

TWO FORMS OF EPISTEMOLOGICAL CONTEXTUALISM

Duncan PRITCHARD
University of Stirling, Scotland

Summary

The recent popularity of contextualist treatments of the key epistemic concepts has tended to obscure the differences that exist between the various kinds of contextualist theses on offer. The aim of this paper is to contribute towards rectifying this problem by exploring two of the main formulations of the contextualist position currently on offer in the literature—the ‘semantic’ contextualist thesis put forward by Keith DeRose and David Lewis, and the ‘inferential’ contextualist thesis advanced by Michael Williams. It is argued that by evaluating these theses in the light of each other one can gain a deeper understanding of the contextualist position. In particular, it is argued that this relative evaluation highlights one interesting way in which contextualism might be developed.

0. *Introduction*

Contextualist epistemological theories—theories which hold that key epistemic concepts like ‘knowledge’ and ‘justification’ are context-relative—are, without doubt, the new vogue in epistemology. The popularity of this general approach belies, however, the tremendous degree of difference that exists between various formulations of this thesis. For example, the contextualist story found in the contemporary

literature is firmly lodged within the debate concerning radical scepticism that arose in the decades after 1970. These discussions take as their reference point the so-called ‘relevant alternatives’ (RA) account of knowledge offered by Fred Dretske (1970; 1971; 1981) and the later more explicitly modal account of knowledge put forward by Robert Nozick (1981). Earlier forms of contextualism, in contrast, make little or no appeal to the contextualist response to radical scepticism, much less express their position in contrast to the anti-sceptical proposals made by Dretske and Nozick.¹ Another axis of difference within the contextualist camp is the commitment to an externalist epistemology. Most forms of contextualism take this commitment for granted, but more recent versions of the view have challenged this assumption.²

Although a taxonomy of all the main contextualist theses would be useful (and long overdue), I will not be undertaking this mammoth task here.³ Instead, I will be examining a third axis of difference between contextualist theories, one that I think is highly significant but which has been almost completely overlooked. This contrast is between those contextualist theses that focus on the role of conversational contexts in the determination of epistemic contexts, and those that instead individuate epistemic contexts in terms of inferential structure. One finds versions of the first view, which I here call ‘semantic’ contextualism, in work by Keith DeRose (1995) and David Lewis (1996). The most thorough presentation of the second view, which I here call ‘inferential’ contextualism, is that proposed by Michael Williams (1991). It is my contention that by contrasting these two views one can gain a deeper appreciation of the contextualist thesis and thus secure a better sense of both the advantages and the difficulties posed by this type of epistemological theory.⁴

1. The most obvious example in this respect is Annis (1978). For discussion, see Airaksinen (1982; cf. Annis 1982) and Henderson (1994).

2. See, in particular, Cohen (1999). Space does not allow a thorough discussion of Cohen’s view, interesting though it is. This is especially unfortunate given that the position which he develops is, in essence, simply a variation on the semantic model that I discuss below. Nevertheless, it is sufficiently different to merit a separate treatment and so I won’t be discussing it here. For the main papers in which Cohen expresses his view, see Cohen (1986; 1988; 1991; 1998a; 1999; 2000a). For discussion, see the recent critical notices by Hawthorne (2000); Klein (2000); and Prades (2000); and the reply by Cohen (2000b).

3. Norman (1999) comes closest to offering a survey of the main types of contextualist accounts, both historically and in the contemporary literature.

4. For the sake of brevity, I will not be discussing here the broader question of whether

In §1, I discuss the RA account of knowledge and explain how such a view can be regarded as a ‘proto-contextualist’ thesis. In §2, I offer an exposition of the DeRose-Lewis semantic contextualist theory and, in §3, I describe, in the light of the DeRose-Lewis model, the opposing inferential contextualist thesis put forward by Williams. In §4, I offer a relative evaluation of the two theories, arguing that Williams’ view has advantages over the semantic contextualist view. In §5, I draw upon this relative evaluation to outline one potential way in which inferential contextualism could incorporate core semantic contextualist claims. I further contend that such a ‘two-component’ contextualist position would possess a key advantage that is lacking in the inferential theory taken by itself. Finally, in §6, I offer some concluding remarks.

1. *From Relevant Alternatives To Contextualism*

Following DeRose (1995), Ernest Sosa (1999), Jonathan Vogel (1999) *et al*, I will understand the radical sceptical paradox to consist in the joint incompatibility of three claims, each of which appears, on the surface of things and taken individually, to be perfectly in order.⁵ Take ‘K [φ]’ to say that an agent knows the proposition φ;⁶ take ‘SH’ to refer to some suitable sceptical hypothesis, such as the hypothesis that one might be brain-in-a-vat (BIV) being ‘fed’ one’s experiences by computers; and understand ‘O’ to be some ‘ordinary’ proposition that one would typically take oneself to know and which entails the falsity of the sceptical hypothesis in question (such as that one has two hands). We can then formulate these three incompatible claims as follows:

- (1) K [O]
- (2) ¬ K [¬ SH]
- (3) ¬ K [¬ SH] → ¬ K [O]

we should endorse *any* contextualist thesis, semantic, inferential, or otherwise.

5. There are, of course, other ways of formulating the radical sceptical challenge. Nevertheless, since the kind of anti-sceptical proposals that I will be concerned with here tend to have this type of formulation in mind, I will concentrate on this version of the challenge. For a survey of recent work on radical scepticism, see Pritchard (2002d).

6. 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 in our formalisations without loss.

Clearly, these three propositions cannot all be true, and yet they do seem to possess at least an initial degree of plausibility. If one knows anything, then one ought to be able to know the ordinary proposition at issue in line (1). Furthermore, as line (2) makes explicit, the very sort of propositions which one seems unable, in principle, to know are the denials of radical sceptical hypotheses, since these hypotheses concern scenarios that are phenomenologically indistinguishable from everyday life. Finally, it seems relatively uncontroversial to argue, with line (3), that in order to know these ordinary propositions one must be able to know that the relevant sceptical hypotheses are false, since they seem to constitute defeaters for our everyday knowledge. One cannot consistently endorse all three of these claims, however. We are thus in a bind.

Of course, the sceptic offers a very simple way out of this puzzle, which is to deny, on the basis of lines (2) and (3), that we ever have knowledge of the ordinary propositions at issue in line (1). That is, the sceptic argues as follows:

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

But this is less of a proposal than a *reductio* of epistemological theorising. After all, since one can repeat such an argument with any agent and just about any ordinary proposition (though one might have to vary the sceptical hypothesis at issue to suit), it would follow that we could never know anything much at all, and this seems like an intolerable conclusion to draw.

Accordingly, since it has been largely taken for granted that line (S1) must be correct, the focus of attention has tended to fall on line (S2). After all, this does seem to be the weakest element of the sceptical argument since, although it is at first pass intuitive, on reflection it is far from immediately obvious that everyday knowledge should be dependent upon anti-sceptical knowledge in this fashion. One response to the problem of scepticism has thus been to deny this premise in the sceptical argument by arguing that one can perfectly well know everyday propositions whilst failing to know the denials of anti-sceptical hypotheses such as the BIV hypothesis.

Matters are not quite so straightforward, however, because the sceptic can motivate this premise by using a highly plausible epistemic principle. In particular, the suggestion has been that what motivates line (S2) is an implicit appeal to the principle that knowledge is ‘closed’ under known entailments, or the ‘closure’ principle for short. This is the principle that if an agent knows a proposition ϕ (such as that she is currently seated), and knows that this proposition entails a second proposition ψ (such as that she is not a BIV), then she also knows ψ . More formally, we can express this principle as follows:

$$(CK) \{K [\phi] \ \& \ K [\phi \rightarrow \psi]\} \rightarrow K [\psi]^7$$

And note that, unlike (S2), the plausibility of closure is not merely *prima facie*. After all, we reason in conformity with closure all the time in cases where we gain knowledge of previously unknown propositions via knowledge of other propositions and the relevant entailment. For example, one could imagine a detective reasoning as follows: I know that the Butler was in the conservatory at the time of the murder and I know that if the Butler was in the conservatory at the time of the murder then he wasn’t at the scene of the crime (the kitchen, say); hence, I know that the Butler wasn’t at the scene of the crime.

The manner in which closure is thought to support the sceptical argument ought to be clear to see. Given the uncontentious assumption that the agent in question knows that the ordinary proposition at issue entails the denial of the relevant sceptical hypothesis (i.e., that O entails \neg BIV), closure will licence the inference from a lack of knowledge that the sceptical hypothesis is false to the conclusion that one lacks knowledge of the ordinary proposition as well, just as line (S2) says. In effect, what closure achieves is to forge the required connection between knowledge of everyday propositions and knowledge of the denials of sceptical hypotheses. And since it is this principle that motivates the sceptical argument, it follows that one can only reject the second premise of the sceptical argument (and thus the third of our ‘intuitions’ listed above) provided that one is able to support a denial of closure.⁸

7. One could, of course, be much more precise about how one is understanding this principle. This rough characterisation should suffice for our purposes here, however.

8. For a more in-depth discussion of the role of closure in radical sceptical arguments,

The standard proposal put forward to support this move has been some variation of the original RA model advanced by Dretske (1970; 1971; 1981). Essentially, the idea is to claim that knowledge only transmits across known entailments where the entailments in question are ‘relevant’.⁹ Thus knowledge that one is sitting down may transmit across a known entailment to the relevant proposition that one is not standing up, but it won’t transmit across a known entailment to the irrelevant proposition that one is not a BIV. Accordingly, the link between ordinary knowledge and anti-sceptical knowledge required by the sceptic is severed and ordinary knowledge is saved.

As it stands, of course, this is merely a promissory note for a theory rather than a fully-fledged anti-sceptical thesis. For one thing, we have still to be told what makes sceptical error-scenarios, in contrast to ‘everyday’ error-scenarios, irrelevant. It is for this reason that Dretske—and, later on, Nozick (1981)—extend this RA treatment of closure by offering the following counterfactual account of knowledge, where ‘ $[\] \rightarrow$ ’ indicates the subjunctive conditional, and ‘B $[\varphi]$ ’ says that the agent believes the contingent proposition φ :

$$(DK) \neg \varphi [\] \rightarrow B [\varphi]$$

In words, DK says that an agent knows a contingent proposition φ , if, and only if, had φ been false, the agent would not have believed it.¹⁰

see Pritchard (2002b; 2002c; 2002e). Although it is commonly accepted that it is closure that underlies the core radical sceptical argument, there have been some suggestions in the literature that what motivates the sceptical reasoning is in fact a different epistemic principle, what has been called an ‘underdetermination’ principle. For discussion of this view, see Yalçın (1992); Brueckner (1994); Cohen (1998); and Pritchard (2002f). Nevertheless, although there is a substantial difference between these commentators regarding the issue of the nature and role of the underdetermination principle in sceptical arguments, there is still a general consensus that denying closure would suffice to meet the sceptical argument.

In a similar fashion, the claim that what drives scepticism is the infallibilist thesis that knowledge requires the elimination of every error-possibility likewise adds nothing substantial in terms of the debate conducted here to the contention that scepticism presupposes closure. After all, since infallibilism would demand that one at least rule-out all known error-possibilities that are incompatible with the proposition putatively known, it follows that infallibilism entails closure anyhow (though not *vice versa*). Accordingly, if one is able to adequately motivate a denial of the closure principle, then one gains a rejection of the infallibilist thesis by default. For a discussion of infallibilism and its role in sceptical arguments, see Unger (1971; 1975); Craig (1990, chapter XII and appendix); and Pritchard (2000b).

9. Dretske (1970) calls epistemic operators “semi-penetrating”.

10. I’m simplifying somewhat here. Dretske adds an extra story about the means by

This claim is typically cashed-out in terms of the modal thesis that knowledge requires not only a true belief in the actual world, but also a belief which is sensitive to the truth such that, in the nearest possible world in which what is (actually) believed is false, the agent does not believe it.

This account of knowledge offers a neat explanation of the failure of closure. To begin with, it explains our knowledge of ‘ordinary’ propositions. I know that I am currently seated, for example, because not only is this true, but in the nearest possible world in which it is false (the world in which, for instance, I am standing next to my computer now), I no longer believe that I am seated. My belief thus ‘tracks’ the truth in an appropriate fashion. In contrast, I fail to know that I am not a BIV, on this account, because although this is (say) true in the actual world, in the nearest possible world in which this is false, the *BIV*-world, I continue to believe that I am not a BIV. Accordingly, my belief that I am not a BIV fails to adequately track the truth and thus I lack knowledge of this proposition. As a result, I can know that I am sitting here now, know that this entails that I am not a BIV, and yet fail to know that I am not a BIV. The denial of closure is thus motivated in terms of a neat, and intuitive, modal account of knowledge.

Essentially, this RA line explains the failure of closure in sceptical arguments by adverting to the fact that sceptical possible worlds are (at least if we know anything much at all) significantly further away from the actual world than the near-by worlds at issue in the evaluation of everyday knowledge. Accordingly, it is no surprise that closure fails because a number of logical principles are likewise dubious in modal contexts where variable sets of possible worlds are at issue.

Nevertheless, the obvious question that arises at this point is why the core RA thought should be spelt-out in this particular fashion. After all, the Dretske-Nozick line takes far-off sceptical possible worlds to be relevant to the determination of knowledge (albeit only knowledge of anti-sceptical propositions), whereas the basic RA idea seemed to be that such far-off worlds were manifestly *irrelevant* to the determination

which the agent forms the belief, the “epistemic credential” that gives rise to the belief. Nozick goes further, relativising his account of knowledge to a method, and adding an extra condition on knowledge. Nevertheless, despite these simplifications, DK represents the core of the proposal and, moreover, it is this element of the view that results in the denial of closure. Accordingly, nothing should be lost by focussing upon this formulation of the key Dretske-Nozick thesis.

of knowledge (*any* knowledge). There thus seems to be an ambiguity in the RA thesis. Either we take it as meaning that relevance is determined by the nearest not-P world (no matter how far-out that might be), and thus end up with the Dretske-Nozick line, or else we construe it as simply demanding that only near-by possible worlds are relevant worlds. This is no mere technical dispute either, since a great deal hangs upon which alternative we adopt.

The first commentator to give clear expression to this worry (though not precisely in these terms) was Gail C. Stine (1976). She noted that so long as one sticks to the latter conception of relevance then the proper conclusion to be extracted from the core RA thought is precisely that we *shouldn't* concede a lack of knowledge of the denials of radical sceptical hypotheses, and thus reject closure. Instead, the conclusion we should draw is that, since such sceptical error-possibilities are indeed modally far-off, and thus irrelevant, it follows that we do know their denials after all and thus that closure remains intact.

It is in the interplay between these two RA models of knowledge that we find the first inklings of the two types of contextualist thesis that concern us here. For although there are considerations that speak in favour of both of these conceptions of the core RA thesis, they each face their problems. The difficulty with the Dretske-Nozick line is that it is forced to deny the highly intuitive principle of closure, and this is a hard result to swallow. As DeRose (1995, §5) has pointed out, dropping closure means allowing what he calls “abominable conjunctions”, such as that one knows that one has two hands but one does not know that one is not a BIV. In contrast, the difficulty facing the opposing RA camp is to explain how it can be that closure holds and thus that we do know the denials of anti-sceptical hypotheses after all. As Stine herself admits, such a conclusion does indeed “sound odd”. In defence of her position she argues that this ‘oddness’ is not due to what is said being false, however (as the Dretske-Nozick line would suggest), but is rather a result of the kind of false conversational implicatures that such a claim to know generates.

In effect, one could view the kind of contextualist theses that I am concerned with here as being a compromise between these two conceptions of the RA thesis. For both of these contextualist theories evade this *impasse* by arguing that, in effect, the standards of relevance here are highly context-sensitive. The basic claim is thus that closure holds but that this does not lead inexorably to scepticism because,

relative to certain standards of relevance, ascriptions of knowledge to an agent of both everyday propositions and anti-sceptical propositions can be true. Furthermore, both camps try to explain the “oddness” of Stine’s conclusion by arguing that there is something about the very drawing of this conclusion that puts one in a context in which sceptical standards of relevance operate and thus which make that conclusion false.

2. *Derose-Lewis Semantic Contextualism*

This contextualist reading of the RA thesis has been given one of its most subtle presentations in a paper by DeRose (1995). This paper, in turn, draws upon a suggestion made by Lewis (1979) and which is also developed along similar lines by Lewis (1996) himself. What both of these authors advocate is a form of *attributer* contextualism where (at least primarily) what they have in mind when they talk of the notion of a ‘context’ is the *attributer’s conversational* context, in a sense that will become apparent in a moment. Henceforth, I will call the DeRose-Lewis view ‘semantic’ contextualism. Since DeRose’s view is the most detailed, I will concentrate my attentions on his presentation of the position. Nevertheless, nearly all of what I say regarding DeRose will equally apply, *mutatis mutandis*, to Lewis’ account.

Before we consider DeRose’s view in detail, however, it is first necessary to be clearer about what attributer contextualism is. Attributer contextualists maintain that the verb ‘knows’ is context-sensitive in the sense that sentences of the form “S knows that p” (call this the ‘ascription sentence’) will express different propositions, and thus have different truth-values, depending upon the context of utterance (which will, of course, be the *attributer’s* context). Typically, attributer contextualists argue that relative to *everyday* contextual standards where sceptical error-possibilities are not relevant, ascription sentences will (*ceteris paribus*) express truths, whilst relative to *sceptical* contextual standards in which such error possibilities *are* relevant, those same ascription sentences will express falsehoods.

Stewart Cohen has been one of the foremost advocates of an attributer contextualist position and offers a neat overview of what it involves in the following passage:

I [...] defend the view that ascriptions of knowledge are context sensitive. According to this view, the truth value of sentences containing the words “know”, and its cognates will depend on contextually determined standards. Because of this, such a sentence can have different truth-values in different contexts. Now when I say “contexts”, I mean “contexts of ascription”. [...] This view has the consequence that, given a fixed set of circumstances, a subject S, and a proposition p, two speakers may say “S knows that p”, and only of them thereby say something true. For the same reason, one speaker may say “S knows p”, and another say “S does not know p”, (relative to the same circumstances), and both speakers thereby say something true. (Cohen 2000, 94)

Note this last sentence. It is an immediate consequence of attributer contextualism as it is usually understood that the very same ascription sentence can simultaneously express a truth and a falsehood relative to two different contexts of utterance. Interestingly, it is for this reason that Dretske at one point explicitly decries an attributer contextualist reading of the RA thesis. For instance, *contra* Cohen’s (1991) characterisation of his position along attributer contextualist lines, he writes:

Knowledge is relative, yes, but relative to the extra-evidential circumstances *of the knower* and those who, like the knower, have the same stake in what is true in the matter in question. Knowledge is context sensitive, according to this view, but it is not indexical. If two people disagree about what is known, they have a genuine disagreement. They can’t both be right. (Dretske 1991, 191)

So whereas Dretske holds that any particular ascription sentence will have a determinate truth-value that stays fixed in all contexts of ascription, the attributer contextualist demurs. Her claim, in contrast, is that whether or not an ascription sentence expresses a truth can essentially depend upon the epistemic standards employed by those who are (or would be) making that ascription.

With this in mind, let us return to DeRose’s version of this thesis. DeRose’s 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:

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

Intuitively, these three claims are in conflict because, or so the standard non-contextualist thought runs, *one* of these judgements must be wrong. That is, since changes in the (*ascriber's*) conversational topic (unlike, for example, changes in the subject's informational state) have no obvious bearing on the epistemic status of a subject's beliefs, an ascription sentence should either simultaneously express a truth in every conversational context of ascription, or else simultaneously express a falsehood in every conversational context of ascription. Semantic contextualism opposes this thought with the suggestion that what is actually occurring is a responsiveness, on the part of the attributor, to a fluctuation in the epistemic standards caused by a change in the conversational context. In this way, DeRose can explain the “oddness” of Stine's conclusion that she ‘knows’ the denials of anti-sceptical hypotheses by showing how drawing such a conclusion takes one into a different conversational context, one where such a self-ascription sentence (“I know that \neg SH”) is no longer appropriate because, due to a shift in the standard of relevance, it no longer expresses a truth. The idea is thus to find a way of reconciling each of the three ‘intuitions’ that we noted in §1 whilst also offering a cogent explanation of each of these supposed features of the phenomenology of our engagement with scepticism.

As regards motivating an explanation of the first element of the phenomenology of scepticism (P1), DeRose turns to the Dretske-

Nozick modal account of knowledge that we saw above expressed as DK which explains the fact that one is unable to know the denials of sceptical hypotheses (that, for instance, I am not a BIV) in terms of the *counterfactual insensitivity* of the beliefs involved. DeRose notes that:

[...] we have a very strong general, though not exceptionless, inclination to think that we don't know that P when we think that our belief that P is a belief that we would hold even if P were false. Let's say that S's belief that P is *insensitive* if S would believe P even if P were false. [...] We tend to judge that S doesn't know that P when we think S's belief that P is insensitive. (DeRose, 1995, 18)

He thus holds that DK explains, at least in part, why we are inclined to regard the sceptic's claim that we do not know the denials of radical sceptical hypotheses as being true.

DeRose cannot leave the story there, of course, because accepting counterfactual sensitivity (so described) involves committing oneself to denying the closure principle (and with it, line (S2) of the paradox presented above), an anti-sceptical strategy distinct from contextualism. As we saw in §1, if counterfactual sensitivity is a necessary condition for knowledge, then it follows that one can know that one has two hands whilst simultaneously *failing* to know that one is not a BIV, regardless of whether one knows that having hands entails that one is not a BIV. So if DeRose is to accept line (S2) of the sceptical paradox, then he is going to have to find some construal of counterfactual sensitivity that does not entail the denial of closure.¹¹ And this is not the only reason why DeRose's acceptance of counterfactual sensitivity must be limited. After all, doesn't this notion of sensitivity show us that sentences of the form "S knows that \neg SH" *never* express truths in *any* conversational context?

Given such considerations it is perhaps unsurprising that DeRose

11. Or, at least, he is going to have to find some construal of counterfactual sensitivity that is sufficient to support some *restricted* version of the closure principle that is able to support line (S2). After all, it is not inconsistent with contextualism *per se* that the universal validity of closure be denied since one could accept, on a principled basis, line (S2) without thereby endorsing an unrestricted version of the closure principle. That said, however, very few contextualists actually make this distinction, perhaps the only exception being Heller (1999a) who both denies scepticism on contextualist grounds *and* offers independent reasons for thinking that although closure holds in the sceptical context, it nevertheless fails in other contexts and thus lacks universal application. For discussion of this proposal, see Pritchard (2000a).

does indeed make the acceptance of the closure principle a pivotal part of his argument for contextualism. He focuses upon the acceptability of line (S2) in this respect, pointing out that one's epistemic position as regards the ordinary proposition, O, can be no better than one's epistemic position *vis-à-vis* the denial of the sceptical hypothesis, \neg SH. Accordingly, it is a rather strange state of affairs to argue, with those who deny closure, that one can simultaneously know the former whilst lacking knowledge of the latter. As DeRose rather neatly expresses the matter, this sort of reasoning goes against the "intuitively compelling realisation that it would be no wiser to bet one's immortal soul on O's being true than to bet it on not-SH's being true." (DeRose 1995, 32)

But how are we to understand "epistemic position" here if it is not to be expressed in terms of counterfactual sensitivity? The distinction runs as follows. Recall that counterfactual sensitivity requires the agent to lack the belief that P in the nearest possible worlds in which P is false, no matter how 'far-out', modally speaking, these worlds are. Hence, counterfactual sensitivity to an everyday proposition such as O will (one would hope at any rate), only require the agent to consider very nearby possible worlds, unlike counterfactual sensitivity to the BIV hypothesis, which would require the agent to consider far-off BIV-worlds. In contrast, epistemic position is determined 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)

Unlike counterfactual sensitivity, then, which is entirely dependent upon one's ability to track the truth in the nearest *not-P* worlds (no matter how far-out they may be), epistemic position is simply a matter of tracking the truth *simpliciter*. Accordingly, one can be in a strong

epistemic position with respect to one's belief that one is not a BIV even though that belief is never sensitive in the required manner. The out-right denial, made by both the sceptic and those who support the Dretske-Nozick line, that an agent can know \neg SH on grounds of counterfactual sensitivity thus fails to acknowledge this fact.

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 that P*] to be true." Although counterfactual sensitivity may be an important factor in determining whether an ascription sentence is true, the agent's strength of epistemic position will also be relevant. DeRose characterises the mechanism that brings about a 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)

The semantic contextualist thought is thus that for an ascription sentence to be true only requires that the agent's beliefs be counterfactually sensitive when the propositions believed are explicitly at issue within the ascriber's conversational context (we will consider an example in a moment). When the propositions believed are not explicitly at issue in that context, then strength of epistemic position can suffice where this is characterised in terms of the most demanding proposition raised in that context (i.e., the proposition which, in order for the agent's beliefs to be counterfactually sensitive to it, requires the agent to consider the furthest possible worlds). If an agent's epistemic position as regards belief in a non-explicit proposition is such that it matches the truth as to whether this proposition is true in the range of possible worlds demanded by counterfactual sensitivity to that context's explicit propositions, then the relevant ascription sentence will be true even if the belief in question is not itself counterfactually sensitive.

What motivates this general approach is the fact that, as Lewis (1979; cf. Lewis 1996) 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. Here is DeRose:

According to Lewis, “rules of accommodation” operate in many spheres of discourse that contain context-sensitive terms. Such rules specify that when a statement is made containing such a term, then [...] the “conversational score” tends to change, if need be, so as to make that statement true. For example, ‘flat’, according to Lewis, is a context-sensitive term: how flat a surface must be in order for a sentence describing it as “flat” to be true is a variable matter that is determined by conversational context. And one way to change the conversational score with respect to the standards in place for flatness is to say something that would require for its truth such a change in standards. Suppose, for example, that in a certain conversation the standards for flatness are relaxed enough that my desktop counts as being flat. If I were then to say, “My desktop is not flat”, what I say would be false if it were evaluated according to the standards for flatness in place immediately before this is said. But the Rule of Accommodation specifies that in such a situation [...] the standards for flatness are raised so as to make my statement true. (DeRose 1996, 5)

DeRose and Lewis treat ‘knowledge’ in an analogous fashion, such that the truth of an ascription sentence is dependent upon the epistemic standards employed in the context of ascription.¹²

An example will help clarify matters here. Imagine that the epistemic credentials of John’s beliefs are being evaluated relative to a quotidian conversational context in which only everyday propositions, such as whether or not he has hands¹³, are at issue.¹⁴ Counterfactual sensitivity

12. This Lewis line on context-sensitive terms is prefigured in work by Unger (1971, 1975). Crucially, however, Unger’s original observation of this context-sensitivity was used to motivate scepticism, not contextualism. In later work (e.g. Unger 1984), he was more willing to concede, albeit equivocally, the contextualist reading. For discussion of Unger’s later contextualist view, see Craig (1990, appendix).

13. Contrary to the contextualist line in this respect, it is plausible to think that any context which brought so obvious a proposition as ‘I have hands’ into question would almost certainly not be quotidian. As Wittgenstein (1969) famously argued, these ‘obvious’ propositions are not, at least not in normal circumstances, up for coherent doubt because they perform a role akin to a “framework” assumption around which everyday inquiry functions. Accordingly, doubt of them would either indicate an abnormal context (one in which we could not take this proposition for granted), or, in the extremal case, a context in which, either directly or indirectly, a sceptical hypothesis was at issue. For the sake of brevity, however, I shall let this complicating factor pass. For discussion of this point, see Pritchard (2001b).

14. What is involved in making a proposition “at issue” in a context? Contextualists rarely offer much in the way of detail on this point, but the paradigm case of a proposition being at issue is clearly meant to be that the question of its truth is under discussion in that context. This is all the detail we need for our purposes, since what is peculiar about sceptical

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. Relative to such contexts one would thus expect an ascription sentence such as “John knows that he has two hands” to be true. Accordingly, since being in an epistemic position which suffices to make what is expressed by this sentence in this context true entails that being a BIV is a relatively far-fetched possibility, hence John will be in an epistemic position that will suffice to ensure that what is expressed by the ascription sentence “John knows that he is not a BIV” in that context would also be true even though John’s belief that he is not a BIV is counterfactually insensitive.

But now suppose that a sceptical possibility—such as the BIV hypothesis—is raised in the context of ascription. This proposition will then be relevant to that conversational context and thus the strength of epistemic position demanded in that context will rise accordingly. In order for the ascription sentence “John knows that he is not a BIV” to be true in this context, John’s belief that he is not a BIV must now exhibit counterfactual sensitivity, and the possible worlds relevant to the determination of that counterfactual sensitivity will be relevant to the corresponding ascription sentence that is concerned with whether or not he has two hands. As a result, *both* ascription sentences will now express falsehoods in this context. The latter (“John knows that he is not a BIV”) because John’s belief in this respect is not counterfactually sensitive, and the former (“John knows that he has two hands”) because even though John’s belief in this case is counterfactually sensitive, he is not (and never could be) in an epistemic position that is strong enough to track the truth of everyday propositions in far-off BIV worlds.

Most of the ascription sentences that we would be inclined to assert would thus come out as expressing truths in everyday conversational contexts, whilst as nearly always expressing falsehoods in sceptical conversational contexts, and hence semantic contextualism claims to have solved the paradox presented above in a way that explains both the phenomenology of scepticism and the intuitiveness of each of the three claims involved. The intuitions are saved because there is a clear sense in which if an ascription sentence involving an everyday proposition expresses a truth then (at least provided one keeps the context of

conversational contexts—such as the epistemology seminar—is that they are conversational contexts in which we seriously discuss the truth or otherwise of sceptical hypotheses.

ascription fixed and it is an everyday context), what is expressed by an ascription sentence that is concerned with the denials of sceptical hypotheses will also be true.¹⁵ This is so even though both types of ascription sentences will express falsehoods in every context in which sceptical hypotheses are at issue. Closure is not contravened, however, because there will be no single conversational context where an ascription sentence involving an everyday proposition and an ascription sentence involving the denial of a relevant sceptical hypothesis will have different truth-values. The apparent breakdown of closure in sceptical contexts identified by Dretske and Nozick is thus accounted for in terms of a failure to appreciate that such cases simply involve an equivocation between two conversational contexts. Moreover, the phenomenology of scepticism is explained by showing how, contrary to first intuitions, whether or not an ascription sentence expresses a truth *can* be entirely dependent upon conversational factors.

3. Williams' Inferential Contextualism

Whereas the DeRose-Lewis variety of semantic contextualism has received a great deal of discussion,¹⁶ Williams' (1991) closely related contextualist theory has not been accorded quite the same degree of attention, and where there has been an examination it has tended to

15. Of course, when it comes to knowledge of the denials of sceptical hypotheses, once one makes the ascription one thereby makes sceptical hypotheses relevant and so raises the epistemic standards in such a way as to make that ascription sentence express a falsehood. In this sense, then, there is always something inherently 'ineffable' about this knowledge of the denials of sceptical hypotheses. In general, contextualism faces an expository difficulty in that the context in which we are considering the theory is one in which scepticism is at issue and thus where sceptical epistemic standards apply. Here is Lewis:

But wait. Don't you smell a rat? Haven't I, by my own lights, been saying what cannot be said? (Or whistled either.) If the story I told was true, how have I managed to tell it? In trendyspeak, is there not a problem of reflexivity? Does the story deconstruct itself?

He goes on to defend himself by arguing that in order

[...] to get my message across, I bent the rules. If I tried to whistle what cannot be said, what of it? I relied on the cardinal principle of pragmatics [...]: interpret the message to make it make sense—to make it consistent, and sensible to say. (Lewis 1996, 566)

16. For the main discussions of semantic contextualism, see Schiffer (1996); Feldman (1999); Heller (1999b); Sosa (1999; 2000); Vogel (1999); Fogelin (2000); and Pritchard (2001a).

group his theory in with the DeRose-Lewis camp.¹⁷ Nevertheless, Williams' theory is very different from the semantic account, and deserves an individual treatment.

Perhaps the best place to begin an exposition of this theory is by identifying those areas where one finds broad agreement with the semantic contextualist theory. As might be expected, what identifies Williams' inferential thesis as a contextualist account is that it incorporates the claim that contextual factors are relevant to the determination of an agent's knowledge. Like the DeRose-Lewis view, Williams argues that certain error-possibilities are only epistemically relevant in certain contexts. Moreover, Williams is also keen to retain the closure principle. He does so on contextualist grounds, arguing that provided one keeps to the one context then closure will hold. As with the DeRose-Lewis position, then, apparent failures of closure are simply due to equivocations between different contexts.

The Williams line thus shares a central core of claims with the DeRose-Lewis view. Nevertheless, their disagreements are significant. The three main areas of disagreement between the two theories are as follows. First, Williams' view is *not* a version of attributer contextualism. Second, Williams does not individuate contexts along a 'conversational' axis, but rather in terms of the inferential structure of that context (hence the name, *inferential* contextualism). Third (and as we shall see this point is closely related to the inferential thesis), Williams does not allow a context-independent hierarchy of contexts.

That Williams' view is not a version of attributer contextualism is perhaps the most radical difference between the two theses. Although Williams holds that contextual factors are relevant to the issue of whether or not an agent knows (whether she is aware of those contextual factors or not), he nevertheless maintains that the contextual factors that are relevant to an agent's knowledge concern the *agent's* context alone, and thus not also (or only) the attributer's context. His position is thus a version of what is known as *subject* contextualism, and accordingly lacks the 'indexicality' of the attributer contextualist thesis (an ascription sentence, on this view, will express a proposition with the same truth value whatever context it is asserted in).¹⁸

17. One of the main exceptions in this respect is Putnam (1998). See also Brady (1998).

18. Williams' early remarks on contextualism were ambiguous in this respect (see

Given this feature of the position, one might wonder what it is about the view that ensures that it is a specifically contextualist thesis at all. As we saw above, Dretske seems perfectly happy with a ‘contextualist’ construal of his position in *this* sense, in that the epistemological rubric that he offers should be applied in a way that is “relative to the extra-evidential circumstances of *the knower* and those who, like the knower, have the same stake in what is true in the matter in question”, but no one would ordinarily describe the Dretskean view as a ‘contextualist’ position. Indeed, in this limited sense of the term ‘contextualist’, one could imagine any number of extant non-contextualist positions being given a ‘contextualist’ flavour in this sense. One could hold, for example, that meeting a reliability condition is necessary for knowledge but nevertheless also demand that the exact manner in which one meets that condition in each case is sensitive to the contextual factors in play.

What marks out the Williams view as a thesis that is contextualist in a significant sense is thus not its commitment to subject contextualism (still less any commitment to attributer contextualism), but rather the particular way in which Williams formulates his subject contextualist thesis. The two factors that are relevant here are the other main differences between Williams’ position and the semantic contextualist view that were just listed—namely, his rejection of a ‘hierarchical’ conception of contexts, and his insistence that contexts be individuated in terms of their inferential structure rather than in conversational terms. I will take each in turn, though they are inter-related features of his position.

That the semantic contextualist view is committed to such a hierarchical conception of contextual standards should be apparent by considering how DeRose characterises his thesis in terms of strength of epistemic position. One of the consequences of this is that the sceptic does indeed impose a more demanding epistemic standard than that which is at issue in everyday contexts, with the scale of standards measured in terms of the strength of epistemic position needed in

especially Williams 1991, chapter 8), possibly because at this time attributer contextualism had not yet been fully developed and so there was no stimulus to contrast his view with this alternative. In more recent work, however, Williams has been more explicit about how his position is a subject contextualist thesis. See, for example, Williams (1999; 2000; 2001a; 2001b; 2003). For more on the distinction between subject and attributer contextualism, see Brower (1998) and DeRose (1999, §4).

order for ascription sentences to express truths. Accordingly, a change in context for the semantic contextualist consists of either a raising, or a lowering, of the epistemic standards. When sceptical scenarios are made explicit to the conversational context, for example, then one epistemically ‘ascends’ to a more demanding epistemic context. When one returns to more quotidian concerns, on the other hand, the context changes back again and one epistemically ‘descends’ once more to the old epistemic standard.

In contrast, Williams contends that a change in context does not involve a raising or a lowering of the contextual standards, but rather consists of merely a change in the inferential structure at play. As a consequence, the sceptic’s context is not more demanding, but merely *different*. Indeed, Williams goes so far as to diagnose the spectre of scepticism in terms of the view that there is such a context-independent manner of evaluating the relative merits of different contextual standards. Williams’ thesis thus represents a marked expansion of the standard contextualist line in a quite radical direction.

Indeed, on Williams’ view, contextualism entails that there is no such hierarchy of epistemic contexts—instead, each context is, epistemically speaking, autonomous. To think otherwise is, he argues, to fall victim to the doctrine of “epistemological realism”—the view that the objects of epistemological inquiry have an inherent, and thus context-independent, structure. In contrast, the inferential contextualism that he advances is defined as the denial of this thesis. It holds that:

[...] the epistemic status of a given proposition is liable to shift with situational, disciplinary and other contextually variable factors: it is to hold that, independently of such influences, a proposition has no epistemic status whatsoever. (Williams 1991, 119)

And Williams is quite clear that this last phrase “has no epistemic status whatsoever” is meant to indicate that there is no context-independent means by which we can comparatively evaluate the standards in operation in different contexts. He describes his view as a “deflationary” theory of ‘knowledge’, in that it holds that there need be nothing that ties all instances of ‘knowledge’ together other than the fact that they are instances. He writes:

A deflationary account of “know” may show how the word is embedded in a teachable and useful linguistic practice, without supposing that

“being known to be true” denotes a property that groups propositions into a theoretically significant kind. We can have an account of the use and utility of “know” without supposing that there is such a thing as human knowledge. (Williams 1991, 113)

One might think that one could abstract away from particular contextually variable factors in order to discover the underlying epistemological essence of certain propositions (e.g., that propositions concerned with the external world are necessarily only known in an inferential way). Williams rejects such a picture, and argues that as a consequence the very validity of the epistemological enterprise, at least as it is commonly understood, is called into question:

If we give up the idea of pervasive, underlying epistemological constraints; if we start to see the plurality of constraints that inform the various special disciplines, never mind ordinary, unsystematic factual discourse, as genuinely irreducible; if we become suspicious of the idea that “our powers and faculties” can be evaluated independently of everything having to do with the world and our place in it: then we lose our grip on the idea of “human knowledge” as an object of theory. (Williams 1991, 106)

Although it looks here as if Williams is dangerously close to endorsing the sceptical contention that knowledge is impossible, he is in fact arguing for what he regards as being the *anti*-sceptical contextualist claim that there is no epistemological analysis to be conducted outside of contextual parameters and, accordingly, that there is no context-independent hierarchy of epistemic standards either, as the semantic contextualist model would suggest.

This is the key feature of Williams’ view that differentiates it from other possible subject contextualist positions that do not appear to be contextualist in any significant way, such as the ‘contextualized’ version of reliabilism noted earlier. Whereas the subject contextualist thesis by itself leaves it as an open possibility that there might be a general (context-independent) epistemological rubric that determines what qualifies as knowledge, albeit with the proviso that this rubric should be applied in a context-sensitive manner, Williams’ radical version of subject contextualism closes-off this possibility. For him, there is no general epistemic condition the satisfaction of which is a requirement of all instances of knowledge, since what is (epistemically) relevant to

knowledge possession is *entirely* a context-dependent manner.¹⁹

It is as a consequence of this stance that Williams is forced to concede that the sceptic is perfectly correct in her assessment of the epistemic status of our beliefs, but only because she is operating within a context that utilises a different epistemic standard than that in play in everyday contexts. That is, the sceptic's conclusions are correct, but, by being confined to a specific epistemic context, they lack the hegemony that she requires in order to cause the intended epistemic harm. It does not follow from the truth of scepticism in this restricted sense that the ascription sentences employed within everyday contexts express falsehoods, or even that what is being ascribed here is inferior to that which is at issue in the sceptical context (which would, in line with the semantic contextualist view, presuppose a hierarchy). That is:

The sceptic takes himself to have discovered, under the conditions of philosophical reflection, that knowledge of the world is impossible. But in fact, the most he has discovered is that knowledge of the world is *impossible under the conditions of philosophical reflection*. (Williams 1991, 130)

So whereas the 'hierarchical' camp of contextualists tend to individuate contexts in terms of a context-transcendent criterion of rigour, Williams' schema allows no such ordering of contexts. For him, a context is individuated purely in terms of the epistemic structure it endorses—in terms of the inferential relations that obtain between the types of beliefs at issue in that context. And since no context employs universal standards, this contextual epistemic structure is also identified in terms of what it takes for granted—which propositions it regards as being immune from doubt in terms of that context. Williams calls the defining assumptions of a context of inquiry its "methodological necessities".²⁰

An example will be helpful here. When we do history we take the general veracity of historical documentation for granted, as well as the denials of certain sceptical scenarios such as that the world came into existence five minutes ago replete with the 'traces' of a distant ancestry

19. Williams' rejects other externalist positions in the theory of knowledge that would qualify as subject contextualist on these grounds (though he does not use this terminology) in Williams (1991, chapter 8). For further discussion of this 'deflationary' aspect of Williams' position, see Pritchard (2002a).

20. Williams' thesis in this respect is meant to be explicitly Wittgensteinian. See especially Wittgenstein's (1969, e.g. §§341–3) remarks on "hinge" propositions and Williams (2003).

(the so-called ‘Russellian Hypothesis’). To doubt such methodological necessities is not, Williams argues, to conduct our historical investigations in a more exacting fashion, but rather to engage in a different sort of investigation altogether, one that is guided by traditional epistemological concerns. So in this context where the beliefs at issue concern particular historical claims, agents are entitled to hold various other beliefs that are concerned with the methodological necessities of that context even if they lack specific evidential grounds in support of those beliefs. And since these methodological necessities are in play, it follows that, for example, one item of historical data can legitimately serve as evidential support for an historical claim even though, had the relevant methodological necessities been false, this evidence would not serve this supporting function, and the agent may have no evidence which suffices to indicate that these methodological necessities are false. For example, that items found on a certain battlefield are dated as being from the eleventh century can provide evidential support for the claim that a battle was fought on that field at this time even though, had the universe come into being five seconds previously replete with the traces of a distant ancestry, then the ‘evidence’ found would provide no evidential support for this claim at all, and the agent has no evidence which indicates that the universe did not come into existence five seconds ago.

It is this inferential structure (and, in particular, what stands fast relative to what) that determines the context and thereby determines the epistemic status of particular beliefs evaluated relative to that context. In the case of the historical context just mentioned, for example, as participants of this context we would no doubt be willing, all other things being equal of course, to regard the relevant ascription sentence (e.g., “S knows that a battle was fought on this field in the eleventh century”) as expressing a truth. A change in the features of the context can undermine this, however, and this reflects what Williams (1991, chapter eight) calls the “instability of knowledge”. In particular, if our historian were to actively participate in an epistemology seminar and continue to assert the same ascription sentences, then what she would express with those sentences would now be false. This is because in this context a different inferential structure is in play which does not employ the same set of methodological necessities as in the historical context—specifically, this context will not take the falsity of the Russellian Hypothesis for granted. In this context, then, the grounds

that the agent is able to adduce in favour of her belief are worthless because she now needs evidence in favour of her belief in the denial of the Russellian Hypothesis, evidence that she does not have. So, relative to two different contexts, the same ascription sentence expresses both a truth and a falsehood.

Three points regarding this example are worthy of emphasis. First, note that the change in context brought about by our historian actively participating in the epistemology seminar does not reflect merely a change in the conversational context. On Williams' view, the mere raising of a sceptical possibility will have no effect in itself on the truth-values of ascription sentences asserted in that context. In this example, for instance, it is not that the historian has begun discussing scepticism that has changed the context, but rather the fact that she has stopped doing history and begun a different sort of inquiry with a different set of methodological necessities. Accordingly, on this view, the sceptic needs to do more than merely raise sceptical error-possibilities if she wishes to change the context in her favour—what she needs to do is persuade agents to join her in a specific type of theoretical investigation.

Second, since the sceptic's context merely reflects a *different* inferential structure, the epistemic evaluations made by the sceptic do not reflect more austere epistemic standards but simply different ones. One does not do history in a more rigorous or exacting fashion by considering the truth of the Russellian Hypothesis. Rather, one ceases to do history altogether and begins a different type of inquiry.

Third, the contextual standards that are relevant to the truth of an ascription sentence are the *subject's* contextual standards. Accordingly, you do not get the same indexicality on this view as you do with attributer contextualism. It is true *in all contexts* that the agent in the historical context knows that a battle was fought on the field in question in the eleventh century, and it is true *in all contexts* that the same agent lacks this knowledge once she enters the epistemological context where a different inferential structure is in play.

The methodological necessities of the traditional epistemological project which, Williams claims, spawns the sceptical threat, are meant to involve a commitment to this false doctrine of epistemological realism. This leads, he argues, to an antiquated foundationalism,²¹

21. Williams (1991, chapters 7 & 8), puts forward a sophisticated set of arguments to the effect that even supposedly non-foundationalist epistemologies such as coherence theory or

one manifestation of which is the traditional epistemologist's concern with the problem of the external world. This problem is meant to reflect the inadequacy of beliefs of one type—concerning immediate experience—at serving the purpose of epistemically supporting beliefs of another type—concerning material objects in the external world. With the problem so characterised, Williams maintains that it should come as no surprise to find that it is without a solution (there are no beliefs concerning immediate experience that are able to act as epistemic guarantors for beliefs concerning objects in the material world). But, he contends, this need not result in a general external world scepticism because this conception of the inferential ordering needed for warranted beliefs about 'external' objects is far from obligatory. In terms of another type of inquiry, such as psychological investigations of perception for example, we may legitimately begin with beliefs about material objects and draw inferences about immediate experience. And since no context has any epistemological ascendancy over any other, the project of justifying psychological beliefs with reference to external world beliefs is just as valid as any epistemological theory which demanded that the inferential relations should point in the opposite direction.

Williams' inferentialist contextualist view is thus both more radical and more demanding than its semantic counterpart. On the one hand, it is more radical because Williams does not, unlike DeRose and Lewis, concede to the sceptic that her scepticism functions at a higher epistemic standard. Instead, he argues that such scepticism merely reflects a different, and faulty, conception of the epistemological landscape. On the other, it is more demanding because, for related reasons, Williams does not believe, *contra* DeRose and Lewis, that mere changes in the conversational context can suffice to bring about a different epistemic context. Instead, there must be an actual difference in the inferential structure that is employed, and thus, given his contextualism, a difference in what is being taken for granted relative to what.

4. *Comparing the Theories*

Both theories face the usual kinds of complaints that are levelled against accounts of this sort. In particular, they both have to accommodate the

reliabilism involve an implicit commitment to the basic foundationalist structure.

anti-contextualist intuition that context plays a negligible role in the determination of knowledge such that one cannot employ contextualist intuitions in order to motivate a response to scepticism. Moreover, they each have to explain how it is that ascription sentences concerning the denials of radical sceptical hypotheses can express truths in *any* context, since this is far from obvious (even at low standards).²² Nevertheless, there are difficulties that are specific to the semantic contextualist thesis, and which thus put a greater theoretical burden upon proponents of this view. I will here focus in turn on three inter-related problems facing theories of this sort which concern the following specific features of the position. First, that it involves a commitment to the thesis that it is the context of the attributer that is relevant to the determination of knowledge. Second, that it places contexts into a hierarchy of epistemic standards. And finally, third, that it allows changes in context to reflect mere conversational changes.

The issue as regards the first point is primarily more a question of methodology than of theoretical principle. Attributer contextualism is clearly a more radical thesis than its subject contextualist counterpart since it consists of not only taking the subject's context seriously in the determination of knowledge, *but also* of taking the attributer's context seriously. The main problem is thus to see why we should endorse the stronger thesis when a weaker version of contextualism seems to do just as well. After all, the subject contextualist view can accommodate the contextualist intuition that assertions of ascription sentences *within* sceptical contexts (i.e., where the agent under consideration is also a participant of that context) usually express falsehoods, whilst assertions of ascription sentences *within* quotidian contexts usually express truths. What attributer contextualism adds to this is the further claim that assertions of ascription sentences in sceptical contexts that concern agents in quotidian contexts will also usually express falsehoods, but it is hard to see why this extra consequences is desirable given that it amounts to a significant concession to the sceptic.²³

22. See DeRose (2000) for discussion of this last problem. Williamson (2000) offers a response.

23. Recall that semantic contextualism does not straightforwardly generate the related consequence that assertions of ascription sentences in quotidian contexts that concern agents in sceptical contexts typically express truths, since on this view to be concerned with sceptical issues is usually sufficient to change the context into a sceptical context. These assertions, once made, will thus ordinarily change the context in such a way that they now express falsehoods.

The second problem facing semantic contextualist concerns the fact that if we are to take seriously the idea that sceptical standards are at the more demanding end of a scale of rigour, then it is difficult to see how we are to avoid the conclusion that, strictly speaking, the sceptic is right. That is, the proper conclusion does not seem to be that lax everyday epistemic standards are equally as legitimate as austere sceptical epistemic standards, but rather that they are *sub*-standard and we should prefer the rigour introduced by the sceptic. After all, it is not as if the semantic contextualist (unlike the inferential contextualist) examines the sceptical context and finds something amiss with it—it is meant to be a set of epistemic standards that the sceptic is perfectly entitled to employ. Supposing that ‘knowledge’ is an absolute concept thus seems to play into the hands of the infallibilist (and thus the sceptic) rather than the contextualist, since it seems to show that knowledge is never *really* possessed (because we can never completely rule-out all possibilities of error).²⁴ Indeed, on Williams’ view, the sceptic’s conclusions are correct only when they are confined to the sceptical context—they do not have any influence on the truth of ascription sentences concerning agents who are not in that context. Semantic contextualism thus concedes a great deal to the sceptic via, first, its concession that sceptical standards are indeed more demanding than everyday standards and then, second, its claim that the sceptical standards can affect whether or not the assertion of an ascription sentence regarding an agent outside of that context expresses a truth.

The support offered by the supposed analogy with the term ‘flat’ is also ambiguous in this respect. If we regard this as an absolute concept then it invites the response that the everyday usage of this term is somehow ultimately illegitimate, in that it fails to reflect the more demanding standards employed by the scientist. Indeed, don’t we ordinarily allow scientific standards to ‘trump’ our everyday standards? On this interpretation, then, the analogy provides support for an infallibilism-based scepticism rather than semantic contextualism. Alternatively, one might argue that the scientist is not using the term ‘flat’ in a more demanding way but rather in a different way, and thus that our everyday ‘flat’ talk need not be trumped by scientific ‘flat’ talk

24. This is the kind of conclusion that Unger (1971; 1975) famously drew from the observation that ‘knowledge’ is an absolute concept. Although he tempers this conclusion somewhat in more recent writings (e.g., Unger 1984), he is still sympathetic to its general tenor.

(such that a scientist in a scientific context will regard the assertion of an ascription sentence involving the term ‘flat’ which concerns an agent in a non-scientific context as expressing, all other things being equal of course, a truth). The problem with this construal, however, is that it seems to provide grounds for Williams’ inferential contextualist model if anything, and certainly does not offer any comfort to the semantic contextualist. The support presented by this analogy is thus suspect at best.

Williams’ inferential model does not experience this problem of hierarchy because it straightforwardly denies that sceptical contexts are on a continuum of rigour with everyday contexts. Instead, it both contends that such contexts merely reflect different epistemic standards and that there is in any case something problematic about the standards employed by the sceptic. Accordingly, this view provides no cause to be tempted by infallibilism. Moreover, since the view is not a form of attributer contextualism, neither does William’s position allow that the sceptic’s standards have application when it comes to evaluating the truth of an ascription sentence regarding an agent in a non-sceptical context, since in this case it is *non*-sceptical standards that are to be applied.

The third problem facing semantic contextualism is how it allows that changes in the context, and thus in the epistemic standards at issue, can be brought about merely by *conversational* changes. In particular, the mere raising of a sceptical error-possibility will, at least sometimes, be sufficient to raise the epistemic standards to those employed by the sceptic. The problem here, however, is that, intuitively at least, mere changes in the conversational context only affect the *propriety* of claims made in that context, where this is understood in Gricean terms such that an inappropriate claim is one that generates false conversational implicatures.²⁵

This seems to be true of all statements, but let’s take a specific example of the assertion of an ascription sentence in order to get straight to the point. Suppose that a group of people are standing around

25. Following Grice (1989), take a conversational implicature to be (roughly) an inference that one is entitled to draw upon hearing an assertion given that one is allowed to make certain assumptions about the agent making the assertion—that she is, for example, honest, co-operative and (at least otherwise) rational. Such implicatures can always be ‘cancelled’ of course, by the agent offering an explicit *caveat* to her assertion, but I here assume that this is not the case.

the water fountain discussing who the murderer is going to turn out to be in the long-running television thriller that they have been watching. Moreover, suppose that David, who is not present at this discussion, knows who the killer is because he noticed the crucial clue that was disclosed in an earlier episode. Two salient counterpart situations can now be distinguished. First, the situation in which one of the people talking—let’s call him Keith—maintains that “David knows who the murderer is”. Second, the situation that is exactly the same in every other respect except that there is more detail to the conversation at issue. Let’s suppose that those whom Keith is talking to agree that there have not been any clues given away in previous episodes such that if one were to know the identity of the murderer then this can only be because they have some special inside information. Suppose now that Keith goes on to say that “David knows who the murderer is”. Clearly, there is something inappropriate about the second claim that is specific to that claim, which is that it generates a false conversational implicature to the effect that David has special inside information about the identity of the murderer. Here we thus have a mere conversational difference affecting the propriety of a claim. Crucially, however, such cases give us no cause to think that mere conversational changes can make a sentence that once stated something true now state something false, since by hypothesis what is expressed by the sentence in question is true in both conversational contexts. Moreover, notice that in one sense the second claim does seem to incorporate more demanding epistemic standards, in that in order for the claim to be appropriate more is demanded of the subject at issue in the claim (i.e., not just that he knows, but that he gained that knowledge in a very specific way). Nevertheless, there is no obvious need here to accommodate this ‘raising’ of the epistemic standards by construing ‘know’ as a context-relative term.

In general, mere changes in the conversation will only affect what it is appropriate to say. Why then should we think that the raising of a sceptical error-possibility would make an ascription sentence that previously expressed a truth now express a falsehood as opposed to merely making that claim inappropriate?²⁶ Again, the analogy with our

26. Indeed, as we saw above, Stine (1976) seems to be expressing a similar point as regards the Dretskean position. Recall that she claimed that if we know everyday propositions then we must know the denials of anti-sceptical hypotheses, and that that the “oddness” of this latter claim can be explained in terms of the impropriety of the assertions in question, rather than in terms of their falsity.

use of the term 'flat' is ambiguous in this respect. If we suppose that all that need be changing when we move from everyday usage to scientific usage of this term are conversational factors, then the temptation will be to regard our reluctance to call certain surfaces 'flat' that we would otherwise happily claim are 'flat' is that in this context such a claim is inappropriate because it carries false conversational implicatures (for example, that the table is 'flat' enough for the use of a certain scientific instrument). Conversely, insofar as we think that the sentence in question states something false in the scientific context, then the temptation will be to think that this is because there is more at issue here than mere conversational changes. Indeed, this is precisely the kind of case where we would tend to automatically suppose that more than mere conversational changes are taking place and would look instead to changes in the disciplinary parameters. On this construal, however, the analogy provides support for Williams' inferential subject contextualist model rather than its semantic attributer contextualist counterpart. Again, then, we find this analogy only offering ambiguous support at best for the semantic contextualist thesis.

Moreover, notice that Williams' view doesn't face this problem because it demands more of context change than the semantic contextualist thesis, requiring a change of, for example, discipline. Whilst changes in the conversation typically only affect the propriety conditions of claims, different disciplines may well use terms in different ways such that the very same sentence could be used to make different claims in each context, one of which was true and the other false.

5. A 'Two-component' Epistemological Contextualism

Of course, the acceptability of the inferential view is entirely dependent upon the correctness of Williams' critique of epistemological realism, something which I have only given the barest of outlines here. Accordingly, it may well be that the DeRose-Lewis variety of contextualism could function as a kind of weaker default contextualist position should the inferential project be seen to fail. Given the problems facing the semantic contextualist thesis that have been canvassed above, however, before we become pessimistic about the prospects for the inferentialist variety of contextualism it is incumbent

upon us to first examine ways in which one might give this view the best run for its money. To this end, I will conclude by outlining one way in which the inferential model can draw upon the work done by semantic contextualists in order to strengthen its account. In effect, the theory that I propose is a kind of ‘two-component’ contextualist thesis with inferential contextualism at its heart.

Williams is sensitive to the fact that although what is expressed by an ascription sentence may well be true this does not always mean that one can properly assert that sentence (see, for example, Williams 1991, 352). Nevertheless, he actually says very little about how a pragmatic account of assertions of ascription sentences which systematically explains this fact is to function. It is in this regard that the semantic contextualist thesis may be put into use in the service of inferential contextualism. For as I also noted above, there are strong reasons for construing the semantic contextualist thesis as a pragmatic thesis rather than as a semantic thesis. That is, there are good reasons available for thinking that, properly understood, the kind of specifically *conversational* context-variability that the semantic contextualist focuses upon is properly diagnosed as a variability in the *propriety* conditions for assertions of ascription sentences. The reason why such an account would represent an important addition to the inferential model is that it could go a long way to explaining the elements of the phenomenology of our engagement with scepticism that were noted above. Recall that these features were primarily concerned with the context-variability of the propriety of assertions of ascription sentences, and since the inferential model contains no general thesis regarding the propriety conditions for these assertions it does not present an adequate account of this phenomenology at all.

Indeed, consider again the specifics of the semantic contextualist model that DeRose describes. One of the central claims that DeRose makes is that the truth of an ascription sentence depends upon whether the agent in question is able to track the truth of the proposition at issue in the nearest possible world in which that proposition is false (what I called ‘counterfactual sensitivity’). Now consider this claim as a thesis about the propriety of an assertion of a *self*-ascription sentence, where this understood as being a sentence of the form “I know that p”. Plausibly, it does seem to be a conversational implicature of asserting a self-ascription sentences that one at least takes oneself to be in an epistemic position such that, had what is claimed been false, one would

not believe it. For example, if I vouch for someone's honesty by claiming "I know that X is of good character", then it would be inferred from this that I have grounds in support of this claim that I take to be sufficient to indicate that were it not true that this person was of good character, then I wouldn't believe that she was. Accordingly, if it turned out that I was aware that this person was hiding something relevant from me then, *ceteris paribus*, we would regard this claim as being improper since, from my point of view, there was (at the very least) a near-by possible world where this agent was not honest but where I continued to believe that she was honest nonetheless.

That claims which involve self-ascriptions of this sort generate this kind of conversational implicature would explain our reluctance to make similar claims where it is the denial of a radical sceptical hypothesis that is at issue (i.e., "I know that \neg SH"), even though, scepticism aside, we do tend to think that such a claim would express a truth. After all, as noted above, one can never have a counterfactually sensitive belief in the denial of a radical sceptical hypothesis and so such a claim (even if true) would always generate a false conversational implicature and so be deemed improper as a result.

Moreover, if one combines this thesis with an account of epistemic position in the manner that DeRose suggests, then one could further explain why it is that in contexts in which sceptical hypotheses are in play it seems inappropriate to assert *any* self-ascription sentences, whether they involve the denials of sceptical hypothesis or just everyday propositions. This would be accounted for in terms of how the standard for legitimate self-ascriptions (the extent of possible worlds that one should be sensitive to) has been raised by the presentation of sceptical error-possibilities. For example, it could be that asserting a self-ascription sentence that involves an everyday proposition in a sceptical context carries the false conversational implicature that one takes oneself to be counterfactually sensitive to the denials of sceptical hypotheses. After all, if the truth of scepticism is at issue, then to assert any self-ascription sentence would seem to imply that one takes oneself to be in a position to dismiss the sceptical doubt, and this would ordinarily mean that one takes oneself to be in a position to properly assert a self-ascription sentence involving the denial of a sceptical hypothesis. If one can never properly assert such a self-ascription sentence, however, because it concerns a proposition that one can never believe in a counterfactually sensitive way, then it follows that neither

can one properly assert a self-ascription sentence regarding everyday propositions in contexts in which the truth of scepticism is at issue.

What applies to claims involving self-ascription sentences will apply with equal weight to assertions of ascription sentences in general. In any conversational context, asserting an ascription sentence that is concerned with the denial of a sceptical hypothesis will carry the conversational implicature that the person in question has special grounds in support of her belief in this proposition which suffice to make that belief counterfactually sensitive. Since this is never the case, assertions of this sort will always carry false conversational implicatures and thus be inappropriate. Similarly, in sceptical conversational contexts, asserting ascription sentences involving everyday propositions will be inappropriate because such assertions will generate the false conversational implicature that the person in question has counterfactually sensitive beliefs in the denials of sceptical hypotheses.

We thus have an explanation of why it is that asserting ascription sentences regarding anti-sceptical propositions seems inappropriate in all conversational contexts and why asserting ascription sentences regarding everyday propositions seems inappropriate in specifically sceptical conversational contexts. Moreover, note that the impropriety of these assertions is entirely consistent with what is being ascribed in each case being true. We can thus explain the ‘contextualist’ phenomenology that the semantic contextualist adverts to without adducing the semantic contextualist thesis by paying closer attention to the context-sensitivity of the propriety conditions for assertions of ascription sentences. In effect, we diagnose the appearance of context-change in terms of the context-sensitivity of the conversational implicatures that are generated by the assertion of these ascription sentences in different contexts rather than in terms of the context-sensitivity of ‘knowledge’.²⁷ And without

27. DeRose (1999, §§8–11) does briefly consider the possibility that semantic contextualism is best understood as a pragmatic thesis about the propriety of assertions of ascription sentences. In response, he lays down two general constraints on the use of objections of this sort (what he calls “warranted assertibility manoeuvres”). First, that both the assertion of the proposition in question and the assertion of the negation of that proposition is improper in that conversational context. And, second, that the impropriety of the claim in question can be explained in terms of the generation of false conversational implicatures in a way that is not *ad hoc* (for example, the conversational rules at issue should not be specific to assertions of that proposition). The strategy discussed here meets both these constraints since the conversational rules appealed to are clearly not *ad hoc*, and in the relevant cases

the conversational motivation for semantic contextualism, the appeal of such a view in the light of its inferentialist alternative wanes. This is especially so once one considers that one cannot diagnose the kind of context-variability that is at issue in the inferentialist model purely in terms of the context-sensitivity of the conversational implicatures generated by assertions of ascription sentences since it is not part of Williams' thesis that mere conversational changes could alter the context. Williams is thus in a position to explain away the attraction of the semantic contextualist view whilst keeping the motivation for his own position intact.

By combining an inferential contextualist thesis with a semantic contextualist thesis (albeit one construed merely at a pragmatic level), the proponent of the inferentialist view would thus be able to explain the phenomenology of our engagement with scepticism, and thereby account for the fact that part of that phenomenology concerns how our use of epistemic terms seems to be sensitive merely to changes in the course of the conversation. We thus get a form of conversational contextualism without the associated claim that we found to be contentious above that the truth of ascription sentences can be dependent upon mere conversational changes in the ascriber's context.

6. *Concluding Remarks*

What this comparative discussion of the two main contextualist theses has highlighted is that, despite the superficial similarities between the semantic and the inferential versions of contextualism, these views in fact approach scepticism in very disparate ways. For whilst the semantic view takes the sceptical threat to be one that only arises at a certain epistemic standard, and which thus must be responded to by 'insulating' one's undemanding everyday 'knowledge' from the demanding sceptical standards, the inferentialist view construes the threat as being itself based upon an erroneous conception of the epistemological landscape. In this sense, inferential contextualism is not so much a response to scepticism as a *product* of a rejection of scepticism. If my remarks concerning the viability of a two-component theory are right, however, then the proper focus for attention in the contextualist's debate with

assertions of both the proposition in question and its negation will be improper.

the sceptic should be on this latter approach, with the semantic theory relegated to a secondary partner in the project, one whose task is merely to explain the phenomenology of our engagement with scepticism.²⁸

REFERENCES

- Airaksinen, T. (1982). 'Contextualism: A New Theory of Epistemic Justification?', *Philosophia* 12, 51-5.
- Annis, D. B. (1978). 'A Contextualist Theory of Justification', *American Philosophical Quarterly* 15, 213-19.
- Annis, D. B. (1982). 'The Social and Cultural Component of Epistemic Justification: A Reply', *Philosophia* 12, 55-60.
- Brady, M. S. (1998). 'Can Epistemic Contextualism Avoid the Regress Problem?', *The Southern Journal of Philosophy* 36, 317-28.
- Brower, B. W. (1998). 'Contextualism, Epistemological', *Routledge Encyclopaedia of Philosophy*, volume 2, E. Craig (editor), 646-50, Routledge, London.
- Brueckner, A. (1994). 'The Structure of the Skeptical Argument', *Philosophy and Phenomenological Research* 54, 827-35.
- Cohen, S. (1986). 'Knowledge and Context', *Journal of Philosophy* 83, 574-83.
- Cohen, S. (1988). 'How to be a Fallibilist', *Philosophical Perspectives* 2, 91-123.
- Cohen, S. (1991). 'Skepticism, Relevance, and Relativity', *Dretske and his Critics*, (ed.) B. McLaughlin, 17-37, Basil Blackwell, Oxford.
- Cohen, S. (1998a). 'Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery', *Australasian Journal of Philosophy* 76, 289-306.
- Cohen, S. (1998b). 'Two Kinds of Skeptical Argument', *Philosophy and Phenomenological Research* 58, 143-59.
- Cohen, S. (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* 13, 57-90.
- Cohen, S. (2000a). 'Contextualism and Skepticism', *Philosophical Issues* 10, 94-107.

28. I am grateful to Michael Brady and an anonymous referee from *Grazer Philosophische Studien* for their subtle and perceptive comments on an earlier version of this paper.

- Cohen, S. (2000b). 'Replies', *Philosophical Issues* 10, 132–9.
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Clarendon Press, Oxford.
- DeRose, K. (1995). 'Solving the Sceptical Problem', *Philosophical Review* 104, 1–52.
- DeRose, K. (1999). 'Contextualism: An Explanation and Defense', *Epistemology*, (eds.) J. Greco & E. Sosa, 187–205, Basil Blackwell, Oxford.
- 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.
- Dretske, F. (1981). 'The Pragmatic Dimension of Knowledge', *Philosophical Studies* 40, 363–78.
- Dretske, F. (1991). 'Knowledge: Sanford and Cohen', *Dretske and his Critics*, (ed.) B. McLaughlin, 185–96, Basil Blackwell, Oxford.
- 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.
- Grice, H. P. (1989). 'Logic and Conversation', reprinted in his *Studies in the Way of Words*, Harvard University Press, Cambridge, Massachusetts.
- Hawthorne, J. (2000). 'Reply to Cohen', *Philosophical Issues* 10, 117–20.
- 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, 115–130.
- Henderson, D. K. (1994). 'Epistemic Competence and Contextualist Epistemology: Why Contextualism is not Just the Poor Person's Coherentism', *Journal of Philosophy* 91, 627–49.
- Klein, P. (2000). 'Contextualism and the Real Nature of Academic Skepticism', *Philosophical Issues* 10, 108–16.
- 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.
- Norman, A. P. (1999). 'Contextualism: Its Past, Present, and Prospects', *Philosophia* 27, 1–28.

- Nozick, R. (1981). *Philosophical Explanations*, Oxford University Press, Oxford.
- Prades, J. L. (2000). 'Scepticism, Contextualism and Closure', *Philosophical Issues* 10, 121–31.
- 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). 'Radical Scepticism, Epistemological Externalism, and "Hinge" Propositions', *Wittgenstein-Studien*, 81–105.
- Pritchard, D. H. (2002a). 'Epistemic Deflationism', *typescript*.
- Pritchard, D. H. (2002b). 'McKinsey Paradoxes, Radical Scepticism, and the Transmission of Knowledge across Known Entailments', *Synthese* 130, 279–302.
- Pritchard, D. H. (2002c). 'Radical Scepticism, Epistemological Externalism and Closure', *Theoria* 69, 129–61.
- Pritchard, D. H. (2002d). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* 39, 215–57.
- Pritchard, D. H. (2002e). 'Resurrecting the Moorean Response to Scepticism', *International Journal of Philosophical Studies* 10, 1–25.
- Pritchard, D. H. (2002f). 'The Source of Skepticism', *typescript*.
- Putnam, H. (1998). 'Skepticism', *Philosophie in Synthetischer Absicht* (Synthesis in Mind), (ed.) M. Stamm, 239–68, Klett-Cotta, Stuttgart, Germany.
- 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, 141–54.
- Sosa, E. (2000). 'Skepticism and Contextualism', *Philosophical Issues* 10, 1–18.
- Stine, G. C. (1976). 'Scepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies* 29, 249–61.
- Unger, P. (1971). 'A Defence of Skepticism', *Philosophical Review* 80, 198–219.
- Unger, P. (1975). *Ignorance—The Case for Scepticism*, Clarendon Press, Oxford.
- Unger, P. (1984). *Philosophical Relativity*, Basil Blackwell, Oxford.
- Vogel, J. (1999). 'The New Relevant Alternatives Theory', *Philosophical Per-*

- spectives* 13, 155–80.
- Williams, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Basil Blackwell, Oxford.
- Williams, M. (1999). ‘Skepticism’, *Epistemology*, (eds.) J. Greco & E. Sosa, 35–69, Basil Blackwell, Oxford.
- Williams, M. (2000). ‘Is Contextualism Statable?’, *Philosophical Issues* 10, 80–5.
- Williams, M. (2001a). ‘Contextualism, Externalism and Epistemic Standards’, *Philosophical Studies* 103, 1–23.
- Williams, M. (2001b). *Problems of Knowledge: A Critical Introduction to Epistemology*, Oxford University Press, Oxford.
- Williams, M. (2003). ‘Wittgenstein’s Refutation of Idealism’, *Wittgenstein and Scepticism*, (ed.) D. McManus, Routledge, London.
- Williamson, T. (2000). ‘Scepticism, Semantic Externalism and Keith’s Mom’, *The Southern Journal of Philosophy* 38, 148–57.
- Wittgenstein, L. (1969). *On Certainty*, (eds.) G. E. M. Anscombe & G. H. von Wright, (tr.) D. Paul & G. E. M. Anscombe, Basil Blackwell.
- Yalçın, U. D. (1992). ‘Skeptical Arguments from Underdetermination’, *Philosophical Studies* 68, 1–34.