

4

How to be a Neo-Moorean

Duncan Pritchard

Moore's mistake lies in this—countering the assertion that one cannot know [. . .] by saying 'I do know it'.

(Wittgenstein 1969: § 521)

1. MOOREANISM

My aim here is a somewhat ambitious one: to sketch the contours of a research project. In particular, what I'm interested in is whether we can add some flesh to a certain style of response to the problem of scepticism, one that has historically been the source of much derision but which, as we will see, derives support from a number of theses that one can find being defended in the current literature—so long, that is, as one integrates these theses into a coherent whole.

The response to scepticism that I have in mind mirrors in key respects the 'commonsense' proposal often ascribed to G. E. Moore. As the story goes, Moore would respond to radical scepticism as it is typically conceived of in the contemporary literature by offering the following sort of argument.¹ First, he would claim that he knew a paradigm 'everyday' proposition—i.e., a proposition

Earlier versions of the material in this chapter have been presented on a number of occasions, including the following conferences: 'Contextualism', Bled, Slovenia (June 2004); 'Contextualism', Amsterdam Free University (October 2004); 'Externalism and Internalism in Semantics and Epistemology', University of Kentucky (April 2005); and 'Disjunctivism', University of Glasgow (June 2005). Earlier versions of this material have also been presented at talks at Faculties of Philosophy at the Universities of York, Birmingham, Copenhagen, Aarhus (Denmark), Rijeka (Croatia), and Hull during 2004 and 2005. Special thanks to Jessica Brown, Tony Brueckner, Earl Conee, Gary Ebbs, Sandy Goldberg, Adrian Haddock, Marie McGinn, Alan Millar, Ram Neta, Jim Pryor, Finn Spicer, and two anonymous referees. Finally, I am grateful to the Arts and Humanities Research Council for the award of a research leave grant which enabled me to conduct work in this area.

¹ I say 'as the story goes', since my interest here is not with what Moore actually said, or meant to say, but rather with what he is typically supposed to have said. Indeed, I think that a close examination of the key articles in this regard (Moore 1925, 1939) indicates that there is good reason to think that he might not respond to BIV scepticism in the manner outlined here. For an overview of some of the issues, see Baldwin (1993).

that all of us would think we knew, in those circumstances—such as, famously, that he has two hands. A claim of this sort is surely intuitive. Second, he would note that since having two hands is inconsistent with the relevant sceptical hypothesis, such as the hypothesis that he is a (handless) brain-in-a-vat (BIV), it follows that if he knows that he has two hands then he also knows that he's not a BIV. Again, this claim is also intuitive. Finally, he would conclude that he knows that he's not a BIV. Here, then, is a run-down of the argument:

Mooreanism

- M1. I know that I have two hands.
- M2. If I know that I have two hands, then I know that I'm not a BIV.
- MC. I know that I'm not a BIV.

An anti-sceptical conclusion is thus validly derived from two intuitive premisses. In what follows, we will treat this argument as representative of the Moorean response to scepticism.

Note that the whole point of Moorean anti-scepticism is that there is nothing more to the stance than the presentation of an argument of this sort. The sceptic has called our knowledge into question, via the presentation of the sceptical hypothesis, and the Moorean, via his opposing argument, has rebutted the sceptic's claims. Thus, there is no case to answer, and hence nothing more that needs to be said. In this sense, then, the Moorean stance is a *pre-theoretical* proposal, in that it attempts to deal with the sceptical challenge in an entirely commonsense way which avoids the need for a theoretical response to the problem.

I think that many who engage with scepticism for the first time find a response of this sort compelling; indeed obvious. Interestingly, though, this conviction usually doesn't last very long. Instead, many initial proponents of Moorean-style anti-sceptical arguments quickly abandon their claims once they realize that, on closer inspection, such an argument is not as intuitive as it first appears. With Mooreanism abandoned, the scene is set for various other anti-sceptical theses to come to the fore which are all far more complex and theoretical than the Moorean response—such as arguments for non-closure, contextualism, and contrastivism, to name but three of the big anti-sceptical theories that have come to prominence in recent years.

Nevertheless, I think it is true to say that many epistemologists secretly look back on the abandoned Moorean response to scepticism with a rueful sigh, and still view it as being the, alas unrealizable, apogee of anti-sceptical endeavour. Roughly speaking, the (unspoken) thought here is this: *if only we could have salvaged the Moorean response, then we wouldn't have had to engage with all these myriad anti-sceptical theories, all highly theoretical and all beset by their own fundamental problems.*

It is not my aim here to show how one might salvage Mooreanism as such, since I share the conventional wisdom that this view is unsalvageable. That said, however, I do think that there is a sister view available—what I have termed 'neo-Mooreanism'—which shares many of the key features of Mooreanism but

FN:2 which *can* be made to work.² As we will see, such a view is certainly more complex and theoretical than its Moorean counterpart, but it is still—or so I claim—less revisionary than other competing anti-sceptical views, and thus it is in this sense more Moorean than anti-Moorean.

2. THE FAULTS OF MOOREANISM

Before we can resurrect the Moorean proposal, we first need to be clear just what is wrong with it. While most epistemologists agree that Mooreanism is a non-starter, there is very little in the way of consensus about exactly why the view is problematic. I want to suggest that there are six key problems with the view, many of them interrelated, which thus identifies six problems that a neo-Moorean position has to overcome.

Perhaps the most common complaint levelled at the Moorean argument is that there is something question-begging about responding to the sceptical problem in this way, in that it simply takes as an unquestioned premiss in its argument the denial of the very claim that the sceptic will want to motivate as a conclusion of her argument. Call this the *dialectal impropriety* objection.³

FN:3 We can get a grip on what the problem is here by considering the sceptical argument that the Moorean stance is supposed to be a response to. Keeping to the first-person, as Moore does, and sticking with the examples used above of having two hands and not being a BIV, we can formulate the opposing sceptical argument as follows:

Scepticism

- S1. I don't know that I'm not a BIV.
- S2. If I know that I have two hands, then I know that I'm not a BIV.
- SC. I don't know that I have two hands.

Both of these premisses are intuitive. The first is intuitive because it seems that whether or not we are BIVs is just something that we could never know because sceptical scenarios are defined such that there is nothing in our experiences that could offer us any definitive indication one way or another as to whether we are the victim of such a scenario. The second premiss is exactly the same as the second premiss in the Moorean argument, which we saw was intuitive above. This premiss can be further motivated in terms of the closure principle:

Closure

FN:4 For all S , φ , ψ , if S knows φ , and knows that φ entails ψ , then S also knows ψ .⁴

² I first coined the phrase 'neo-Mooreanism' in Pritchard (2002d).

³ See, for example, Wright (2002) for a complaint against Mooreanism of this sort.

⁴ One might want to modify this principle in a number of ways in order to deal with potential counterexamples of a trivial sort (such as possible cases where the agent doesn't even believe the entailed proposition), but this unembellished version of closure should suffice for our purposes here.

Closure certainly seems plausible, since it is hard to see how this principle could fail. How could one know one proposition, know that it entailed a second proposition, and yet fail to know the second proposition? Crucially, however, with closure in play—and given that one knows the relevant entailment, as presumably one does—it follows that if one knows that one has two hands then one also knows that one is not a BIV. We thus get the second premiss, motivated in terms of the highly intuitive closure principle.

With these two premisses in hand, however, the sceptical conclusion immediately follows. And since this argument can be repeated with any number of everyday propositions (one would just have to vary the sceptical hypothesis to suit), so the full intellectually devastating radical sceptical conclusion—that we are unable to know most of the empirical propositions which we typically think we know—is in the offing.⁵

FN5

With the sceptical argument and the Moorean argument set side by side, one can see that the debate here encapsulates that old philosophical chestnut that one philosopher's *modus ponens* is another philosopher's *modus tollens*. Whereas the Moorean takes his everyday knowledge as secure and argues on this basis that he also has the required anti-sceptical knowledge; the sceptic begins by highlighting the implausibility of anti-sceptical knowledge and argues on this basis that we also lack everyday knowledge. Indeed, notice that the second premiss of the Moorean argument can be regarded as resting on closure just as much as the second premiss of the sceptical argument. It is clearly not the closure principle, then, that is at issue in this debate.

With the debate so construed, however, one can see why the Moorean strategy can seem so dialectically inappropriate. The sceptic has given us an apparently compelling argument for thinking that we lack everyday knowledge. In response, the Moorean simply helps himself to the denial of the contested conclusion and reasons on this basis to the negation of the premiss of the sceptical argument. Given that the Moorean argument begins and ends with this strategy, it is little wonder that few find it persuasive. Imagine, for example, an atheist responding to an apparently compelling proof for the existence of God—let us suppose that such a proof could exist—by citing as a premiss God's non-existence! Clearly, we wouldn't think that such a person had engaged with his opponent, and the same seems true of the Moorean.

A second, and related, difficulty with the Moorean response is that it seems to offer us, at most, a *draw* with the sceptic, rather than a resolution of the sceptical problem. After all, given that the sceptical argument is just the *modus tollens* to Moore's *modus ponens*, and since both arguments have intuitive premisses, it appears that the dialectical situation is that we are faced with two opposing arguments of equal force. If this is right, then even despite the Moorean argument we still have just as much reason to be sceptics as to be Mooreans. Put another

⁵ For more on the contemporary discussion of scepticism, see Pritchard (2002c).

way, it is still the case on the Moorean view that we have no good reason *not* to be sceptics. This is a kind of second-order scepticism which, while not obviously reducible to its first-order cousin (which would hold that we have reason to be sceptics), is still enough to make Moorean anti-scepticism not nearly as intellectually satisfying as it might at first appear. Call this the *impasse* objection.⁶

FN:6

The third problem with Mooreanism is that the Moorean argument seems to consist of a series of assertions which strike one as conversationally inappropriate, if not just plain absurd or contentless. As a number of commentators have noted—most trenchantly Wittgenstein (1969) in his final notebooks—the assertions in question in the Moorean argument seem to offend against our usual usage of the term ‘know’. In particular, we very rarely use the phrase ‘I know that *p*’ in order to convey the fact that we have knowledge of *p*—instead, we just assert ‘*p*’. This phrase thus plays a very special role in our practices of knowledge self-ascription, but, crucially, not one that seems applicable to the kind of anti-sceptical assertions that the Moorean makes. We never normally say we know the denials of sceptical hypotheses, and neither do we usually say we know everyday propositions which are just plain obvious to everyone in that conversational context. Call this the *conversational impropriety* objection.

The fourth problem with the Moorean stance is that it isn’t backed up by a plausible epistemological theory about how it is that we might come to know that we are not BIVs, still less does it engage, for obvious reasons, with the contemporary literature on this topic which features, as we will see, a number of different proposals in this regard. The problem here is that, as we saw when we looked at the opposing sceptical argument, it is very intuitive to suppose that this is the sort of proposition which one could never come to know. Since it is an empirical proposition it isn’t obviously in the market to be known *a priori*, but neither does it seem to be the kind of proposition that one could know empirically either since, *ex hypothesi*, there is no empirical investigation that one could undertake which would indicate to you that this hypothesis did not obtain. Accordingly, to be told on pre-theoretical grounds that one can know such a proposition is unlikely to offer any intellectual satisfaction, since one would also want to be told just how this could be so. Call this the *no supporting epistemology* objection.

This complaint about a lack of a supporting epistemology connects with a fifth worry about the Moorean proposal, which is that it does not support the anti-sceptical line with a *diagnostic* story which explains why we ever found scepticism to be so plausible in the first place. By the lights of this approach, it seems just plain *mysterious* why we were ever taken in by the sceptical problem. After all, if the proper way to deal with scepticism is this straightforward, then why wasn’t it recognized all along? Why did we expend so much effort engaging

⁶ Wright (1991) discusses this ‘impasse’ objection to Mooreanism, and offers an argument which shows that the second-order scepticism which results collapses into first-order scepticism.

with the problem in a myriad of complex ways when we could have dealt with it simply by offering the Moorean argument? Call this the *no diagnosis* objection.

There is also a sixth and final difficulty facing the Moorean response to scepticism that I want to mention, but which requires a little more explaining. In short, this problem is that Mooreanism doesn't really engage at all with a related form of scepticism which targets the *evidential basis* of our knowledge. Call this the *evidential scepticism* objection.

Consider the following argument:

Evidential Scepticism

- ES1. My evidence for E does not favour E over the known to be incompatible BIV hypothesis.
- ES2. If I know E, then my evidence for E favours E over the known to be incompatible BIV hypothesis.
- ESC. I don't know E.

Whereas the sceptical argument we considered above, like the Moorean counterargument, traded on the closure principle for knowledge, this argument rests on the logically weaker 'underdetermination' principle, which we can roughly express as follows:

Underdetermination

For all S, φ, ψ , if S knows φ , and S knows that φ entails ψ , then S 's evidence for believing φ favours φ over $\neg\psi$.⁷

FN:7

In essence, this principle—a version of which can plausibly be found in ancient Pyrrhonian sceptical writings—demands that one's knowledge be evidentially supported, where evidential support here means support which *favours* what is believed over known to be incompatible alternatives (i.e., which provides *more* support for what is believed than it does for the known to be incompatible alternatives). For example, if you know that you are at present in the town's Odeon cinema, and you know that if you are in the Odeon then you are not in the town's other cinema, the Grand, then the evidence which supports your belief that you are at present in the Odeon must prefer this belief over the known to be incompatible alternative that you are at present in the Grand.

So construed, the principle seems entirely uncontentious, since it is hard to see how one's evidence could be genuinely supporting evidence if it did not perform this 'favouring' function. The trouble is, however, that once one feeds sceptical hypotheses into this principle then one immediately generates the sceptical problem. This is because if you know that, for example, you have two hands, and you also know that having hands is inconsistent with being a BIV, then it follows, via underdetermination, that your evidence for believing that you have hands must favour your belief that you have hands over the alternative

⁷ For a thorough discussion of the underdetermination principle, including some of the different ways in which it can be formulated, see Pritchard (2005d).

hypothesis that you are a BIV. *Ex hypothesi*, however, this is impossible, and thus you are unable to know that you have two hands (and much else besides).⁸

[FN:8]

The main point to note about the underdetermination principle is that it is more fundamental to the sceptical problem than the closure principle in at least two ways. To begin with, notice that the underdetermination principle is logically weaker than closure. I have not the space to elucidate this point at length here—see Pritchard (2005d) for the full argument—but one can get a sense of why this is so simply by looking at the kind of sceptical inferences that the two principles license. Taking the knowledge of the relevant entailment as given, the closure principle licenses an inference from knowledge of an everyday proposition, such as that one has two hands, to knowledge of an anti-sceptical proposition, such as that one is not a BIV. In contrast, the underdetermination principle only licenses an inference from knowledge of an everyday proposition, let's say two hands again, to the claim that one's evidence favours one's belief that one has two hands over sceptical alternatives, such as the BIV hypothesis. Since the antecedent is the same in both cases, we just need to focus on the consequent. Intuitively, the claim that one knows that one is not a BIV is much stronger than the claim that one's evidence favours having two hands over being a BIV. Indeed, it seems that all the latter immediately entails is that one does *not* know that one *is* a BIV, and that is certainly weaker than the claim that one *knows* that one is *not* a BIV. *Prima facie*, then, the underdetermination principle is logically weaker than closure, and this means that merely denying the closure principle will not suffice by itself to block the sceptical challenge.

There is also a second—and, I think, more important—sense in which the underdetermination principle is more fundamental to the sceptical problem than the closure principle. This is that the sceptical argument is best thought of as primarily attacking the evidential basis of our knowledge. Think, for example, of how the sceptic motivates the first premiss of her argument, (S1). Clearly, the claim here is that one is unable to know that one is not a BIV because one, perforce, lacks evidence for a belief in this proposition. An evidential claim is thus right at the heart of the standard sceptical argument, even though there is no explicit mention of the agent's evidence in that argument. Making this evidential aspect of the sceptical argument explicit is important because it highlights just how counterintuitive the Moorean conclusion, (MC), is. To say that we can know the denials of sceptical hypotheses is one thing; to say (as seemingly the

⁸ Now one might want to qualify this principle in order to accommodate the intuition that *every* belief needs to be evidentially grounded in order to be an instance of knowledge. Notice, though, that such an amendment to the principle wouldn't necessarily make any real difference to its ability to generate sceptical conclusions. After all, it is surely the case when it comes to *most* of our beliefs in empirical propositions that they are only evidentially grounded to a degree that would support knowledge provided that such evidence is able to play this 'favouring' role. And since the sceptic only needs to call the epistemic status of most of our beliefs in empirical matters into question in order to motivate her sceptical doubt, this weaker construal of the underdetermination principle would, it seems, serve the sceptic's purposes equally well.

Moorean must) that one has adequate evidence for one's belief in the denials of sceptical hypotheses is quite another, and far more problematic.

Evidential scepticism gets right to the heart of this issue by focusing on an evidential thesis and highlighting how this thesis can itself engender the sceptical problem. Moreover, note that attempting a Moorean response to the evidential sceptical challenge just does not seem possible. Suppose, for example, that the Moorean responded by 'inverting' the sceptic's reasoning in the way attempted above as regards closure-based scepticism. The evidential Moorean argument would then go something like as follows:

Evidential Mooreanism

ME1. I know E.

ME2. If I know E, then my evidence for E favours E over the known to be incompatible BIV hypothesis.

MEC. My evidence for E favours E over the known to be incompatible BIV hypothesis.

Clearly, unless this argument is supplemented with further philosophical support then it is extremely suspect. It is one thing to say, on 'commonsense' grounds, that since I know E it must also be the case that I know that the BIV hypothesis is false; quite another to specifically claim that I have evidence which favours E over the BIV hypothesis. The underlying problem here is that to respond to scepticism in this way is not, it seems, to offer a 'pre-theoretical' response to the sceptic at all since it directly issues in a very counterintuitive and highly theoretical claim about the evidential basis of our anti-sceptical knowledge. Given the oddity of the conclusion, the Moorean strategy of refusing to engage further with the sceptical problem once the relevant Moorean argument has been offered seems just plain dogmatic.

We have thus identified six problems with Moorean anti-scepticism, and this means that we have thereby set out what work needs to be done to get a neo-Moorean response on a sound footing. Essentially, what we are looking for is a neo-Moorean anti-sceptical account which shares as many of the central features of the Moorean response to scepticism as possible consistent with it no longer being subject to these complaints. This is what I will be attempting in what follows, with one further constraint. This is that we want the neo-Moorean response to scepticism to have dialectical advantages over its main competitor anti-sceptical theories, and so we need to evaluate the view in this light. As we will see, bringing in the competition in this way is useful in that neo-Mooreanism can exploit aspects of other anti-sceptical views to its own advantage.

3. SCEPTICISM AND CLOSURE

Before we get down to the nitty-gritty of the task before us, I want to make one simplifying assumption explicit. This is that in what follows I will take it

as given that closure holds (at least in some slightly modified form). I take it that most contemporary epistemologists no longer have the inclination to deny this principle outright. Indeed, I suspect that Fred Dretske (2005a; cf. Dretske 1970; Nozick 1981) is perhaps the only prominent epistemologist left who is still defending the outright denial of closure. Part of the reason for this is that epistemologists are now more aware of the alternatives to denying closure here—especially contextualism, which we will consider in a moment. But the main reason is still that rejecting such a principle seems to be itself a form of intellectual self-harm, and so unable to offer us any comfort in our dealings with scepticism. Moreover, recent work on closure has made explicit just how intellectually disastrous the rejection of this principle would be.⁹ In any case, as we have already noted, denying closure is of little help when it comes to the sceptical problem unless one is also in a position to deny the underdetermination principle as well, and that fact in itself diminishes the importance of responses to scepticism that rest on the rejection of closure alone.

FN:9

Accepting closure means accepting that one's response to the sceptic must be, as Jonathan Schaffer (2004: 22) has put it, *immodest*, in the sense that it must allow that we can know the denials of sceptical hypotheses. The primary task for the neo-Moorean, as with any proponent of an anti-sceptical strategy which is immodest in this way, is to account for such knowledge.

One final coda is in order here, which is that in moving directly from the claim that we should retain closure to the claim that we should therefore adopt an immodest response to the sceptic, one might object that I am ignoring the contrastivist response to scepticism, as developed, for example, by Schaffer (e.g., 2005b). On this view, knowledge is always to be understood as a contrastive notion, such that one never knows that *p*, but always knows that *p* rather than a set of contrasts, *Q*. Knowledge is thus to be understood as a three-place relation between a subject, a proposition, and a contrast class, rather than in the usual way as a two-place relation between a subject and a proposition.

It is indeed true that I am ignoring the contrastivist option here, and this is primarily because I am not convinced that contrastivists can retain closure, at least not in any form that we would recognize as closure. After all, contrastivists like Schaffer explicitly grant that there is no contrast relative to which we are able to know the denials of sceptical hypotheses, and thus the *prima facie* tension between closure and contrastivism is manifest. Schaffer (2005a) himself has gone to great lengths to show how one could 'contrastivize' the closure principle in order to reconcile this tension with closure intact, but it is hard to see why the astonishingly complex set of principles that results should be thought to model closure. That is, given the extent of the revision in play, it is hard to see why Schaffer isn't simply denying the principle rather than merely modifying it.

⁹ See, for example, Hawthorne (2005), to which Dretske (2005b) responds. I explicitly discuss anti-sceptical arguments that are predicated on the rejection of closure in Pritchard (2002a, 2002b).

In any case, there is an overriding reason why we can set both of these forms of immodest anti-scepticism to one side, which is that if we can make neo-Mooreanism a palatable anti-sceptical strategy, as I think we can, then there is no longer any obvious need to deny (or, if you prefer, radically modify) closure in order to deal with the sceptical problem, and thus the impetus to endorse anti-sceptical theories of this sort subsides accordingly. Moreover, as we will see, neo-Mooreanism as it will be developed here can accommodate a substantial part of the motivation for contrastivism anyway. With this in mind, we will now focus on those immodest anti-sceptical theories that remain under consideration.¹⁰

FN:10

4. SCEPTICISM, CONTEXTUALISM, AND NEO-MOOREANISM

As I just noted, what all immodest anti-sceptical views have in common is that they claim that we can have knowledge of the denials of sceptical hypotheses. The chief anti-sceptical theory in this regard—and the neo-Moorean’s main competition—is attributer contextualism, as defended, for example, by such figures as Keith DeRose (1995), David Lewis (1996), and Stewart Cohen (e.g., 2000). There are other forms of contextualism of course, such as the variety of subject contextualism—confusingly known as *subject-sensitive invariantism*—recently defended by John Hawthorne (2004) and Jason Stanley (2004, 2005). If, however, we can show that neo-Mooreanism has distinct advantages over attributer contextualism, then this will constitute a substantial part of the task of motivating the view over other contextualist theories as well, especially since, as we will see, the key advantage that neo-Mooreanism has over its main contextualist rival is that it demonstrates that we do not need to suppose that there is any substantial epistemic context-relativity in order to deal with the sceptical problem. Thus we will focus our attentions on attributer contextualism.

Attributer contextualism (henceforth just *contextualism*) holds that ‘knows’ is a context-sensitive term in the sense that the truth of ascriptions of knowledge will depend upon the context in which those ascriptions are made (and thus the context of the *attributer*). In particular, contextualists maintain that different contexts employ different epistemic standards, such that while an agent might meet the epistemic standards in operation in one context of ascription—so that relative to this context an ascription of knowledge to this agent would express a truth—this is consistent with that agent failing to meet the more demanding set of epistemic standards in operation in another context of ascription—so that relative to this context an ascription of knowledge to this agent would express a falsehood. Given the broadly indexical nature of ‘knows’, however, there is no

¹⁰ I consider the contrastivist response to scepticism in its own right in Pritchard (forthcoming a).

conflict between these two claims, since the proposition that is being expressed in the one context is not the same proposition that is being expressed in the other, and thus the one can be true while the other is false without contradiction.

The basic contextualist idea is that while the epistemic standards are high enough in sceptical contexts to ensure that any ascription of knowledge of either everyday propositions or the denials of sceptical hypotheses would be false; nevertheless the epistemic standards are low enough in quotidian contexts to ensure that any ascription of knowledge of either everyday propositions or the denials of sceptical hypotheses would (normally) express a truth. Thus, relative to the standards in operation in quotidian contexts at any rate—which are presumably the contexts that are of most interest to us—we do come out as having the everyday knowledge that we typically take ourselves to have and also the anti-sceptical knowledge which we must have if closure is accepted and scepticism is false.

So construed, contextualism is primarily a linguistic thesis which must be backed up by an appropriate epistemological theory. We do not need to get into the details of the different formulations of the contextualist epistemological thesis in order to see that there is a Moorean aspect to this type of anti-scepticism (indeed, some contextualists have explicitly referred to the view as ‘Moorean’). This aspect relates to the immodesty of the position in allowing that we can know the denials of sceptical hypotheses. Notice, however, that this Moorean commitment on the part of contextualism also points to why a neo-Moorean view ought not to be a contextualist thesis if this is at all possible. After all, the contextualist is committed to *two* revisionary theses here, when it seems that just one of them would suffice. That is, the contextualist is denying not only the intuition that we cannot know the denials of sceptical hypotheses but also the invariantist (i.e., non-contextualist) intuition that ‘knows’ is not a substantively context-sensitive term (i.e., not context-sensitive in the kind of radical way that it would need to be if it is to enable us to resolve key epistemological problems). Interestingly, however, the first claim alone would suffice by itself to block the closure-based sceptical argument, since this is the denial of one of the two key premisses in that argument, (S1). Furthermore, notice that going contextualist without also denying this premiss is of little appeal, since it would mean that one would have all the disadvantages of being a contextualist along with all the disadvantages of rejecting closure.¹¹ Given that one should always minimize one’s revisionism where possible, and bearing in mind also the general nature of the Moorean response which tries to evade the need for complex theorizing, it should be clear that the natural way for the neo-Moorean to go is to aim to simply deny the first premiss of the sceptical argument without also advancing a form of contextualism.

[FN:11]

¹¹ For more on this point, see my discussion of Heller (1999) in Pritchard (2000).

There is also a further reason to prefer a form of neo-Mooreanism run along these lines over a contextualist account, and this concerns the inherently concessive nature of the contextualist thesis. After all, it is part of the contextualist response to scepticism to allow that there are some contexts of ascription in which it would be true to say that we lack everyday knowledge. Given that sceptical contexts are by their nature contexts in which more demanding epistemic standards are being employed than is usual, the problem with this concession is that it prompts the natural thought that, strictly speaking, we do not have the knowledge which we typically ascribe to ourselves. That is, that while, speaking loosely, it is in some sense correct to ascribe knowledge to ourselves in quotidian contexts where the epistemic standards are not very exacting, once one tightens up one's use of the language one realizes that those ascriptions are, strictly speaking, false. Put simply, the worry here is that contextualism seems to validate a train of infallibilist thinking which leads us right back into the sceptical problem. If we can avoid making such a concession, as neo-Mooreanism, as we are now understanding that view, would, then this is all to the good.

A related issue here is the unusual way in which the contextualist is understanding the relevant alternatives thesis. Relevant alternatives epistemology is a form of fallibilism, in that it demands that not all error-possibilities must be eliminated in order for an agent to have knowledge. Instead, just a relevant sub-set of the total class of error-possibilities will need to be eliminated. I take it that the core relevant alternatives intuition is that those error-possibilities which only obtain in far-off possible worlds are automatically irrelevant, in that modal closeness is a prerequisite for relevance. Crucially, however, it is the sceptic's claim that empirical knowledge (most of it at any rate) is impossible, and thus the sceptical argument should go through however we understand the actual world to be. With that in mind, we are entitled to suppose that the actual world is much as we take it to be, and thus that sceptical hypotheses are indeed modally far off, and then ask the question of whether we have any empirical knowledge. The reply given by those who subscribe to the core relevant alternatives intuition is that on this supposition we do have knowledge of much of what we take ourselves to know—because in this case sceptical hypotheses are irrelevant to our knowledge—and thus that the sceptic is wrong to say that it is impossible to possess the empirical knowledge that we typically take ourselves to possess.

There is a strong rationale for the core relevant alternatives intuition, and it arises out of the platitude that knowledge is, at root, non-lucky true belief. For so long as sceptical error-possibilities are indeed modally far off, then the mere fact that one would not be able to tell a sceptical and a counterpart non-sceptical scenario apart should not by itself suffice to indicate that one's true beliefs in everyday propositions—which entail the denials of sceptical hypotheses—are thereby only luckily true. Put another way, if sceptical error-possibilities are indeed far-fetched, then it is not the case that one's everyday beliefs—and

one's anti-sceptical beliefs for that matter—could have been in this regard very easily false.

It is an anti-luck intuition of this sort that, I take it, lies behind safety-based theories of knowledge—as defended, for example, by Ernest Sosa (1999)—where safety means something like ‘could not have easily been false’. There are various ways of formulating this principle, not all of them plausible, but the basic formulation has it that for a true belief to be safe it must be the case that, across a wide range of nearby possible worlds, where the agent believes the target proposition (on the same basis), that belief continues to be true. Note that the anti-luck intuition encapsulated in safety is also what lies behind the contextualist's claim that we can know everyday propositions and the denials of sceptical hypotheses relative to quotidian contexts of ascription (indeed, DeRose 1995 explicitly appeals to a safety-type principle in this regard). What is important for our purposes, however, is that this intuition can be read as motivating the core relevant alternatives thesis that far-off error-possibilities are by default irrelevant to knowledge possession. The core relevant alternatives intuition and the anti-luck intuition thus go hand in hand.¹²

FN:12

Contextualists are fallibilists, and thus relevant alternatives theorists, in the sense that they allow that an ascription of knowledge to an agent can be true even though that agent is unable to rule out all possibilities of error.¹³ This will certainly be the case in quotidian contexts of ascription, where the epistemic standards are low. Notice, however, that contextualists do not subscribe to the core relevant alternatives intuition just noted, since they allow that even scenarios which are, *ex hypothesi*, modally far off can still be relevant to knowledge ascription if the epistemic standards in the context of ascription demand that one take them into consideration. In sceptical contexts, for example, where sceptical hypotheses are explicitly at issue, whether or not an agent is truly ascribed knowledge will be dependent upon whether that agent is in a position to rule out the sceptical hypothesis, even if, as a matter of fact, that hypothesis is indeed modally far off.

FN:13

This points to a further juncture at which the neo-Moorean and the contextualist should part company, in that the neo-Moorean should, it seems, cling on to that core relevant alternatives intuition—and the associated anti-luck intuition—and insist that just so long as the actual world is roughly as we take it to be then we are able to have knowledge even despite our inability to eliminate modally far-off sceptical hypotheses. Of course, the neo-Moorean will have to explain how, given closure, one is able to know the denials of sceptical hypotheses given that this is the case, but that is a problem shared by the

¹² I discuss this point at length in Pritchard (2005b, *passim*). See also Pritchard (2004, 2005c, 2006).

¹³ Notice that this formulation of fallibilism is consistent with Lewis's (1996) famous claim to be an infallibilist. This is because it is an *explicit* fallibilism that Lewis rejects, such that an ineliminable error-possibility is taken into consideration and yet the agent is ascribed knowledge nonetheless.

contextualist anyway, and so not relevant to the specific issue of theory choice when it is just these two theories under consideration.

There are thus strong *prima facie* grounds for preferring a neo-Moorean anti-sceptical strategy over a contextualist one. There is, however, one key feature of the contextualist position which speaks in its favour in this respect, and that is that contextualism seems able to meet some of the problems we saw facing the Moorean response above. To begin with, the contextualist can account for why the assertions made by the Moorean seem so improper, the simple explanation being that in the sceptical context in which these assertions are, perforce, made, what is asserted is false (because very demanding sceptical epistemic standards will apply). Furthermore, the contextualist can diagnose the attraction of scepticism, since on this view the problem emerges out of a failure to realize that ‘knows’ is a context-sensitive term. Finally, the contextualist has the supporting epistemology to back all of this up. It is thus essential to neo-Mooreanism that it is able to supply a similar response to the problems that beset the Moorean strategy, since without it the view loses much of its appeal relative to its competitor contextualist theory.

5. EVIDENTIAL CONTEXTUALISM AND EVIDENTIAL NEO-MOOREANISM

This discussion of contextualism—and a potential neo-Moorean rival theory—is all by-the-by, however, since, as we noted above, it is essential that any anti-sceptical theory be able to deal with the threat posed by evidential scepticism, and it is far from obvious how the views just set out (which make no mention of evidence) would do that. There is, however, a contextualist theory in the literature—due to Ram Neta (2002, 2003)—which also tries to deal with the evidential sceptical problem, so it is worthwhile turning to this view for inspiration in this regard.

Neta’s proposal is that we should treat what counts as evidence as being context-sensitive, and thus offer a contextualist account of knowledge in virtue of the context-sensitivity of evidence. In particular, Neta’s claim is that in different contexts of ascription what counts as an agent’s evidence can change, so that the evidential support that the agent has for her belief relative to one context of ascription might be strong (strong enough to make an ascription of knowledge in that context true, say), even though that same agent has very weak evidential support for her belief relative to another context of ascription which employs different epistemic standards (and so may not have sufficient evidential support to make an ascription of knowledge in that context true). Thus, the truth of ascriptions of knowledge is still a variable matter as it is on the standard contextualist account, it is just that the variability is explicitly accounted for in terms of the shifting standards of what counts as evidential support.

In essence, Neta's idea is that an agent's evidence is determined by the error-possibilities at issue in the context of ascription in the sense that the agent's evidence is the evidence the agent would have were the relevant error-possibilities to be true. In a context in which sceptical hypotheses are at issue, therefore, an agent's evidence can only be the evidence she would have were she to actually be, say, a BIV, and that will of course severely limit the scope of the agent's evidence. As far as perceptual evidence goes, for example, her evidence for her perception-based belief that there is, say, a table before her can only consist of the way the world seems to her (that it looks to her as if there is a table in front of her). With this in mind, the underdetermination-based sceptical argument gets a grip, since relative to this conception of what counts as the agent's evidence it will not favour her belief in the everyday proposition (in this case that there is a table before her) over the known sceptical alternative, and thus she will be unable to have knowledge of this proposition.

In quotidian contexts, in contrast, the error-possibilities at issue will not include sceptical hypotheses, and thus what counts as an agent's evidence relative to this context will inevitably be much broader. In particular, Neta claims that in quotidian contexts in which sceptical hypotheses are not at issue an agent's perceptual evidence can include factive evidence, such as that the agent sees that such-and-such is the case, where *seeing that p* entails *p*. If this is right, then relative to quotidian contexts an agent's evidence can be such that it favours the agent's belief in an everyday proposition, such as that there is a table before her, over a sceptical alternative, since relative to these quotidian contexts the evidence for the former could logically exclude the latter alternative.

Notice that this view is most naturally construed along content externalist lines as being allied to a form of disjunctivism. This is because the factive nature of one's perceptual evidence in quotidian contexts intuitively reflects the factive nature of one's perceptual experiences in those contexts, such that the content of one's perceptual experience can be such that it is determined, at least in part, by facts in the agent's environment (for instance, in the case just described, the fact regarding *p*), facts that would not obtain in corresponding cases of illusion or delusion where as a result the content of one's experience, and thus the nature of one's evidence on this view, would be different.

This is an intriguing suggestion, but note that, when evaluated relative to neo-Mooreanism, it faces similar problems to those we noted above as regards the standard non-evidential version of contextualism. In particular, the natural question to ask at this juncture is that if we can indeed make sense of the idea not just of an agent knowing that she is not the victim of a sceptical hypothesis, but also that this knowledge is evidentially supported within the constraint imposed by the underdetermination principle, then why should we be attracted at all to an evidential contextualism? Why not instead simply deny, on content externalist grounds, the key premiss in the underdetermination-based sceptical argument that we are unable to have evidence which favours our beliefs

in everyday propositions over sceptical alternatives and leave the matter at that? Indeed, once the role of the disjunctivist thesis in the arguments for evidential contextualism is made explicit, then the idea that one lacks the factive perceptual evidence in sceptical contexts that one possessed in quotidian contexts starts to look rather suspect. After all, how can the content of one's perceptual experience vary in line with purely conversational matters? And if it doesn't vary, then why should we concede that what counts as one's perceptual evidence changes with the shift of context?

Moreover, adopting a neo-Moorean form of evidentialism ensures that nothing of substance is conceded to the sceptic, since, as with standard forms of contextualism, evidential contextualism leaves one with the nagging thought that, strictly speaking, one lacks evidential support for one's beliefs in everyday propositions, it is just that, loosely speaking in quotidian contexts, it is acceptable to regard such beliefs as appropriately evidentially supported. An evidential neo-Moorean view that eschewed contextualism would not face this problem.

A further motivation for the neo-Moorean stance in this regard comes from the core relevant alternatives intuition that we noted earlier. So long as sceptical scenarios are indeed far-fetched, then why should they be relevant to the determination of knowledge in *any* context? More specifically for our present purposes, why should they be relevant to the determination of the scope of an agent's evidence in any context? And since we can connect this thought to the more general, and widely held, intuition that knowledge is at root non-lucky true belief, the support this consideration provides for an evidential neo-Mooreanism is quite strong.

Of course, the major advantage that evidential contextualism has over evidential neo-Mooreanism is the same one that non-evidential contextualism has over non-evidential neo-Mooreanism, which is that it has a story to tell about why we get taken in by sceptical arguments, and why Moorean assertions can seem so plain odd. Again, then, we need to remember that neo-Mooreanism, evidentially construed or otherwise, is in serious trouble unless it can incorporate the necessary diagnostic story.

6. VARIETIES OF EVIDENTIAL NEO-MOOREANISM AND THE CLASSICAL EPISTEMIC EXTERNALISM/INTERNALISM DISTINCTION

So far we have characterized neo-Mooreanism mostly in negative terms; in terms of which theses the neo-Moorean rejects. In particular, we have noted that a neo-Moorean will argue that we can know the denials of sceptical hypotheses, but will not contextualize this thesis by claiming that such knowledge is only possessed relative to certain contexts of ascription. Moreover, we have seen that

such a view would need to be given an evidential spin such that one's evidence can favour one's beliefs in everyday propositions over sceptical alternatives, thereby ensuring that one's beliefs in the known to be entailed anti-sceptical propositions, such as that one is not a BIV, can also be appropriately evidentially grounded. Again, it is important to the view that this evidential claim is not contextualized. Finally—and this is the only positive aspect of the view so far—we have started to see one way of motivating the neo-Moorean position in terms of a certain construal of the core relative alternatives intuition which we have noted is closely associated with the anti-luck requirement common to most theories of knowledge, and which can be encapsulated in some formulation of the safety principle.

These theses alone clearly do not yet represent an anti-sceptical theory, nor do they suffice to motivate neo-Mooreanism over contextualism unless they are supplemented with further claims that can match the diagnostic appeal of contextualism. Before we can make any headway at developing this view, however, it is important to first factor in the classical externalism/internalism distinction in epistemology since, as we will see, this has important ramifications for how such a view will be developed.

By classical epistemic internalism here, I mean *access* internalism such that what makes an epistemic condition (i.e., a condition which, perhaps in conjunction with other epistemic conditions, can turn true belief into knowledge) an internal epistemic condition is that the agent concerned is able to know by reflection alone those facts which determine that this condition has been met. Meeting the justification condition, for example, at least as it is standardly conceived, involves the possession of grounds in support of the target belief, where these grounds—and the fact that they are supporting grounds—are reflectively accessible to the subject. Classical epistemic externalism denies this, and so holds that there are epistemic conditions which do not demand reflective access on the part of the subject of this sort. I will understand classical internalism about knowledge as being the view that meeting a substantive internal epistemic condition is *necessary* for knowledge possession, with externalism about knowledge as the denial of this thesis.

There are other ways of drawing the internalism/externalism distinction of course. One could put the point in terms of supervenience rather than access—see, for example, Earl Conee and Richard Feldman (2000)—or one might weaken the internalist requirement by saying that one only needs reflective access to the supporting grounds for one's belief and not also to the fact that they are supporting grounds, as William Alston (1988) suggests (though note that he doesn't regard this view as an internalist thesis as such). The account of the distinction offered here is fairly standard, however, and, I believe, it also gets to the heart of what is at issue in the standard debates about epistemic externalism and internalism. Think, for example, about the normal cases over which internalists and externalists diverge, such as the chicken sexer case. Here we have an agent

who is exhibiting a highly reliable cognitive ability (to distinguish between the sexes of chicks), and yet who has false beliefs about how she is doing what she is doing (she thinks she is using her senses of sight and touch, when actually it is her sense of smell) and who typically lacks good grounds for thinking that she is reliable in this respect. Clearly such an agent does not meet a substantive internal epistemic condition as we have just defined that notion. Nevertheless, that, for the externalist at least, there is still an issue about whether or not the agent has knowledge indicates that what is in question here is the *necessity* to knowledge of meeting such an internal epistemic condition, which is just as it should be given how I have just characterized the classical externalism/internalism contrast. We will consider a non-classical form of internalism below, but for now we will focus on the distinction as it has just been drawn.¹⁴

FN14

With this distinction in mind it ought to be clear that motivating a classical internalist version of neo-Mooreanism is not going to be at all easy. There are two reasons for this. The first is that the primary intuition behind the idea that we are unable to know the denials of sceptical hypotheses is an internalist intuition to the effect that since there is nothing in our present experiences that we can, as it were, introspectively point to in order to indicate that we are not the victim of a radical sceptical hypothesis, hence we are unable to have adequate reflectively accessible grounds to support our belief in this respect. Accordingly, on an internalist account, lack of knowledge of the denials of sceptical hypotheses seems to follow fairly quickly.

The second reason why a classical internalist version of neo-Mooreanism would be tricky to sustain concerns the evidential basis of our putative anti-sceptical knowledge on this view. By the lights of an internalist epistemology one would expect one's evidence to be individuated on internalist grounds in terms of facts which are accessible to the agent. If this is right, however, then one's evidence even in cases where one is not being radically deceived can be no better than one's evidence in counterpart cases in which one is being radically deceived. After all, if the deception is such that there is nothing that is reflectively accessible to the agent which could indicate which of these situations is the one that obtains, it follows that one's evidence reduces to the lowest common denominator—or,

¹⁴ For further discussion of the chicken-sexer example, see Foley (1987: 168–9), Lewis (1996), Zagzebski (1996: § 2.1 and § 4.1), and Brandom (1998). For more on the classical epistemic externalism/internalism distinction in general, see Kornblith (2001). Note that the chief problem with supervenience internalism is that it seems it can only accommodate the traditional epistemic internalist intuitions provided one does not combine the view with content externalism. After all, most internalists take as a datum to be explained by the theory that if an agent has a justified belief then so too does his envatted counterpart. On most content externalist views, however, there will tend to be a difference in the content of the mental states of the two agents, and thus this claim will not obviously go through. It may be, of course, that all formulations of the epistemic externalism/internalism distinction are hostage to what conclusions are derived from the debate about content externalism/internalism, but it is odd that such a basic incompatibility between epistemic internalist intuitions and content externalism should be so immediately apparent on this way of drawing the epistemic externalism/internalism distinction.

if you prefer, the highest common factor—between the two cases. As a result, an evidential version of neo-Mooreanism would be unable to account for the evidential basis for our anti-sceptical knowledge and would also be immediately susceptible to the underdetermination-based sceptical argument.

It is unsurprising, then, that there are very few proto-neo-Moorean stances in the literature that are conceived along classically internalist lines. Indeed, perhaps the only such account is due to Crispin Wright (e.g., 2004). It is part of Wright's view to grant that no warrant can, as he puts it, be 'earned' for beliefs in anti-sceptical propositions. Nevertheless, he claims that it does not follow that such warrant is not possessed. In essence, Wright's idea is that what the sceptic highlights is that it is essential to our cognitive projects that one has a standing entitlement to certain anti-sceptical propositions, such that the warrant in question is 'unearned'. That is, the idea is that given that it is essential to enquiry that one believe certain anti-sceptical propositions—and since without engaging in enquiry one is unable even in principle to earn a warrant for one's beliefs, hence the moral of scepticism is that not all warrants are earned.

This view faces the immediate problem of accounting for how, by classical internalist lights, an unearned warrant is any warrant at all, but we do not need to get into this issue here.¹⁵ This is because even if the classical internalist neo-Moorean can make sense of the idea of an unearned warrant—perhaps in terms of the requirements of an epistemic rationality—it is still not clear how such a notion would enable the proponent of a view of this sort to deal with the sceptical challenge posed by the underdetermination-based argument. How could it be that by internalist lights one has evidence which favours one's belief in everyday propositions over sceptical alternatives? Indeed, Wright's claim that the epistemic standing of our anti-sceptical beliefs is 'unearned' is surely meant to imply that they have no evidential support at all.¹⁶

¹⁵ I expand upon this objection at length in Pritchard (2005e).

¹⁶ Another position in the literature which might plausibly qualify as a classical internalist version of neo-Mooreanism is that set out by Pryor (2000). Roughly speaking, Pryor's idea is that we have a default warrant—and in this sense 'unearned' and, thus non-evidential, warrant—for our basic perceptual beliefs on account of how forming beliefs in this way is not epistemically blameworthy. The thought is then that given this default standing epistemic status of our basic perceptual beliefs, we are entitled to make inferences to (known to be entailed) anti-sceptical beliefs, where this default epistemic standing transmits to the inferred conclusion. Our beliefs in the denials of sceptical hypotheses thus have a default epistemic standing in virtue of the default epistemic standing of our basic perceptual beliefs and the associated (correct) inference. Moreover, Pryor wishes to make this anti-sceptical move from within a classical internalist epistemology, so that the warrant in question is to be understood as satisfying classical internalist criteria. There is certainly a tradition of understanding classical internalism deontologically in terms of obedience to epistemic norms (where it is essential, of course, that one is in a position to determine by reflection alone what these norms are, and what one needs to do in order to follow them). As Pryor (2001) himself elsewhere concedes, however, it is not plausible to suppose that the minimal epistemic standing that accrues on this picture simply in virtue of not having contravened any epistemic norm will itself suffice, with true belief, for knowledge, even if we further suppose that the belief in question also meets a further anti-Gettier external condition. Thus, even if we can assuage our worries about how a

In any case, what is perhaps more important for our purposes is that the classical externalist is clearly in a stronger position than the classical internalist when it comes to motivating a neo-Moorean position. After all, that one lacks adequate reflectively accessible grounds for believing that one is not a BIV will not on this view decide the issue of whether or not one can have knowledge of this proposition. Similarly, such a lack of reflectively accessible grounds need not mean that one's evidence is restricted to merely the evidence that one would have were one to be in the counterpart deceived case. This is especially the case if one allies the view to a form of disjunctivism which allows that the content of one's perceptual experience can be different in non-deceived and counterpart deceived cases, since there is surely a close connection between the content of one's perceptual experience and one's perceptual evidence. Thus, the path is cleared to allowing one's anti-sceptical knowledge to be evidentially grounded and thus to blocking the underdetermination-based sceptical argument.

7. KNOWING AND SAYING THAT ONE KNOWS

This aspect of the envisaged neo-Moorean view—namely, the combination of content externalism and epistemic externalism to meet underdetermination-style sceptical arguments—is, of course, found in recent work by Timothy Williamson (2000a, 2000b), though it is not advanced under this description. Significantly for our purposes, however, this view is not set within the kind of diagnostic account that we saw above was essential to any plausible rendering of the neo-Moorean proposal. It is thus vital that we do not end the anti-sceptical story at this point, but continue to consider the further issue of diagnosis. What is key here, I believe, is understanding just why the Moorean assertions, while (on this view) true, are nevertheless conversationally inappropriate.

To begin with, note that the classical externalist is again on stronger ground than the classical internalist in this regard. This is because of the connections between the conditions for knowledge possession and the conditions for appropriate assertions of explicit knowledge claims (i.e., explicit self-ascriptions of knowledge of the form 'I know that p '). On the classical internalist view, there will be a close connection between these two conditions. This is because one ordinarily needs good reflectively accessible grounds in order to properly make an assertion—especially assertions which involve explicit self-ascriptions of knowledge—and the possession of grounds of this sort will also be a prerequisite for

default warrant for propositions of this sort is to be maintained on the classical internalist account, and similarly overcome the concern we raised for Wright above regarding how such a default conception of warrant could respond to the specific challenge posed by evidential scepticism, it would still remain that this type of anti-scepticism is not directly relevant for our purposes, where it is scepticism about knowledge that is our focus.

knowledge on the classical internalist account. On the classical externalist account, in contrast, the connection between these two conditions will be relatively weak in that knowledge might be possessed even in the absence of good reflectively accessible grounds, and thus be possessed in cases where the agent would not be in a position to properly assert the proposition that she knows. Think, for example, of the chicken sexer case described above. Here we have an agent who has knowledge by classical externalist lights, but who is clearly not in a position to properly assert what she knows—much less assert that she knows it—since she lacks any good reflectively accessible grounds to back up that assertion.

With this in mind, it is going to be a lot easier for the classical externalist to account for the apparent impropriety of Moorean assertions in such a way as to retain the thought that what Moore is asserting is nevertheless true. It could be, for example, that one's anti-sceptical knowledge that, say, one is not a BIV is like the 'brute' knowledge possessed in the chicken sexer case, and if this is so then this would explain why it cannot be properly asserted. What makes this claim plausible is the fact that, as just noted, what does seem to be a clear consequence of the sceptical reasoning is that we are unable to possess good reflectively accessible grounds for believing that we are not the victims of sceptical hypotheses. On this view, however, the lack of such grounds will only affect the propriety of one's assertions in this respect, and need not undermine one's anti-sceptical knowledge.

Further reflection on the peculiar role of 'I know' in our linguistic practices also supports this sort of contention. Typically, one conveys one's knowledge of a proposition simply by asserting the proposition in question. Adding the further phrase 'I know' is rare, and standardly reflects not just emphasis but also an ability to resolve a particular challenge that has been raised. For example, one might initially convey one's knowledge of what the time is by simply asserting, say, 'It's 10.22 a.m.', but then be prompted into the further explicit claim to know this proposition by a challenge to one's original assertion. There are two main ways in which these challenges could be issued. The first concerns the presentation of an error-possibility which is held to be salient. Call this an *epistemic challenge*. The second concerns those occasions where it is pointed out that a lot hangs on the correctness of the assertion in question. Call this a *standards challenge*.¹⁷

FN:17

In responding to either of these types of challenge with an explicit knowledge claim one is representing oneself as being in possession of stronger reflectively accessible grounds in support of one's assertion than would be implied simply by making the assertion itself. Notice, however, that the kinds of additional grounds required in each case can be very different. When it comes to standards challenges, for example, all that is normally required is stronger grounds *simpliciter*. In contrast, when it comes to epistemic challenges, the additional grounds have

¹⁷ One might have challenges which are a mix of the two, of course, but we will bracket this possibility here in order to keep matters as simple as possible.

to speak specifically to the error-possibilities raised. In particular—and this is a point which, I think, has often been overlooked in this regard, despite its epistemic importance—the grounds one needs available to one in making a claim to know in response to this second sort of challenge must be such as to *discriminate* between what is asserted and the relevant error-possibility. This claim is important because, intuitively, the kind of evidential support one needs in order to have knowledge is weaker than this. All that is required here is the sort of ‘favouring’ evidence that we saw above in our discussion of the underdetermination principle (which is still a strong requirement on knowledge, as we also saw).¹⁸

FN:18

In order to make this point clear, consider the famous ‘zebra’ case offered by Dretske (1970). Here we are asked to imagine someone who is at the zoo in normal circumstances and sees what looks like a zebra in the zebra enclosure. Clearly, such an agent would normally be attributed knowledge that the creature before him is a zebra, and we would be perfectly happy with any assertion he might make to the effect that there is a zebra before him—which would represent him as having knowledge of this proposition—since he has adequate reflectively accessible grounds to back up that assertion. Similarly, an explicit claim to know that he sees a zebra would also be deemed appropriate in this context, if the circumstances were right. For example, if the original claim that the creature before him is a zebra is challenged in some mundane fashion—perhaps by someone short-sighted who wonders out loud why, since they were expecting to be near the gorilla enclosure, there should be zebras here—then it would be unproblematic for our agent to respond to this challenge by saying that he knows that this creature is a zebra.

It is important to recognize why such an assertion would be entirely appropriate, given how we have described the situation. The reason for this is not just that the agent is in a position to offer very good reflectively accessible grounds in favour of what he asserts, nor even that he has good reflectively accessible grounds in favour of what he asserts which prefer what he asserts over the target error-possibility (that it is a gorilla rather than a zebra), but more specifically that he has good reflectively accessible grounds which discriminate between the target proposition and the target error-possibility—i.e., between creatures that are zebras and (non-zebra) creatures that are gorillas. That is, explicitly claiming knowledge in this context will generate the conversational implicature that one is able to offer grounds in support of the proposition claimed as known which would suffice to distinguish the scenario described by this proposition from the specified error-possibility. In this case, however, the agent does have the required grounds. After all, our agent is aware, presumably, that zebras and gorillas have very different shapes and gaits, and this will suffice to enable such discrimination

¹⁸ I think a failure to make this distinction is what makes contrastivism such an initially plausible conception of knowledge. For more on this point, see Pritchard (forthcoming a).

to take place. That such a discriminative ability is required in order to claim knowledge in this case should not, however, lead the neo-Moorean into thinking that it is thereby required for knowledge possession, since, intuitively, mere favouring evidence will suffice in this regard.

This last point is important because there are cases where one has the required favouring evidence but where one lacks the relevant discriminative capacity. Accordingly, if one fails to pay due attention to this point then one will be led into denying knowledge to the agent even though on the neo-Moorean view it is possessed. Imagine, for example, that the error-possibility that the zebra may in fact be a cleverly disguised mule is raised and taken seriously in that conversational context. Since the original assertion that the creature is a zebra has been challenged, it would ordinarily be appropriate for the agent, if he knows this proposition, to explicitly say so, just as he did in the case just described where the objector wonders why he isn't at present looking at a gorilla. If the agent now claims to know that there is a zebra before him in the light of this error-possibility being raised, however, then this will generate the implicature not just that the agent has reflectively accessible grounds which prefer the proposition claimed as known over the target error-possibility, but also that the agent has reflectively accessible grounds which could serve to discriminate between the proposition claimed as known and the target error possibility. That is, the agent is representing himself as having grounds which would suffice to enable him to tell the difference between these two creatures (i.e., a zebra and a cleverly disguised mule). Such grounds might be, for example, that he has examined the creature at close range and been able to determine that it is not painted. Typically, of course, the agent will not have grounds of this sort available to him, and so his claim to know will be inappropriate because it generates a false conversational implicature.

Notice, however, that this fact alone does not suffice to indicate that the agent lacks knowledge of the target proposition. After all, in the standard case at least, the agent will have evidence which favours the hypothesis that the animal before him is a zebra over the alternative hypothesis that it is a cleverly disguised mule. Think, for example, of the grounds he has regarding the implausibility of a zookeeper going to such lengths to deceive patrons, and the penalties that would be imposed were such a deception to come to light, as presumably it would eventually. Moreover, this evidence will typically be reflectively accessible to the subject, and so this point stands alone from any general considerations regarding the relevance of the classical internalism/externalism distinction in epistemology here. Thus, by both externalist and internalist lights, the agent has evidentially grounded knowledge that there is a zebra before him. The issue is solely whether or not he can properly claim to possess that knowledge in these circumstances; not whether it is possessed.

What is different about these cases and the sceptical case is that we can make sense of an agent having evidence—even reflectively accessible evidence—which

favours the believed hypothesis over the alternative hypothesis. Moreover, we can also make sense of there being agents who are in better epistemic positions relative to the believed proposition—such as, in the zebra example, zoologists—who are able to possess the reflectively accessible discriminating grounds required for an appropriate knowledge claim. In sceptical cases, in contrast, matters are very different. This is because we have difficulty comprehending evidence that can play the required favouring role (it is essential that the evidence not be understood along classical internalist lines if it is to play this role); and we can make no sense at all of the idea that one has adequate reflectively accessible grounds which could serve to discriminate the target hypothesis from the sceptical alternative. On the neo-Moorean account sketched here, then, we can only account for knowledge possession in these cases relative to a classical externalist epistemology, and we can make no sense at all of appropriate knowledge claims in sceptical contexts.

Consider again the Moorean anti-sceptical assertions. If we grant that knowledge possession requires evidence which favours one's everyday beliefs over their sceptical alternatives, but also that such evidence can be possessed by externalist lights, then we are in a position to maintain that the Moorean assertions could well be true. Nevertheless, they cannot be properly made, since they are entered in a context in which sceptical alternatives are explicitly at issue, and thus they will generate false conversational implicatures. In particular, the claim that, for example, one knows one has two hands will generate the false conversational implicature that one has reflectively accessible grounds which would suffice to distinguish between the scenario in which one has hands and the alternative sceptical scenario in which one is, say, a BIV who merely seems to have hands.¹⁹ Similarly, the claim to know that one is not a BIV will generate the false conversational implicature that one has reflectively accessible grounds which would suffice to indicate that one can distinguish between the scenario in which one is not a BIV from the alternative scenario in which one is a BIV. It is little wonder then, on this view, why such knowledge can never be properly claimed even in cases when it is possessed. Once one factors in the further consideration of the dialectical impropriety of these assertions, it becomes manifest why a neo-Moorean stance must not try to deal with the sceptical problem head-on by making anti-sceptical assertions in this way.

It is interesting to note that Williamson's (1996) own account of assertion—in terms of the overarching rule that one should only assert what one knows—does not seem to offer any diagnosis of what is wrong with the Moorean assertions. Presumably, this rule demands in this case that one should know that one knows the relevant propositions (i.e., $E, \neg \text{BIV}$) if one is to properly assert them.

¹⁹ This consideration also explains, at least in part, why such claims to know what is (taken to be) plainly obvious *always* seem problematic. This is because it is hard to imagine a non-sceptical context in which such an assertion would be entered.

Accordingly, one might think, on broadly epistemically externalist grounds, that the problem with the Moorean assertions is that while one has the knowledge in question (such that what is asserted is true), one lacks the corresponding second-order knowledge and it is this fact that ensures that the rule of assertion has been broken and thus that the assertion is inappropriate.

It is far from clear, however, why the neo-Moorean view as it is described here should ally itself with a move of this sort. After all, if one does know the relevant everyday and anti-sceptical propositions, then intuitively one's beliefs that one knows these propositions could well be just as safe as one's beliefs in the propositions themselves. Thus, there seems no reason, on this view, for denying the second-order knowledge in this case.²⁰ The trouble is, of course, that the assertion would, intuitively, still be inappropriate even if the second-order knowledge were possessed. If this is right, then on the face of it the Williamsonian view contains an important lacuna, since it cannot offer the required diagnostic story that distinguishes the neo-Moorean view from its problematic Moorean ancestor.

FN:20

In any case, with the neo-Moorean picture of the conversational propriety of explicit self-ascriptions of knowledge just set out in mind, it is hardly surprising that there will be a context-sensitivity in the propriety conditions for making such assertions. One type of context-sensitivity is that just considered regarding claims to know in response to epistemic challenges. In these cases, the evidential demands on appropriate assertion shift in response to features of the conversational context, such that the very same assertion can be appropriate in one context and yet inappropriate in another context even though all that has changed has been the introduction of an error-possibility to that conversational context. A similar sort of context-sensitivity is also in play when it comes to standards challenges, though, as noted above, the type of evidential demand made by a standards challenge is usually subtly different in that responding to such a challenge merely demands of the asserter a greater degree of reflectively accessible evidence, rather than reflectively accessible evidence which can specifically serve to discriminate between the asserted proposition and the target error-possibility. Repeating a claim to know in response to someone simply pointing out how important it is that one is right will typically involve merely representing oneself as having very strong reflectively accessible grounds in favour of the proposition claimed as known; it will not usually involve representing oneself as having specific discriminating grounds because there is no error-possibility at issue for which such grounds would be relevant in this case.

²⁰ This is especially so on the Williamsonian picture because of its commitment to the thesis that one's knowledge is identical to one's evidence (see Williamson 2000a). Accordingly, simply in virtue of having the first-order knowledge one thereby has excellent evidence to support any second-order belief that one might hold to the effect that one has this knowledge, and thus one is by default well on the way to having second-order knowledge as a result.

Contextualists often talk as if their focus is simply standards challenges, but it should be clear that it is epistemic challenges that are really what is at issue here. In any case, what is important is that this pragmatic account of the shifting propriety conditions for explicit knowledge claims removes much of the impetus for contextualism, in that one can accommodate the apparent context-sensitivity of our ‘knows’ talk without thereby treating ‘knows’ as a context-sensitive term. More needs to be done to complete the pragmatic story of course, since one needs to extend the view so that it deals with explicit knowledge ascriptions more generally, rather than just explicit self-ascriptions of knowledge, but the beginnings of such a view are clear to see.²¹

FN:21

Furthermore, such a pragmatic story also goes a long way towards offering the kind of diagnostic support for neo-Mooreanism that we noted above was so lacking, especially in light of the strong diagnostic support that contextualists can offer in favour of their view. In particular, one can now account for the impropriety of the Moorean assertions, and thereby explain the intuitive pull behind sceptical arguments, without conceding anything substantial to the sceptic. With this diagnostic story in place, the remaining theoretical advantage held by the contextualist over the neo-Moorean is finally undermined.

8. Mc DOWELLIAN NEO-MOOREANISM

Interestingly, there is a very different neo-Moorean position that has been proposed in the literature which deserves our consideration, one that does not comfortably fit into either of the classical internalist or classical externalist camps. This position is due to John McDowell (e.g., 1995), and what is distinctive about it (amongst other things) is that it incorporates two theses which, collectively, set the view apart from classical versions of epistemic externalism and internalism. These theses are: (i) a claim in the spirit of epistemic internalism which demands of a knower that she be in a position to know by reflection alone what the reasons which support her knowledge are; and (ii) a content externalist claim of the disjunctivist sort noted above which allows that one’s reasons can be both empirical and factive—i.e., can be reasons for believing an empirical proposition and entail what it is that they are a reason for. This last claim entails content externalism since on McDowell’s view one’s experiences can function as one’s reasons, and yet the content of one’s experiences will clearly be sometimes—i.e., in those cases where the reasons are factive—determined by facts concerning one’s environment (such that in corresponding cases of illusion or delusion where the fact does not obtain the content of the experience will be different, and hence the target reason will be absent). These two theses, when conjoined,

²¹ I offer the extended account in Pritchard (2005a).

pose a problem for the classical way of understanding the internalism/externalism debate because it is standardly thought that what one has reflective access to cannot extend beyond the ‘inner’ to take in factive empirical reasons in this way. According to McDowell, however, this conventional wisdom of contemporary philosophy is false and leads to a philosophical picture which invites the sceptical challenge. Moreover, the source of the problem with this conventional wisdom, according to McDowell, is the failure to endorse the kind of content externalism that he has in mind.

For example, for McDowell one’s reason for believing an empirical perception-based proposition—say, that John is in my office—could simply be the factive empirical reason that one *sees that* John is in my office. Moreover, since one’s reasons are reflectively available to one (since otherwise they would not be one’s reasons at all), it follows that one has reflective access to this factive empirical reason. That is, one’s reason for believing that John is in one’s office could be that one sees that he is in one’s office, where one is able to know by reflection alone that one possesses this reason.

We saw above that there is an everyday conception of evidence according to which we can legitimately cite factive grounds in favour of our empirical beliefs in this way. According to the classical internalist, however, this everyday practice of offering factive grounds cannot be taken at face value. Instead, the classical internalist will claim that the evidential grounding of one’s belief must be, strictly speaking, the non-factive counterpart of the factive claim, such as that *it seems to one as if* John is in one’s office. Part of the reason for this restriction, presumably, is that it is only the non-factive counterpart that could be reflectively accessible to the subject. Any straightforward accommodation of this everyday conception of evidence is thus automatically placed in the classical externalist camp. It is this classical internalist orthodoxy that McDowell’s view challenges.

On the face of it, the McDowellian line is susceptible to a straightforward problem, one that mirrors the famous ‘McKinsey’ puzzle that concerns the supposed compatibility of first-person authority and content externalism.²² For it seems that one can use one’s reflective access to one’s factive empirical reasons, along with one’s reflective knowledge of the relevant entailment (i.e., from the factive empirical reason to the empirical fact), to acquire reflective, and thus non-empirical, knowledge of the empirical proposition which the reason is a reason for (in this case that John is in one’s office). Intuitively, this is just a *reductio* of the view, and this in part explains, I think, why very few commentators have taken McDowell seriously on this point and have stuck, instead, to the conventional wisdom on reflective access to reasons that is implicit in the classical way of drawing the externalism/internalism contrast in epistemology.²³

²² For more on this problem, see Nuccetelli (2003).

²³ For example, Greco (2004) takes the factivity of reasons claim seriously, but doesn’t take the internalist reflective access claim seriously as a result, and therefore regards McDowell’s view

It is not obvious, however, that McDowell's position is subject to a problem of this sort. The reason for this is that the difficulty only emerges if one *acquires* non-empirical knowledge by running through an inference of this type, and it is far from clear why, on the McDowellian view, this would be so. In particular, the McDowellian conception of perceptual knowledge in these cases is such that one could only be in a position to run through such an inference provided one already has the empirical knowledge in question, and thus the problem concerning the non-empirical acquisition of empirical knowledge does not arise.

In order to see this, one only needs to note that for McDowell a factive reason for p and a relevant belief that p which is based on that factive reason will suffice for knowledge that p . Note, however, that reflective accessibility is itself factive, in that if one is in a position to reflectively access that one has a factive reason for believing p then it must be the case that one has that factive reason for believing p . Furthermore, even though the possession of the factive reason is consistent with non-belief in the target proposition on the McDowellian view, it is obviously not going to be possible for someone to run through the reasoning described above without in the process acquiring the reason-based belief in the target proposition, thereby meeting all the conditions required on this view for empirical knowledge in this case. Thus, given the further trivial claim that if one has empirical knowledge of p then one cannot also have reflective (i.e., non-empirical) knowledge of p , it follows that there is no prospect of acquiring non-empirical knowledge in this case, and thus McDowell's McKinsey-style difficulty disappears.

On the face of it, then, there is a position available here, one that retains that aspect of the classical internalist thesis that insists on reflectively accessible grounds in favour of one's beliefs if one is to count as a knower, but which also allows that one's grounds, so construed, can be robust enough to meet the underdetermination-based sceptical argument. Accordingly, one could combine the contextualist pragmatic thesis outlined above with this line of argument to derive an internalist formulation of the view which does not concede that our knowledge of these anti-sceptical propositions is necessarily brute. Instead, one can have adequate reflectively accessible grounds in favour of such beliefs, such as one's factive perception-based reasons which entail the denials of sceptical hypotheses.

FN:24 This view clearly merits further exploration.²⁴ Notice, however, that while it changes the shape of the dialectic here by opening up a new direction of research,

as simply being a version of classical epistemic externalism. Wright (2002), in contrast, takes the internalist reflective access claim seriously but as a consequence does not take the factivity claim seriously. On his reading, what one can have reflective access to is not the factive empirical reason but rather a 'disjunctive' reason—i.e., a reason for believing that either one is in the factive state or one is not in the factive state and deluded in some way.

²⁴ For more detailed discussion of the view, see Neta and Pritchard (forthcoming) and Pritchard (forthcoming b).

it does not make any substantial difference to the key points raised earlier. For while it is now true on this view that one does have reflectively accessible grounds in favour of one's anti-sceptical beliefs, this does not mean that one is in a position to, for example, properly cite them in a sceptical conversational context in which we are taking sceptical error-possibilities seriously. Even if it is true that one possesses a factive empirical reason which entails the denial of, say, the BIV sceptical hypothesis, it can hardly be thought that such a reason—that I see that I have two hands, for example—would represent grounds which speak to this particular contrast, and yet this assertion in this context will certainly generate this conversational implicature. We have seen above just why this is so, since to explicitly claim knowledge in a sceptical context is to represent oneself as having reflectively accessible evidence which *discriminates* between the proposition asserted as known and sceptical alternatives. Crucially, however, one's reflectively accessible factive empirical reasons will not serve *this* role.²⁵

[FN:25]

I think this is part of the reason why McDowell—at least in his more careful remarks on this subject—only claims to be showing how one can legitimately *ignore* the sceptical argument, rather than claiming to have actually *responded* to it. The thought is that once one recognizes that one's reflectively accessible grounds can be factive then one should not feel the pressure to conversationally engage with the sceptical problem with anti-sceptical assertions any more, not that one thereby has a response which should silent even the sceptic (which is what the Moorean strategy seems to aspire to). (I take it that the classical externalist neo-Moorean would be inclined to agree with this claim, and thus on this score at least the two views are very similar.) Thus, the McDowellian line is very much in the spirit of neo-Mooreanism.

I here leave it open whether one should be a McDowellian neo-Moorean or simply a classical externalist neo-Moorean, since either will suffice to meet the sceptical problem. I do, however, want to make one small remark on this issue before I close, which points to how we should go about exploring this issue. This is that on the face of it the McDowellian view is plausibly in a better position to be counted as the true heir to the Moorean tradition on account of how it retains core internalist intuitions. After all, part of the desiderata of neo-Moorean positions is that they are able, where possible, to accommodate our pre-theoretical intuitions, and internalist intuitions are surely highly embedded within folk epistemology. Moreover, since both views appear to endorse some form of disjunctivism, it seems that one could similarly argue that what the McDowellian picture highlights is how disjunctivism, properly understood, enables one to evade the sceptical problematic *without* having to resort to the revisionism of classical epistemic externalism. Such fine-grained issues of which sort of neo-Mooreanism one should endorse once one has dealt with the sceptical problem can, however, be left for another occasion.

²⁵ For more on this point, see Pritchard (2003).

9. CONCLUDING REMARKS

We have thus seen that there is a viable neo-Moorean view available which can deal with the sceptical problem in a more satisfactory way than other competing anti-sceptical theories, and which avoids the problems facing Mooreanism. In particular, it is an anti-luck epistemology that retains closure while avoiding a contextualization of our understanding of ‘knows’, but which also accommodates contextualist intuitions by incorporating a view about the context-sensitivity of the propriety conditions for explicit knowledge self-ascriptions. Furthermore, this pragmatic account also serves to demarcate the neo-Moorean view from its Moorean ancestor, since it explains, in part, why the Moorean assertions were problematic (and thus why scepticism can seem so plausible) even though they were true. Finally, we have seen that one can resolve the evidential sceptical problem without resorting to an evidential form of contextualism by adopting a version of content externalism and either opting for a classical externalist epistemology or a McDowellian version of epistemic internalism.

There is one further element of the neo-Moorean account that I think is required, though I will not explore it at length here. This is that the neo-Moorean would be wise, I think, to concede something to the sceptic; to say that there is something right about the sceptical challenge. The key to this concession lies in the fact that the neo-Moorean anti-sceptical response is in a certain way necessarily *mute*, since it is part of the view that one cannot properly respond to the sceptical challenge by repeating one’s everyday claims to know. I think there is a deep point here about the limits of our cognitive responsibility, a point that I have explored at length elsewhere.²⁶ Making such a concession to the sceptic does not undermine the view so long as one steadfastly retains the core claim that knowledge is nonetheless possessed in such cases, so that the epistemic lack at issue here, if that is the right way to characterize it, is not an epistemic lack that would undermine knowledge. Such a concession would thus only strengthen the view by accounting for our visceral attraction to sceptical arguments, even despite our strong anti-sceptical intuitions. The point would be that there is a deep truth in sceptical arguments, though not the deep truth that the sceptic advertises. Neo-Mooreanism, while obviously a particularly robust anti-sceptical theory, need not be a view that dismisses the sceptical problem entirely.

FN:26

REFERENCES

- Alston, W. 1988: ‘An Internalist Externalism,’ *Synthese* 74: 265–83.
 Baldwin, T. 1993: ‘G. E. Moore.’ In J. Dancy and E. Sosa (eds.), *A Companion to Epistemology* (Oxford: Blackwell), 283–5.

²⁶ See, for example, Pritchard (2005b, *passim*; 2005c).

- Brandom, R. 1998: 'Insights and Blindspots of Reliabilism.' *The Monist* 81: 371–92.
- Cohen, S. 2000: 'Contextualism and Skepticism.' *Philosophical Issues* 10: 94–107.
- Conee, E. and Feldman, R. 2000: 'Internalism Defended.' In H. Kornblith (ed.), *Epistemology: Internalism and Externalism*, (Oxford: Blackwell), 231–60.
- DeRose, K. 1995: 'Solving the Skeptical Problem.' *Philosophical Review* 104: 1–52.
- Dretske, F. 1970: 'Epistemic Operators.' *Journal of Philosophy* 67: 1007–23.
- 2005a: 'The Case against Closure.' In M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Oxford: Blackwell), 13–26.
- 2005b: 'Reply to Hawthorne.' In M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*, (Oxford: Blackwell) 43–6.
- Foley, R. 1987: *A Theory of Epistemic Rationality* (Cambridge, Mass.: Harvard University Press).
- Greco, J. 2004: 'Externalism and Skepticism.' In R. Shantz (ed.), *The Externalist Challenge: New Studies on Cognition and Intentionality*, New York, de Gruyter.
- Hawthorne, J. 2004: *Knowledge and Lotteries* (Oxford: Oxford University Press).
- 2005: 'The Case for Closure.' In M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Oxford: Blackwell) 26–43.
- Heller, M. 1999: 'Relevant Alternatives and Closure.' *Australasian Journal of Philosophy* 77: 196–208.
- Kornblith, H. (ed.) 2001: *Epistemology: Internalism and Externalism* (Oxford: Blackwell).
- Lewis, D. 1996: 'Elusive Knowledge.' *Australasian Journal of Philosophy* 74: 549–67.
- McDowell, J. 1995: 'Knowledge and the Internal.' *Philosophy and Phenomenological Research* 55: 877–93.
- Moore, G. E. 1925: 'A Defence of Common Sense.' In J. H. Muirhead (ed.), *Contemporary British Philosophy* (2nd series) (London: Allen and Unwin).
- 1939: 'Proof of an External World.' *Proceedings of the British Academy* 25: 273–300.
- Neta, R. 2002: 'S Knows that P.' *Nous* 36: 663–81.
- 2003: 'Contextualism and the Problem of the External World.' *Philosophy and Phenomenological Research* 66: 1–32.
- and Pritchard, D. H. Forthcoming: 'McDowell and the New Evil Genius.' *Philosophy and Phenomenological Research*.
- Nozick, R. 1981: *Philosophical Explanations* (Oxford: Oxford University Press).
- Nuccetelli, S. (ed.) 2003: *New Essays on Semantic Externalism and Self-Knowledge* (Cambridge, Mass.: MIT Press).
- Pritchard, D. H. 2000: 'Closure and Context.' *Australasian Journal of Philosophy* 78: 275–80.
- 2002a: 'McKinsey Paradoxes, Radical Scepticism, and the Transmission of Knowledge across Known Entailments.' *Synthese* 130: 279–302.
- 2002b: 'Radical Scepticism, Epistemological Externalism and Closure.' *Theoria* 69: 129–61.
- 2002c: 'Recent Work on Radical Skepticism.' *American Philosophical Quarterly* 39: 215–57.
- 2002d: 'Resurrecting the Moorean Response to Scepticism.' *International Journal of Philosophical Studies* 10: 283–307.
- 2003: 'McDowell on Reasons, Externalism and Scepticism.' *European Journal of Philosophy* 11: 273–94.

- 2004: 'Epistemic Luck.' *Journal of Philosophical Research* 29: 193–222.
- 2005a: 'Contextualism, Scepticism and Warranted Assertibility Manoeuvres.' In J. Keim-Campbell, M. O'Rourke, and H. Silverstein (eds.), *Knowledge and Scepticism* (Cambridge, Mass.: MIT Press).
- 2005b: *Epistemic Luck* (Oxford: Oxford University Press).
- 2005c: 'Scepticism, Epistemic Luck and Epistemic Angst.' *Australasian Journal of Philosophy* 83: 185–206.
- 2005d: 'The Structure of Skeptical Arguments.' *Philosophical Quarterly* 55: 37–52.
- 2005e: 'Wittgenstein's *On Certainty* and Contemporary Anti-Scepticism'. In D. Moyal-Sharrock and W. H. Brenner (eds.), *Investigating On Certainty: Essays on Wittgenstein's Last Work* (London: Palgrave Macmillan).
- 2006: 'Knowledge, Luck, and Lotteries.' In V. F. Hendricks and D. H. Pritchard (eds.), *New Waves in Epistemology* (Aldershot: Ashgate).
- forthcoming a: 'Contrastivism, Evidence, and Scepticism.' *Social Epistemology* (special issue on *Contrastivism*) 20.
- forthcomingb: 'McDowellian Neo-Mooreanism.' In A. Haddock and F. Macpherson (eds.), *Disjunctivism: Perception, Action, Knowledge* (Oxford: Oxford University Press).
- Pryor, J. 2000: 'The Skeptic and the Dogmatist.' *Nous* 34: 517–49.
- 2001: 'Highlights of Recent Epistemology.' *British Journal for the Philosophy of Science* 52: 95–124.
- Shaffer, J. 2004: 'From Contextualism to Contrastivism in Epistemology.' *Philosophical Studies* 119: 73–103.
- 2005a: 'Closure, Contrast, and Answer.' Manuscript.
- 2005b: 'Contrastive Knowledge.' In T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology* (Oxford: Oxford University Press).
- Sosa, E. 1999: 'How to Defeat Opposition to Moore.' *Philosophical Perspectives* 13: 141–54.
- Stanley, J. 2004: 'On the Linguistic Basis for Contextualism.' *Philosophical Studies* 119: 119–46.
- 2005: *Knowledge and Practical Interests* (Oxford: Oxford University Press).
- Williamson, T. 1996: 'Knowing and Asserting.' *Philosophical Review* 105: 489–523.
- 2000a: 'Scepticism and Evidence.' *Philosophy and Phenomenological Research* 60: 613–28.
- 2000b: *Knowledge and its Limits* (Oxford: Oxford University Press).
- Wittgenstein, L. 1969: *On Certainty*, ed. G. E. M. Anscombe and G. H. von Wright, tr. D. Paul and G. E. M. Anscombe (Oxford: Blackwell).
- Wright, C. 1991: 'Scepticism and Dreaming: Imploding the Demon.' *Mind* 397: 87–115.
- 2002: '(Anti-)Skeptics Simple and Subtle: G. E. Moore and John McDowell.' *Philosophy and Phenomenological Research* 65: 330–48.
- 2004: 'Warrant for Nothing (and Foundations for Free)?' *Proceedings of the Aristotelian Society* 78 (suppl. vol.): 167–212.
- Zagzebski, L. 1996: *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge: Cambridge University Press).