

## RESURRECTING THE MOOREAN RESPONSE TO THE SCEPTIC

Duncan Pritchard  
University of Stirling

ABSTRACT. G. E. Moore famously offered a strikingly straightforward response to the radical sceptic which simply consisted of the claim that one could know, on the basis of one's knowledge that one has hands, that there exists an external world. In general, the Moorean response to scepticism maintains that we can know the denials of sceptical hypotheses on the basis of our knowledge of everyday propositions. In the recent literature two proposals have been put forward to try to accommodate, to varying extents, this Moorean thesis. On the one hand, there are those who endorse an externalist version of contextualism, such as Keith DeRose, who have claimed that there must be some contexts in which Moore is right. More radically still, Ernest Sosa has expanded on this externalist thesis by arguing that, *contra* DeRose's contextualism, Moore may be right in all contexts. In this paper I evaluate these claims and argue that, suitably modified, one can resurrect the main elements of the Moorean anti-sceptical thesis.

## I

G. E. Moore (1939) famously argued against radical scepticism about our knowledge of the external world in the following 'common-sense' fashion:

- (1) I know that I have two hands.
  - (2) I know that if I have two hands then there must be an external world.<sup>1</sup>
- Hence:
- (C) I know that there is an external world.

Moore regarded this as being a perfectly "rigorous" proof, reminding us that

[...] we all of us do constantly take proofs of this sort as absolutely conclusive proofs of certain conclusions—as finally settling certain questions, as to which we were previously in doubt. (Moore 1939, 147)

To illustrate this point he offers the example of proving that there are at least three misprints on a page. To settle this question we simply look for one, then another, and then another. As Moore points out, if this counts as an adequate 'proof' of the contested proposition in this context, then it is far from clear why the displaying of one's hands should be thought to be a deficient response to the external world sceptic.<sup>2</sup>

In general, we might take any 'Moorean' response to scepticism to have the following form.<sup>3</sup> Given that we know lots of everyday propositions, and given that we also know that the truth of these propositions entails the denials of radical sceptical hypotheses, we must know the denials of radical

sceptical hypotheses. More formally, we might express the argument as follows, where ‘K [ $\varphi$ ]’ says that the subject knows the proposition  $\varphi$ ;<sup>4</sup> ‘O’ stands for any suitable ordinary proposition which we standardly take ourselves to know; and ‘ $\neg$  SH’ stands for the denial of a sceptical hypothesis which is inconsistent with the truth of the ordinary proposition:

*The Moorean Template*

- (M1) K [O]  
 (M2) K [O  $\rightarrow$   $\neg$  SH]  
 Hence:  
 (MC) K [ $\neg$  SH]

So, for example, one could view the following response to brain-in-a-vat (BIV) scepticism as conforming to this Moorean template, where ‘O’ says that the agent is, for example, currently seated at her computer; and ‘ $\neg$  BIV’ says that the agent is not currently a BIV:

- (M1\*) K [O]  
 (M2\*) K [O  $\rightarrow$   $\neg$  BIV]  
 Hence:  
 (MC\*) K [ $\neg$  BIV]

In a sense, the Moorean response to scepticism simply runs the sceptical argument in reverse, in that it takes the negation of the sceptical conclusion as its main premise. In contrast, a standard radical sceptical argument will tend to begin by appropriating as its main premise the denial of the conclusion of the Moorean template, (MC). The sceptic thus takes as her starting-point the claim that the agent in question does not (indeed, *cannot*) know that she is not the victim of a sceptical hypothesis. Given that the sceptic accepts the (relatively uncontentious) second premise, (M2), of the Moorean template, however, she is then in a position to derive the denial of the first premise, (M1), of the Moorean argument, and therefore conclude that the agent lacks knowledge of the ordinary proposition in question:

*The Sceptical Template*

- (S1)  $\neg$  K [ $\neg$  SH]  
 (S2) K [O  $\rightarrow$   $\neg$  SH]  
 Hence:  
 (SC)  $\neg$  K [O]

And since one could repeat this argument with any agent and with a wide class of ordinary propositions, so radical scepticism ensues.

One might be tempted to express this difference between the two arguments by saying that the sceptic simply argues *modus tollens* in response to the Moorean’s *modus ponens*, but that is not quite right. For it is not deductive closure as such that both arguments pivot upon, but rather an epistemic

version of deductive closure that concerns whether *knowledge* transfers across, and is therefore ‘closed’ under, known entailment. This principle states that, roughly, if one knows a proposition  $\varphi$ , and one knows that this proposition entails a second proposition  $\psi$ , then one also knows  $\psi$ . Formally:

*Epistemic Deductive Closure*  
 $\{K[\varphi] \ \& \ K[\varphi \rightarrow \psi]\} \rightarrow K[\psi]$ <sup>5</sup>

Indeed, like *modus ponens* itself, this is a highly intuitive principle. It is at least odd to say that an agent knows a proposition, knows that this proposition entails a second proposition, and yet does not know the second proposition. Following convention, we shall henceforth refer to epistemic deductive closure as simply ‘closure’.

It has seemed to some, however, that this principle must be false—at least if radical scepticism is to be rejected—for the simple reason that although we do appear to know lots of everyday propositions (which we know entail the denials of sceptical hypotheses), we *cannot*, intuitively at least, know the denials of sceptical hypotheses. After all, *ex hypothesi*, there is (or, at least, need be) no phenomenological difference between everyday life and, say, being a BIV who is ‘fed’ her experiences by neuroscientists. And if there is no phenomenological difference then, so the argument goes, there can be no evidence available to the subject which could support any belief she might have to the effect that she is not a BIV. Accordingly, since it is entirely natural to suppose that where there is no evidence then, *a fortiori*, there can be no knowledge, it follows that we cannot know the denials of sceptical hypotheses.

Given that it endorses closure, the Moorean response to the sceptic is thus contentious, first and foremost, in its acceptance of the claim that we can know the denials of sceptical hypotheses. It is for this reason that the Moorean argument is dialectically at a disadvantage when it comes to scepticism. After all, if one has reason to believe that the conclusion of an apparently valid argument must be false, then one is entitled to doubt the veracity of one of the premises. And since premise two of the Moorean argument, (M2), is above reproach in this respect (along with the closure principle itself, by Moorean lights), it would thus appear that it is the first premise, (M1), that must be denied. Such, in any case, is the enduring attraction of radical sceptical arguments of this form.<sup>6</sup>

On this account of the sceptical debate, we are thus faced with a dilemma. Either we confront the sceptic head-on by denying the highly intuitive closure principle, or else we retain this principle at the expense of leaving us in a dangerous *impasse* with the sceptic, such that there seems just as much reason to draw the sceptical conclusion of denying (M1) as there is to endorse the anti-sceptical Moorean conclusion, (MC). As we shall see, these are not the only two dialectical options available here, but they represent a good place to start our discussion because they focus our evaluation of the

Moorean proposal. After all, if the Moorean argument is such that its conclusion seems to be, on reflection at least, false, then the prospects for a Moorean revival are slim indeed.

## II

Given the above remarks, one might think that the natural *anti*-sceptical response to the sceptical argument is precisely *not* to deny the sceptical premise, (S1)—and therefore argue, with Moore, that we do know the denials of sceptical hypotheses after all—but rather to query the sceptic’s use of the epistemic principle of closure. One could imagine how an informal argument along these lines might go. Typically, knowing an ‘ordinary’ proposition requires that I rule-out a wide range of relevant possibilities of error—possibilities that are inconsistent with the truth of the proposition putatively known. What it does not require, however, is that I rule-out *every* possibility of error, even manifestly far-fetched and irrelevant ones. Accordingly, I can know a proposition, know that it entails the denial of some irrelevant possibility of error, and yet consistently *fail* to know that the irrelevant possibility of error does not obtain.

Of course, part of the task of motivating this informal thought is showing how one could erect a plausible conception of knowledge on the basis of it. To this end, the two main proponents of the view that closure should be rejected, Fred Dretske (1970; 1971) and Robert Nozick (1981), have argued that a necessary requirement of an agent’s knowledge of a contingent proposition is that the agent’s belief in that proposition should be *sensitive*, where this, in turn, means that, in the nearest possible world in which that proposition is false, the agent does not believe it.<sup>7</sup> More formally, where ‘ $\Rightarrow$ ’ expresses the subjunctive conditional, and ‘B [ $\varphi$ ]’ says that the agent believes the contingent proposition  $\varphi$ :

*Sensitivity*

$\neg \varphi \Rightarrow \neg B [\varphi]$

If this were a necessary condition on knowledge, then both the impossibility of knowing that one is not the victim of a sceptical hypothesis and the failure of closure would be explained. In the nearest possible world in which it is false that one is not a BIV (i.e., the BIV-world), one continues to believe that one is not a BIV (because one is being ‘fed’ everyday experiences by the neuroscientists). Accordingly, one can never know that one is not a BIV (or, indeed, not the victim of any radical sceptical hypothesis). In contrast, if the actual world is pretty much as we take it to be (so that sceptical hypotheses are indeed far-fetched), then it will follow that one’s belief in everyday

propositions *will* be sensitive. Take my ‘Moorean’ belief that I have two hands. Intuitively, in the nearest possible world in which this is false, such as the world in which my hands got caught in the garden shredder I used last weekend, I don’t believe that I have hands because I look down and see two bandaged stumps before me. As a result, one’s belief in a proposition might be sensitive, and thus one might be in a position to know that proposition (depending upon whatever other conditions one applies to one’s theory of knowledge), and yet one might *lack* a sensitive belief in (and so lack knowledge of), the denial of a known consequence of that proposition—such as the denial of the sceptical hypothesis. Hence, closure fails, and our everyday knowledge is now ‘insulated’ from sceptical attack.

This is certainly an interesting proposal. In particular, it is important to note that it is an anti-sceptical strategy that is entirely *externalist* in spirit in that it points to factual conditions that must be met in order for an agent to know whilst eschewing any claim that the agent should have the kind of reflective access to those facts characteristically demanded by the internalist. Provided that one understands internalism as essentially incorporating a commitment to the idea that knowledge involves, at the very least, reflective access to the factors that make it such that one knows (rather than just truly believes), then this proposal breaks with the internalist paradigm.<sup>8</sup> Moreover, if this view is to have any anti-sceptical plausibility, then this commitment to externalism is not an incidental feature of the position. After all, as we just noted above, there need be no reflectively accessible difference between, say, sitting at one’s computer and being a BIV being fed experiences by neuroscientists. And since, intuitively, one only has ‘evidence’ (where this is given an internalist reading) for one’s everyday beliefs in the first place on the assumption that one is not being systematically deceived, it would appear to follow that the reflective access necessary for knowledge of even the ordinary propositions in question is lacking. So even with closure gone, it still remains that one can only use this strategy as a means of answering the sceptic provided one also endorses the externalist epistemology of which it forms part.<sup>9</sup>

It is only by moving away from the internalist paradigm, then, that one finds such a proposal plausible. Indeed, this may be no bad thing, since it does seem to be more central to our everyday conception of knowledge—as opposed to say, justification—that one regards subjects as exhibiting this kind of ‘tracking’. For example, it is well documented that we are quite happy to ascribe knowledge to agents who have no substantive reflective awareness of what factors support that knowledge, just so long as they do in fact have beliefs that are sufficiently responsive to the truth.<sup>10</sup>

It is not my intention to argue for an externalist epistemology here, however, but merely to note what the consequences of endorsing such an epistemology might be. For example, an immediate repercussion of the externalism inherent in the Dretske-Nozick position is that it calls into question

the very anti-sceptical move that lies at the heart of that proposal. If one is happy with the thought that an agent might know a proposition simply by meeting such an external condition, then the motivation for accepting the sceptical premise, (S1)—that we do not know the denials of radical sceptical hypotheses—diminishes accordingly, and, with it, the motivation for denying closure.

Recall that the reason we gave above for finding this premise so compelling was the thought that evidence was lacking in such cases, and thus, *a fortiori*, that knowledge must be lacking as well. Whilst this line of argument might be uncontentious when set within an internalist conception of knowledge, however, it fails to be compelling when set within an externalist thesis. After all, according to the latter view, the absence of such evidence need have no impact upon whether or not I do in fact meet a purely external condition. Granted, I cannot meet the specific requirement of sensitivity that Dretske and Nozick propose, but *that* external condition was placed on knowledge in order to *explain* how it could be that I could know everyday propositions whilst failing to know the denials of radical sceptical hypotheses (even when I know that the former entails the latter). If, on an externalist account, it is no longer a datum that one does lack knowledge of the denials of sceptical hypotheses, then this casts the adoption of sensitivity as a requirement on knowledge into a new light. Moreover, since it seems intuitively plausible to contend that, all other things being equal, one should accept an epistemological proposal that retains closure over a more revisionistic one that doesn't, then we have grounds for being sceptical about the Dretske-Nozick argument for non-closure.

Of course, this claim that we lack evidence to support our belief in the denials of radical scepticism hypotheses is not the only motivation for accepting the sceptical premise, and we shall explore other arguments to this end in due course. The interesting possibility raised by the Nozick-Dretske line, however, is that it may re-open the door to the Moorean proposal, a central plank of which is the contention that we can indeed know the denials of radical sceptical hypotheses. In what follows, I shall consider three ways in which one might develop this thought: via the particular brand of contextualism adopted by Keith DeRose; via Ernest Sosa's 'neo-Mooreanism'; and via my own approach which aims to expand upon the best points of the other two theories.

### III

For DeRose (1995), the basic contextualist strategy pivots upon the acceptability, and appropriate use, of the following contextualist thesis:

Suppose a speaker A (for "attributor") says, "S knows that P", of a subject S's true belief that P. According to contextualist theories of knowledge attributions, how strong an epistemic position S must be in with

respect to P for A's assertion to be true can vary according to features of A's conversational context. (DeRose 1995, 4)

DeRose employs this thesis as a means of explaining the following supposed features of the 'phenomenology' of our engagement with scepticism. First, that ascriptions of knowledge to subjects in conversational contexts in which sceptical error-possibilities have been raised seems wholly inappropriate. Second, that in conversational contexts in which no sceptical error-possibilities are in play it seems perfectly appropriate to ascribe knowledge to subjects. And, third, that all that may change when one moves from a non-sceptical conversational context to a sceptical context are mere conversational factors. Intuitively, these three 'intuitions' are in conflict because, or so the standard non-contextualist thought runs, *one* of these judgements must be wrong. That is, since conversational topic has no obvious bearing on the epistemic status of a subject's beliefs, it ought to be universally true (i.e., *whatever* the conversational context) that the subject either does or does not know the propositions in question. Contextualism opposes this thought with the suggestion that what is actually occurring is not a contradiction but a responsiveness, on the part of the attributor of knowledge, to a fluctuation in the epistemic standards (and with them the subject's possession of knowledge) caused by a change in the conversational context.<sup>11</sup>

The first thing that DeRose tries to capture is the intuition that as one moves from one conversational context to another one's epistemic situation (one's total informational state for instance) could remain exactly the same. DeRose accommodates this intuition in conjunction with the contextualist picture by arguing, as the above quotation indicates, that although one's "epistemic position" is constant at any one time, the epistemic position that one *needs to be in* so as to count as possessing knowledge can be variable. Strength of "epistemic position" is characterised by DeRose as follows:

[...] being in a strong epistemic position with respect to P is to have a belief as to whether P is true match the fact of the matter as to whether P is true, not only in the actual world, but also at the worlds sufficiently close to the actual world. That is, one's belief should not only be true, but also should be non-accidentally true, where this requires one's belief as to whether P is true to match the fact of the matter at nearby worlds. The further away one gets from the actual world, while still having it be the case that one's belief matches the fact at worlds that far away and closer, the stronger a position one is in with respect to P. (DeRose 1995, 34)

In order to see this, imagine that Lars believes that his car is outside on the basis of a certain fixed informational state (which involves, perhaps, his memory of the car being there a few hours ago, his grounds for believing that no-one would steal it, and so forth). Now imagine an (almost) exact counterpart of Lars—Lars\*—who is in exactly the same cognitive state except that he has the extra piece of information that the car was there a minute ago (perhaps he looked). Clearly, Lars\* will be in

a better epistemic position with respect to his belief that his car is outside than Lars. Although they will, in general, track the truth across the same set of possible worlds, Lars\* will track the truth in a few extra possible worlds, such as the possible worlds in which the car was stolen ten minutes ago.

We are now a little nearer to understanding DeRose's initial formulation of the contextualist position in terms of the variability of "how strong an epistemic position S must be in with respect to P for A's assertion [*that S knows P*] to be true." DeRose's idea is that the notion of sensitivity which we saw Dretske and Nozick advocating above—whilst not essential to knowledge possession itself—will play a part in the mechanism that changes a conversational context in such a way as to raise (or lower) the strength of epistemic position that one needs in order to count as knowing. Recall that for an agent to have a belief in P that is sensitive, the agent must not only have a true belief in P in the actual world, but also *not* believe P in the nearest possible world in which P is false. DeRose's thought is that in any particular conversational context there is a certain set of propositions that are explicitly at issue and the agent must, at the very least, be sensitive to all these 'explicit' propositions if she is to know them. Moreover, the most demanding of these propositions—the proposition which has a negation that occupies the furthest-out possible world—will set the standard for that conversational context since this not-P world will determine the extent of possible worlds that one's beliefs must be able to track if one is to be truly said to know a proposition in that context. Knowing a proposition thus involves being in an epistemic position sufficient to track the truth across the range of possible worlds determined by the most demanding proposition explicit to that context. Crucially, however (and I shall be expanding upon this detail in a moment), this point also applies to propositions which are *implicit* to a conversational context (i.e., propositions which one believes but which are not explicit to that conversational context). In order to know such a proposition—even if one's belief in that proposition is not sensitive—one need only be in a sufficient epistemic position to meet the standards of that context. (The importance of this point shall soon become apparent).

DeRose then characterises the mechanism that brings about an upward shift in epistemic standards as follows:

When it is asserted that some subject S knows (or does not know) some proposition P, the standards for knowledge (the standards for how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require S's belief in that particular P to be sensitive for it to count as knowledge. (DeRose 1995, 36)

That is, what changes a conversational context is when a new proposition is made explicit to that context which is more demanding than any of the propositions currently explicit in that context. This will thus increase the range of possible worlds at issue in the determination of knowledge, and thereby increase the strength of epistemic position required in order to be truly said to know.

What motivates this claim is the fact that, as David Lewis (1979) famously argued, when it comes to ‘context-sensitive’ terms like ‘flat’ or ‘knowledge’, the conversational ‘score’ tends to change depending upon the assertions of that context. We may all agree that the table in front of us is ‘flat’ in an everyday context, but, *ceteris paribus*, if someone enters the room and denies that it is flat we do not thereby disagree with her. Instead, we take it that she means ‘flat’ in some more demanding sense and so raise the standards for ‘flatness’ so as to make her assertion true (this is what Lewis (1996, 559) calls a “rule of accommodation”). That is, we take it that the new participant of our conversational context means flat in some more restricted sense so that the barely perceptible bumps on the table before us are sufficient to make the claim “This table is flat” false. DeRose considers the Lewis line to have captured something intuitive about the pragmatics of how we use our ‘context-sensitive’ terms and, moreover, believes that epistemic terms such as ‘knowledge’ behave in a similar way.<sup>12</sup>

An example will help clarify matters here. Imagine an agent in a quotidian context in which only everyday propositions, such as whether or not one is currently having dinner with one’s brother (P), or whether or not the garden gate has been closed (Q), are at issue. Sensitivity to these everyday propositions will only require the consideration of nearby possible worlds and thus the strength of epistemic position demanded will be very weak. Let us say, plausibly, that the possible world in which one is not having dinner with one’s brother is ‘further-out’ than the possible world in which the garden gate is not closed. This proposition will thus determine the range of possible worlds at issue in the determination of knowledge in that conversational context. Let us further suppose that the agent in question does have a sensitive belief in this proposition. The issue of what other propositions the subject knows will now be decided by whether the agent’s belief in those propositions will track the truth across the range of possible worlds determined by not-P. If, for instance, the agent’s belief that the garden gate is closed matches the truth as to whether Q in all of the possible worlds within that range, then she will know Q. Equally, however, if the subject’s belief in a proposition which is implicit to that conversational context tracks the truth across this range of possible worlds then that proposition will be known also, even if the agent’s belief in that proposition is not sensitive.

Consider, for instance, the agent’s belief that she is not a BIV, a proposition which, as we saw above, one cannot be sensitive to because in the nearest BIV-world one still believes that one is not a BIV. On the contextualist model, however, if one is in a conversational context in which such a proposition is not explicit, then one can know this proposition just so long as one has a belief as to whether this proposition is true which matches the facts as to whether it is true within the range of possible worlds at issue. And, clearly, this demand will be trivially satisfied in the above scenario where the subject has a sensitive belief in the ordinary proposition, P. After all, insofar as one has

such a sensitive belief in P, then it must be the case that the BIV sceptical world is, modally speaking, far-out, for if it weren't, then this would affect the sensitivity of the subject's beliefs in ordinary propositions like P. Accordingly, on this view, all the subject needs in this context is a stubborn belief that she is not a BIV in order to be truly said to know this proposition in this conversational context. The contextualist can thus capture the second element of the 'phenomenology' of our engagement with scepticism that we noted above—that, in quotidian conversational contexts, we are perfectly willing to ascribe knowledge of everyday propositions and also feel that we must know the denials of sceptical hypotheses as well.

One might wonder why I use the word "feel" here. Well, the reason is that, on the contextualist account, if one were to explicitly mention these anti-sceptical propositions (as one would if one were to verbally ascribe knowledge of them to oneself), then one would thereby make that proposition explicit to the conversational context and so change the epistemic standards needed for knowledge accordingly.<sup>13</sup> In order to have knowledge that one is not a BIV within the new conversational context, one's belief that one is not a BIV must now exhibit sensitivity (which, as we saw above, is impossible), and the possible worlds relevant to the determination of that sensitivity will be relevant to one's knowledge of even everyday propositions. Accordingly, one will now lack knowledge both of the denial of the sceptical hypothesis (because one's belief in this respect is not sensitive), *and* of the everyday propositions (since even though one's beliefs in these propositions are sensitive, one can never be in an epistemic position that would support knowledge of them which would be strong enough to track the truth in far-off BIV-worlds). The contextualist thus claims to have captured the other two aspects of the 'phenomenology' of our engagement with scepticism—that we are extremely reluctant to ascribe knowledge in sceptical conversational contexts, and that this is even so when the only thing that may have changed from the non-sceptical conversational context in which we were willing to ascribe knowledge is the course of the conversation. Moreover, the contextualist has done this without either conceding the universal truth of scepticism (since scepticism is false in everyday contexts), or denying closure (since there is no single context in which one both knows an everyday proposition whilst lacking knowledge of the denial of a sceptical hypothesis).<sup>14</sup> Accordingly, DeRose claims to have 'solved' the sceptical paradox in an entirely intuitive manner.

What is particularly interesting about this proposal from our point of view, however, is that it also endorses the Moorean claim that we can know the denials of sceptical hypotheses, at least provided they are not made relevant to our current conversational context.<sup>15</sup> Moreover, DeRose effects this result by employing the insight gained in the last section that if we endorse an externalist account of knowledge then the conceptual space is left open for us to allow knowledge of the denials

of radical sceptical hypotheses and thus retain closure *and* everyday knowledge. In effect, DeRose weakens the blow of allowing this knowledge by making it such that the knowledge can never be put to any conversational use, since in making it explicit to a conversational context we thereby raise the epistemic standards and destroy the knowledge. In making this claim it is thus incumbent upon DeRose to reject sensitivity as a necessary condition on knowledge (although it does retain an important epistemic role), and cash-out knowledge instead in terms of the sufficiency of an agent's epistemic position. If, as a matter of fact, one's belief that one is not a BIV is such that it matches the truth across the range of possible worlds relevant to that context, then one knows that proposition.

It is important to emphasise at this point the necessity that this be an *externalist* proposal since if it were demanded that the subject should have reflective access to the relevant factual conditions that supports her knowledge then the thesis ceases to be at all plausible. After all, as we saw above, it is only by incorporating externalism that *any* epistemological theory, contextualism included, can make sense of an agent's knowledge of the denials of sceptical hypotheses (in *any* context). Moreover, provided closure is to be retained, without this knowledge agents would lose their knowledge of even everyday propositions. The contextualist anti-sceptical thesis thus goes hand-in-hand with externalism.<sup>16</sup>

There are, however, a number of striking problems that this account faces, the most pressing of which is that, on this view, knowledge possession can be dependent upon purely conversational factors.<sup>17</sup> Intuitively, whether or not one knows is a matter fixed by concrete environmental circumstances, not by mere changes in the conversation. Indeed, one might think that, all other things being equal, one should prefer a thesis which views knowledge possession as independent of conversational factors over one that doesn't. This point is important because what does seem to be influenced by conversational change is the propriety of *claims* to know, and so we have grounds to wonder whether or not it would be better to consider contextualism as a thesis about how it is that claims to know, rather than the possession of knowledge itself, are affected by conversational factors.

This concern is significant because it impacts upon the Moorean aspect of the thesis—that it allows knowledge of the denials of sceptical hypotheses (at least in some contexts). Recall, that the Dretske-Nozick argument for non-closure was undermined by the fact that it endorsed an externalist thesis that called into question one of the central claims that drove that brand of epistemological revisionism—that we could not know the denials of radical sceptical hypotheses. The DeRose brand of revisionism takes this worry further by actually endorsing an account of knowledge wherein one might know the denials of these propositions, but only in particular conversational contexts where that knowledge is never explicitly considered (and therefore never claimed). If one is willing to allow that knowledge of this sort is possible, however, then the motivation for denying it in some

conversational contexts starts to wane. The thought might run that although it may well be true that one could not, for example, *claim* to know that one was not a BIV, still, if one really does know everyday propositions (so that sceptical possible worlds really are far-fetched), then one could indeed be in a position to know the denial of a sceptical hypothesis in *any* conversational context. That is, if one is not independently convinced of the thesis that mere conversational factors can influence knowledge possession, then the temptation is to take the contextualist model but apply it to knowledge claims rather than to knowledge possession and thereby use it to resurrect the Moorean proposal. I shall consider how such a strategy might run in more detail below. First, I want to consider a neo-Moorean argument, due to Sosa (1999), which marks the beginnings of just this thought.

#### IV

Sosa's response to the Nozick-Dretske line and the DeRose variant on contextualism is to reject the notion of sensitivity, both as a necessary condition on knowledge and as a mechanism for raising and lowering contextual standards. Instead, he proposes the following notion of "safety" which he does take to be necessary for knowledge (Sosa 1999, 142):

*Safety*  
 $B[\varphi] \Rightarrow \varphi$

In words, this says that a subject's belief in a contingent proposition  $\varphi$  is safe if, and only, if, for a wide-range of near-by possible worlds, if the subject believes  $\varphi$ , then  $\varphi$  is true.<sup>18</sup> The anti-sceptical advantage that this proposal offers over both the Dretske-Nozick line and the DeRose variant on contextualism is that it allows Sosa to endorse a fully-fledged version of the Moorean proposal which neither issues in the denial of closure nor results in contextualism.<sup>19</sup>

Take the former point first. What prompted the denial of closure was the fact that, if we take sensitivity as a necessary condition on knowledge, then we must allow that a subject can know everyday propositions whilst being unable to know all the known consequences of those everyday propositions—i.e., the denials of sceptical hypotheses. In contrast, on this view we can allow both that agents have knowledge of everyday propositions and that they can know the denials of radical sceptical hypotheses. Suppose, for example, that the agent does have a safe belief in an everyday proposition. Not only is this belief true in the actual world, but, across the range of near-by possible worlds where she believes this proposition, it is true there as well. Insofar as this belief really is safe,

however, then there will not be any sceptical possible worlds in the realm of near-by possible worlds which determine that safety (henceforth, the “realm of safety”). For if there were such worlds, then the agent would have a belief in an everyday proposition, such as that she is currently sitting down at her computer, which, because she is in fact a BIV, is false. But since sceptical possible worlds are now excluded from the realm of safety, it follows that the agent must also have a safe belief that she is not a BIV (or, indeed, the victim of any sceptical hypothesis). The reason for this is that there will be no possible world within the realm of safety in which this proposition is false, and thus, in every world in the realm of safety in which she believes this proposition (which, I take it, is all of them), her belief is true. Accordingly, scepticism is evaded and closure, as least as it functions in sceptical and anti-sceptical reasoning, is retained.

Moreover, the adoption of safety as a necessary condition on knowledge is also able to speak to the informal worry that we raised about closure above. As expressed there, the concern was that I should be able to know everyday propositions without having to rule-out the possibility of a far-fetched sceptical scenario. The Dretske-Nozick construal of this worry is that I should be able to know everyday propositions whilst simultaneously being unable to know the denials of sceptical hypotheses, but this is not the only way to read this concern. Alternatively, the thought might be that whatever influences one’s knowledge is only the relevant range of near-by possible worlds, not also worlds far away. Accordingly, if the sceptical possibility is indeed far-fetched then it ought to be unable to influence my knowledge of everyday propositions or, indeed, my knowledge of the denials of sceptical hypotheses. Safety captures this concern by allowing agents to have knowledge in both cases provided sceptical possible worlds do not feature in the realm of safety. Sensitivity, in contrast, violates this intuition by making knowledge dependent not just on the relevant circumstances in near-by worlds but also on the circumstances that obtain in far-off worlds (such as sceptical worlds) where the target proposition is false.

Furthermore, since the realm of safety does not vary in response to mere conversational factors, it follows that this is not a contextualist thesis. If the agent does indeed know everyday propositions then, in line with the central Moorean contention, she will also know the denials of radical sceptical hypotheses, and this will be so no matter what conversational context the agent is in. We thus have a Moorean variety of anti-scepticism which, whilst keeping to the externalist spirit of both the Nozick-Dretske and DeRose proposals, lacks the epistemological revisionism of either.

There are problems with this view, however, the most pressing of which is how it explains why it is that it seems so implausible to actually contend that we do know the denials of radical sceptical hypotheses. Note that this is not necessarily the worry about lacking good reflectively accessible evidence raised above, because, I take it, we all have the brute intuition, regardless of our broader epistemological prejudices, that there is at least *something* wrong about claiming knowledge of the denials of sceptical hypotheses. Sosa (1999, 147-8) himself tries to evade this difficulty by claiming that people just tend to mix-up sensitivity with safety. The disadvantage of this proposal is that it seems to imply that people *do* view (albeit wrongly) sensitivity as a condition of knowledge, and this can only serve to undermine the general style of anti-scepticism that Sosa is proposing. More work thus needs to be done to make this proposal appealing by showing how it can account for our Moorean intuition that we must know these propositions whilst also accommodating our sceptical intuition that such knowledge is, at best, atypical because it can never be legitimately claimed.

The first thing to note is that we can strengthen the notion of safety that Sosa proposes whilst retaining the dialectical gain that it offers. Accordingly, I propose the following rough characterisation of “super-safety”, which, following Sosa, I take to be a necessary condition of knowledge:

*Super-safety*

An agent has a *super-safe* belief in the contingent proposition  $\varphi$  if, and only if that agent has a true belief in  $\varphi$  and her belief in  $\varphi$  ‘tracks’ the truth as to whether  $\varphi$  across a wide-range of near-by possible worlds (i.e., where  $\varphi$  is true, she believes it; where  $\varphi$  is false, she does not believe it).<sup>20</sup>

What this formulation adds to Sosa’s notion of safety is the contention that knowledge does not just require an agent’s belief to track the truth across the realm of safety, but also that the agent’s *non*-belief should do likewise. On Sosa’s account, one might have a safe belief in a proposition even though there is a near-by possible world in which, whilst the proposition in question is true, one does not believe it. In contrast, the notion of super-safety rules out such a possibility, and therefore constitutes a more demanding account of knowledge.

I take it that Sosa did not incorporate this extra requirement into his notion of safety because it looks dangerously similar to the sensitivity requirement proposed by Nozick and Dretske. That condition was only problematic, however, because it demanded that we must consider the nearest possible world in which the proposition the agent believes is false in order to determine whether or not that agent knows, no matter how far-out, modally speaking, that world was. In contrast, this notion of super-safety evades this difficulty by only demanding that the subject’s non-belief be

responsive to the truth in a fixed range of near-by possible worlds.

Indeed, the notion of super-safety incorporates the best elements of both the Sosa and the DeRose proposals. On the one hand, it incorporates Sosa's notion of safety. On the other, it strengthens that notion by making the 'tracking' relation in question 'two-way' in the manner that DeRose envisaged, so that the agent's belief has to be responsive both to the truth and to the falsity of the proposition believed across a range of possible worlds. Where it departs from the DeRose model, however, is by following Sosa in taking the realm of possible worlds at issue as fixed to near-by worlds. Accordingly, if one really does know a proposition, then one knows it *period*, and changes in mere conversational factors cannot alter this fact.

We are thus able to employ a more robust non-contextualist conception of knowledge in our presentation of the Moorean thesis without thereby falling foul of sceptical worries or being committed, *à la* Dretske and Nozick, to arguing for non-closure. And if we have a more demanding epistemic notion in play, then we are in a better position to regard this externalist thesis as being able to actually capture our *knowledge* of the denials of radical sceptical hypotheses.

Nevertheless, it might be thought that even with this strengthened notion of safety in mind there is still something problematic about the claim that we can know the denials of anti-sceptical propositions. Recall how the phenomenology of our engagement with scepticism was characterised above in terms of the following three features:

- (P1) Ascriptions of knowledge to subjects in conversational contexts in which sceptical error-possibilities have been raised is wholly inappropriate.
- (P2) Ascriptions of knowledge to subjects in conversational contexts in which no sceptical error-possibilities are in play can be entirely appropriate.
- (P3) All that may change when one moves from a 'non-sceptical' conversational context to a 'sceptical' context are mere conversational factors.

Insofar as these three claims do indeed capture the phenomenology of our engagement with scepticism, then they seem to be, *prima facie* at least, in conflict with the Moorean approach sketched here. For although the Moorean line (*contra* DeRose), holds to (P3), it seems to lack any means of explaining (P1) and (P2). Why is it that we don't tend to ascribe knowledge in sceptical contexts that we are happy to ascribe in everyday contexts? Moreover, why is that we are disinclined to ascribe knowledge of the denials of radical sceptical hypotheses in *any* conversational context?

I think that the answer to these questions lies in appreciating how DeRose's contextualism can offer an account of how it is that certain knowledge ascriptions can *seem* wrong in particular conversational contexts even though the ascription is fact correct. This requires us to reflect on what is involved in making a legitimate *claim* to know, and the manner in which the conversational implicatures generated by a claim to know can be sensitive to changing features in one's

conversational context. In what follows, we will take a conversational implicature to be any inference that one is entitled to draw upon hearing an assertion (but which is not entailed by the assertion itself), given that one may legitimately make certain assumptions about the agent making that assertion—that she is, for instance, honest, co-operative and (at least otherwise), rational.<sup>21</sup> In what follows I will simply take it that any assertion that carries a false conversational implicature is improper (even though the assertion itself may of course be true).

One conversational feature of a claim to know which holds in all conversational contexts is that it implies a readiness to offer grounds to support one's claim—it puts one's own personal stamp on the proposition and indicates a willingness to support the truth of that proposition if necessary. One consequence of this feature is that any proposition that cannot be evidentially grounded cannot be properly claimed to be known since it would generate a false conversational implicature. As Wittgenstein puts the point:

One says “I know” when one is ready to give compelling grounds. “I know” relates to a possibility of demonstrating the truth. Whether someone knows something can come to light, assuming that he is convinced of it.

But if what he believes is of such a kind that the grounds he can give are no surer than his assertion, then he cannot say that he knows what he believes. (Wittgenstein 1969, §243)<sup>22</sup>

A claim to know carries with it the implicature that one is able to offer evidence to support that claim. Accordingly, if no evidence is available to the subject (or if the evidence is “no surer than his assertion”), then the claim to know is improper. And if this thesis is accepted then it follows that any claim to know the denial of a radical sceptical hypothesis will always be improper. Note, however, that this fact need not prejudice the agent's *possession* of knowledge of these propositions, since a lack of adducible evidence does not, on an externalist account at least, entail a lack of knowledge. It is thus entirely plausible that one could know the denials of radical sceptical hypotheses whilst never being in a position to properly claim knowledge of them.<sup>23</sup>

This feature of our conversational practices can therefore explain why, in part, it is that we are so uncomfortable with the Moorean *claim* that we know the denials of sceptical hypotheses, even if we do feel that these are propositions which, if we know anything much at all, we must know. However, although it helps to account for why the Moorean approach to scepticism can seem so counterintuitive, it does not fully explain the first two features of the phenomenology of scepticism noted above. After all, in itself this observation does nothing to account for the apparent ‘context-relativity’ of knowledge.

Nevertheless, it does provide us with the means to lever a Moorean interpretation of these phenomena into place. To begin with, it needs to be noted that this observation that one can never

claim to know the denials of sceptical hypotheses presents us with a case in which the second feature of the phenomenology of scepticism, (P2), does not apply. For although it is true that we are typically content to allow ascriptions of knowledge in everyday conversational contexts, this does not extend to ascriptions of knowledge that sceptical hypotheses have not obtained. And note, this does not just apply to *self*-ascriptions of knowledge, but also to ascriptions of knowledge to others. Since there is, in principle, an absence of adducible evidence available in this regard, an ascription of knowledge of this sort of proposition to a third person would carry false implicatures in just the same way as a self-ascription would. After all, it is not as if anyone else could be in a significantly better epistemic position regarding this hypothesis than we are.

What makes this point important is the manner in which it inter-connects with the further claim that the range of conversational implicatures that are generated by a claim to know can be influenced by the conversational context and, in particular, by whether or not sceptical hypotheses are at issue in that context.

Consider the following scenario. Bill knows that Jack is out of town today because he saw him off at the airport this morning. Moreover, let us take it that this is a genuine case of knowledge, and thus that it is super-safe in the manner outlined above. *Ceteris paribus*, then, Bill can claim to know that Jack is out of town with impunity. In such everyday contexts, however, making this claim to know does not carry with it very strenuous demands. It demands that Bill be able to adduce evidence which is sufficient to rule-out certain mundane possibilities of error, such as that it was Jack, and not his brother Barney, that he took to the airport, but little more than that. And since Bill has excellent grounds to support his claim to know, such demands do not present him with any difficulty.

Suppose, however, that we raise the conversational standards. Imagine that the following conversational context arises. Bill and his work colleagues James and David are standing by the coffee machine idly discussing the movie from last night. The element of the film that they are talking about concerns a scene in which an everyday guy with murderous intent fakes his departure onto a plane by having a body-double hide in the airport lounge and swap places with him as he walks down the departures corridor. Suppose now that, with this possibility in mind, James and David ask whether Bill really knows that Jack flew out of town this morning (as opposed to a body-double that Jack hired to stand in his place). In this conversational context, can Bill still legitimately claim to know that Jack is out of town?

The problem with this assertion is that, in this conversational context, it appears to carry the conversational implicature that Bill is able to offer evidence that not only rules out mundane possibilities of error (such as that it was Barney that he dropped off at the airport), but which also

rules out quasi-sceptical possibilities of error, such as that it was a body-double that got on that plane. The trouble is, however, this seems like the very sort of evidence that Bill lacks. So even though he does know the proposition in question, in this conversational context he cannot properly claim that knowledge because such a claim carries with it a false conversational implicature. Note that, at least on the account sketched here, this does not mean that Bill does not know this proposition in this conversational context, only that he cannot properly *claim* such knowledge in this context.

Moreover, what is true of faked departures will be even more applicable in the conversational contexts in which radical sceptical hypotheses have been raised. If one makes an everyday claim to know in a conversational context in which sceptical hypotheses are under discussion, then that claim to know carries with it the conversational implicature that one is able to adduce evidence which would rule-out the hypotheses in question. And since such evidence is, in principle, unavailable, it follows that one is unable to properly claim knowledge of even everyday propositions in ‘sceptical’ conversational contexts. We can thus account for the two aspects of the phenomenology of scepticism noted above. First, that, at least in general, we are willing to ascribe knowledge to agents in everyday contexts. And, second, that we are unwilling to ascribe such knowledge in ‘sceptical’ conversational contexts. Furthermore, we can account for these features of the phenomenology of scepticism without endorsing the epistemological revisionism of contextualism.

Moreover, this feature of our conversational practices also explains the weakest element of the Moorean stance, which is the oddity of Moore’s own response to scepticism. After all, what Moore did was straightforwardly *claim* that he knew lots of everyday (and even philosophical) propositions in a conversational context in which radical scepticism was being explicitly discussed. As a result, his claims to know carried false implicatures and therefore were improper even though they were, one would hope at any rate, entirely true.

## VI

We can thus do a great deal of work to make the Moorean anti-sceptical proposal palatable by making two key adjustments to the version of the thesis advocated by Sosa. First, by enhancing the formulation of knowledge at issue to make it more robust. And, second, by paying due attention to certain salient features of our conversational practices so as to account for the aspects of our use of epistemic terms that the contextualist focuses upon. With these adjustments in play, the Moorean anti-sceptical strategy, far from being a relic of a distant epistemological debate, is still a live option.<sup>24</sup>

## REFERENCES

- Austin, J. L. (1961). 'Other Minds', reprinted in his *Philosophical Papers*, (eds.) J. O. Urmson & G. J. Warnock, 44-84, Clarendon Press, Oxford.
- Baldwin, T. (1990). *G. E. Moore*, Routledge, London.
- Baldwin, T. (1993). 'G. E. Moore', *A Companion to Epistemology*, (eds.) J. Dancy & E. Sosa, 283-5, Basil Blackwell, Oxford.
- Bonjour, L. (1980). 'Externalist Theories of Justification', *Midwest Studies in Philosophy* 5, 27-52.
- Bonjour, L. (1987). 'Nozick, Externalism, and Skepticism', *The Possibility of Knowledge: Nozick and his Critics*, (ed.) S. Luper-Foy, chapter 11, Rowman & Littlefield, Totowa, New Jersey.
- Brandom, R. (1998). 'Insights and Blindspots of Reliabilism', *The Monist* 81, 371-92.
- Brueckner, A. (1991). 'Unfair to Nozick', *Analysis* 51, 61-4.
- Chisholm, R. M. (1989). *Theory of Knowledge*, (3rd ed.), Prentice Hall, Englewood Cliffs, New Jersey.
- Cohen, S. (1986). 'Knowledge and Context', *Journal of Philosophy* 83, 574-83.
- Cohen, S. (1987). 'Knowledge, Context, and Social Standards', *Synthese* 73, 3-26.
- Cohen, S. (1998). 'Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery', *Australasian Journal of Philosophy* 76, 289-306.
- Cohen, S. (1999). 'Contextualism, Scepticism, and the Structure of Reasons', *Philosophical Perspectives* 13, 57-90.
- Craig, E. (1989). 'Nozick and the Sceptic: The Thumbnail Version', *Analysis* 49, 161-2.
- DeRose, K. (1995). 'Solving the Sceptical Puzzle', *Philosophical Review* 104, 1-52.
- DeRose, K. (2000). 'How Can We Know That We're Not Brains in Vats?', *The Southern Journal of Philosophy* 38, 121-48.
- Dretske, F. (1970). 'Epistemic Operators', *Journal of Philosophy* 67, 1007-23.
- Dretske, F. (1971). 'Conclusive Reasons', *Australasian Journal of Philosophy* 49, 1-22.
- Feldman, R. (1999). 'Contextualism and Skepticism', *Philosophical Perspectives* 13, 91-114.
- Fogelin, R. (2000). 'Contextualism and Externalism: Trading in One Form of Skepticism for Another', *Philosophical Issues* 10, 43-57.
- Fumerton, R. (1988). 'The Internalism/Externalism Controversy', *Philosophical Perspectives* 2, 443-59.
- Grice, H. P. (1989). 'Logic and Conversation', reprinted in *Studies in the Way of Words*, Harvard University Press, Cambridge, Massachusetts.
- Goldman, A. (1980). 'The Internalist Conception of Justification', *Midwest Studies in Philosophy* 5, 27-51.
- Heller, M. (1999a). 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy* 77, 196-208.
- Heller, M. (1999b). 'The Proper Role for Contextualism in an Anti-Luck Epistemology', *Philosophical Perspectives* 13, 115-29.
- Kornblith, H. (2000). 'The Contextualist Evasion of Epistemology', *Philosophical Issues* 10, 24-32.
- Lehrer, K. (2000). 'Sensitivity, Indiscernibility and Knowledge', *Philosophical Issues* 10, 33-7.
- Lewis, D. (1979). 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* 8, 339-59.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.
- Moore, G. E. (1925). 'A Defence of Common Sense', *Contemporary British Philosophy* (2nd series), (ed.) J. H. Muirhead, Allen and Unwin, London.
- Moore, G. E. (1939). 'Proof of an External World', *Proceedings of the British Academy* 25.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford University Press, Oxford.
- Pritchard, D. H. (2000a). 'Closure and Context', *Australasian Journal of Philosophy* 78, 275-80.
- Pritchard, D. H. (2000b). 'Understanding Scepticism', *Sats—Nordic Journal for Philosophy* 1, 107-24.

- Pritchard, D. H. (2001a). 'Contextualism, Scepticism, and the Problem of Epistemic Descent', *Dialectica* 55, 327-49.
- Pritchard, D. H. (2001b). 'Meta-Epistemological Constraints on Anti-Sceptical Theories', *Facta Philosophica* 3, 101-26.
- Pritchard, D. H. (2001c). 'Radical Scepticism, Epistemological Externalism, and "Hinge" Propositions', *Wittgenstein-Studien*, 81-105.
- Pritchard, D. H. (2001d). 'The Opacity of Knowledge', *Essays in Philosophy, Vol. 2: The Internalism/Externalism Debate in Epistemology*, (ed.) H. Benjamin Schaeffer, 1-10, Humboldt University Press, Humboldt, California.
- Pritchard, D. H. (2002a). 'McKinsey Paradoxes, Radical Scepticism, and the Transmission of Knowledge across Known Entailments', *Synthese* 130, 279-302.
- Pritchard, D. H. (2002b). 'Radical Scepticism, Epistemological Externalism and Closure', *Theoria* 68, 129-61.
- Pritchard, D. H. (2002c). 'Two Forms of Epistemological Contextualism', *Grazer Philosophische Studien* 64, 97-134.
- Pritchard, D. H. (2002d). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* 39, 215-57.
- Pritchard, D. H. (2004). 'The Structure of Sceptical Arguments', *The Philosophical Quarterly* 54.
- Pryor, J. (2001). 'Epistemology c. 1988-2000', *The British Journal for the Philosophy of Science* 52, 95-124.
- Schiffer, S. (1996). 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society* 96, 317- 33.
- Sosa, E. (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-53.
- Sosa, E. (2000a). 'Skepticism and Contextualism', *Philosophical Issues* 10, 1-18.
- Sosa, E. (2000b). 'Replies', *Philosophical Issues* 10, 38-42.
- Stroll, A. (1994). *Moore and Wittgenstein on Certainty*, Oxford University Press, Oxford.
- Stroud, B. (1984). *The Significance of Philosophical Scepticism*, Oxford University Press, Oxford.
- Tomberlin, J. (2000). 'Skepticism, Tracking, and Warrant', *Philosophical Issues* 10, 19-23.
- Vogel, J. (1999). 'The New Relevant Alternatives Theory', *Philosophical Perspectives* 13, 155-80.
- Williams, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Basil Blackwell, Oxford.
- Wittgenstein, L. (1969). *On Certainty*, (eds.) G. E. M. Anscombe & G. H. von Wright, (tr.) Denis Paul & G. E. M. Anscombe, Basil Blackwell, Oxford.
- Wright, C. (1985). 'Facts and Certainty', *Proceedings of the British Academy* 71, 429-72.

## NOTES

<sup>1</sup> More precisely: I know that if I have two hands then there are two 'external' objects in the world, and thus that there is an external world.

<sup>2</sup> It is actually a moot point whether Moore saw himself as answering the sceptic at all, let alone offering an *argument* against scepticism (cf. Baldwin 1990; 1993), but many commentators, including, most famously, Wittgenstein (1969), have conceived of Moore's remarks in this way. Of the two cited papers upon which Moore's reputation as an anti-sceptic was largely established (Moore 1925; 1939), the first was written in response to an invitation to provide a *personal statement* (and therefore not necessarily a *validation*) of his position, whilst the second is regarded by some as being an attempt at a refutation of *idealism*, not scepticism, in that Moore was only aiming to prove the *existence* of an external world, not also that we have any *knowledge* of it. (I leave it to the reader to decide whether this distinction is entirely plausible). To add to the confusion, in the latter paper Moore notes that although he knows the premises of his 'proof', he cannot prove them (Moore 1939, 150). Nevertheless, the interpretation of Moore offered above is widely endorsed. See Stroud (1984, chapter 3); Wright (1985); Williams (1991, 40-5); Stroll (1994, *passim*); DeRose (1995; 2000); and Sosa (1999).

<sup>3</sup> Note, however, that in calling the response ‘Moorean’ it is not thereby being claimed that it would be a response that Moore himself would endorse, only that it is sufficiently similar to the basic anti-sceptical line that Moore offers.

<sup>4</sup> For the sake of simplicity, the reader should take every instance of the ‘K’ operator to be indexed to a particular agent and a particular time. Since the sceptical argument we will be considering is meant to work with (just about) any agent and (just about) any time, we can ignore these complications without loss.

<sup>5</sup> In order to avoid philosophically uninteresting objections to this principle which rest upon the subject’s lack of belief in the consequent proposition, I shall simply assume that the agent does form the requisite beliefs. For further discussion of closure and its role in sceptical arguments, see Pritchard (2002b; 2004).

<sup>6</sup> For more on the theoretical burdens incumbent upon anti-sceptical theses, see Pritchard (2001b). For a survey of recent work on the sceptical problem, see Pritchard (2002d).

<sup>7</sup> It is of course important that sensitivity, as a requirement on knowledge, should only apply to contingent propositions since necessary propositions are not false in *any* possible world. The reader should also note that I’m simplifying somewhat in my account of the Dretske-Nozick conception of sensitivity. Dretske adds an extra story about the means by which the agent forms the belief, the “epistemic credential” that gives rise to the belief. Nozick goes further, relativizing his account of knowledge to a method, and adding an extra condition on knowledge. Nevertheless, despite these simplifications, the principle described is at the core of the proposal and, furthermore, it is this element of the view that results in the denial of closure. Accordingly, nothing should be lost by focussing upon this formulation.

<sup>8</sup> This is a widely held view of the internalist thesis. In a recent survey article on contemporary epistemology, for example, Pryor (2001) argues that this conception of internalism constitutes the “core internalist position”. For an example of an internalist view which proceeds along these lines, see Chisholm (1989), chapter one. For more on the internalism/externalism distinction itself, see Bonjour (1980); Goldman (1980); Fumerton (1988); and Pritchard (2001d).

<sup>9</sup> Craig (1989; cf. Bonjour 1987) pursues this line of argument as it applies to Nozick’s use of the denial of closure for anti-sceptical ends. For discussion, see Brueckner (1991) and Pritchard (2000b).

<sup>10</sup> For more on this claim, see Brandom (1998).

<sup>11</sup> For the main texts in the recent literature which endorse this general contextualist anti-sceptical strategy, see Cohen (1986; 1987; 1998; 1999); Lewis (1996); and Heller (1999b). Note that this is not the only way to cash-out the basic contextualist thesis. For example, in Pritchard (2002c), I distinguish between this contextualist view, which I label ‘semantic contextualism’, and a different model developed by Williams (1991) that I call ‘inferential’ contextualism. I offer a critique of the former version of contextualism in Pritchard (2001a).

<sup>12</sup> Lewis agrees. In Lewis (1979) he briefly applies this line of argument to epistemic terms (and thus to the sceptical problem), and expands upon this thought in Lewis (1996).

<sup>13</sup> The issue of what is involved in bringing about a change of conversational context is still somewhat controversial. Most contextualists agree, however, that merely mentioning a sceptical hypothesis will, at least sometimes, suffice.

<sup>14</sup> For more on the relationship between arguments for contextualism and arguments for non-closure, see the exchange between Heller (1999a) and Pritchard (2000a).

<sup>15</sup> Indeed, in more recent work, DeRose (2000) refers to his position as ‘Moorean contextualism’ precisely because it allows, albeit only in everyday conversational contexts, that one can know the denials of radical sceptical hypotheses. Interestingly, in this paper DeRose tries to defend this move by arguing that such knowledge is *a priori*, buttressing this contention by arguing that a similar conclusion follows from some versions of semantic externalism. If anything, however, this only makes the situation worse since it ties the fate of a contentious epistemological theory to an even more contentious putative consequence of a semantic thesis. Moreover, if one endorses an externalist epistemology, and DeRose certainly does, then it is hard to see why such knowledge can only be accounted for in an *a priori* fashion. That one’s knowledge lacks reflectively accessible empirical grounds does thereby not make it *a priori* on the externalist account, since there is still the possibility that it rests on *non*-reflectively accessible empirical grounds. For more on the relationship between the sort of radical sceptical arguments under consideration here and semantic externalism, see Pritchard (2002a).

<sup>16</sup> See Heller (1999b) for a development of this claim. Although most contextualists do endorse some version of externalism (see, for example, Lewis (1996)), one prominent defender of the contextualist position, Cohen (1999), has argued for an internalist reading of this thesis. The union is, at best, awkward, since it depends upon a particular construal of internalism that allows one’s evidential knowledge to rest upon non-evidentially known foundations. Indeed, ultimately, Cohen defends the contextualist claim that one can know that one is not the victim of a sceptical hypothesis on broadly *pragmatic*, rather than *epistemic*, grounds. It is thus at least questionable that the epistemological position on offer, if coherent at all, bears any of the standard hallmarks of epistemological internalism.

<sup>17</sup> A burgeoning industry has sprung up of critical appraisals of the contextualist position. For some of the main texts in this regard, see Schiffer (1996); Feldman (1999); Vogel (1999); Fogelin (2000); and Pritchard (2001a).

<sup>18</sup> Sosa (1999, 142) himself actually offers three different characterisations of this notion. I take the formulation offered here to capture the spirit of all three. For (slightly) more on Sosa’s view in this respect, see Sosa (2000a). For critique of this latter paper, and Sosa’s response, see Kornblith (2000); Lehrer (2000); Sosa (2000b); and Tomberlin (2000).

<sup>19</sup> Or, at least, it does not lead to either contextualism or the denial of closure insofar as these theses are understood as anti-sceptical views. There may be reasons independent of the sceptical problem for conjoining a safety thesis with either the denial of closure or a contextualist view, but they are not relevant here. The central point is that closure, insofar as it figures in sceptical arguments, is consistent with safety, as is the sceptic's denial of contextualism.

<sup>20</sup> Even construed merely as a necessary condition for knowledge, this formulation needs some work. To take one problem that this formulation faces, consider the following scenario (adapted from one used by Nozick (1981)). Suppose that I am in my office and, as luck would have it, I happen to look up and see a burglar, whom I recognise, exiting the window opposite. As a result, intuitively at least, I come to know who burgled the office across the courtyard. There is, however, a nearby possible world where I don't happen to look up at that moment, and thus where I don't form the required belief. As things stand, it would be a consequence of the formulation of super-safety under consideration here that I fail to know who the burglar is because there is a near-by possible world where I lack the belief even though it is true. Accordingly, we need to do some tinkering to the formulation of super-safety in order account for these cases. There are a number of ways of doing this (e.g., one could, following Dretske (1970), restrict the range of relevant possible worlds to those where the "epistemic credential" that gave rise to the belief is the same as in the actual world), though I will spare the reader such an involved and technical discussion. After all, my primary aim in advancing this notion is to convince the reader of the plausibility of the proposal in general, not to present a detailed account of knowledge.

<sup>21</sup> The *locus classicus* for this account of conversational implicature is, of course, Grice (1989). As Grice notes, an agent can always 'cancel' a conversational implicature of her assertion by simply accompanying that assertion with a suitable *caveat*. In what follows, I shall only be considering claims to know the conversational implicatures of which have not been cancelled in this way.

<sup>22</sup> One finds a similar contention in Austin (1961). Consider, for example, the following passages:

When I say, 'I know', I give others my word: I *give* others *my authority for saying* that 'S is P'. [...]

If you say you know something, the most immediate challenge takes the form of asking, 'Are you in a position to know?': that is, you must undertake to show not merely that you are sure of it, but that it is within your cognisance. (Austin 1961, 99-100, *italics in the original*)

<sup>23</sup> It is, of course, an interesting question why a legitimate claim to know can demand more of an agent than simply knowing does. Why, for example, can it be true that an agent has knowledge which she is unable to legitimately claim to possess? One possible explanation of this can be found by examining the social transmission of knowledge. For example, many of us are quite happy to ascribe basic perceptual knowledge to small children even though they lack the cognitive capacities to reconstruct the grounds on which that knowledge is based. For this reason, however, we also tend to regard their claims to know as being illegitimate. Since children lack the cognitive resources to reconstruct the grounds for their knowledge, it follows that, from their point of view, they lack good reasons for thinking that they do know. Accordingly, whilst children might be *de facto* reliable transmitters of true belief, their claims to know are, in general, unreliable. Wittgenstein (1969) offers a subtle account of the rules that govern claims to know. For discussion, see Pritchard (2001c).

<sup>24</sup> Thanks to three anonymous referees from this journal for useful comments, and to Patrick Greenough, Jesper Kallestrup, Patrice Philie, Michael Williams and Crispin Wright for discussions on previous versions of this material. An earlier version of this paper was presented both at the University of St. Andrews and the University of Glasgow, and I am grateful to the audiences on those occasions, especially Bob Hale, Jimmy Lenman, Paul Markwick, Andrew McGonigal, Philip Percival and Stewart Shapiro.