

7. What the Skeptic Still Can't Learn from How We Use the Word 'Know'

Wai-hung Wong

Not only rules, but also examples are needed for establishing a practice.
Our rules leave loop-holes open, and the practice has to speak for itself.
One is often bewitched by a word. For example, by the word 'know'.

(Ludwig Wittgenstein, *On Certainty*, §§139 & 435)

1. Philosophical *Déjà Vu*?

Since its publication more than twenty years ago, Barry Stroud's *The Significance of Philosophical Scepticism* has been widely read and discussed by philosophers who are interested in skepticism about our knowledge of the external world.¹ Some of his later writings on the topic (such as Stroud (1989) and (1994)) are considered essential reading too. This does not, however, mean that what Stroud says about skepticism² has as much impact on the discussion of skepticism as it deserves. It seems that his insights into the nature of skepticism have been largely misunderstood or missed. Although Stroud has never argued for skepticism or claimed that skepticism is true, he has been referred to as a "contemporary skeptic" (Huemer, 2001, p.37). The main theme of *The Significance of Philosophical Scepticism* is, as the title says so clearly, that skepticism is philosophically significant and should be considered seriously, but skepticism is still felt by most philosophers, including those who write about the topic and related issues, "to be even less worthy of serious consideration than, say Ptolemaic astronomy or

¹ In what follows I will use 'the external world' and 'the world' interchangeably.

² Hereafter 'skepticism' refers to skepticism about our knowledge of the external world.

the account of creation in the book of Genesis” (Stroud, 1977, p.39).³

One way of taking skepticism seriously is to see it as a real threat to human knowledge of the world⁴ in the sense that the Darwinian theory of evolution is a real threat to the belief in divine creation. Philosophers who take skepticism seriously in this way think that because of the force of the skeptical argument⁵ we need to show that we do have knowledge of the world, and that we cannot show that unless we can refute skepticism. That is, they think we have to conclude that we do not know anything about the world if we cannot refute skepticism. Many philosophers who discuss skepticism, however, do not take skepticism seriously in this way.⁶ A widespread attitude among philosophers towards skepticism is that it presents us with *no more than* a paradox: the skeptical argument is an apparently valid argument with true premises but a conclusion that is obviously false. For them, it is intellectually challenging to solve the paradox in the way that a logical puzzle is solved, but human knowledge is not really threatened by skepticism even if we don't know how to solve the paradox. Achilles can easily overtake the Tortoise even if Zeno's paradox cannot be solved; likewise, these philosophers think, it is obvious that human beings have knowledge of the world even if the skeptical paradox cannot be solved.

Stroud can be read as *not* taking skepticism seriously in this way, for he has never

³ Stroud wrote this in an article published in 1977, but what he said remains true today. Stewart Cohen, for example, remarks in an article published in 2001 that “[a]fter all, in the end, skepticism is crazy” (Cohen, 2001, p.96). Indeed, Cohen's remark suggests that skepticism is even less worth taken seriously than Ptolemaic astronomy or the account of creation in the book of Genesis, for the latter, though false, is at least not crazy.

⁴ For simplicity I will use 'human knowledge' to refer to human knowledge of the external world.

⁵ There is no one single argument that can be referred to as *the* skeptical argument. I am speaking loosely of the kind of argument that is well-known in the discussion of skepticism, such as the dream argument, the evil demon argument, and the brain-in-a-vat argument.

⁶ One can see skepticism as a real threat to human knowledge without conceding that one will accept it as true if it is not refuted. As Stroud remarks, “philosophical scepticism is not something we should seriously consider adopting or accepting” (Stroud, 1984b, p.1). I argue in Wong (2005) for the stronger thesis that no one can live a human life as we understand it without believing that she herself and others know many things about the world.

attempted to refute skepticism or argued that we have to conclude that we do not know anything about the world if we cannot refute skepticism. But this does not mean that he does not think skepticism needs to be answered in some way. As a matter of fact, he does offer an anti-skeptical strategy in some of his more recent writings, such as Stroud (1999) and (2003), which consists in showing that the possibility that our beliefs about the world are mostly false “is not a possibility that we could consistently believe to be actual” (Stroud, 1999, p.199).⁷ And although Stroud thinks skepticism is comparable to Zeno’s paradox, the liar paradox, and Russell’s paradox,⁸ he does not think skepticism is no more than a paradox in the sense that the study of skepticism will not reveal anything important about human knowledge (just as the study of Zeno’s paradox will not reveal anything important about motion).

Indeed, another way of taking skepticism seriously is to have the understanding that the study of skepticism can “reveal something deep or important about human knowledge or human nature or the urge to understand them philosophically” (Stroud, 1984a, p.ix). This is clearly Stroud’s way of taking skepticism seriously. When we take skepticism seriously in this way, we do not have to think we have to conclude that we do not know anything about the world if we cannot refute skepticism. We may even think that although we do not know how to refute skepticism, we do have good reason to believe that we have knowledge of the world. But we may still see skepticism as threatening. Here is how Stroud describes that threat:

We want an account of our knowledge of the world that would make all of it intelligible to us all at once. We want to see how knowledge of the world could come to be out of something that is not knowledge of the world. Without that, we will not have the kind of doubly general explanation we seek. I think skepticism in epistemology now represents,

⁷ For a detailed discussion of Stroud’s anti-skeptical strategy, see Ellis (2010).

⁸ See Stroud (1984b).

and perhaps always did represent, the possibility that such an explanation is impossible; that we cannot consider all our knowledge of the world all at once and still see it as knowledge. Given that project, the threat is that skepticism will be the only answer. (Stroud, 1984b, p.8).

On this understanding of the threat of skepticism, even if skepticism does not threaten our knowledge of the world, it still threatens “a certain kind of understanding of our position in the world” (ibid.), which can be referred to as *the philosophical understanding of knowledge*.

Stroud wrote in the 1970s that one of the reasons why skepticism was not taken seriously⁹ was “the rise of ‘linguistic’ philosophy in the 1950s, as a result of which scepticism and the arguments thought to lead to it came to be regarded as little more than a mess of false analogies, definite errors, and even identifiable fallacies which had bewitched the intelligence of earlier philosophers through insufficient attention to the complexities of language and to the general conditions for the significant functioning of our actual conceptual scheme” (Stroud, 1977, p.38). The heyday of ordinary language philosophy has long been gone. There are philosophers who still think linguistic or conceptual analysis is useful and important, but even among them there are not many who think that a philosophical problem can be solved (or dissolved, as it was put in the old days) simply by considering carefully how we use some particular words. Recently, however, there have been revived attempts to solve the problem of skepticism about our knowledge of the external world by analyzing how we use the word ‘know’ (and its cognates).¹⁰ This is part of what has been referred to as ‘the new linguistic turn in

⁹ Apparently Stroud spoke of the second way of taking skepticism seriously here, but what he said was also true with respect to the first way of taking skepticism seriously.

¹⁰ Philosophers who offer such an analysis presumably do not intend it to be an analysis only of the English word ‘know’, but intend it to apply to the counterparts of ‘know’ in other languages.

epistemology'.¹¹ Those who witnessed the flourishing of ordinary language philosophy and particularly the influence of J. L. Austin may feel a sort of philosophical *déjà vu* when they see again so many philosophers writing and debating so spiritedly about the semantics of the word 'know'. Although some philosophers who participate in this new linguistic turn certainly have said something new about the word 'know' and contributed much, and probably will contribute even more, to a better understanding of how we talk about knowledge, it is not clear that they have done a better job of illuminating the nature of human knowledge than ordinary language philosophers in the past did.

In his discussion of Austin's way of dealing with skepticism, Stroud concludes that "no anti-sceptical conclusion to the contrary could be drawn simply from the fact that we use the expressions 'I know ...', 'He knows ...', etc., as we do in fact use them" (Stroud, 1985, p.75). This conclusion still holds with respect to the new linguistic turn in epistemology. In what follows I will substantiate this claim by examining how contextualism about knowledge attributions, which has been the main impetus for the new linguistic turn, fails to solve the problem of skepticism. Although I focus on contextualism here, most of the points I am going to make are applicable, *mutatis mutandis*, to other attempts to solve the problem that make use of the idea that epistemic standards (i.e. standards for true knowledge attributions) vary in different particular cases. Sensitive invariantism, for example, shares with contextualism the idea that epistemic standards vary, though it disagrees with contextualism on *how* epistemic standards vary, or more precisely, on what the mechanism is by means of which epistemic standards are raised or lowered in different cases; and an anti-skeptical strategy similar to the one offered by

¹¹ See Ludlow (2005).

contextualists can be formulated in terms of sensitive invariantism.¹² I focus on contextualism because “[c]ontextualist accounts of knowledge attributions have been almost invariably developed with an eye toward providing a response to philosophical skepticism” (DeRose, 2002, p.168) and because the anti-skeptical strategy based on contextualism has been quite fully developed by Stewart Cohen and Keith DeRose. Besides, it seems to me that contextualism, particularly Keith DeRose’s formulation of it, is the most convincing among the theories concerning the semantics of knowledge-attributing sentences that have been suggested.¹³ Since what I try to show is that we won’t know how to solve the problem of skepticism even if we have a *correct* understanding of the semantics of ‘know’ or of knowledge-attributing sentences, I had better choose as my target of criticism a theory that is true or most likely true.

I am certainly not the first to criticize the contextualist anti-skeptical strategy. Earl Conee, Richard Feldman, Peter Klein, Hilary Kornblith, and Ernest Sosa, for instance, have made good points about how contextualism fails to solve the problem of skepticism;¹⁴ what I am going to say may overlap some of their points. What I think is new in my criticism is that it is based on two of Stroud’s insights into skepticism: first, skepticism “is a *general* theory of human knowledge” (Stroud, 1994, p.141, italics added) and “[a]ll of [our] knowledge of the external world is supposed to have been brought into question at one fell swoop” (Stroud, 1984a, p.118, original italics); and second, “the considerations that can make that conclusion [i.e. the skeptical conclusion] look plausible, perhaps unavoidable, always depend on certain ideas about sense-

¹² For sensitive invariantism, see, for example, Hawthorne (2004) and Stanley (2005). MacFarlane (2005) offers a helpful categorization of views on the semantics of ‘know’ that develop the idea that epistemic standards vary in different particular cases. MacFarlane criticizes both contextualism and sensitive invariantism, but he develops that same idea in his own way and calls his view ‘relativism’. The difference between MacFarlane’s relativism and the other two views has, again, to do with how epistemic standards vary.

¹³ The semantics of knowledge-attributing sentences suggested by contextualists has been seriously criticized, such as by Schiffer (1996) and Stanley (2004), but DeRose does a great job of clarifying and defending contextualism, most impressively in DeRose (2002) and (2005).

¹⁴ See Conee (2005), Feldman (1999) and (2001), Klein (2000), Kornblith (2000), and Sosa (2000).

perception and its role in knowledge of the world” (Stroud, 2009, p.559). I will argue that it is because contextualists lose sight of these two important aspects of skepticism that they have misconstrued skepticism in their anti-skeptical strategy.

The problem with the contextualist anti-skeptical strategy is not that it fails to *refute* skepticism, for contextualists do not aim at refuting skepticism. The problem is that they have changed the subject without even noticing it. Contextualists may have solved ‘the skeptical paradox’ in their own terms, but such a solution cannot in any way make skepticism, that is, skepticism correctly construed, less threatening to human knowledge or to the philosophical understanding of human knowledge.

2. The Contextualist Anti-Skeptical Strategy: DeRose’s Version

It may not be fair to contextualists if we consider the contextualist anti-skeptical strategy only in completely general terms, for how forceful such a strategy is depends on how it is formulated in detail. I think the best formulation of the strategy is DeRose’s in DeRose (1995), which we will look at in this section. Since DeRose’s formulation has all the features of contextualism that I want to criticize, and these features are shared by other formulations, my argument should be understood as an argument against the contextualist anti-skeptical strategy generally rather than against only DeRose’s version of it.¹⁵

The central idea of contextualism is that epistemic standards vary in different contexts of knowledge attribution.¹⁶ According to DeRose, epistemic standards determine how strong one’s

¹⁵ And not just against the contextualist anti-skeptical strategy generally, but also, as I have pointed out, against other attempts to solve the problem that make use of the idea that epistemic standards vary in different particular cases.

¹⁶ Note that the contextualist idea is not just that epistemic standards vary in different contexts, which most epistemologists would grant (that is, except strict invariantists; see MacFarlane (2005)), but that epistemic standards

epistemic position with respect to a proposition has to be in order for one to be truly described as knowing that proposition. He explains the strength of one's epistemic position in terms of how well one's belief tracks the truth. To track the truth, one's belief concerning p has to match the fact as to whether p not only in the actual world, but also in some of the nearby possible worlds; the further away the possible worlds are in which one's belief concerning p still matches the fact as to whether p , the better one's belief tracks the truth. The strength of one's epistemic position with respect to p is thus a matter of how remote (from the actual world) the remotest possible worlds are to which one's belief concerning p tracks the truth — the remoter these possible worlds are, the stronger one's epistemic position is with respect to p .¹⁷

One's epistemic position (under the same circumstances) with respect to different propositions can have different strength, and DeRose suggests that such differences can be expressed by comparative conditionals. That is, the fact that S 's epistemic position with respect to q is at least as strong as his epistemic position with respect to p can be expressed as follows: if S knows that p , then S knows that q ; and if S does not know that q , then S does not know that p . For example, my epistemic position with respect to 'I am not a brain-in-a-vat' is at least as strong as my epistemic position with respect to 'I have hands', for if I know that I have hands, then I know that I am not a brain-in-a-vat; and if I do not know that I am not a brain-in-a-vat, then I do not know that I have hands.

Another central notion in DeRose's account is the *sensitivity* of beliefs, which can also be

vary in different contexts of *knowledge attributions*. That is, it is the knowledge attributor's context that determines the epistemic standards used.

¹⁷ This account of epistemic standards is externalist in the sense that one does not have to know what epistemic position one is in to be in that epistemic position and to satisfy or fail to satisfy the epistemic standards concerned. Cohen suggests, by contrast, an internalist account of epistemic standards: to have different epistemic standards is to be more or less demanding on how good S 's *evidence* to believe that p has to be in order for ' S knows that p ' to be true (see, for example, Cohen (1987), (1988), and (1999)). In other places than DeRose (1995), however, DeRose seems to adopt an internalist account of epistemic standards (see, for example, the bank cases discussed in DeRose (1992) and (2002)).

explained in terms of truth-tracking. One's belief that p is sensitive if it tracks the truth as to whether p to at least the closest possible worlds in which p is false; or more intuitively, one's belief that p is sensitive if it is not the case that one would still have that belief even if p were false when not much else had changed. One is presumably in quite a strong epistemic position with respect to p when one's belief that p is sensitive, and one is always in a stronger epistemic position with respect to p when one's belief that p is sensitive than when it is insensitive.

It is important to note, however, that how strong one's epistemic position has to be for one's belief to be sensitive depends on what the proposition is that is the object of one's belief, for what the proposition is determines how far away from the actual world the closest possible worlds are in which the proposition is false. It is thus possible for one to be in a stronger epistemic position with respect to q than with respect to p even if one's belief that q is *insensitive* while one's belief that p is sensitive. To use the standard pair of examples again, I am in a stronger epistemic position with respect to 'I am not a brain-in-a-vat' than with respect to 'I have hands', but this does not imply that my belief that I am not a brain-in-a-vat must also be sensitive if my belief that I have hands is sensitive.

Let us now look at how DeRose uses the above account of epistemic standards to 'solve' the problem of skepticism. The work is done by what he calls the 'Rule of Sensitivity', which is a conversational rule governing the change of epistemic standards. Here is the rule: when it is asserted that S knows that p , epistemic standards will be raised to such a level that S 's belief that p has to be sensitive for ' S knows that p ' to be true.¹⁸ In a philosophical context in which the skeptic proposes a radical skeptical hypothesis h and asserts that we do not know that not- h , the Rule of Sensitivity is applicable; no one in that context can then know that not- h unless his belief

¹⁸ See DeRose (1995), pp.205-206.

that not-*h* is sensitive.¹⁹ Since our belief that not-*h* is insensitive (we would still believe that we are not brains-in-vats even if that belief were false), the application of the Rule of Sensitivity to it results in the skeptical conclusion that we do not know that not-*h*. This skeptical conclusion can then be used, in conjunction with the principle that knowledge is closed under known entailment, to show that we do not know a particular proposition about the world provided that the proposition is incompatible with *h*. If *h* is radical enough, such as the hypothesis that we are brains-in-vats, then most of the propositions that we believe about the world are incompatible with *h*; the final skeptical conclusion the skeptic arrives at, a conclusion we have to rationally accept in such a context, is that we do not know most of the things we believe we know about the world, including those that we strongly believe we know, such as that we have hands.²⁰

But contextualism is no skepticism. The Rule of Sensitivity makes the skeptical conclusion true *only in philosophical contexts*, for it is in such contexts alone that a radical skeptical hypothesis is proposed and claimed not to be known by us to be false. And when the Rule of Sensitivity is applied to our belief that a certain radical skeptical hypothesis *h* is false, the epistemic standards it induces are indeed extraordinarily high: in order for us to know that not-*h*, our epistemic position with respect to not-*h* has to be so strong that none of us can ever be in

¹⁹ In DeRose (2004), DeRose argues that the skeptic does not *always* succeed in raising epistemic standards, such as in a case in which “our skeptic has executed a maneuver [...] that has a tendency to raise the epistemic standards, and that her opponent has responded by executing a maneuver that has at least some tendency to keep lower, ordinary standards in place” (p.4). His view is that neither of them succeeds. It seems that this view does not accord with his formulation of the Rule of Sensitivity in DeRose (1995), but I am not going to pursue the issue here.

²⁰ If this is the skeptic’s argument, and if the skeptical conclusion is that we do not know *anything* about the external world, then, as DeRose points out, the skeptic may need a principle stronger than the principle that knowledge is closed under known entailment. The proposition that I am in Houston, for example, does not contradict the skeptical hypothesis that I am a brain-in-a-vat. See DeRose (1995), note 33. For a related discussion, see Stroud (1984a), pp.25-30. As we will see, there is a way of understanding how skepticism attacks *all* our knowledge of the external world that is independent of understanding how this skeptical argument should be refined. It should be noted that the evil demon argument does attack all our knowledge of the external world, for whatever we believe about the external world is incompatible with the skeptical hypothesis that we are being deceived by the evil demon. I will explain in section 4 why the evil demon argument may not be the best argument for bringing out the force of skepticism.

such an epistemic position. However, in everyday contexts when nothing like radical skeptical hypotheses are considered, our epistemic standards are set much lower (i.e. our beliefs don't have to track the truth to remote possible worlds), so low so that we often meet them — our epistemic position with respect to most propositions about the world is strong enough for our knowing those propositions.

The Rule of Sensitivity is also applicable to our everyday knowledge claims or knowledge attributions, but most of these claims and attributions survive the test of the rule, for most of our beliefs about the world *are* sensitive. I do not, for example, believe that I have hands in the closest possible worlds in which I do not have hands, such as the possible worlds in which I lost my hands in a car accident. In the possible worlds in which I am a brain-in-a-vat, I do (falsely) believe that I have hands, but, DeRose suggests, my belief as to whether I have hands does not have to track the truth to these extremely remote possible worlds in order to be sensitive. So our knowledge claims or attributions are mostly true in everyday contexts. Indeed, since our knowing things about the world implies our knowing that radical skeptical hypotheses are false, in everyday contexts we can be taken to know that radical skeptical hypotheses are false (though it seems that the Rule of Sensitivity makes it impossible to *assert* that we have such knowledge without making the assertion false²¹).

Contextualists like DeRose are quite willing to surrender the philosophical territory to the

²¹ David Lewis makes a similar point when he compares the epistemologist's situation with that of a person who wants to say of a silent group of which he is a member that 'All of us are silent' (see Lewis, 1996, p.238). DeRose remarks, however, that he "would find it a bit embarrassing if we could never *claim* to have such knowledge [i.e. knowledge that radical skeptical hypotheses are false] by means of simple knowledge attributions"; he thinks that "in special conversational circumstances, it seems we *can* truthfully claim to know that not-H, despite the fact that our belief that not-H is insensitive" (DeRose, 1995, note 36). On this point he gets a helpful interpretation of his view from Hilary Kornblith: "It is important to recognize that on DeRose's view, the mentioning of skeptical possibilities does not automatically raise the standards for knowledge attribution; it merely creates some pressure in that direction, to which the otherwise anti-skeptical partner may or may not accede... Attempts to raise the standards for knowledge, like attempts to change the topic of conversation, require the engagement of both conversational partners" (Kornblith, 2000, p.28).

skeptic and fight only for (re)claiming the non-philosophical territory. This is because they admit that they do not know how to refute skepticism; indeed, they find skepticism persuasive and want to “explain the persuasiveness of the skeptic’s attack” (DeRose, 2002, p.168). On the other hand, they do not think that we should accept skepticism, which is why they attempt to explain the persuasiveness of skepticism “in a way that makes it unthreatening to the truth of our ordinary claims to know” (ibid.).²² Contextualists see having knowledge as, in many cases, a piece of cake (such as when you look at your hands and know that you have hands); but since they do not know how to refute skepticism, they try to have the cake and eat it too — try to show that we do have knowledge of the world even if skepticism is true.

3. Non-Threatening Skepticism

What makes skepticism a threat to human knowledge and our understanding of it is its generality: *all* our (putative) knowledge of the external world is subject to the attack of the skeptical argument. It is impossible for us to make use of any particular piece of knowledge of the world to answer skepticism or to explain the possibility of human knowledge because that very piece of knowledge is subject to the attack of the skeptical argument *in the same way* as the rest of our knowledge. This is a point that Stroud emphasizes over and over again in his work on skepticism. In the light of this understanding of the threat of skepticism, we should see that any ‘skepticism’ that attacks only particular pieces of knowledge rather than all our knowledge of the external world at one fell swoop is not a real threat to human knowledge. Let us call such skepticism *non-threatening skepticism* in order to distinguish it from the philosophical

²² Cohen expresses a similar understanding of the contextualist anti-skeptical strategy when he remarks that by means of contextualism we can have a solution to the problem of skepticism that “preserves our belief that we know things” and at the same time “explain[s] the undeniable appeal of skeptical argument” (Cohen, 2000, p.100).

skepticism that was first clearly formulated by Descartes and has been a worry of epistemologists since then.

Here is an example of non-threatening skepticism. I believe I know that my car is parked in my driveway. If I know that, then I also know that my car was not stolen two minutes ago. But I do not know that my car was not stolen two minutes ago. Therefore, I do not know that my car is parked in my driveway. Following John Hawthorne, let us call ‘My car was not stolen two minutes ago’ a *lottery proposition*, which is “a proposition of the sort that, while highly likely, is a proposition that we would be intuitively disinclined to take ourselves to know” (Hawthorne, 2004, p.5). Obviously this argument is only about my knowledge that my car is parked in my driveway and has no implication for my other knowledge of the world. Of course, what makes the argument philosophically interesting is that it can be generalized: for many ordinary propositions like ‘My car is parked in my driveway’ we can find some lottery propositions and form arguments like the one above (call them *lottery arguments*). This is why Hawthorne claims that “[t]hese considerations generate powerful pressure towards a skepticism that claims that we know little of what we ordinarily claim to know” (ibid.). Even with such generalization, however, the skepticism engendered is still non-threatening.

In most lottery arguments, the subject does not know the lottery proposition concerned simply because of the particular epistemic position she is in with respect to that proposition. Her not knowing it does not rule out the possibility that another subject in a different epistemic position may know the proposition — a lottery proposition to one subject may not be a lottery proposition to another subject. Yes, I do not know that my car was not stolen two minutes ago because I am in my study and cannot see my car from there, but my neighbor Chris may be standing outside his house and looking at my car, and hence know that my car was not stolen two

minutes ago. Indeed, if I go out now and see my car, I will also know that my car is still parked in my driveway. A lottery proposition to a subject at a particular time may not be a lottery proposition to her at another time because her epistemic position with respect to the proposition may change.

For those who think that lottery arguments are the basis of skepticism, lottery arguments threaten human knowledge in general because for *every* ordinary proposition about the world there is at least one lottery proposition, and hence at least one lottery argument by virtue of which a person, no matter what circumstances she is in, can be shown not to know that ordinary proposition. Even if the same proposition p is a lottery proposition to person A but not to person B, it may, as the suggestion goes, still be the case that for B and p , there is a different lottery proposition and a lottery argument that can show that B does not know that p . It is difficult to see how this suggestion can be established. Perhaps it can be argued that Chris does not know that my car was not stolen two minutes ago because he does not know that my car was not stolen two minutes ago and was replaced with one that looks exactly like my car. But what if my brother has been fixing my car for me the whole day? Isn't it clear that he knows that my car was not stolen two minutes ago?

Skepticism as construed by contextualists is non-threatening for a similar reason. Although contextualists focus on cases in which a person's epistemic position remains the same with respect to a particular proposition in order to explain how epistemic standards determine whether a knowledge attribution is true or false, they assume that people can be in different (i.e. stronger or weaker) epistemic positions under different circumstances, whether with respect to the same proposition or different propositions. When contextualism is used to explain why a knowledge attribution is false, that is, why a subject does not know a particular proposition, it

leaves open the possibility that another subject under different circumstances may know the proposition, either because the subject's epistemic position is different, or because the epistemic standards are different, or because both are different.

It may appear that according to DeRose's contextualist explanation of the persuasiveness of the skeptical argument, we are all in the same epistemic position with respect to a proposition that expresses the negation of a radical skeptical hypothesis (such as 'I am not a brain-in-a-vat'), and the same extremely high epistemic standards are used for assessing whether we know that proposition, the result being that none of us knows it. But the very idea of contextualism is that assessment of knowledge attributions depends on the particular contexts or circumstances in which the attributions are made. Even if there are contexts in which each of us is correctly assessed as not knowing 'I am not a brain-in-a-vat', contextualism leaves open the possibility that there are other contexts in which we may be correctly assessed as knowing it. Indeed, as I mentioned above, contextualists can argue that since in some contexts I can be correctly assessed as knowing that I have hands, and since 'I have hands' entails 'I am not a brain-in-a-vat', in those context I can be correctly assessed as knowing that I am not a brain-in-a-vat.

I am not simply recapitulating contextualism here; my point is that the contextualist construal of skepticism has the built-in feature of being non-threatening because contextualism explains only particular assessments of knowledge attributions under particular circumstances. Skepticism, by contrast, assesses all our putative knowledge of the external world *at one fell swoop*. When the skeptic considers a particular knowledge claim or knowledge attribution, he uses it only as a representative example of our putative knowledge of the world, such as Descartes' example of his putative knowledge that he is sitting in his room by the fire in his

dressing gown with a piece of paper in his hand,²³ which is, as Stroud points out, “a best-possible case” that “serves as the basis for a completely general assessment of the senses as a source of knowledge about the world around us” (Stroud, 1984a, p.9). No such completely general assessment can be made in contextualist terms.

As we will see in the next section, the generality of the threat of skepticism consists not only in its attacking all our beliefs about, and all our putative knowledge of, the external world at one fell swoop, but also in its showing that we are in the *exact same* poor epistemic position with respect to *any* proposition about the external world (at least any proposition about the world that we believe on the basis of sense-perception). Since contextualism assumes that we can be in different epistemic positions with respect to the same proposition or different propositions, any skepticism put in contextualist terms cannot be threatening in the way in which what Stroud calls philosophical skepticism is. The way contextualists make skepticism “unthreatening to the truth of our ordinary claims to know” (DeRose, 2002, p.168) is by replacing it with a straw man, which is, of course, not very threatening.

It should be noted that this criticism of the contextualist construal of skepticism is not a criticism of the contextualist semantics of knowledge-attributing sentences. Even if the semantics is correct, the criticism still stands. Although the skeptical conclusion as construed by contextualists is in a sense general, that is, the content of it is about all our putative knowledge of the world, the truth of the conclusion is supposed to be always tied to a particular context in which a particular human being has made a particular knowledge attribution. The skeptic, by contrast, does not understand the truth of the skeptical conclusion as being tied to any particular context. Contextualists may think this means that the skeptic does not understand the semantics

²³ See Descartes’s *First Meditation*.

of knowledge-attributing sentences. But the point here is, as far as his skepticism is concerned, the skeptic is simply not in the business of assessing particular knowledge attributions made in particular contexts. Contextualists may insist that the skeptic cannot arrive at the skeptical conclusion without being constrained by the semantics of knowledge-attributing sentences. But this has to be established by an argument other than a mere formulation of the semantics, an argument which contextualists have not offered. Unless there is such an argument, comparing the skeptic's argument for the skeptical conclusion to assessment of particular knowledge attributions made in particular contexts is like comparing an atheist's argument against miracles to the Catholic Church's assessment of particular miracle claims.

There is another way of seeing why skepticism as construed by contextualists is non-threatening. It is clear that DeRose's contextualism assumes that the world is in the main how we believe it is, that is, that most of our beliefs about the world are true. Without such an assumption, neither our different epistemic positions under different circumstances nor the sensitivity of our beliefs can be determined in the way DeRose suggests — to determine how well our beliefs track the truth, we have to assume where the tracking starts. And it is not only DeRose's version of contextualism that needs the assumption that most of our beliefs about the world are true; all versions of contextualism need it. Contextualism appeals to the idea that in some contexts we can truly be said to have knowledge of the world because epistemic standards are low enough in those contexts, but this idea makes sense only on the assumption that the beliefs in question are true. For if our beliefs are false, then no matter how low epistemic standards are, meeting those standards will not turn our beliefs into knowledge — false beliefs cannot be knowledge; and no matter how high epistemic standards are, it is not *because of* our not meeting those standards that our beliefs are not knowledge — false beliefs are not knowledge

anyway.

DeRose admits that he assumes in his argument “certain things that we believe but that the skeptic claims we can’t know” (DeRose, 1995, p.215). He does not think this is problematic:

I’m ready to admit to the skeptic that if I am a BIV [i.e. a brain-in-a-vat], then I don’t know I have hands, according to any standards for knowledge. But, of course, as I firmly believe, I’m not a BIV. Is it legitimate for me to use this conviction in a debate against the skeptic? Not if we’re playing King of the Mountain. But if the skeptic is marshalling deeply felt intuitions of ours in an attempt to give us good reasons for accepting his skepticism, it’s legitimate to point out that other of our beliefs militate against his position, and ask why we should give credence to just those that favor him. (ibid.)

This may sound reasonable given that DeRose (and other contextualists) is not trying to refute skepticism. Since skepticism does not imply that all our beliefs about the world are false or that we are brains-in-vats, contextualists cannot be accused of begging the question against skepticism when they assume that some or even most of our beliefs about the world are true.

This does not, however, mean that the assumption is unproblematic. Although the skeptic does not argue that our beliefs about the world are false, his conclusion that we do not know anything about the world does imply that our beliefs about the world are *doubtful*. Now the problem is not merely that contextualists assume that some of our beliefs about the world are true, but also that their account of knowledge attribution applies exclusively to cases in which “a speaker A (or ‘attributor’) says, ‘S knows that P,’ of a subject S’s *true* belief that P” (ibid., p.185; italics added). Accordingly, when contextualists say that in philosophical contexts the skeptic is right that we do not know anything about the world, what they really mean is that in those contexts the skeptic is right that even our true beliefs about the world are not knowledge. But the

skeptic's question is never 'Are our *true* beliefs about the world knowledge?'; it is simply 'Do we know anything about the world?' or 'Are our beliefs about the world knowledge?'. The skeptic's negative answer to the question, namely, that we do not know anything about the world, does not mean 'Even our true beliefs about the world are not knowledge'. 'Even our true beliefs about the world are not knowledge' cannot express the doubts that the skeptic argues we have good reasons to have. What the skeptic thinks we have good reasons to doubt is not that we have knowledge of the world; for him this is not a matter of doubt — he believes he has shown that we do *not* have knowledge of the world. What he thinks we have good reasons to doubt, reasons he believes he has given us by showing that we do not know anything about the world, is that our beliefs about the world are mostly true. Any skepticism that does not imply such doubts about our beliefs is non-threatening.

4. Perceptual Access to the World

One serious misunderstanding of skepticism by contextualists is that a radical skeptical hypothesis is *merely* a remote possibility that we cannot rule out. From the perspective of what we already believe, or believe that we know, about the world, a radical skeptical hypothesis is indeed a very remote possibility, but what makes the skeptical argument so persuasive is not that the radical skeptical hypothesis used in it is a possibility that cannot be ruled out. A radical skeptical hypothesis that brings forth the threat of skepticism is one that effectively reveals to us the difficulty of understanding, as Stroud puts it, "how we could get any knowledge of things around us on the basis of sense-perception, given certain apparently undeniable facts about sense-perception" (Stroud, 1984b, p.5). Whether it is the dream hypothesis, the brain-in-a-vat hypothesis, or the evil demon hypothesis, the difficulty revealed is the same. Here is how Stroud

explains the difficulty:

There have been many versions of that fundamental idea. But whether it is expressed in terms of ‘ideas’ or ‘experiences’ or ‘sense data’ or ‘appearances’ or ‘takings’ or ‘sensory stimulations’, or whatever it might be, the basic idea could be put by saying our knowledge of the world is ‘underdetermined’ by whatever it is that we get through that source of knowledge known as ‘the senses’ or ‘experience’. Given the events or experiences or whatever they might be that serve as the sensory ‘basis’ of our knowledge, it does not follow that something we believe about the world around us is true. (ibid., p.6)

What we are led by a radical skeptical hypothesis to see is, in a nutshell, that we do not actually perceive that the world is the way we believe it is and that, as a result, we are in the exact same poor epistemic position with respect to *any* proposition about the world that we believe on the basis of sense-perception.

I am not going to rehearse the details of any particular skeptical argument here; the point I want to make is merely that the persuasiveness of skepticism has everything to do with the problem of sense-perception I have just mentioned and has very little to do with a radical skeptical hypothesis being a (remote) possibility that we cannot rule out. Even if it is possible to solve the problem of skepticism without refuting the skeptical argument, any ‘solution’ to it that does not address the problem of sense-perception, such as the contextualist ‘solution’,²⁴ is no real solution; this is because without addressing the problem of sense-perception, such a ‘solution’ cannot help us understand how human knowledge is possible, and it is such understanding that the problem of skepticism forces us to admit we need to show that we have.

The reason why the problem of sense-perception revealed by a radical skeptical

²⁴ The title of DeRose (1995) is ‘Solving the Skeptical Problem’.

hypothesis makes skepticism so persuasive is that it is based on some “apparently undeniable facts about sense-perception” (ibid., p.5). One such apparently undeniable fact is that in hallucination a person can have sensory experience of something that is not really there. Another is that a dream can be so vivid and realistic that the sensory content of it is indistinguishable from what we would experience if it was not a dream but was what really happened in our waking life. The latter fact (if it is a fact) is, of course, the basis for the famous dream argument offered by Descartes.

In order to lead us to see that there is a *general* problem of sense-perception, the problem that our knowledge of the world is underdetermined by our sensory experience, a skeptical hypothesis has to be radical enough to make it impossible for us to rely on any particular sensory experience. However, if the skeptical hypothesis is too radical, it may lead us so far afield that we do not see that it is about our actual situation. This may be why Stroud focuses on the dream hypothesis rather than the evil demon hypothesis when he discusses Descartes’s skeptical argument, though the dream hypothesis is, as Penelope Maddy points out, “the functional equivalent of the Evil Demon hypothesis” (Maddy, 2007, p.23): we all have the experience of dreaming, but we do not believe that we have the experience of being deceived by the evil demon (even though the very point of the hypothesis is that all our experiences are the evil demon’s tricks). The brain-in-a-vat hypothesis, which DeRose focuses on, is in this respect more like the evil demon hypothesis than the dream hypothesis — we do not believe that we have the experience of being a brain-in-a-vat; at least none of us has seen a brain-in-a-vat.

Contextualists are presumably familiar with the problem of sense-perception, for it is a major traditional problem in epistemology and most contextualists are epistemologists who have studied the problem and its history. So why don’t they see that a radical skeptical hypothesis

points to precisely this problem? Why do they treat it instead as nothing more than a remote possibility? The answer cannot simply be that they focus on skeptical hypotheses that are too radical, such as the brain-in-a-vat hypothesis, for some of them do discuss the dream hypothesis. I am going to suggest an explanation of this oversight that, I think, may tell us something about an aspect of the difficulty of philosophy.

We have seen in the last section that contextualists assume that most of our beliefs about the world are true. The word ‘assume’ may, however, be too weak to convey the fact that contextualists examine skepticism from the view of the world that they actually have — the view of the world that we all have, the view that is constituted by the beliefs we share about the world, such as that there are trees and mountains, that most of us have hands, that we are not brains-in-vats, etc. This view of the world would not be possible if we did not have the following central belief that constitutes part of the view: our most direct and reliable access to specifics of the world is through our senses, our perceptual experiences. As P. F. Strawson puts it, “we think of perception as a way, indeed the basic way, of informing ourselves about the world of independently existing things: we assume, that is to say, the general reliability of our perceptual experiences” (Strawson, 1979, p.103); and more importantly, such a belief about the relation between perceptual experiences and the world “so thoroughly permeates our consciousness that even those who are intellectually convinced of its falsity remain subject to its power” (ibid., p.106). Whether contextualists are (and they probably are not) intellectually convinced of the falsity of the belief that we have perceptual access to the world, they are, like the rest of us, subject to the power of the belief. Unless we adopt persistently a *detached* attitude towards the belief, we cannot help seeing things through it. A plausible answer to the question of why contextualists lose sight of the connection between the problem of skepticism and the problem of

sense-perception even though they are familiar with the problem of sense-perception is thus this: when they deal with skepticism, they are not able to, as it were, free themselves from the grip of the belief that we have perceptual access to the world.

The difficulty of philosophy that this explanation reveals is the difficulty of maintaining the detached attitude required for doing philosophy, that is, the detached attitude towards the beliefs that we have and that we are also philosophically examining.²⁵ On the one hand, we do have the beliefs in question; on the other hand, our task is to examine the beliefs without employing or being affected by the beliefs. The detachment required is particularly difficult in the case of the belief that we have perceptual access to the world because the belief “thoroughly permeates our consciousness” and because all other beliefs about the world that we acquire on the basis of this very belief — and they are numerous — fortify it.

5. The Practice of Knowledge Attribution

When contextualists discuss knowledge attribution, they focus on epistemic standards and suggest a mechanism by means of which epistemic standards are determined in different contexts. Suppose they are right about the mechanism; this does not, however, mean that contextualism is all there is to say about knowledge attribution. Knowledge attribution is a complex human *practice* that is intertwined with other practices we have, with how we interact and communicate with one another, and with how we make evaluations about ourselves as well as people and things around us. In this section I will discuss some important aspects of the practice of knowledge attribution that can help explain further why contextualists can neither take skepticism seriously nor solve the problem of skepticism. I will try to establish two points:

²⁵ For a detailed discussion of philosophical detachment, see Wong (2002).

first, the more we can make sense of particular knowledge attributions, the less we can take skepticism seriously; and second, the practice of knowledge attribution as we understand and engage in it presupposes that we have knowledge of the world.

The two points are, of course, related. Let me begin with the second point. When contextualists speak of contexts, they mean primarily *conversational* contexts in which knowledge-attributing sentences are uttered. But the practice of knowledge attribution is more than utterances of knowledge-attributing sentences. I am not suggesting that contextualists are not aware of the fact that utterances of knowledge-attributing sentences are closely related to practical concerns that we have in our everyday life. DeRose, for example, is clearly aware of that when he suggests, in his initial formulation of contextualism, that “requirements for making a knowledge attribution true go up as *the stakes* go up” (DeRose, 1992, p.915, italics added). He makes this remark with respect to the bank cases he discusses (*ibid.*, p.913); he refers to these cases again in DeRose (2005) and relates them to “the practical concerns involved” (DeRose, 2005, p.176) and “the actual practical situation that a speaker faces” (*ibid.*, p.177). Nevertheless, when contextualists do speak of practical concerns, they mean no more than practical concerns that *an individual* has under some particular circumstances. They overlook the fact that the practice of knowledge attribution is a *collective* practice in the sense that its existence requires that, first, most of us engage in the practice; second, we have some common understanding of the practice; and third, we have some common purposes when we engage in the practice.

In the *Philosophical Investigations* Wittgenstein makes the following insightful remarks about some human practices:

It is not possible that there should have been only one occasion on which one obeyed a rule. It is not possible that there should have been only one occasion on which a report

was made, an order given or understood; and so on. — To obey a rule, to make a report, to give an order, to play a game of chess, are *customs* (uses, institutions). (Wittgenstein, 2001, §199, original italics)

The practice of knowledge attribution is, I think, in the same sense a custom or an institution: it is not possible that there should have been only one occasion on which a knowledge attribution was made. The development of the practice of knowledge attribution took a long time and depended on many factors; it is, like “[c]ommanding, questioning, storytelling, chatting”, “as much a part of our natural history as walking, eating, drinking, playing” (ibid., §25).

The collective practice of knowledge attribution would not have been developed if there had never been cases in which human beings agreed that what they did in attributing knowledge to themselves and others was done *sincerely* and *successfully*. Indeed, when we learn to engage in the practice of knowledge attribution, part of what we have to learn to do is to distinguish between sincere and insincere attributions as well as between successful and unsuccessful attributions.

In sincere knowledge attributions, the knowledge attributor typically understands herself as playing the role of a knowledge acquirer or the role of a knowledge examiner.²⁶ To be a knowledge acquirer is to acquire knowledge by finding someone who knows what one wants to know, someone who can provide one with the knowledge concerned;²⁷ to be a knowledge examiner is to examine whether someone who claims to know something really knows it. For a knowledge acquirer, positive knowledge attributions are more important than negative ones; a

²⁶ This distinction between knowledge acquirer and knowledge examiner is inspired by Bernard Williams’s discussion of what he calls ‘the examiner situation’ in Williams (1970) and Williams (1972).

²⁷ We sometimes say that we can acquire knowledge by finding *something*, such as a book or a fossil, that contains the knowledge that we want, but it is clear that this is a derivative use of the word ‘knowledge’ — a book does not know anything, only the author does; and a fossil cannot contain any knowledge unless it contains something that has already been known by someone.

positive knowledge attribution is an expression of his judgment that the knowledge provider has been found. By contrast, a knowledge examiner may, so to speak, pass or fail someone in a knowledge examination by making, respectively, a positive or a negative knowledge attribution. None of this implies that there cannot be particular contexts of knowledge attribution in which there is no knowledge acquirer or knowledge examiner, but these are atypical cases which would not be possible without typical cases in which there is a knowledge acquirer or a knowledge examiner.

Moreover, in most cases of knowledge attribution, the knowledge acquirer is *successful* or the knowledge examiner is *qualified*, that is, the knowledge acquirer finds her knowledge provider or the knowledge examiner possesses the knowledge about which she tests others. Again, this does not imply that there cannot be particular contexts in which a knowledge acquirer fails to find her knowledge provider. A knowledge acquirer may fail for various reasons; sometimes it is because the knowledge provider looked for simply does not exist (that is, no one has the knowledge). But if there had never been any knowledge provider, there would not have been any knowledge acquirer either. There may also be particular contexts in which a knowledge examiner does not possess the knowledge about which she tests others. Sometimes we can tell whether someone knows something even if we ourselves do not have the knowledge. But if no one had ever known anything about the world, it would not make sense for any of us to test whether someone knows something in particular under certain circumstances.

The practice of knowledge attribution as we understand and engage in it thus presupposes that we have knowledge of the world. Some may object that what it takes for there to be such a practice is that we *believe* that we have knowledge of the world rather than that we do have such knowledge. Yes, such a belief is required, but it is important to note that I am speaking of the

practice of knowledge attribution *as we understand and engage in it*. There is no such thing as the practice of knowledge attribution as we understand and engage in it if none of us knows anything about the world. It is in this sense that the practice presupposes our knowledge of the world, and it is from the perspective as *practitioners* of knowledge attribution that contextualists are unable to take skepticism seriously. An analogy may help here. The practice of prayer as religious believers understand and engage in it presupposes the existence of a personal God. We could say that what it takes for there to be the practice of prayer is that religious believers believe that a personal God exists and listens to their prayers, but this would not be the practice of prayer as religious believers understand and practice it. There is a sense in which there is no such thing as prayer if a personal God does not exist.

To make particular knowledge attributions is to engage in the collective practice of knowledge attribution. It is because of our having learnt to attribute knowledge under different circumstances, and because of our understanding of the practice as a result of such learning, that we are able to make sense of particular knowledge attributions by ourselves and others. The more we are able to make sense of particular knowledge attributions under different circumstances, the more we are part of the practice — or custom — of knowledge attribution, and the more difficult it is for us to take skepticism seriously. Since contextualism applies only to particular knowledge attributions under particular circumstances, it is very difficult for contextualists to detach themselves from the practice of knowledge attribution when they devote their attention to how contextualism works in particular knowledge attributions, and hence very difficult for them to take skepticism seriously.

Should we then say that contextualists beg the question against skepticism by assuming, *qua* practitioners of knowledge attribution, that we have knowledge of the world? The answer is

more complicated than a straightforward ‘Yes’. For one thing, contextualists do not attempt to refute skepticism, so there may not be a problem of begging the question even though they do assume that we have knowledge of the world; for another, contextualists may be willing to grant that skepticism is true in philosophical contexts. Now even if contextualists do not beg the question against skepticism, they are in no position to solve the problem of skepticism if solving it involves explaining *in the most general way* how we have knowledge of the world. As for philosophical contexts, it is not clear how they should be understood with respect to the practice of knowledge attribution. If they are contexts in which knowledge attributions are made in the normal way, then they cannot be contexts in which we can understand ourselves as knowing nothing about the world. But if they are contexts in which knowledge attributions are supposed to be made independently of the normal practice, then we may not be able to make sense of such knowledge attributions — philosophical contexts of knowledge attribution may simply be a fantasy of contextualists.²⁸

6. What the Skeptic May Learn from How We Use the Word ‘Know’

Contextualists have not shown that skepticism is not a real threat to human knowledge of the world, nor have they shown how human knowledge is possible. The skeptic thus has nothing to learn from them. Although Stroud has never criticized the contextualist anti-skeptical strategy directly, what he says below is close to my conclusion that the skeptic has nothing to learn from contextualists’ view on how we use the word ‘know’:

Identifying the assumptions about sense-perception that are responsible for the ease with

²⁸ For a discussion of the difficulty of understanding the notion of a philosophical context, see Wong (2006), sections V-VII.

which a sceptical conclusion can be reached, and learning why those assumptions are unacceptable and what should be put in their place, would be a rich reward for taking the traditional sceptical reasoning seriously and trying to get to the bottom of it. We could perhaps even discover how knowledge by perception is to be correctly understood. That would be a substantial, positive payoff: much better than what is offered by those who see in the sceptical reasoning nothing more than arbitrary insistence on impossibly high standards for knowledge or one or another form of confusion, equivocation, or legerdemain. (Stroud, 2009, p.559)

This does not, however, mean that the skeptic cannot learn anything from how we use the word 'know'. How we use the word 'know' is not just a matter of how truth conditions of knowledge-attributing sentences are determined. By considering the important roles that the practice of knowledge attribution plays in our lives and in our understanding of the world around us, the skeptic may realize a grave difficulty in being a skeptic, a difficulty that he has not been aware of so far. The difficulty is that the skeptic cannot avoid engaging in the practice of knowledge attribution even though he believes that skepticism is true. The more general lesson the skeptic may learn is that, as Stroud puts it, "we can understand ourselves as perceiving and believing and knowing things only from within whatever position we already occupy in understanding the world, and so only as part of understanding everything else that is so in the very world we all perceive and believe in" (ibid., p.568).²⁹

²⁹ I am grateful to Jason Bridges and Niko Kolodny for very helpful comments.

References

- Cohen, S. (1987). Knowledge, context, and social standards. *Synthese*, 73, 3-26.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspective*, 2, 91-123.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57-89.
- Cohen, S. (2000). Contextualism and skepticism. *Philosophical Issues*, 10, 94-107.
- Cohen, S. (2001). Contextualism defended: Comments on Richard Feldman's 'Skeptical problems, contextualist solutions'. *Philosophical Studies*, 103, 87-98.
- Conee, E. (2005). Contextualism contested. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp.47-56). Oxford: Blackwell.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913-929.
- DeRose, K. (1995). Solving the skeptical problem. Reprinted in K. DeRose & T. A. Warfield (Eds.), *Skepticism: A contemporary reader* (pp.183-219). Oxford: Oxford University Press, 1999.
- DeRose, K. (1999). Contextualism: An explanation and defense. In J. Greco & E. Sosa (Eds.), *The Blackwell guide to epistemology* (pp.187-205). Oxford: Blackwell, 1999.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review*, 111, 167-203.
- DeRose, K. (2004). Single scoreboard semantics. *Philosophical Studies*, 119, 1-21.
- DeRose, K. (2005). The ordinary language basis for contextualism, and the new invariantism. *Philosophical Quarterly*, 55, 172-198.
- Descartes, R. (1988). *Selected philosophical writings* (J. Cottingham, R. Stoothoff & D. Murdoch, Trans.). Cambridge: Cambridge University Press.

- Ellis, J. (2010). Stroud's proposal for removing the threat of skepticism. In this volume.
- Feldman, R. (1999). Contextualism and skepticism. *Philosophical Perspectives 1: Epistemology*, 91-114.
- Feldman, R. (2001). Skeptical problems, contextualist solutions. *Philosophical Studies*, 103, 61-85.
- Feldman, R. (2004). Comments on DeRose's 'Single scoreboard semantics'. *Philosophical Studies*, 119, 23-33.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Clarendon Press.
- Huemer, M. (2001). *Skepticism and the veil of perception*. Lanham, MD: Rowman & Littlefield.
- Klein, P. (2000). Contextualism and the real nature of academic skepticism. *Philosophical Issues*, 10, 108-16.
- Kornblith, H. (2000). The contextualist evasion of epistemology. *Philosophical Issues*, 10, 24-32.
- Lewis, D. (1996). Elusive knowledge. Reprinted in K. DeRose & T. A. Warfield (Eds.), *Skepticism: A contemporary reader* (pp.220-219). Oxford: Oxford University Press, 1999.
- Ludlow, P. (2005). Contextualism and the new linguistic turn in epistemology. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp.11-50). Oxford: Clarendon Press.
- MacFarlane, J. (2005). The assessment sensitivity of knowledge attributions. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology*, volume 1 (pp.197-233). Oxford: Oxford University Press.
- Maddy, P. (2007). *Second philosophy: A naturalistic method*. Oxford: Oxford University Press.
- Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*,

96, 317-333.

Sosa, E. (2000). Skepticism and contextualism. *Philosophical Issues*, 10, 1-18.

Stanley, J. (2004). On the linguistic basis for contextualism. *Philosophical Studies*, 119, 119-146.

Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Clarendon Press.

Strawson, P. F. (1979). Perception and its objects. Reprinted in J. Dancy (Ed.), *Perceptual knowledge* (pp.92-112). Oxford: Oxford University Press, 1988.

Stroud, B. (1977). Taking scepticism seriously. Reprinted in B. Stroud, *Understanding human knowledge* (pp.38-50). Oxford: Oxford University Press, 2000.

Stroud, B. (1984a). *The significance of philosophical scepticism*. Oxford: Oxford University Press.

Stroud, B. (1984b). Skepticism and the possibility of knowledge. Reprinted in B. Stroud, *Understanding human knowledge* (pp.1-8). Oxford: Oxford University Press, 2000.

Stroud, B. (1989). Understanding human knowledge in general. Reprinted in B. Stroud, *Understanding human knowledge* (pp.99-121). Oxford: Oxford University Press, 2000.

Stroud, B. (1994). Scepticism, 'externalism', and epistemology. Reprinted in B. Stroud, *Understanding human knowledge* (pp.139-154). Oxford: Oxford University Press, 2000.

Stroud, B. (1999). Radical interpretation and philosophical scepticism. Reprinted in B. Stroud, *Understanding human knowledge* (pp.177-202). Oxford: Oxford University Press, 2000.

Stroud, B. (2003). Anti-individualism and skepticism. In M. Hahn and B. Ramberg (Eds.), *Reflections and replies: Essays on the philosophy of Tyler Burge* (pp.15-27). Cambridge, MA: MIT Press.

Stroud, B. (2009). Scepticism and the senses. *European Journal of Philosophy*, 17, 559-570.

- Williams, B. (1970). Deciding to believe. Reprinted in B. Williams, *Problems of the self* (pp. 136-151). Cambridge: Cambridge University Press, 1973.
- Williams, B. (1972). Knowledge and reasons. In G. H. von Wright (Ed.), *Problems in the theory of knowledge* (pp.1-11). The Hague: Martinus Nijhoff, 1972.
- Wittgenstein, L. (1969). *On certainty* (G. E. M. Anscombe & G. H. von Wright, Eds.) (D. Paul & G. E. M. Anscombe, Trans.). Oxford: Basil Blackwell.
- Wittgenstein, L. (2001). *Philosophical investigations*, third edition. (G. E. M. Anscombe, Trans.). Oxford: Blackwell Publishing.
- Wong, W. (2002). The problem of insulation. *Philosophy*, 77, 349-373.
- Wong, W. (2005). The skeptical paradox and the indispensability of knowledge-beliefs. *Synthèse*, 143, 273-290.
- Wong, W. (2006). Moore, the skeptic, and the philosophical context. *Pacific Philosophical Quarterly*, 87, 271-287.