisions in one's views that philosophy might legitimately inspire? When (if ever) is one justified in refusing to be swayed by an argument which is flawless as far as one can

tell, and when is such refusal simply dogmatism? These are large questions, and ones which I will not attempt to fully answer here. My more modest ambition is to attempt to make progress with respect to these and related issues by way of examining the notion of a Moorean fact that I find in Lewis and others.

To place my own cards on the table: my sympathies lie with the Moorean. I believe that there are very substantial limits on how radical a change in our views philosophy might legitimately inspire. For example, in epistemology—the domain on which I'll focus in what follows—I suspect that, ultimately, the skeptic simply cannot win.⁵ The sense in which the skeptic cannot win is not that he will inevitably fail to persuade us of his conclusion—that, after all, might be a matter of mere psychological stubbornness on our part, which would, I think, be of rather limited philosophical interest. Rather, the sense in which the skeptic cannot win is that it would never be reasonable to be persuaded by the skeptic's argument. Moreover, I think that this is something that we can know even in advance of attending to the specifics of the skeptic's argument: in a sense, the skeptic has lost before the game begins. I concede that it is not easy to see what distinguishes this stance from simple dogmatism (if anything does). Indeed, much of what follows is devoted to criticizing various ways of developing the Moorean response to skepticism that I believe are unsuccessful. Towards the end of the paper, I sketch a way of understanding that response on which it has, I believe, a great deal of force.

Of late, there has been a strong resurgence of interest in Moore's response to the skeptic. In particular, the status of his famed 'proof of an external world'—'Here is one hand; Here is another; Therefore, the external world exists'—has been vigorously

debated.⁶ This focus on the virtues and vices of Moore's proof naturally encourages a certain picture of the dialectic between Moore and his opponents. According to the picture in question, Moore is essentially *playing offense* against the skeptic. That is, Moore has taken up the burden of proof, a burden that he can discharge only by providing a successful argument for an anti-skeptical conclusion. On this reading of Moore, he is primarily of interest to the philosopher concerned with skepticism as one among countless others who has taken the bait and risen to meet the challenge of providing an anti-skeptical argument that fulfills the criteria of argumentative goodness (whatever exactly those criteria are). In aim and ambition, Moore is comparable, perhaps, to one who offers an inference to the best explanation argument against the skeptic.

I think that this debate over the status of Moore's proof has proven illuminating and worthwhile. However, there is a danger that a fixation on Moore's proof will lead us to neglect what is arguably the deepest and most distinctive aspect of his critique of radical skepticism. For there is an alternative reading of Moore available, according to which Moore is ultimately not playing offense against the skeptic but rather *playing defense*. On this reading of Moore, the onus is on the skeptic to provide a compelling argument for *his* conclusion, and Moore is providing reasons for thinking that such a project will inevitably end in failure. Thus, even when Moore is presenting his own proof, that presentation is, in the vocabulary of Scott Soames (2003a: 23), 'ironic': Moore is really calling attention to the relative weakness of the skeptic's own premises compared to those which he himself employs. Whether or not this is ultimately correct as a matter of textual exegesis, I believe that much of Moore's influence on contemporary philosophy consists in the belief that he succeeded in providing an effective recipe for playing

defense against the skeptic. (Notice, for example, that in each of the passages from Lewis, Armstrong, and Fine quoted above, the focus is on the alleged impotence of the *skeptic's* argument.) Among contemporary expositors and defenders of Moore, this perspective is most well-developed by William Lycan (2001) and Soames (2003a). Although I depart from both Lycan and Soames in important respects in what follows, the general approach adopted here is most akin to theirs.

2. Predictions and Policies

The skeptic presents us with an argument for a conclusion radically at odds with common sense. How should we respond? Best of all, of course, would be to identify some false premise or fallacious step in the argument. But suppose that we find that we can do neither. What then? As we have seen, Fine suggests that, in such circumstances, we should conclude that *something* must be wrong with the skeptic's argument and simply retain our original beliefs. But how can a stubborn refusal to be moved by an argument for which one can admittedly find no flaw be anything other than dogmatic?

Perhaps the picture is something like this. In any case in which one scrutinizes an argument and fails to find any flaw in that argument, there are two competing potential explanations of one's failure. First, one's failure might be due to the flawlessness of the argument. Alternatively, it might be that the argument is in fact flawed, and one's failure is due to one's own cognitive limitations. (If one were better informed, one would recognize one of the premises as false; if one were more sophisticated or insightful, one would detect some subtle fallacy.) In deciding how to respond to any argument which

appears to be flawless, one is in effect in the position of performing an inference to the best explanation, where the explanandum is one's inability to identify any flaw despite having attempted to do so. If the better explanation of this fact is the flawlessness of the argument, then one should come to believe its conclusion and revise one's other beliefs accordingly. If, on the other hand, the better explanation of one's failure is one's own cognitive limitations, then one should remain unmoved in the face of the argument. We can view the Moorean as someone who holds that, for arguments aimed at overturning Moorean facts, 'hidden flaw' explanations will inevitably trump 'no flaw' explanations.

Notice that, if this is dogmatism, there is a respect in which it is an unusually modest variety. For when one reasons in this way, one's refusal to change one's beliefs is due to the weight that one gives to one's own cognitive limitations. In marked contrast, the skeptic will insist that one treat the fact that his argument *seems* or *appears* to be flawless as a reliable indication that it *is* flawless. Here--but not elsewhere--the skeptic will insist that one treat appearances as a reliable guide to reality.

Still, this does nothing to answer the question of why Moorean facts might have the relevant status. I'll consider some answers to this question in the next section. But first, I want to further explore some related Moorean themes.

Recall Lewis' informal gloss of a Moorean fact as 'one of those things that we know better than we know the premises of any philosophical argument to the contrary'. This gloss admits of weaker and stronger readings. A relatively weak reading is the following: a Moorean fact is one of those things that we know better than we know the premises of any philosophical argument to the contrary *which has yet been offered*. On this reading, to declare that F is a Moorean fact is to declare that none of the known philosophical

arguments for not-F is rationally compelling. It is to say nothing about the possibility, or even the likelihood, that there is some compelling argument which we have yet to encounter. (One might, after all, hold that while as a matter of fact none of the known arguments for not-F is strong enough to undermine one's belief that F is true, this situation might very well change with the publication of the next volume of *The Philosophical Review* or *Nous*.)

The claim that some fact is Moorean in this weak sense is not a trivial one. There are, after all, some who suspect that we *do* presently possess arguments that are sufficiently strong to undermine our confidence even in those beliefs that we ordinarily take to be among the most certain. Nevertheless, I think that to adopt this interpretation would be to seriously underestimate what those who appeal to Moorean facts often have in mind.

When Armstrong declares

...certainly, Zeno should not persuade us that things do not move. Neither should anybody else

I don't think that he is plausibly interpreted as *noncommittal* on the question of what the next volume of *The Philosophical Review* might bring, or what our epistemic position with respect to the proposition that *Things move* will be after we've fully absorbed its contents.

One who declares that F is a Moorean fact is not simply making a claim about the relationship between F and those arguments for not-F that have been offered thus far, I think. Rather, he or she is also making a claim about the relationship between F and all of those arguments for not-F that might be offered in the future. But what, exactly, is the relevant relationship? In declaring something a Moorean fact, what is one doing?

Here there are at least two possibilities worth considering:

- (i) *To declare something a Moorean fact is to make a prediction.* In particular, to declare that F is a Moorean fact is to claim not only that none of the presently known arguments suffices to rationally undermine one's belief that F is true, but that we will not encounter such an argument in the future, either.

Perhaps one simply takes the epistemic standing of a particular proposition to be such as to license the relevant prediction straightaway. On the other hand, the prediction might also be construed as the conclusion of an inductive inference. Consider the claim that We know a lot. For thousands of years, philosophers have been offering skeptical arguments for the denial of this claim. Suppose that one judges that even the most formidable of these arguments is insufficient to undermine one's belief that *We know a lot*. Given such an estimation of the relevant inductive base, one might naturally infer that we will not encounter a sufficiently strong argument in the future, either.

Alternatively, it might be that

- (ii) *To declare something a Moorean fact is to endorse a policy.* In particular, to declare that F is a Moorean fact is to endorse the following policy for evaluating arguments: any argument that has as its conclusion not-F should be judged a bad argument.

In general, one's judgement as to the probative force of a given argument is not independent of one's judgement as to the credibility of its conclusion. As we've noted, if one takes oneself to have strong reasons to believe that F is true, then one will take oneself to have strong reasons to believe that a given argument for not-F is a flawed argument, even if one finds oneself unable to identify any particular flaw that the

argument contains. Given this, adoption of the following policy might seem the most reasonable course: one should conclude that any argument which has as its conclusion the denial of a Moorean fact is a flawed argument. In effect, one who adopts such a policy resolves to treat the fact that an argument has the denial of a Moorean fact as its conclusion as a *reductio ad absurdum* of the (conjunction of) that argument's premises.

Are those who appeal to Moorean facts in the context of philosophical argument better understood as making a prediction or as endorsing a policy? I believe that they are best understood as doing both. On the one hand, they are endorsing a policy of evaluating skeptical arguments negatively. On the other hand, their willingness to endorse this particular policy is presumably not independent of their confident belief that we will never encounter a skeptical argument that is sufficiently strong to make it reasonable to believe its conclusion. That such individuals mean to be endorsing a policy is seen, I believe, in the contexts in which appeals to Moorean facts are typically made. The relevant contexts are ones in which methodological advice is on offer: what is up for discussion is the proper way to respond to skeptical arguments, what one could and could not learn from skepticism, and so on.⁷ On the other hand, it is not as though one's willingness to endorse such a policy is independent of one's willingness to predict that we will not encounter a compelling skeptical argument in the future. For of course, the goodness of the policy would seem to depend directly on the accuracy of the prediction: if it's in fact the case that we will never encounter a compelling skeptical argument, then consistently adhering to a policy of negatively evaluating skeptical arguments will invariably lead one to classify such arguments correctly.

Still, even if one is completely convinced that we will never encounter a compelling skeptical argument, it is not obvious why one would want to endorse any general policy at all. Why not judge each argument by the content of its character, rather than engaging in what would seem to be a kind of invidious group discrimination? After all, some skeptical arguments are more formidable than others: this much is not in dispute. (Presumably, even those who think that no skeptical argument is or could be compelling will admit that some skeptical arguments are *transparently bad* in a way that others are not.) Given this, why would one think it a good idea to have some general policy for classifying arguments that might have nothing in common other than their conclusion (e.g. ‘We don’t know a lot’)?

Here is a *possible* reason why one might think that it is a good idea (I make no claim that this consideration is in fact what drives those who appeal to Moorean facts): one thinks that one is more likely to make mistakes in particular cases if one judges each individual without recourse to a general policy. For example, it’s sometimes suggested that a central intrapersonal function of adopting general policies is the role that doing so plays in enabling us to overcome particularly tempting mistakes (Ainslie 1975, 1986; Nozick 1993: 17-18). One’s considered, reflective judgement is that it is best to do A1 when in circumstances C; however, one knows that when one is actually in circumstances C, one will be strongly tempted to do A2 instead. One thus decides to adopt a general policy of always doing A1 when in circumstances C, and one’s explicit adoption of this policy makes it easier to resist the temptation to do A2 rather than A1 when one later finds oneself in circumstances C. The adoption of a policy by one who will resolutely adhere to the policy which she has adopted thus involves a certain trade-off. On the one

hand, one formulates the policy without the benefit of relevant information that might become available later. On the other hand, one has a certain psychological bulwark against local temptations that ought to be ignored. Applied to the present case, the analysis yields the following result. One confidently believes that it will never be rational to abandon one's belief that *We know a lot* on the basis of a skeptical argument. However, one knows that skeptics are a crafty lot and capable of great ingenuity in the service of their cause: the skeptic will labor to construct transparently valid arguments whose premises seem intuitively beyond reproach. Faced with such a skeptical argument, and finding oneself unable to identify any particular flaw which it contains, one might be strongly tempted to conclude--incorrectly, by one's present lights--that we don't know a lot after all. By adopting a general policy of evaluating skeptical arguments negatively, one fortifies oneself against the possibility of being taken in by the skeptic. By adopting such a policy, one helps to remind oneself that, even in these circumstances, one has (what one now takes to be) a decisive reason for rejecting such an argument: its conclusion.

On this way of understanding things, the prediction is prior to the policy. One confidently predicts that one will never encounter a compelling skeptical argument; because of this, one thinks that if, one did change one's mind in response to a skeptical argument in the future, one would surely be making a mistake in doing so. One thus adopts the policy as an attempt to avoid making such a mistake.⁸

But there is, I think, an alternative way of understanding things that better captures the intent of the Moorean. On this alternative, it is the policy which is prior to the prediction. That is, the Moorean thinks that if we possess a sufficiently rich understanding of what

are in fact the correct norms of belief revision, we will see that these norms effectively guarantee that it would never be reasonable to abandon one's belief in a Moorean fact in response to a skeptical argument. Thus, the relevant prediction is not some potentially precarious prediction to the effect that, say, we will never encounter a sufficiently ingenious skeptic in the future. It is, rather, a trivial consequence of the correct application of the norms of belief revision that we ought to employ.

Consider an analogy drawn from the philosophy of mathematics. In the heyday of conventionalist accounts of the a priori, a staple of conventionalist manifestos was the claim that we will never make any observations which will falsify (or even disconfirm) a proposition of arithmetic or geometry.⁹ In the mouth of the conventionalist, of course, this claim was not some potentially precarious prediction about the future course of natural science or about the contents of the experiences that human beings will in fact undergo in the future. Rather, the conventionalist is impressed with what is, arguably, an important aspect of our mathematical practice: our adhering to a general policy of refusing to treat empirical considerations as *the kind of thing* which might count as evidence against a select class of propositions. For the conventionalist, the knowledge that no future experimental outcome will disconfirm a proposition of arithmetic is readily available to anyone with an adequate grasp of the relevant bits of epistemology. We should, I think, view the Moorean in a parallel way: as one who thinks that, according to what are in fact the correct norms of belief revision, philosophical considerations are simply not the kind of thing which could undermine another select class of propositions, 'the Moorean facts'. That we will never encounter a compelling skeptical argument is a

piece of knowledge that is readily available to anyone with a sufficiently rich understanding of those norms.

But what would the norms of belief revision have to be like, in order for this picture to be correct? I turn to this question next.

3. Some Norms of Belief Revision

Suppose that I believe that F is true. Attempting to convince me otherwise, you offer a philosophical argument for the contrary conclusion. I remain unmoved. You invite me to point out some false premise or fallacious step in your argument. I decline the invitation. I assure you that there must, of course, be *something* wrong with your argument, but I insist that it is not a condition of my knowing this that I am able to identify some specific flaw. Determined, you signal your intention to offer further arguments for the same conclusion. I advise you not to bother. For F is a *Moorean fact*, and thus, the rationality of my continuing to believe F is simply not susceptible to being undermined by your arguments. Infuriated, you accuse me of dogmatism. I deny the charge—at least, if the charge of dogmatism carries with it the suggestion that the dogmatist is being unreasonable. Indeed, far from being unreasonable, I respond, my unwillingness to abandon my belief in the face of your arguments is the uniquely reasonable response in the circumstances.¹⁰

In carrying on in this way, I might have either one of two pictures in mind. First, I might think that it is simply a fundamental epistemic norm that

MOORE One should never abandon one's belief in a Moorean fact on the basis of a philosophical argument.

One way of thinking about Moorean facts then, would be this: Moorean facts make up a class of special, privileged propositions, and it is simply a fundamental norm of belief revision that one should never stop believing a member of this class in response to a philosophical argument.

This approach involves reifying the notion of a Moorean fact in a particularly strong way. Moorean facts make up a category of epistemically special entities, things to which the normal rules do not apply. Of course, much of the history of epistemology involves such reification: Cartesian foundations, empiricist sense data reports, Kantian synthetic a priori principles, Carnapian linguistic rules and Wittgensteinian hinge propositions were all thought of as in some respect standing outside the rules that apply to more mundane propositions. According to this way of thinking about Moorean facts, Moorean facts are simply among the more recent additions to the epistemologist's bestiary.¹¹

Perhaps some have thought of Moorean facts in this way, or along similar lines.¹² But there is an alternative—and, I think, preferable- way to attempt to make sense of the notion of a Moorean fact. On this way of thinking about Moorean facts, MOORE is not itself a fundamental norm. While it's true that one should never abandon one's belief in a Moorean fact in response to a philosophical argument, the fact that one should never do so falls out of higher-level, more fundamental norms. Again, Lewis writes of Moorean facts as things that we 'know better' than the skeptic's premises. Consider then the following norm of belief revision:

KNOWN BETTER One should never abandon a belief in response to an argument when that belief is known better than (at least one of) the premises of the argument.

The present picture then, is one according to which propositions can be ranked along some relevant dimension (perhaps: ‘the known better’ dimension). It is the position of a proposition along this dimension that determines which other propositions might in principle be employed to rationally undermine one’s belief that it is true. A Moorean fact simply falls so far along the relevant dimension that there is simply no place to stand from which one might hope to dislodge it.

On this picture, it is not that Moorean facts differ in kind from more mundane propositions and inherit their relative immunity by dint of possessing some special property or feature which they and they alone possess. Rather, their relative immunity is a de facto matter, and consists in their scoring highest along that dimension—whatever it is—which determines the relative vulnerability of *any* proposition.¹³

In what follows, I’ll proceed on the assumption that this is in fact the correct way to think about Moorean facts: the framework is adopted as a working hypothesis, in the hope that proceeding in this way might shed light not only the notion of a Moorean fact but also on the relevant norms of belief revision themselves. The immediate task will be to figure out exactly what the relevant dimension might be.

Why not simply settle for KNOWN BETTER? My own primary reason for dissatisfaction with this norm is not so much a conviction that it is false but rather that it is extremely obscure. What is it, exactly, for one proposition to be ‘known better’ than another? One difficulty, perhaps, is that the locution ‘known better’ or ‘better known’ strongly suggests knowledge by acquaintance rather than propositional knowledge. That

is, the ‘known better’ locution seems difficult to disentangle from its strong associations with the idea of *greater familiarity*—which, presumably, is not what is at issue here. We should, I think, insist upon a less enigmatic formulation of the relevant norm.

A more popular candidate for the relevant dimension is *plausibility*. As we’ve seen, Fine holds that it is the plausibility of our ordinary beliefs that justifies us in supposing that skeptical arguments must be flawed. Plausibility also seems to be the central notion for Lycan (2001). Consider then, the following norm:

MORE PLAUSIBLE: One should never abandon a belief in response to an argument when the proposition believed is more plausible than (at least one of) the premises of the argument.

One immediate attraction of MORE PLAUSIBLE for the Moorean is the following: the sorts of common sense propositions which are his stock in trade really do seem more plausible than the kinds of principles that are typically employed by the skeptic in attempting to cast doubt on them. I take this point to have been successfully established by Lycan, who drives it home with great gusto. Given this fact about comparative plausibility, the truth of MORE PLAUSIBLE would seem to deliver a quick vindication of the Moorean response to skepticism.

Unfortunately for the Moorean, MORE PLAUSIBLE is false—at least, it’s false if we understand ‘plausibility’ in its literal sense. For strictly speaking, the plausibility of a proposition concerns, not its all-things-considered worthiness of belief, but rather its *apparent* or *seeming* worthiness of belief, or its worthiness of belief upon preliminary examination. Roughly: a proposition is plausible to the extent that it seems to be true to one who considers it. However, as Earl Conee has noted (2001: 57), plausibility in this

sense is not a good candidate for being that which determines normative facts about what one ought to believe all things considered. Indeed, a given proposition's being extremely plausible seems to be consistent with its being known to be false: Frege's Unrestricted Comprehension Principle does not cease to be plausible when one learns of its falsity. Given that plausibility is consistent with known falsity, it's clear that comparative plausibility is not the correct guide to belief revision. Thus, MORE PLAUSIBLE is itself an example of a plausible principle that turns out to be false.¹⁴

Perhaps the most popular candidate for the relevant dimension among Mooreans has been that of **certainty**. Moore himself employed the vocabulary of certainty throughout his corpus, and others have followed him in this. The following is a characteristic passage from Moore:

Russell's view that I do not know that this is a pencil or that you are conscious rests, if I am right, on no less than four distinct assumptions...And what I can't help asking myself is this: Is it, in fact, as certain that all these four assumptions are true, as that I do know that this is a pencil and that you are conscious? I cannot help answering: It seems to me *more* certain that I *do* know that this is a pencil and that you are conscious, than that any single one of these four assumptions is true, let alone all four...of no one of these...do I feel as certain as that I do know for certain that this is a pencil (2000: 29, all emphases his).

Compare Armstrong on 'the less certain' and 'the more certain':

It is the bedrock of our beliefs that G.E. Moore defended in his vindication of common sense...One of the problems involved in casting doubt upon such beliefs is that the doubt-casting arguments require premises, but it is not easy to see where the premises can be collected. To use premises which are not drawn from the bed-rock of our beliefs is to bring the less certain as a reason for doubting the more certain (1983: 53-54).

Consider then the following norm:

MORE CERTAIN: One should never abandon a belief in response to an argument when one is more certain of that belief than one is of (at least one of) the premises of the argument.

How should we understand MORE CERTAIN? It is a familiar fact that talk of ‘certainty’ is often ambiguous. On the one hand, ‘certainty’ might mean *psychological certainty* or confidence. In this sense of certainty, to say that one is more certain of p than of q is to say that one is *more confident* that p is true than that q is true. It is thus to report on one’s present psychological state. Alternatively, ‘certainty’ might mean *evidential certainty*. The notion of evidential certainty concerns, not one’s actual level of confidence that some proposition is true, but rather the level of confidence that it is rational for one to have that the proposition is true given one’s epistemic situation. This is the sense of certainty which is in play in statements such as ‘Given what we know now, it is certain that there is no intelligent life at the center of the sun’. When used in this way, certainty concerns what it is *reasonable to believe* given the evidence and arguments with which one is acquainted.¹⁵

Some Mooreans suggest that it is psychological certainty which is the key notion. The idea that it is psychological certainty or confidence which is crucial seems to be the view of Pollock and Cruz:

If we reflect upon our beliefs, we will find that we are more confident of some than of others. It is reasonable to place more reliance on those beliefs in which we have greater confidence, and when beliefs come in conflict we decide which to reject by considering which we are least certain of...In typical skeptical arguments, we invariably find that we are more confident that we have the knowledge seemingly denied us than we are of some of the premises. Thus it is not reasonable to adopt the skeptical conclusion that we do not have that knowledge. The rational stance is

instead to deny one or more of the premises (1999: 6-7).

Moore himself is sometimes interpreted as having held that what is crucial are psychological facts about how confident one is that various propositions are true.

Consider, for example, the interpretation offered by Soames:

As Moore saw it, conflicts between speculative philosophical principles and the most basic convictions of common sense confront one with a choice. In any such case, one must give up either one's common sense convictions or the speculative philosophical principle. Of course, one ought to give up whichever one has the least confidence in. But how, Moore wondered, could anyone have more confidence in the truth of a general philosophical principle than one has in the truth of one's most fundamental convictions—convictions such as one's belief that there are many different objects, and many different people, that exist independently of oneself? In the end, Moore came to think that one's confidence in a general principle of philosophy never could outweigh one's confidence in convictions such as these...As a result philosophers have nothing that could be used to undermine the most central and fundamental parts of what we take ourselves to know (2003a: 8-9).

Here is a seemingly straightforward model of how we should resolve conflicts among our beliefs. Moreover, if this model is correct, it would seem to constitute a quick vindication of the Moorean response to skepticism. Consider then the following norm of belief revision:

MORE CONFIDENT In resolving conflicts among one's beliefs, one should always favor those beliefs of which one is more confident over those beliefs of which one is less confident.

The suggestion that, in revising one's beliefs, one should resolve conflicts in favor of those beliefs of which one is more confident has an undeniable ring of plausibility to it. Indeed, the suggestion that one should favor those beliefs of which one is more confident might seem to be simply common sense. (Might one claim that this suggestion is itself a

Moorean fact?) What, after all, is the alternative—favoring those beliefs of which one is *less* confident? Moreover, the idea that the Moorean facts are just those propositions in which we invest the greatest confidence is itself a plausible view about what ultimately distinguishes Moorean facts from other propositions. Finally, as we've just noted, MORE CONFIDENT is explicitly embraced by some Mooreans and is attributed to Moore himself by Soames. For these reasons, I want to consider it at some length.

Despite its plausibility, I don't believe that MORE CONFIDENT withstands scrutiny. In what follows, I'll consider what I take to be the two most natural ways of understanding this norm and argue that, on neither interpretation is it a good candidate for being the norm that we seek.

First, let's try to bring MORE CONFIDENT into sharper focus. As Pollock and Cruz emphasize, we are more confident of some of our beliefs than of others. Imagine an ordered list of all of those propositions that you believe at the present moment, time t_0 . The position of a given proposition on the list is determined by how confident you are that that proposition is true: the more confident you are of its truth, the higher its position on the list. Call this **The List of Things You Believe**. At time t_1 , an instant later, the skeptic will present you with an argument that you have never seen before. Of course, a formidable skeptic will choose his premises with great care—unless the skeptic finds premises that you believe--or at least, are strongly inclined to believe--he has no hope of dialectical effectiveness. Imagine that the best case scenario for the skeptic is realized: at time t_1 , the skeptic succeeds in producing an argument that is transparently valid, and the premises of which all appear on The List of Things that You Believe. Of course, the negation of the skeptic's conclusion (e.g., 'We know a lot') will also be on the List. The

skeptic has thus succeeded in at least this much: he has succeeded in identifying a genuine conflict among your beliefs. The List of Things that You Believe is inconsistent. Question: How should you resolve this conflict? Here is a possible decision procedure for doing so. Locate the position of the proposition that is under attack by the skeptic on the List, and compare its position to the position of each of the skeptic's premises. Resolve the conflict by eliminating whichever proposition was lowest on the List at t_0 from the new List of Things that You Believe at time t_1 .

Notice that, if this decision procedure is in fact the one which we should employ, then, given that the Moorean facts are just the highest propositions on the List at time t_0 , it seems as though the Moorean will inevitably win against the skeptic. For at time t_0 , the Moorean will be able to reason as follows: 'No matter what happens at t_1 , I know now that, even if I'm forced to employ the decision procedure then, it will tell me to retain my belief in the Moorean fact and abandon my belief in one of the skeptic's premises'.

Consider then the following interpretation of MORE CONFIDENT:

MORE CONFIDENT (1): In resolving conflicts among one's beliefs, one should abandon one's belief in whichever proposition one was least confident of immediately prior to becoming aware of the conflict.

Unfortunately for the Moorean, the decision procedure in question is not correct. In order to appreciate its inadequacy, consider first the following point: it might very well be that the fact that a given proposition occupies a high position on the List of Things that You Believe at time t_0 depends in part on your assuming that *there are no formidable arguments to be made against that proposition*. Indeed, I suspect that this is the usual case with respect to propositions of which we are extremely confident. Typically, when I

am extremely confident that something is true, I also think that, if someone were to argue that it is not in fact true, he or she wouldn't get very far. For example, I am extremely confident that the next President of the United States will be either a Democrat or a Republican. My confidence that this is so is by no means independent of my belief that if someone were to argue that this is not the case ('No, it won't be a Democrat or a Republican, a third party candidate will win') the case that he or she would be able to offer would be extremely weak. Contrast some belief that I hold with considerably greater diffidence: for example, my belief that capital punishment as presently practiced in the United States does not have a deterrent effect on crime. My relative diffidence here is not at all unrelated to my belief that, if someone were to argue for the opposite conclusion, he or she could adduce some formidable reasons in support of that view.¹⁶ The general moral: how confident one is that something is true is not independent of one's expectations about the quality of the case that might be made against it. This fact might seem obvious. But it's enough, I think, to undermine the decision procedure on offer.

For suppose that my being extremely confident that *p* is true at the earlier time *t*₀ depends on my (perhaps tacit) assumption that there are no formidable arguments for not-*p*. Suppose further that at time *t*₁, I discover first-hand that this assumption is false: someone presents me with a formidable argument for not-*p*. The argument in question is transparently valid and each of its premises appears on the List of Things That I Believe. Perhaps one of its premises falls below *p* on the List. Does it follow from this that I should retain my belief that *p* in these circumstances? It does not follow. For the superior position of *p* on the original List depended in part on an assumption that I *now*

know to be false: viz., that there are no formidable arguments for not-p. When I'm subsequently presented with a formidable argument for not-p, I in effect come into possession of a new piece of relevant information, viz. that there is such an argument. To rely solely on how confident I was that various propositions are true before I came into possession of this new piece of relevant information would be in effect to neglect part of what is now my total evidence.¹⁷

Of course, in response to this kind of objection, someone might offer a somewhat different model of how to resolve conflicts in one's beliefs, a model on which confidence remains the central notion. In particular, someone might propose the following: 'Look, what matters isn't how confident you *were* that various things are true back at t0, prior to being presented with the skeptic's argument. Rather, what matters is how confident you *are* at time t1, once you have seen the argument, or (better yet) at some later time t2, after you have had an opportunity to thoroughly scrutinize and digest the argument. That is, what matters is how confident you are that the relevant propositions are true *after the dust has settled*.'

Consider then

MORE CONFIDENT (2): In resolving conflicts among one's beliefs, one should abandon one's belief in whichever proposition one is least confident of once one becomes aware of the conflict.

However, unlike the previously considered interpretation of MORE CONFIDENT, the view on offer simply could not be a general recipe or decision procedure for resolving conflicts among one's beliefs. For consider again my situation, immediately after a newly-encountered argument has called my attention to an inconsistency among my

beliefs. Suppose that I'm in the process of actively *deliberating* about how to resolve this inconsistency; that is, suppose that I have not yet made up my mind which of the conflicting beliefs I will abandon and which I will retain. In these circumstances, the recommendation of the previously-considered norm--'Abandon whichever belief you were least confident of, *prior* to being presented with the argument'--is at least applicable advice (although for the reasons provided above, it is not, I think, the correct view). In contrast, the present recommendation--'Abandon whichever belief you are least confident of now that you have seen the argument'-- is simply not advice that is applicable in these circumstances. For, ex hypothesi, what I am in the process of deciding is exactly how confident I will be that the various propositions are true, now that I have seen the argument. In these circumstances, one simply cannot appeal to how confident one will be after the dust settles, for what is at issue in one's deliberations is precisely how the dust *ought* to settle.

In deliberating as to how I should resolve a newly-discovered inconsistency among my beliefs, what, exactly, am I attempting to figure out? Simply this: which of the beliefs it is more reasonable for me to retain, given the totality of relevant evidence and arguments to which I have been exposed—including the argument with which I have just been presented. That is, the true norm here, I think, is simply this:

***MORE REASONABLE** In resolving conflicts among one's beliefs, one should always favor those beliefs that it is more reasonable for one to think are true given the totality of evidence and arguments to which one has been exposed.

Indeed, I think that *this* norm is trivial, or close to trivial. Thus, the sense of 'certainty' in which it's true that we should favor beliefs that are 'more certain' over those that are 'less

certain' is evidential certainty, not psychological certainty. We can view potentially more informative norms such as MORE CONFIDENT as attempts to specify *what makes it the case* that it is more reasonable to retain some beliefs rather than others in cases of conflict. However, if the arguments provided above are sound, none of the other norms that we have canvassed is even extensionally equivalent to MORE REASONABLE in its recommendations. What's the significance of this?

If I'm correct in thinking that MORE REASONABLE is the operative norm, and that this norm is not equivalent to norms such as MORE CONFIDENT or MORE PLAUSIBLE, then this is, I think, a disappointing result for the Moorean. For even if the Moorean is correct in thinking that MORE REASONABLE will in fact always favor Moorean facts over skeptical premises, it seems hopeless to *show* that this is so, or to *explain why* this is so, by appeal to MORE REASONABLE itself. Here a contrast with the discarded norm MORE CONFIDENT is instructive. On the picture suggested by MORE CONFIDENT, normative facts about how one ought to revise one's beliefs in cases of conflict are in effect constituted or determined by psychological facts about how confident one is that the propositions in question are true. Thus, if MORE CONFIDENT were true, the Moorean would have a potentially satisfying answer to the question of why, as a general rule, one should retain one's belief in a Moorean fact when it conflicts with a skeptical principle. Namely: one should retain one's belief in the Moorean fact because one is more confident that it is true than that the skeptic's premises are true, and these are the facts which determine how one ought to respond to the conflict. This, at least, has the form of an acceptable answer to what is the most pressing question for the Moorean. In contrast, the attempt to answer the same question by appealing in a parallel

way to MORE REASONABLE would yield the following: ‘One should always retain one’s belief in a Moorean fact rather than the skeptic’s premises because this is what it is more reasonable to do’. But this, of course, is a mere restatement of the Moorean thesis.

No doubt, the Moorean is convinced that the norm MORE REASONABLE will always favor retaining our common sense beliefs when correctly applied, and that, when push comes to shove, this norm will dictate abandoning the philosophical principles on which the skeptic relies. But again: why should that be? In the last section, I want to sketch what I take to be the Moorean’s most promising answer to this question.

4. Moore and Metaphilosophy

Let’s briefly recapitulate. In section 2, I suggested that we should view the Moorean as someone who both (i) confidently predicts that we will never encounter a compelling skeptical argument and (ii) endorses a policy of concluding that any particular skeptical argument must be flawed in some way (even if we are unable to identify the flaw). I also suggested that, for the Moorean, the endorsement of the policy is more fundamental than the prediction: *that* it will never be reasonable to abandon one’s belief in a Moorean fact is in effect guaranteed by a correct application of what are in fact the true norms of belief revision. In section 3, I distinguished two different ways of understanding this suggestion. According to the first way, Moorean facts differ in kind from other propositions, and it is simply a fundamental epistemic norm that one should never abandon one’s belief in a Moorean fact. According to the second way, while it’s true that one should never abandon one’s belief in a Moorean fact, this is not itself a fundamental

epistemic norm; rather, this fact falls out of more fundamental norms that do not themselves make mention of Moorean facts but rather govern how one ought to revise one's beliefs more generally. I suggested that this latter way of thinking about Moorean facts is preferable. I thus canvassed a number of proposed norms of belief revision that looked initially promising for the Moorean. I provided reasons for doubting that any of these norms was correct. I ultimately endorsed the relatively trivial norm MORE REASONABLE but suggested that, in part because of its triviality, it seems not to advance the Moorean's cause. Given that MORE REASONABLE is in fact the correct norm for resolving conflicts among one's beliefs, the Moorean still needs to supply some reason for thinking that it will be more reasonable to retain our common sense beliefs when these turn out to be inconsistent with the skeptic's premises. How might the Moorean do this?

A recurrent theme among interpreters sympathetic to Moore is that his deepest point against the skeptic is a metaphilosophical one. According to this line of thought, the skeptic is ultimately undermined by a correct appreciation of philosophical method. This understanding of Moore has been embraced by interpreters both early and late. Thus, Malcolm (1942) famously interpreted Moore as offering paradigm case arguments on behalf of common sense. According to Malcolm, the reason why it would be unreasonable to adopt the skeptic's claims is that those claims 'go against ordinary language' (349) and this suffices to guarantee their falsity. However, the paradigm case argument is rightly discredited, and Moore declined to endorse Malcolm's interpretation when explicitly offered the chance to do so (Moore 1942)¹⁸

More recently, Soames has also suggested that Moore should be understood as making a point about philosophical method. According to Soames

[Moore's] position regarding the propositions of common sense is that they constitute the starting point for philosophy, and, as such, are not the sorts of claims that can be overturned by philosophical argument (2003a p.5).¹⁹

However, it is at the very least unclear why something's being among the starting points for philosophy would render it as such immune to being subsequently overturned on the basis of philosophical argument. Perhaps on some substantive views of what makes something a proper starting point for philosophy, anything which qualifies as such will in fact enjoy the relevant kind of immunity. Thus, perhaps on a broadly Cartesian view of philosophical inquiry, the qualifications for being a proper starting point are sufficiently demanding that anything which satisfies them is ipso facto not susceptible to being subsequently overturned. At the other end of the spectrum, Harman (1999), following Quine, insists that the proper starting point for philosophy is simply everything that one presently believes. On such a view, the suggestion that the starting points for philosophy cannot be overturned by philosophical argument is equivalent to the suggestion that philosophical argument is powerless to overturn *any* of our pre-philosophical beliefs, no matter how tenuous or tentatively-held—a conclusion that is surely too strong. In general, to take something as a philosophical starting point is not to hold that it is such immune to being overturned by subsequent philosophical argument—any more than to take something as a starting point for inquiry more generally is to hold that it cannot be given up at some later stage of inquiry.

Nevertheless, I believe that the view that Moore's deepest point against the skeptic is a metaphilosophical one is correct. Why can't the skeptic win? The Moorean's best

answer to this question, I think, runs as follows: the skeptic cannot win because the skeptic is implicitly committed to a methodology for philosophical theorizing that does not withstand scrutiny once it is forced out into the open.

In order to see why this might be so, let's begin by considering one last interpretation of Moore, an interpretation on which metaphilosophical considerations remain central: Moore as **particularist**. This interpretation is offered by Roderick Chisholm in his classic discussion of 'the Problem of the Criterion' (Chisholm 1973). There, Chisholm distinguishes between two different ways in which one might theorize about knowledge: as a particularist or as a methodist. Roughly, the particularist takes as data our considered judgements about whether knowledge is present or absent in particular cases (e.g., 'I know that I have hands', 'I know that my name is Thomas Kelly', 'George Bush knows that he is a Republican', 'I do not know the identity of the next President of the United States'). He then uses these judgements about particular cases to evaluate proposed general principles about knowledge. When a general principle conflicts with a considered judgement about a particular case—say, it suggests that I *do* know the identity of the next President of the United States or that I do *not* know that my name is Thomas Kelly--then this counts as at least some evidence against the principle. On the other hand, to the extent that the general principle accommodates our judgements about cases, this counts in its favor. Thus, for the particularist, the tenability of a general principle is ultimately determined by how well it accommodates our judgements about cases. One philosophical theory is better supported than another when it does a superior job overall of accommodating such judgements. For Chisholm, Moore is the example par excellence of a particularist.

In contrast to the particularist, the methodist begins with a commitment to some general philosophical principle or principles about the nature of knowledge and utilizes these principles to arrive at judgements about whether knowledge is present or absent in particular cases. Chisholm's own example of a methodist is David Hume, who begins with a firm commitment to a general philosophical theory about human knowledge, viz. empiricism, and appeals to that theory again and again in drawing conclusions about whether human beings should be credited with this or that bit of putative knowledge. For Hume, our pre-theoretical judgements about whether knowledge is present or absent in this or that particular case should be retained if and only if the judgements in question match those returned by a consistent and rigorous application of empiricism.

Strictly speaking, particularism is not inconsistent with radical forms of skepticism (Cf. Lemos 2004:110). Indeed, at least one prominent argument for radical skepticism takes as its starting point an intuitive judgement about a particular case.²⁰ However, as Chisholm emphasizes, it seems clear that the skeptic will find the methodology of methodism more congenial than that of particularism. For if one evaluates general principles about knowledge by reference to how well they fit with one's pre-theoretic judgements about whether knowledge is present in particular cases, then, given that at least a substantial number of such judgements are to the effect that knowledge is present, it seems as though the kinds of epistemic principles that would generate skeptical conclusions if true will be eliminated from consideration on this basis. To put the point in comparative terms: if the sole standard of evaluation for theories is how well they accommodate our judgements about particular cases, then surely the *best* theory will not overturn all (or almost all) of one's positive attributions of knowledge.

In contrast, suppose that one proceeds as a methodist. In that case, the purpose in seeing how a presently-accepted principle matches one's judgements about cases is to answer questions about which judgements about cases should be retained and which discarded. As a general rule, friction between judgements about principles and judgements about cases are resolved in favor of the former and at the expense of the latter. Given such a method, there seems to be nothing to preclude a rather sweeping revision of our judgements about cases. In short, methodism seems to leave the door open to skepticism—or more broadly, views that are radically revisionary with respect to our pre-theoretical judgements about particular cases-- in a way that particularism does not. Indeed, I think that the point generalizes beyond epistemology and holds for philosophical revisionism in other domains as well.²¹

Chisholm ultimately embraces the particularist methodology that he finds in Moore. Should we follow him? In considering this question, the first thing that we should note is that the choice between particularism and methodism as described above is clearly not exhaustive. The particularist holds that our judgements about cases have priority over our judgements about principles; the methodist holds that our judgements about principles have priority over our judgements about cases. Once posed in these terms, it seems clear that there is a further alternative. For one might hold that neither type of judgement has such priority. Rather, our most fundamental judgements about cases and about principles should both receive substantial weight in our theorizing as we attempt to achieve a stable coherence among judgements at different levels through a process of mutual adjustment. This, in its barest essentials, is the method of **reflective equilibrium**.

We have then three possible methods for theorizing about knowledge: particularism, methodism, and reflective equilibrium. Even this enriched picture of our options, however, is apt to make the range of possible methodologies appear more limited than it in fact is. For I've described an especially pure form of each of these views. Thus, as described above, the particularist is a theorist who gives literally no weight to the intuitive plausibility of general principles; for him, the sole determinant of the tenability of a general principle is how well it accords with our judgements about cases. On this view, if two incompatible principles matched our judgements about cases equally well, there would be nothing to choose between them, even if one of the two principles had a high degree of intuitive plausibility while the other was completely counterintuitive, hopelessly gerrymandered and *ad hoc*. This is *hyper-particularism*. Hyper-particularism is not an attractive view, and I doubt whether it can be charitably attributed to anyone. But surely there are methodologies worthy of the name 'particularism' other than hyper-particularism: methodologies which, while placing a special emphasis on preserving judgements about cases, also give some significant weight to other factors as well. In a similar way, we can distinguish between hyper-methodism—a view on which our judgements about cases are given zero or negligible weight—from more temperate, moderate versions.

Of course, the more weight a particularist is willing to give her judgements about principles in the course of her theorizing, the less her methodology will resemble hyper-particularism and the more it will begin to shade off into a version of reflective equilibrium. In a parallel way, the more weight a methodist is willing to give to his judgements about particular cases, the less his view will resemble hyper-methodism and

the more it will begin to shade off into (another version of) reflective equilibrium. Undoubtedly, the boundaries in question are vague. The tripartite division between particularism, methodism, and reflective equilibrium encourages us to think of the relevant logical space as carved up into three discrete units. But we do better, perhaps, to think of a continuum of methodologies, ordered according to the relative weights that each would have us give to our judgements about cases on the one hand and our judgements about principles on the other.

Having noted this complication, let's return to the simple tripartite division between particularism, methodism, and reflective equilibrium. Above, I suggested that for the purposes of the skeptic, methodism seems to hold out more hope than particularism. How should we expect the skeptic to fare if we proceed according to the method of reflective equilibrium? Despite its considerable popularity, the method of reflective equilibrium is typically not described in much detail by either its proponents or its detractors: presently, we lack a well-worked out account. I will not attempt to rectify that state of affairs here. Nevertheless, even in the absence of such an account, we have good reason to think that the skeptic is unlikely to fare well given this methodology. For the reflective equilibrium theorist resembles the particularist in that she too wants to give substantial weight to our judgements about cases in assessing which principles it is reasonable for us to accept. Again, given that the claims of the skeptic are radically inconsistent with our judgments about cases, the reflective equilibrium theorist will join the particularist in taking this to constitute strong reason to reject the skeptic's principles. The skeptic advocates what is in effect a revolution in our thinking: the Cartesian skeptic, for example, advocates an abandonment of all of our positive

attributions of knowledge of the external world. The reflective equilibrium theorist, however, will insist that many judgements from this targeted class should be given substantial weight in our deliberations about which philosophical principles to accept and which to reject. Because of this, the method of reflective equilibrium will militate against the kind of sweeping revisionism that the skeptic envisions. In this respect, the method of reflective equilibrium seems to be an inherently conservative one. Given this conservatism, the method would not seem to be a promising one for the skeptic.

Admittedly, in the absence of a well-worked out account of reflective equilibrium, the claim that it is an inherently conservative method takes on a certain speculative air. Moreover, given that I have not attempted to spell out in detail how *I* understand the method, one might naturally worry that my claim of conservativeness is based on presupposing some perhaps idiosyncratic conception that has been left off-stage. However, it's noteworthy that reflective equilibrium has often been taken to be a fundamentally conservative method by a wide range of both its advocates and its detractors. Thus, Goodman (1953) originally proposed the method of as a way of effectively 'dissolving' Humean skepticism about the deliverances of inductive reasoning. Stich (1983, ch.4) takes the conservative character of reflective equilibrium as a reason to reject it. Particularly in the wake of Rawls' endorsement of the method in moral and political philosophy (1972), it was widely criticized on the grounds that it is overly conservative (See, e.g., Singer 1974, Copp 1985).²² On the other hand, the alleged conservatism of reflective equilibrium was both acknowledged and claimed as a virtue by the philosopher who has been perhaps its most persistent and epistemologically sophisticated defender (Harman 1994, 2003, 2004).²³ My present purpose is not to side

either with those who see the apparent conservatism of reflective equilibrium as a virtue or with those who see it as a vice. My point is rather that, unless each of these philosophers is rather fundamentally confused about how reflective equilibrium operates, then it is a conservative method. Again, inasmuch as the skeptic advocates what is in fact a genuinely revolutionary change in our thinking, reflective equilibrium would seem to be an unpromising methodology for his purposes.

Chisholm takes Moore's rejection of skepticism to follow directly from his commitment to a particularist methodology. (Indeed, the anti-skeptical promise of particularism seems to be one of the things that recommends it in Chisholm's eyes.) The question that I would like to press, however, is not which methodology best serves the interests of the anti-skeptic, but rather which methodological options are left open to the skeptic. For as we have seen, it is not only particularism which seems inimical to skepticism—the method of reflection equilibrium also seems to be a method on which it's difficult to see how the skeptic can gain the kind of traction requisite to effect the dramatic change in our views that he advocates.

The picture that the skeptic presupposes seems to be something like the following. There are certain philosophical principles that have radically revisionary implications. The intuitive plausibility of these principles renders them worthy of belief at the outset of inquiry. Moreover, this belief worthiness is indefeasible in the following sense: it survives the realization that the principles in question are inconsistent with large numbers of our most fundamental judgements about cases. The Moorean claims that it would be unreasonable to continue to accept a philosophical principle once it has been shown that this principle is inconsistent with the judgement that *Moore knows that he has hands*.

The skeptic denies this; on his view, inconsistency with *Moore knows that he has hands* is insufficient to undermine the reasonableness of continued acceptance of a philosophical principle. Moreover, it's clear that the skeptic's reaction is not a matter of dissatisfaction with Moore's specific choice of example. The skeptic would be no more impressed if Moore had chosen as his example the proposition that *Lloyd George knows that he has hands*, or that *Moore knows that he has feet*. Rather, against the Moorean, the Cartesian skeptic will maintain the following: there are no judgements attributing knowledge of the external world inconsistency with which suffices to undermine the reasonableness of continuing to accept a philosophical principle. This, of course, is a relatively strong metaphilosophical claim. Indeed, for the skeptic to insist that we endorse this claim is in effect to insist that we embrace the methodology of hyper-methodism: the view that no significant weight should be given to one's considered judgements about particular cases in evaluating principles.

At this point, the dispute between the skeptic and the Moorean seems to have devolved into a dispute about which methodology is in fact correct. How might this dispute be resolved? Chisholm held that, in the end, the choice between particularism and methodism could only be made by begging the question (1973: 37). Perhaps that's so. However, in the present context, the following seems to me to be a fairly telling point in favor of the Moorean: the methodology which the skeptic advocates seems utterly at variance with the methodology employed elsewhere in philosophy in the assessment of general principles.²⁴ Indeed, it is at variance with the generally accepted methodology for theorizing about knowledge when skepticism is not at issue. Thus, Gettier (1963) was famously taken to have refuted an extremely plausible general theory of knowledge by

showing how that theory returned verdicts about particular cases that are inconsistent with our judgements about those cases. If hyper-methodism is the correct methodology, however, then this near²⁵ universal reaction to Gettier was simply a mistake. For if hyper-methodism is correct, then to abandon a highly plausible account of knowledge on the basis of its failure to return the intuitively correct verdict about a particular case is to give too much weight to one's judgements about cases in the assessment of principles. One who looks askance at the Moorean response to skepticism should, I think, take a similarly dim view of those who take Gettier to have shown that the traditional justified true belief account of knowledge is incorrect.

For the sake of further comparison, consider also the history of the logical positivists' attempts to formulate an acceptable version of the verifiability criterion of meaning. Successive principles were abandoned on the grounds that they classified what seemed to be clearly meaningful linguistic expressions as meaningless. Here again, inconsistency with judgements about cases was taken as a sufficient reason to abandon a general principle. Moreover, this was taken as a sufficient reason not only by those unsympathetic to the general picture of meaning in play but by the positivists themselves. Again, given a methodology of hyper-methodism, this reaction was simply a mistake.

Notice, moreover, that the positivists' willingness to abandon general principles on this basis does not mean that the positivists were particularists in Chisholm's sense. On the contrary, the positivists were prepared to be quite revisionary with respect to widely-held pre-theoretical judgements about the meaningfulness of various linguistic expressions. (Indeed, an appetite for such revisionism was a primary motivation for the development of the verifiability criterion of meaning in the first place.) Similarly, it

would be unwarranted to think that all of those who took Gettier to have successfully refuted the traditional analysis of knowledge have thereby incurred a commitment to the substantive metaphilosophical view of particularism. One need not be a particularist in order to believe in counterexamples.

A traditional complaint against the skeptic is that he is guilty of importing artificially demanding standards for knowledge. According to this line of thought, the skeptic proceeds as though some condition C is necessary for knowledge (where C is a condition which human beings seldom if ever satisfy); however, in ordinary life, it is claimed, we confidently and unhesitatingly attribute knowledge even in cases in which it is clear that C is not satisfied.²⁶ I think that this charge is made with greater justice when it is leveled against some skeptics than when it is leveled against others. My own charge against the skeptic is a structurally-similar one, entered one level up: in insisting that we give so little weight to our judgements about cases in the evaluation of principles, the skeptic is guilty of importing artificially demanding standards for what it takes to undermine a general philosophical principle. That is, elsewhere in philosophy, we confidently and unhesitatingly take a general principle to have been adequately undermined by showing that it is inconsistent with sufficiently fundamental judgements about cases.

It is natural to interpret one who appeals to Moorean facts as a particularist. The Moorean is thus viewed as one who is committed to an extremely substantive metaphilosophical thesis. Indeed, some proceed as though showing the unsatisfactoriness of particularism would suffice to undermine the Moorean response to skepticism.²⁷ If the present line of thought is correct, however, then this considerably overstates the vulnerability of the Moorean. For to view the dialectic in this way is to overestimate the

Moorean's methodological commitments while underestimating those of the skeptic. One who appeals to Moorean facts does not thereby incur a commitment to particularism. Rather, the methodological commitment incurred is that there are some particular judgements inconsistency with which suffices to undermine the credibility of a general principle to the point that it is reasonable to reject that principle. And this idea is not the exclusive property of the particularist but might also be shared by methodologies that are incompatible with particularism, viz. reflective equilibrium and more moderate forms of methodism. On the contrary: it is not the Moorean but rather the skeptic who seems committed to an extremely strong metaphilosophical view, and an unattractive one at that. Indeed, on the present view of the dialectic, it is the skeptic who would seem to be guilty of the very charge that is often leveled at the Moorean: that he treats his favorite commitments as sacrosanct. For it is the skeptic who demands that we retain belief in his principles even after those principles have failed exactly the sorts of tests that are ordinarily taken to warrant rejection of general principles. In this respect, it is the skeptic, more than the Moorean, who seems to resemble the dogmatist.²⁸

Bibliography

- Ainslie, George (1975). "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control" in the Psychological Bulletin 82: 463-496.
- Ainslie, George (1986). "Beyond Microeconomics". In The Multiple Self, edited by Jon Elster (Cambridge: Cambridge University Press): 133-178.
- Armstrong, David (1983). What is a Law of Nature? (Cambridge: Cambridge University Press).
- Armstrong, David (1999). "A Naturalist Program: Epistemology and Ontology". The Eleventh Annual Romanell Lecture. Reprinted in Proceedings and Addresses of the American Philosophical Association, vol.73, No.2.
- Ayer, A.J. (1952). Language, Truth, and Logic (New York: Dover).
- BonJour, Laurence (1986). The Structure of Empirical Knowledge. (Cambridge, MA: Harvard University Press).
- BonJour, Laurence (2002). Epistemology. (Oxford, England: Rowman and Littlefield Publishers).
- Byrne, Alex (2004). "How Hard Are the Sceptical Paradoxes?" Nous, 38, 299-325.
- Cassell, Paul (2004). "In Defense of the Death Penalty". In Debating the Death Penalty: Should America Have Capital Punishment? The Experts on Both Sides Make Their Best Case. Edited by Hugo Adam Bedau and Paul G. Cassell (Oxford: Oxford University Press).
- Chisholm, Roderick (1973). The Problem of the Criterion. (Milwaukee, WI: Marquette University Press).
- Cohen, Stewart (1988). "How to Be a Fallibilist" in James Tomberlin (ed.) Philosophical Perspectives 2 (Atascadero, CA: Ridgeview Publishing Company).
- Conee, Earl (2001). "Comments on Bill Lycan's Moore Against the New Skeptics", Philosophical Studies 103: 55-59.
- Copp, David (1985). "Considered Judgements and Justification: Conservatism in Moral Theory", in D. Copp and M. Zimmerman (eds.) Morality, Reason, and Truth. (Totowa, NJ: Rowman and Allenheld): 141-169.
- Daniels, Norman (1996). Justice and Justification: Reflective Equilibrium in Theory and

Practice. (Cambridge: Cambridge University Press).

Davies, Martin (1998). "Externalism, Architecturalism, and Epistemic Warrant", in Wright, Crispin, Smith, Michael, and Macdonald, Cynthia (eds.), Knowing Our Own Minds: Essays in Self-Knowledge, (Oxford: Oxford University Press): 321-361.

Davies, Martin (2000). "Externalism and Armchair Knowledge", in Boghossian, Paul and Peacocke, Christopher (eds.), New Essays on the A Priori (Oxford: Oxford University Press): 384-414.

Davies, Martin (2003). "The Problem of Armchair Knowledge", in Nuccetelli, Susana (ed.) News Essays on Semantic Externalism and Self-Knowledge. (Cambridge, MA: MIT Press): 23-55.

Davies, Martin (2004). "Epistemic Entitlement, Warrant Transmission, and Easy Knowledge", Aristotelian Society Supplement, 78, pp. 213-245.

DeRose, Keith (1995). "Solving the Skeptical Problem" Philosophical Review 104: 1-52.

Donellan, Keith (1967). "The Paradigm Case Argument", in Paul Edwards (ed.) The Encyclopedia of Philosophy (New York: MacMillan):39-44.

Edwards, Paul (1951). "Bertrand Russell's Doubts About Induction", in Antony Flew (ed.) Essays in on Logic and Language (Oxford: Oxford University Press).

Feldman, Richard. "Contextualism and Skepticism" in Philosophical Perspectives 13: 91-114.

Fine, Kit (2001). "The Question of Realism". The Philosophers' Imprint vol.1, no.1.

Gettier, Edmund. "Is Justified True Belief Knowledge?" Analysis 23: 121-123.

Goodman, Nelson (1953). Fact, Fiction, and Forecast. (Cambridge, MA: Harvard University Press).

Hanfling, Oswald (1991). "What is Wrong with the Paradigm Case Argument?" in Proceedings of the Aristotelian Society vol.91: 21-38.

Harman, Gilbert (1973). Thought. (Princeton, NJ: Princeton University Press).

Harman, Gilbert (1986). Change in View. (Cambridge, MA: The MIT Press).

Harman, Gilbert (1994). "Epistemology and the Diet Revolution", in Michaelis Michael and John O'Leary-Hawthorne (eds.), Philosophy in Mind. (The Netherlands: Kluwer Academic Publishers).

- Harman, Gilbert (1999). Reasoning, Meaning, and Mind. (Oxford: Oxford University Press).
- Harman, Gilbert (2003). "Skepticism and Foundations", in Steven Luper (ed.) The Sceptics: Contemporary Essays. (Ashgate: Aldershot): 1-11.
- Harman, Gilbert (2004). "Three Trends in Moral and Political Philosophy" in Value Inquiry, 37.3.
- Hetherington, Stephen (2001). Good Knowledge, Bad Knowledge. (Oxford: Oxford University Press).
- Hirsch, Eli (2002). "Against Revisionary Ontology" in Philosophical Topics vol.30, no.1.
- Kagan, Shelly (1989). The Limits of Morality. (Oxford: Oxford University Press).
- Lemos, Noah (2004). Common Sense. (Cambridge: Cambridge University Press).
- Lewis, David (1986). On the Plurality of Worlds (Blackwell Publishers: Cambridge, MA).
- Lewis, David (1999) "Elusive Knowledge", reprinted in his collection Papers in Metaphysics and Epistemology (Cambridge: Cambridge University Press), pp. 418-445.
- Lycan, William (2001). "Moore against the New Sceptics". Philosophical Studies 103 (1):35-53.
- Malcolm, Norman (1942). "Moore and Ordinary Language", in The Philosophy of G.E. Moore, vol.I, ed. P.A. Schilpp (La Salle, IL: Open Court) 343-368.
- Moore, G.E. (1942). "Replies" in The Philosophy of G.E. Moore, vol.I ed. P.A. Schilpp (La Salle, IL: Open Court).
- Moore, G.E (1993). Selected Writings. Edited by Thomas Baldwin. (London: Routledge).
- Moore, G.E. (2000). "Four Forms of Skepticism". Reprinted in Kim and Sosa (eds.) Epistemology: An Anthology (Oxford: Blackwell Publishers): 27-29.
- Nolan, Daniel (2005). David Lewis. (McGill: Queen's University Press).
- Nozick, Robert (1981). Philosophical Explanations. (Cambridge, MA: Harvard University Press).

- Nozick, Robert (1993). The Nature of Rationality. (Princeton: Princeton University Press).
- Pollock, John and Cruz, Joseph (1999). Contemporary Theories of Knowledge (Oxford: Rowman and Littlefield).
- Pryor, James (1996). How to be a Reasonable Dogmatist. Unpublished doctoral dissertation. Princeton University, 1996.
- Pryor, James (2000). "The Skeptic and the Dogmatist", Nous 34: 517-549.
- Pryor, James (2004). "What's Wrong With Moore's Argument?" in Philosophical Issues 14.
- Rawls, John (1972). A Theory of Justice. (Cambridge, MA: Belknap)
- Rosen, Gideon (1994). "Objectivity and Modern Idealism: What is the Question?" in Michaelis Michael and John O'Leary-Hawthorne (eds.), Philosophy in Mind. (The Netherlands: Kluwer Academic Publishers): 277-319
- Scanlon, Thomas (2002). "Rawls on Justification" in The Cambridge Companion to Rawls, edited by Samuel Freeman (Cambridge: Cambridge University Press): 139-167.
- Singer, Peter (1974). "Sidgwick and Reflective Equilibrium" in the Monist 58: 490-517.
- Singer, Peter (1995). Rethinking Life and Death: the Collapse of our Traditional Ethics. (New York: St. Martin's Press).
- Soames, Scott (2003a). Philosophical Analysis in the Twentieth Century. Volume 1: The Dawn of Analysis. (Princeton, NJ: Princeton University Press).
- Soames, Scott (2003b). Philosophical Analysis in the Twentieth Century. Volume 2: The Age of Meaning. (Princeton, NJ: Princeton University Press).
- Stitch, Stephen (1983). The Fragmentation of Reason (Cambridge, MA: The M.I.T. Press).
- Strawson, P.F. (1985). Skepticism and Naturalism: Some Varieties. (New York: Columbia University Press).
- Stroud, Barry (1984). The Significance of Philosophical Skepticism. (Oxford: Oxford University Press).
- Unger, Peter (1975). Ignorance: A Case for Skepticism (Oxford: Oxford University Press).

- Weatherson, Brian (2003). "What Good Are Counterexamples?" Philosophical Studies 15: 1-31.
- Wittgenstein, Ludwig (1972). On Certainty. Edited by G.E.M. Anscombe and G.H. von Wright. (New York: Harper and Row)
- Wright, Crispin (1985). "Facts and Certainty", Proceedings of the British Academy 71: 429-472.
- Wright, Crispin (2000). "Cogency and Question-begging: Some reflections on McKinsey's Paradox and Putnam's Proof" in Philosophical Issues 10: 140-163.
- Wright, Crispin (2002). "(Anti)skeptics simple and subtle: Moore and McDowell". Philosophy and Phenomenological Research 65:.330-348.
- Wright, Crispin (2003). "Some Reflections on the Acquisition of Warrant by Inference", in Nuccetelli, Susana (ed.) New Essays on Semantic Externalism and Self-Knowledge (Cambridge, MA: MIT Press): 57-77.
- Wright, Crispin (2004). "Warrant for Nothing (and foundations for free)? Aristotelian Society Supplement.

¹ A note about the interpretation of Lewis is in order here. One who emphasizes this Moorean aspect of Lewis' thought (as I intend to do) might seem to slight other aspects of his thought which are more congenial to the skeptic. In particular, as is often noted, Lewis' own version of contextualism is in one crucial respect unusually concessive to the skeptic, inasmuch as it makes it a relatively trivial matter for the skeptic to establish extremely demanding standards for knowing in a given conversational context. The upshot of this for Lewis is that any minimally competent skeptic is more or less guaranteed to win any argument with a non-skeptic. Given this, isn't it misleading in the extreme for me to dwell on a few Moorean pronouncements?

Considerable care is needed here, however. A primary motivation for Lewis' particular version of contextualism is his desire to reconcile infallibilism about knowledge with the truth of *We know a lot* (418-420). However, as Lewis makes explicit, his commitment to the truth of *We know a lot* is deeper than his commitment to infallibilism: if forced to choose, he would abandon the latter in order to preserve the former (419). Thus, Lewis' commitment to the truth of many of our everyday attributions of knowledge is more fundamental than his commitment to those aspects of his epistemology that are more congenial to the skeptic. And indeed, similar Moore-style appeals recur throughout the Lewisian corpus. Compare, for example, the following from Lewis (1986): 'Our knowledge of mathematics is ever so much more secure than our knowledge of the epistemology that seeks to cast doubt on mathematics' (p.109).

For more on this Moorean aspect of Lewis' thought, see Nolan (2005), chapter 9.

² Prominent figures here include Bonjour (1985, ch.1), Stroud (1984, ch.3), and Unger (1975).

³ In addition to Lewis, Armstrong, and Fine, I would also include Lycan (2001), Pollock and Cruz (1999, especially pages 6-7), Soames (2003a), and Hirsch (2002) among the ranks of prominent contemporary Mooreans. Finally, the usual allures of historical revisionism notwithstanding, I believe that Moore (1993) himself was a Moorean in the relevant respects.

⁴ For the suggestion that our own philosophical era is more commonsensical than some previous ones, see also Rosen (1994: 277-278) on nineteenth versus twentieth century idealisms, and especially Soames (2003a). Indeed, Soames holds that the adoption of what is an essentially Moorean orientation towards philosophy (or perhaps, a Moorean-cum-Quinean orientation) stands as 'one of the two most important achievements that have emerged from the analytic tradition' (pp.xi-xii). He elaborates on this theme as follows:

One of the recurring themes in the best analytic work...has been the realization that no matter how attractive a philosophical theory might be in the abstract, it can never be more securely supported than the great mass of ordinary, pre-philosophical convictions arising from common sense, science, and other areas of inquiry about which the theory has consequences. All philosophical theories are, to some extent, tested and

constrained by such convictions, and no viable theory can overturn them wholesale. Analytic philosophers are, of course, not the only philosophers to have recognized this; nor...have they always been able to resist the seductions of unrestrained, and sometimes highly counterintuitive, theorizing. Still, the tradition has had a way of correcting such excesses, and returning to firmer foundations (xi-xii).

⁵ Here and below, I use ‘skepticism’ generically, to refer to any sufficiently radical variety of the view (as opposed to, say, skepticism about the existence of God or about the claims of psychical research). If more specificity is wanted, one might take the claims of the Moorean as being directed at skepticism about our knowledge of the external world.

⁶ See, e.g., the papers by Pryor (2004), Wright (1985, 2000, 2002, 2003, 2004), and Davies (2000, 2003, 2004).

⁷ This is true of all of the works by Mooreans mentioned in footnote 3 above.

⁸ Compare Saul Kripke’s ‘dogmatism paradox’, first presented in Harman (1973).

⁹ A locus classicus of the genre is Ayer (1952).

¹⁰ On ‘reasonable dogmatism’, see Pryor (2000) and (1996). We should distinguish carefully, however, between the kind of dogmatism advocated by Pryor and the kind of dogmatism the reasonableness of which is at issue here. Again, the distinction between playing offense and playing defense (Cf. section 1 above) is apposite. According to the dogmatism advocated by Pryor, one can be justified in believing *p* even if one cannot offer a certain kind of non-question begging argument for *p*. Intuitively: one can be justified in believing things even if one is not adept at playing offense. According to the kind of dogmatism presently under consideration, one can be justified in believing propositions even when one is not adept at playing defense. That is, one might be justified in believing a proposition even if one finds oneself unable to identify any objectionable premise or transition in the skeptic’s attempt to provide a non-question-begging argument for the negation of that proposition. It’s natural to think that the task of constructing non-question-begging philosophical arguments is, in general, a more demanding intellectual task than that of raising doubts about someone else’s attempts to do the same. At least at first pass then, the kind of dogmatism at issue here would seem to be considerably more radical than the kind of dogmatism endorsed by Pryor.

¹¹ Although popular within the larger history of philosophy, the claim that a certain class of propositions enjoys an epistemically privileged role runs counter to one of the great trends of late twentieth century epistemology: its leveling tendency. As noted, many of the objects of traditional epistemological concern (Cartesian foundations, empiricist sense data reports, Carnapian linguistic rules) were alleged to enjoy a distinctive epistemic status. Indeed, the properties which allegedly set them off as distinctive (indubitability, incorrigibility, being ‘unrevisable on the basis of experience’

or ‘confirmed come what may’, and so on) often became objects of epistemological reflection in their own right. Much of the second half of twentieth century epistemology might be viewed as a kind of egalitarian reaction to these claims of epistemic privilege. Thus, Quine and his followers held that no statement is in principle immune to revision; their suggestive metaphor of ‘the web of belief’ is one according to which epistemic differences are a matter of degree as opposed to a matter of kind. Contemporary foundationalists go to great lengths to emphasize that the foundations which they posit need not possess any of the special properties (e.g., certainty, indubitability, incorrigibility) that were typically claimed for such foundations by their classical predecessors. In a similar spirit, many contemporary rationalists emphasize that, on their view, a belief’s being *a priori* justified at one time does not preclude that justification from being defeated—indeed, defeated by empirical considerations—at later times. Despite their diversity, all of these trends might be viewed as part of a larger reaction against traditional claims of epistemic privilege. On the present way of thinking about Moorean facts, the friend of such facts seems to swim against this general current.

¹² Notably, Wittgenstein (1972) seems to hold that the status of the propositions to which Moore calls our attention differs in kind, and not merely in degree, from more mundane propositions whose truth might be debated in non-philosophical contexts. Indeed, part of Wittgenstein’s dissatisfaction with Moore seems to be his sense that Moore was insufficiently appreciative of this point. Of course, it would be completely contrary to the spirit of Wittgenstein’s later philosophy to assimilate his notion of a ‘hinge proposition’ too closely to more traditional epistemological categories. My failure to engage with this fascinating though enigmatic text in the present paper does not reflect any doubt about its importance for the topic but rather an awareness of the inadequacy of my present understanding of it.

¹³ The picture adumbrated here is, of course, a Quinean one in its essentials.

¹⁴ Sometimes ‘plausible’ is simply used as a synonym for ‘reasonable to believe, all things considered’. As a terminological matter, I think that this is best avoided. Reading ‘plausible’ as ‘reasonable to believe all things considered’ in interpreting MORE PLAUSIBLE yields a norm equivalent to the norm MORE REASONABLE, which I discuss below.

¹⁵ Notice that in the passage quoted above, Moore himself seems to waver in the way that he uses ‘certainty’: some appearances of the word ‘certain’ seem to call out for a psychological reading while others suggest an evidential reading. Thus, the line ‘of no one of these assumptions do I feel as certain...’ strongly suggests psychological certainty. On the other hand, ‘Is it, in fact, as certain...?’ seems to read more naturally as concerning the evidential notion.

¹⁶ Perhaps along the lines of Cassell (2004). See especially pages 190-197.

¹⁷ Compare: at a young age, I was extremely confident—indeed, I suspect that I was psychologically certain—that the earth was stationary (in an absolute, Newtonian sense of ‘stationary’). I lost this conviction as the end result of a process of reconsideration that was prompted by being told various things by others. No doubt, the things that were told to me were at least somewhat plausible—if they had struck me as ludicrous or absurd I would have dismissed them out of hand. However, given my extreme confidence that the earth is stationary, it’s quite likely that I had greater confidence that the earth is stationary than in at least some of the things that were said to me (and certainly, greater confidence than in the conjunction of the relevant propositions). Nevertheless, the process by which I abandoned my belief that Earth is stationary was not an irrational or even an arational process.

¹⁸ A recent criticism of Malcolm’s argument in particular is Soames (2003b, Chapter 7). An excellent overview of the paradigm case argument and the classical debate over its merits is Donellan (1967); a relatively recent attempt at its rehabilitation is Hanfling (1991)

¹⁹ Compare pages xv-xvi of Soames (2003b) where the starting point idea seems to recur.

²⁰ Here I have in mind arguments of the following form:

- (1) I don’t know that I’m not a Brain-in-a-Vat (BIV).
- (2) If I don’t know that I’m not a BIV, then I don’t know that I have hands.
- (3) Therefore, I don’t know that I have hands.

See, e.g., Nozick (1981: Ch.3), Cohen (1988), and DeRose (1995). Unlike some philosophers, I don’t think that this is an especially strong argument compared to others which the skeptic might offer. In particular, I think that (1) is an *extremely* strong claim to take as an unargued-for premise in an argument that is supposed to establish skepticism. If the skeptic simply asserts (1), then I think that the non-skeptic is well within her rights to simply decline to accept it.

Of course, one who puts forth such an argument might attempt to *motivate* acceptance of premise (1) in various ways (‘But *how* do you know that you’re not a BIV, if you do?’ ‘Wouldn’t everything seem exactly the same to you if were a BIV?’) At this point, I think that the skeptic is best understood as tacitly appealing to some general epistemic principle in order to establish (1). (Perhaps: ‘if one can’t reliably discriminate between H and some incompatible H’, then one doesn’t know that H.’) In that case, I think that it’s at least somewhat misleading to represent such a skeptic as taking a particular judgement about the absence of knowledge as a *starting point*. Rather, I think that the relevant claim is better represented as an intermediate lemma of a more complicated argument for skepticism. (Admittedly, these are large issues which deserve greater scrutiny than I am able to offer here.)

In holding that the skeptic *does* require an argument for (1), I join Byrne (2004), Feldman (1999), and Greco (2000: 52, fn.16). For a strong opinion to the contrary, see especially Nozick (1981, p.201).

²¹ For example, I believe that those who espouse radically revisionary views in ethics (i.e., views that depart in radical ways from ‘common sense morality’) are often best understood as methodists in Chisholm’s sense. Here I have in mind especially those who have championed revisionary forms of consequentialism in the tradition that extends from Bentham to Kagan (1989) and Singer (1995).

²² A recent attempt to defend the use of reflective equilibrium in moral and political philosophy against the charge that the method is overly conservative is Scanlon (2002). See especially pages 145-151.

²³ Harman offers the following characterization:

We correct our considered intuitions about particular cases by making them more coherent with our considered general principles and we correct our general principles by making them more coherent with our judgements about particular cases. We make progress by adjusting our views to each other, pursuing the ideal of reaching a set of particular opinions and general views that are in complete accord with each other. The method is conservative in that we start with our present views and try to make the least change that will best promote the coherence of our whole view (“Three Trends”, p.3).

²⁴ Compare Soames (2003b:160). The second example below is borrowed from this paragraph.

²⁵ Apparent dissenters are Weatherson (2003) and Hetherington (2001).

²⁶ A locus classicus of this charge is Edwards (1951).

²⁷ See, e.g, BonJour (1985, 2002).

²⁸ Earlier versions of this paper were read at Princeton University, the University of Wisconsin at Milwaukee, and a Pacific Division meeting of the American Philosophical Association; I am grateful to the audiences present on those occasions. In addition, I would like to thank the following individuals for helpful feedback: Paul Benaceraf, Mark Johnston, Jonathan Vogel, Jim Pryor, Kelly Jolley, Shelly Goldstein, Roger White, and the participants in my Spring 2005 graduate seminar at Princeton.