MICHAEL WILLIAMS

KNOWLEDGE, REFLECTION AND SCEPTICAL HYPOTHESES

1. TWO KINDS OF CONTEXTUALISM

Keith DeRose defines epistemological contextualism as "the position that the truth-conditions of knowledge ascribing and denying sentences (sentences of the form "*S* knows that *P*" and "*S* doesn't know that *P*" and related variants of such sentences) vary in certain ways according to the context in which they are uttered" (DeRose, 1999, p. 187). As he notes, there are several competing versions of this general view. For example subject contextualism and attributor contextualism divide over whether the shifting standards that a person must meet to count as knowing are set by the context of that person or by that of whoever is describing him as knowing or not knowing. De Rose argues forcefully in favour of the attributor (DeRose, 1999, pp. 190–191). However, subject and attributor are closely related. According to both views, the standards for (truly) saying of a person that (s)he knows that *P* can be more or less severe depending on the (subject's or attributor's) *conversational* context. Thus both are versions of *conversational contextualism*. Indeed, both are versions of what I call *simple conversational contextualism* (SCC). I shall say more about what I mean by this in the next section, where I articulate the position in more detail.

Conversation contextualism contrasts importantly with what DeRose calls *structural contextualism* (DeRose, 1999, p. 190). On this view, hints of which can be found in Austin and Wittgenstein, justification (hence knowledge) presupposes a definite *issue context*. Proponents of structural contextualism tend to see it as an alternative to foundationalism and coherentism. DeRose thinks that it is better seen as a (non-standard) version of foundationalism. Either way, structural contextualism must be distinguished from SCC, since the latter implies *no* view about the structure of knowledge or justification. I have my doubts about this. Accordingly, I shall adopt more neutral terminology and call DeRose's structural contextualism *issue contextualism*.

Erkenntnis **61:** 315–343, 2004.

Do these two forms of contextualism have more than a name in common? I think they do.

First, both sorts of contextualism are committed to the following generic claim about knowledge:

(C) The standards for (truly) attributing (or claiming) knowledge are not fixed but vary (somehow) with the context in which knowledge is attributed (or claimed).

Furthermore, both exploit C in their approach to scepticism. Indeed, they offer competing articulations of the Basic Contextualist Diagnosis:

(BCD) The sceptic reaches his paradoxical results by exploiting the context-sensitivity of epistemic standards. Sceptical conclusions seem plausible because the very practice of sceptical argumentation or "doing epistemology" tends to set epistemic standards so as to make such conclusions true. However, this fact does not invalidate everyday knowledge-claims and attributions, which remain true at everyday ("non-philosophical") standards.

If BCD is correct, the sceptic's mistake is to think that he has discovered, while doing epistemology, that knowledge is impossible. In fact, the *most* that he has discovered is that knowledge is impossible while doing epistemology.

This is an attractive thought. Scepticism is a problem because while, on the one hand, sceptical conclusions are difficult (if not impossible) to accept, sceptical arguments seem (or can be made to seem) intuitively compelling. BCD accounts for this smoothly. As Hume saw, while scepticism leaves us cold in everyday situations, it tends to triumph in the study (the site of "doing epistemology"). But as Hume did not see, the context-bound appeal of sceptical claims reflects the logic of knowledge-attributions, not the psychology of belief. This isn't quite right, for this way of putting things is too concessive for issue contextualists. More of this as my argument proceeds.

So much for common ground. The topic of this paper is the differences in how attributor contextualists and issue contextualists flesh out BCD. I will explore this issue by focussing on a particular question. According to these different articulations of contextualism, why do we take sceptical arguments seriously? In particular, why do we take seriously sceptical hypotheses (that I am a victim of an Evil Deceiver, or a brain in a vat), given that all of us regard them as completely outlandish? I shall argue that issue contextualism has a much better answer to this question. As a result, issue contextualism offers an anti-sceptical strategy that promises much deeper insights into how scepticism arises and how it can be avoided.

2. SCC AND SCEPTICISM

I focus on sceptical hypotheses because SCC is typically applied to the diagnosis of Cartesian scepticism; and Cartesian scepticism is distinguished from Agrippan (regress) scepticism by the use it makes of sceptical hypotheses.

Sceptical hypotheses are a special kind of "defeater" to ordinary knowledge-claims. They are defeaters because they posit situations such that, if they obtained, our ordinary beliefs would not amount to knowledge, either because our beliefs would be false or, even if true, epistemically defective. They are special in that they involve *systematic* error or deception. Thus in both Demon deception and brain-in-vat cases, our experience is manipulated to mimic the experience we have in (what we take to be) our "normal" world. Because such scenarios incorporate systematic deception, it can seem hard to say how we could know that they do not obtain.

A standard deployment of sceptical hypotheses is found in what Keith DeRose calls the Argument from Ignorance (DeRose 1995, p. 183). Let O be some ordinary claim – e.g the claim that I have two hands – and H be some appropriate sceptical hypothesis. The sceptic now argues:

(AI) I don't know that not-*H*.
If I don't know that not-*H*, I don't know that *O*.
So: I don't know that *O*.

Since O could be any ordinary claim about an external object, it seems that we have no knowledge of external things.

Advocates of SCC treat AI, or some obvious variant, as the canonical form of sceptical arguments. Since SCC comes in several versions, reflecting differences over how to understand knowledge, proponents of SCC (SC contextualists) do not all say exactly the same thing about AI. Still, there are certain shared commitments, which I state in terms of SCC's "attributor" version:

(SCC1) "Context" is first and foremost *conversational* or *dialectical* context. That is to say, the standards for attributing (or claiming) knowledge depend on what explicit knowledge-claim has been made, or on what error-possibilities have been brought up, are being attended to, or are otherwise "salient".

(SCC2) The more "remote" (far-fetched, improbable) the error-possibilities ("defeaters" to our knowledge-claim) we feel obliged to consider, the harder it is to have knowledge (or more precisely, the more demanding the standards for truly attributing knowledge).

(SCC3) Everyday epistemic contexts are *restricted*, That is to say, the error-possibilities in play are limited to a restricted range of relevant alternatives to what we claim to be the case, or presuppose in so-claiming. The restricted character of everyday epistemic contexts keeps epistemic standards (comparatively) low.

(SCC4) The epistemic context created by doing epistemology, is *unrestricted*. When reflecting philosophically, we are open to any coherent error-possibility. Thus the effect of "going philosophical" is thus to raise the standards for attributing knowledge to the maximum.

(SCC5) Removing all restrictions on relevant defeaters allows us to bring sceptical hypotheses into play. Alternatively, bringing such hypotheses into play has the effect of lifting all everyday restrictions. Thus, simply presenting AI creates, or tends to create, context in which knowledge-claims about ordinary objects turn false

Of course, what the proponent of this diagnosis gives with one hand he takes away with the other. Knowledge-attributions turn false in the widespread way that sceptics claim only in the peculiar context crated by considering sceptical hypotheses. Ordinary knowledge-attributions made in ordinary contexts are safe from sceptical undermining.

My reasons for talking about simple conversational contextualism should be clear. Calling this form of contextualism "conversational" reflects its stress on the role of conversational context in fixing epistemic standards. Calling it "simple" underlines the idea

that there is a measure of severity for epistemic standards that applies independently of subject-matter. With respect to epistemic standards, context fixes their degree of severity on some *context-independent* scale.

This latter idea is very important. In my view, it is what distinguishes SCC from issue contextualism. Issue contextualists tend to think that epistemic standards are subject-matter sensitive. So, for example, both scientific experiments and historical researches can be conducted according to more or less strict standards. But there need be no answer to the question of whether knowledge in physics is subject to stricter standards than historical knowledge. Physics and history could be too disparate for any such comparison to make good sense. (I am talking here about physics and history as such: in either discipline, standards could deteriorate so that, at a given time, one subject was in a poorer state than the other. A subject could fail to maintain its own standards.) Noticing this feature, some philosophers associate issue contextualism with relativism. I am inclined to demur, though I will not pursue the issue. My interest here is scepticism.[1]

The contrast between SCC and issue contextualism with respect to whether we should see epistemic standards as varying across contexts in ways that go beyond anything that could be captured by a single severity scale goes to the heart of the differences between their approaches to scepticism. SC contextualists think that the sceptic *raises the standards* for knowing. But issue contextualism opens another possibility: that the sceptic changes the subject; that ''doing epistemology'' involves a radical break with ordinary epistemic standards, so that sceptical doubting is *not* an extension of ordinary doubting but at best a partial and deceptive simulacrum of it.

Here is one reason why the difference matters. According to SC contextualists, the sceptic seems right because, in a limited way, he is right. Doing epistemology results in a simple failure of knowledge. The sceptic's *only* mistake is to think that he has shown more than he has. The anti-sceptical strategy on offer is thus one of *pure insulation*. Such a strategy is seriously *concessive*.

A contextualist response to scepticism does not have to be concessive in this way. Identifying scepticism's context-bound character might be only the first step, the second being to call into question the theoretical tenability of the sceptic's context. The point of connecting scepticism with ''doing epistemology'' might be to raise about whether epistemology (so conceived) is worth doing. One way it might prove not to be worth doing is this: doing it may

depend on implicit theoretical presuppositions (presuppositions belonging to what I shall eventually identify as epistemology's disciplinary meta-context) that we have no reason to accept, and even many reasons no to accept. This is what issue contextualists argue, or ought to argue.

SC contextualists do not go down this road. This is not an oversight. The essence of their view is that the sceptic generates his conclusions by exploiting quite ordinary context-shifting mechanisms. To be sure, he exploits them in a surprising way, but there is nothing wrong with that. Sceptical conclusions *seem* intuitive because, within their proper limits, they *are* intuitive. Once the possibility that I am a brain in a vat is visibly on the table, I have no way of knowing that it does not obtain.

Still, a piece of the puzzle is missing. Why do we bring up sceptical hypotheses in the first place? And even if they cross our minds, why do we take them seriously? Certainly, lots of epistemologists take them very seriously indeed, at least in the sense of seeing them as having great theoretical interest. Yet none of us believes that any sceptical hypothesis is true, or even remotely likely to be true. So what is the source of their interest? No diagnosis of scepticism will be satisfactory if it leaves us in the dark on this fundamental point.

Someone might say that taking an interest in scepticism (hence in sceptical hypotheses) just is (an important aspect of) doing epistemology. In epistemology, we try to understand how knowledge is possible. But there is a question about how knowledge is possible only because there are intuitively plausible arguments for the conclusion that knowledge is *im*possible. Of course, no one is under any obligation to be interested in epistemology. If you find epistemology uninteresting, don't do it. Then you can ignore sceptical hypotheses. But if you do take an interest in epistemology, you can't.

I have some sympathy with this: certainly, concern with scepticism has been a driving force behind much epistemological theorizing. But is this really all there is to say about epistemology: that it intrigues some people and not others? I don't think so.

3. IGNORING THE PROBLEM

SC contextualists tend to take interest in sceptical arguments (involving sceptical hypotheses) as a given. Accordingly, while they have a lot to say about how sceptical arguments work, they have little

to say about why sceptical hypotheses are worth taking seriously. This is not an accident: their position lacks the necessary resources. But first let me illustrate the phenomenon.

First David Lewis. According to Lewis, knowledge is infallible in the sense that, for me to know that *P*, my evidence must eliminate every possibility in which not-*P*. (Lewis understands "evidence" and "eliminate" very broadly, so his account of knowledge is externalist in spirit.) However, depending on context, some possibilities may be ignored, so that the standards for "infallibility" are more severe in some contexts than others.

In Lewis's terminology, the standards for knowing depend on our *presuppositions*, where to presuppose proposition *Q* is to ignore all possibilities in which not *Q*. Obviously, if we were free to presuppose, or ignore, whatever we like, knowledge would be far too easy to come by. But we are not free in this way. Rather, there are normative rules governing *proper* presupposition or *proper* ignoring. Lewis suggests a number of such rules. Particularly important is the *Rule of Attention*:

> When we say that a possibility *is* properly ignored, we mean exactly that; we do not mean that it *could have been* properly ignored. Accordingly, a possibility that is not ignored at all is *ipso facto* not properly ignored. (Lewis, 1996, p. 230)

In the light of our current question – why should we take sceptical hypotheses seriously? – this seems a very odd rule. Since Lewis's rules are normative – they govern what it is *permissible* to ignore – we might have thought that the relevant issue is precisely the one Lewis sets aside: i.e. could we have (properly) ignored something, even if we didn't? Not only is the Rule of Attention silent here, it amounts to a license to ignore the question.

Here is another of Lewis's rules, the *Rule of Belief*:

> A possibility that the subject believes to obtain is not properly ignored, whether or not he is right to so believe. Neither is one that he ought to believe to obtain – one that evidence and arguments justify him in believing – whether or not he does so believe. (Lewis, 1996, pp. 226–227)

But again, what about possibilities that we do *not* believe to obtain, or positively believe not to obtain. Absent some reason to take them seriously, are we entitled to ignore them? Not if we notice them apparently.

There is a striking – and completely unexplained – asymmetry in Lewis's attitudes towards ignoring and attending. I *can't* make possibilities irrelevant merely by ignoring them; but I *can* make them relevant merely by paying them some attention. This asymmetry explains Lewis's claim that knowledge is subject to what he calls "the

sotto voce proviso'': *S* knows that *P* iff *S*'s evidence eliminates every possibility in which not *P* – *Psst*, except for those possibilities that conflict with our proper presuppositions (Lewis, 1996, p. 225). The proviso has to be *sotto voce* because, as Lewis interprets the Rule of Attention, a possibility becomes relevant if it so much as crosses one's mind. This is why our knowledge of the world is ''elusive'': it temporarily evaporates every time we so much as think of a sceptical possibility.

This is very implausible. In our normal way of talking, ignoring and noticing are not incompatible. On the contrary, ignoring is often deliberate; and one can deliberately ignore only something (or someone) one has noticed. In explaining a problem in mechanics, a teacher might say, ''In this situation, we can properly ignore resistance due to friction''. It would be silly to reply, ''You just brought it up; so you aren't ignoring it; *a fortiori*, you aren't properly ignoring it''.

What all this suggests is that Lewis's Rule of Attention is not first drawn from reflection on everyday conversational rules and subsequently *applied* to doing epistemology. Rather, from the outset the Rule is designed to set the bar for relevance very low, allowing Lewis to finesse the question of why sceptical hypotheses have any call on our attention.

Robert Fogelin has a more ''justificationist'' (and much more subtle) account of knowledge. But he agrees with Lewis that the standards we impose in attributing knowledge go up and down according to the sorts of defeaters we bring into view. In Fogelin's useful terminology, bringing up new defeaters to a knowledge-claim tends to raise the ''level of scrutiny'' to which that claim is subject. When sceptical defeaters are in play, the level rises so high that we find ourselves reluctant to claim any knowledge at all.

So again, why do we bring sceptical defeaters into play. And even if we bring them into play to the extent that we notice them, what makes them relevant? To his credit, Fogelin acknowledges that there is an issue here. As he notes, Descartes, who seems to be the first to have trafficked in sceptical scenarios, had a methodological interest in them. Wanting an absolutely secure basis for knowledge, he hoped to use sceptical thought-experiments to filter out foundational certainties. But Fogelin asks: Given the failure of Descartes's project, why do other philosophers concern themselves with sceptical scenarios, especially today? An enormous amount has been written on the topic of late: what drives this interest? Fogelin admits that he is not sure

that he knows the answer to this question. But he has a suggestion to make:

> Part of the answer, I think, is this: Dwelling on remote defeaters can itself raise the level of scrutiny. Put differently, defeaters that are not salient in everyday life can be made salient simply through intensely reflecting on them. (Fogelin, 2003, p. 108)

I don't much care for this talk of salience, which Fogelin shares with Stewart Cohen, for it tends to blur the distinction between the psychological question of whether something *is* on our minds with the normative-epistemological question of whether it *ought* to be, or whether we would be offending against any important epistemic norms if, having once noticed it, we decided to ignore it anyway. True, in Fogelin's view, it is not enough merely to notice a sceptical possibility: one has to dwell or reflect intensely on it. But why does mere psychological salience establish epistemic relevance?

Fogelin's account of knowledge helps him here. For Fogelin, an important element in the justification required for knowledge is epistemic responsibility: we are disinclined to count as instances of knowledge beliefs that are held or formed irresponsibly (for example, in the teeth of counter-evidence). So if, for whatever reason, you find yourself taking a defeater to a given belief seriously, then to continue to hold that belief, while unable to cope with the defeater, is to be epistemically irresponsible. This is why we are reluctant to claim knowledge, once sceptical scenarios are (rightly or wrongly) in play.

Fogelin's view is a clear improvement on Lewis's. Even so, it is not even part of the answer to the question he poses. At best, Fogelin explains how sceptical possibilities compromise knowledge, if we take them seriously. But why do we take them seriously? There seems to be no answer to this question. Some people are gripped by scepticism, some aren't. Fogelin again:

> When doing philosophy, one can be made to feel the force of Cartesian doubt. There are, it seems, certain philosophers who do not rise to the bait dangled by sceptical scenarios – either because they see danger in the offing or because they simply do not get it. They, perhaps, are blessed. (Fogelin, 2003, p. 108).

But there is an ambiguity in the phrase "Cartesian doubt", which can refer either to Descartes's standards for doubting – anything that is not absolutely certain is subject to Cartesian doubt – or to the principle target of his doubts in the *First Meditation* – our knowledge of the external world. Fogelin starts from the first use but, in the quotation just given, slides to the second. He has no explanation for why we continue to be interested in Cartesian scenarios – hence external world scepticism – given that we have abandoned

Cartesian aspirations. He says that philosophers feel the force of Cartesian doubt. But if they have given up on Cartesian certainty, the force of the doubt cannot be that nothing is absolutely certain. He says that other philosophers see danger. But what danger: that the quest for certainty is doomed? They know that already. In giving his answer, Fogelin loses track of his question.

For a final example, I turn to Keith DeRose. DeRose develops a variant of Nozick's subjunctive conditional analysis of knowledge. In determining whether *S*'s belief that *P* amounts to knowledge, we have to consider whether it matches the fact of the matter, not just in the actual world, but in all *sufficiently close* possible worlds as well. The greater the distance one can stray from the actual world, such that *S*'s belief continues to match the facts, the *stronger S's epistemic position* with respect to *P*. Knowledge is true belief involving a proposition with respect to which one stands in a sufficiently strong epistemic position.

This notion of strength of epistemic position is related to but distinct from Nozick's idea of sensitivity. A belief that *P* is sensitive given that, if it were not the case that *P*, *S* would not believe that *P*. For mundane propositions, strength and sensitivity go together: if I am in even a minimally strong epistemic position with respect to such propositions, my beliefs regarding them will also be sensitive. Last year, I went to Italy for my family vacation. Naturally, all sorts of familiar mishaps might have disrupted the trip; but if they had, I would not believe myself to have taken an Italian holiday. By contrast, my belief that I am not a brain in a vat, though it will match the facts over a wide range of situations, some quite distant from reality, is not sensitive; for I would continue to believe that I am not a brain in a vat, even if I were. So while I am in a *strong* position with respect to believing that sceptical hypotheses are false, this is only because the possibilities they raise are *so* remote that there aren't any nearby worlds in which such hypotheses are true.

Although DeRose does not build sensitivity into his analysis of knowledge, he appeals to sensitivity to explain the plausibility of the sceptic's claim that I do not know that any sceptical hypothesis is false. We have, he claims a strong, though not exceptionless, reluctance to count insensitive beliefs as knowledge. Accordingly, he proposes a Lewis-type *Rule of Sensitivity*, governing the contextually appropriate standards for knowledge:

> 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, p. 205)

Take my claim to know that I went to Italy for my vacation: the Rule will demand sensitivity for my belief and this requirement will be met across an appropriate range of worlds. Now, since I am not a brain in a vat in any world within that range, my insensitive belief that I am not will also match the facts. Accordingly, I will be in as strong a position with respect to my anti-sceptical opinion as I am with respect to what I believe about my recent travels, so that *both* beliefs amount to knowledge at the standards enforced by a mundane knowledge-claim or attribution. However, if I claim explicitly to know that I am not a brain in a vat, the Rule will require sensitivity for that particular belief, vastly expanding the range of worlds over which, to count as knowledge, *any* belief of mine must match the facts. Thus my explicit anti-sceptical knowledge-claim creates a context in which I must be in a much stronger than normal epistemic position in order truly to claim to know *anything*: even where I went for my holidays. In this way, the AI's second premise – if I know that *O*, then I know that not-*H* – is true at every level. And its first premise – that I do not know that not-*H* – is made true by an explicit knowledge-claim or attribution with respect to not-*H*. The plausibility of the argument is thus accounted for. But sine everyday knowledge claims remain true by everyday standards, a blanket sceptical conclusion is resisted.

This view, too, is much more plausible than Lewis's. For DeRose, sceptical possibilities are brought into play by *explicit knowledge-claims*, not by someone's merely mentioning or noticing them. Even so, DeRose does not really explain why we take sceptical hypotheses seriously. To bring the possibility into conversational play, I need to make an explicitly anti-sceptical knowledge-claim: I have to claim to know that I am not a brain in a vat. But I will only make such a claim if I am *already* disposed to take sceptical hypotheses seriously. The source of this disposition remains to be discovered.

Like Fogelin, DeRose recognizes that not everyone resonates to sceptical arguments. Commenting on Barry Stroud's claim that Descartes's ''dreaming'' argument ''appeals to something deep in our nature and seems to raise a real problem about the human condition'', DeRose notes that some people have a quite different reaction, encapsulated by the exclamation ''Aw, come on''.[2] He takes this reaction to imply a judgement that the sceptical argument is extremely weak. His reply is that it isn't extremely weak but rather strong: it is

clearly valid and, considered individually, its premises "enjoy a good deal of intuitive support". You may think you know that you are not a brain in a vat, but how *could* you know such a thing?[3]

This response anticipates the concessive character of the response to scepticism DeRose favours, but that is about all. Notice that DeRose's rhetorical question implicitly excludes the Moorean response that ordinary knowledge excludes sceptical possibilities: knowing that I have hands, I know that I am not a brain in a vat. It also excludes the Quinean variant of Moore's response: that science excludes them. Maybe it's *logically* possible that I'm a brain in a vat, but it's not *technologically* possible. In posing his apparently innocent question – How *could* you know that you are not a brain in a vat? – DeRose has already projected himself into a context in which ordinary (and scientific) knowledge has been put up for grabs. How does this happen? The answer is: by virtue of our *already* being determined to take sceptical scenarios seriously. Why?

DeRose's answer seems to be Fogelin's: some of us are just built that way. However, for DeRose, who is very sensitive to questions of burden of proof, this is a problematic move. As he notes, his own response to scepticism involves assuming things that the sceptic claims we can't know. For example, in claiming that his belief that he has hands is sensitive, he betrays his conviction that he is not a brain in a vat, in this or in any nearby worlds. Is it legitimate to use this conviction against the sceptic? According to DeRose, the answer is "No", if we are playing King of the Mountain: that is, if we are trying to prove, in terms acceptable to the sceptic, that scepticism is false. But this cuts both ways, so that

> . . .if the sceptic 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 favour him. (DeRose, 1995, p. 215)

While this is a fair question, it can just as well be posed by the "Aw, come on" school of Moore and Quine. If our aim is not to refute the sceptic, once he has been handed all the cards, why take the long way round? Why not use our commonsense convictions to head him off right at the beginning? Particularly, why not do this, given that we have no interest in Cartesian certainty?

One reply is that the aim is not to *refute* the sceptic but to *understand* him. A good response to scepticism should be diagnostic and not merely dialectical. After all, the sceptic is not so much an opponent as a personification of our own tendencies to be swayed by

sceptical arguments. Moore's way with the sceptic is thus not so much wrong as unenlightening. It gives us no insight into why we find scepticism so plausible.

While I am sympathetic to this reply, it takes a lot for granted. DeRose's talk of "deeply felt intuitions" is reminiscent of Stroud's claim that scepticism appeals to something "deep in our nature". DeRose never so much as considers the possibility that our intuitions – which are the intuitions of professional philosophers – are just an artifact of our philosophical education. Still less does he consider that the possibility this education itself involves initiation into a tradition that has long since slipped into a degenerate, scholastic phase. DeRose wants a partially vindicating explanation of scepticism's intuitive appeal, one that reveals the sceptic as subtly misusing intuitions about knowledge that deserve to be endorsed, rather than explained away. There is nothing wrong with looking for such an explanation. But we should remember that such a strategy encourages us to think that the sceptic must be right about something. In consequence, it has an inbuilt tendency to favour responses to scepticism that are to some degree concessive. As we noted SCC is seriously concessive. We should not be surprised.

I will not press this line further. Not that I think that there is nothing to it: quite the contrary. However, there are reasons to suspect that it can't be the whole story. Most philosophers think that philosophical scepticism turns radical. Sceptical arguments seem to show not just that we fail to have knowledge by exalted Cartesian standards, but that we fail even by relaxed, everyday standards. This is the answer to Fogelin's question: "Why are we still interested in Cartesian (external world) scepticism?" How does this radical turn even *seem* to happen?

The issue of diagnostic adequacy is crucial. But that has always been my point. The fact that SC contextualists have so little to say about why sceptical possibilities are interesting (beyond the fact that they are useful in arguing for scepticism) is significant because it points to the inadequacy of their diagnosis.

4. REFLECTION ALONE

SC contextualists think that the sceptic makes an unusual (and perhaps deceptive) use of ordinary context-shifting mechanisms. But as Fogelin insists, sceptical arguments presume that the level of scrutiny can be raised by reflection alone. Does this ever happen in ordinary

cases? Or is the sceptic's attempt to change standards peculiar in and of itself?

Let us consider an ordinary example of raising standards by introducing a new defeater.

> *Timetable*. I have an appointment that I cannot afford to miss; but I also have important things to do beforehand. I therefore ask you if you know when the last train leaves that will get me to the city in time. You say you do. Needing to be reassured, I ask you how you know. You say you happen to have just consulted the timetable and offer to show it to me. I notice that it is last year's edition. Have you looked into whether there have been any revisions? No. So your timetable could be out of date. Perhaps you don't know when the train leaves. Certainly, we have some inclination to judge that you don't.

In this case, which I think is quite typical, we find three factors that contribute to raising the level of scrutiny. Two are epistemic (having to do with reasons and evidence) and one is "economic" (having to do with the costs and benefits of getting things right or wrong).

> *Red Flag*. A specific piece of information suggests a particular error-possibility. Here, the fact that the timetable is not current, suggests that it may not be accurate.
> *Background Information*. An ostensible red flag could be neutralized by knowledge that rules out the error-possibility it points to. Relative to our background information, there must be *some* likelihood of that possibility's being realized: the indicated error-possibility must be *live*. Here, I know that timetables do get revised, and I don't know that it hasn't happened in the case at hand.
> *High* (*enough*) *stakes*. Even flagging a live error-possibility may not be enough to make it worth taking seriously. The possibility may be extremely remote and the costs of error low. But if a mistake has serious consequences, even a rather unlikely error-possibility, once flagged, can be worth considering. This is what is going here. While I accept that the chances of the timetable's having been changed are not high, I simply can't afford to miss my train.

In sum, we raise the level of scrutiny by flagging an error-possibility that, relative to our background information and stakes (accounting for opportunity and information-gathering costs), is probable enough to be worth considering.

> Red flags are not always necessary, as we can see by modifying the timetable case:
>
> *Track repairs*. Your timetable is up-to-date. But it occurs to me that, from time to time, trains are delayed because of track repairs. Do you know that no such repairs are taking place today? No. So perhaps you don't know when the last suitable train leaves.

Here there is no red flag. But given appropriate stakes, our background knowledge alone can certify an error-possibility as sufficiently probable. This will happen readily when the probability is reasonably high. However, the fact that error possibilities can become salient without red flags does not show that the level of scrutiny can be raised by reflection alone.

Can attention alone make error-possibilities relevant, at least in any ordinary case? Consider Gilbert Harman's lottery puzzle:

> *Lottery*. I buy one ticket out of n (a very large number) sold. The probability that I lose is $1 - 1/n$, which in a large lottery is very close to 1. But even though I recognize that the chance of my winning is very remote, it does not seem correct to say that I know that I will lose, if this statistical information is all I have to go on. On the other hand, we do not hesitate to see ourselves as gaining knowledge from sources that are less than 100% reliable: for example, the testimony of a trustworthy informant. Yet the probability that what I come to believe is true, given testimonial evidence, may be lower than the probability that I lose in the lottery.

Stewart Cohen has suggested that we can explain this apparent discrepancy in terms of the ways in which the character of our reasons can call our attention to error-possibilities that, in other circumstances, would remain unnoticed. The explanation for our reluctance to grant knowledge in the lottery case

> . . .lies in the statistical nature of the reasons. Although, as fallibilists, we allow that S can know q, even though there is a chance of error. . ., when the chance of error is salient, we are reluctant to attribute knowledge. Statistical reasons of the sort that S possesses in the lottery case make the chance of error salient. The specification that S's reason is the $n-1/n$ probability that the ticket loses, calls attention to the $1/n$ probability that the ticket wins. Our attention is focused on the alternative that the ticket wins and this creates a context in which we are reluctant to attribute knowledge, unless S has some independent ground sufficient for denying the alternative. (Cohen, 1988, p. 106)

Is this a case in which mere attention creates salience, and hence relevance? I don't think so. The lottery case is the track-repair case all over

again: i.e. a case in which an error-possibility becomes salient *via* the interaction of background information and practical interests? When we enter a lottery, in the hope of winning a large sum of money, we decide then and there that nothing will count as conclusive evidence of losing short of the result's being officially announced. Otherwise, if we know that we are going to lose, why buy the ticket, or why not throw it away? Reflecting on the character of our reasons for thinking we will lose doesn't raise the level of scrutiny. It has been raised already.

Some evidence for this suggestion can be found (at least according to my intuitions) in the asymmetry between the first and third person cases. Imagining myself as a lottery participant, I find myself reluctant to judge that I know that I will lose. But imagining myself reacting to my friend's constantly playing the state lottery, I think I know that he is wasting his money. True, Cohen presents a third-person case. But he does so in a way that invites us to focus on *S*'s reasons. In effect, he invites us to put ourselves in *S*'s shoes. Putting ourselves is *S*'s practical situation is what elicits the reaction that we don't know that *S* will lose.

If this is right, we do not need to get involved with scepticism in order to understand how the level of scrutiny gets raised in ordinary situations. If anything, examination of ordinary context-shifting intensifies our sense of the extraordinary character of the context supposedly created by "doing epistemology". Given the appropriate combination of background information and stakes, we can see how merely thinking of a defeater can raise epistemic standards. But this is a far cry from showing how standards can intelligibly be raised by reflection *alone*.

In sceptical reflections, standards cannot be raised in the ordinary ways just scouted. Since sceptical possibilities are designed to be (supposedly) ineliminable, we cannot have evidence for or against them, and they cannot be flagged. Since, if taken seriously, they eliminate background information, along with specific claims, we have no way of estimating the likelihood of their being realized. True, most of us believe that their probability is vanishingly small. But this belief reflects our common-sense and scientific picture of the world, which the sceptic means to put up for grabs. Finally, the notion of stakes has no clear application. If I am a brain-in-a-vat, I won't catch my train, and I won't make my appointment. But in the image I will "catch my train" and "make my appointment". So it is all the same to me, as far as I will ever know.

To sum up: attributor contextualists want sceptical doubt to involve a natural extension of ordinary doubting: they want it to be no

more than an extreme instance of raising the level of scrutiny by introducing new defeaters. But they also want sceptical defeaters to be ineliminable, so that knowledge-claims really do go false in "sceptical" contexts. This means that sceptical defeaters must be capable of being made relevant in a way that does not seem to be ordinary at all, for in ordinary cases, mere logical possibilities may properly be dismissed. If anything, reflection on ordinary situations suggests that conversational developments induce standard-shifts only in definite issue-contexts (set by stakes and background knowledge). Attributor contextualism starts to look superficial. If it has anything going for it, it is that its friendliness to sceptical hypotheses offers an illuminating diagnosis of scepticism. But as we have seen, this friendliness is not well-motivated. We need to do better and I think we can.

5. KNOWLEDGE AS SUCH

None of the foregoing entails that there could not be a reason for taking sceptical hypotheses seriously. But it would have to be a reason of a different type from any we have so far isolated.

The natural question to ask at this point is: what about purely theoretical inquiry? Here there are no external costs. The cost is simply getting things wrong. This isn't really true: there are always opportunity costs, if only those involved in not following up alternative lines of inquiry. But I will not press this point. Questions about constraints on theoretical inquiry open up a new and fruitful line of investigation.

We can agree that doing epistemology is a form of purely theoretical inquiry. It is purely theoretical inquiry into the nature of knowledge. Not any particular kind of knowledge, but knowledge as such. Does this explain the relevance of sceptical hypotheses? Fogelin thinks that perhaps it does. He writes:

> why should the activity of philosophizing lead us to take cartesian skepticism seriously? Part of the reason might be that in philosophizing we are not interested in knowledge of any particular kind. We are interested in knowledge qua knowledge. Because of this, nothing puts constraints on the range of relevant or salient defeaters. The act of philosophizing done in a certain way makes every possible defeater salient, and, with that, skepticism is inevitable. (Fogelin, 2003, pp. 108–109)

I am not sure that I can imagine a context in which *every* possible defeater to some claim is *salient*. If everything stands out, nothing does. I take it that what Fogelin really means to say is that, in philosophizing about knowledge, there are no constraints on the range

of relevant defeaters, so that any defeater can properly be made salient by our attending to it. But this doesn't seem right either, for it leaves out something that Fogelin is well aware of, namely, the *special* interest that seems to attach to sceptical defeaters. When he first posed the question of why we take scepticism seriously, Fogelin found himself wondering about why philosophers spend so much time on sceptical scenarios. The explanation cannot be that, in philosophizing, all defeaters are equally salient. Or to put the question the other way around: why is it that, in philosophizing about knowledge, ordinary defeaters seem to be irrelevant?

Fogelin's claim that, when we take an interest in knowledge qua knowledge, there are no constraints on the range of relevant defeaters doesn't seem to be right. But I would go farther: I don't think that it *can* be right.

Some well-known remarks of Wittgenstein are very helpful.

> 163. . . .We check the story of Napoleon, but not whether all the reports about him are based on sense-deception, forgery and the like. For whenever we test anything, we are already presupposing something that is not tested. . .
>
> 337. One cannot make experiments if there are not some things that one does not doubt. But that does not mean that one takes them on trust. . .
>
> If I make an experiment I do not doubt the existence of the apparatus. I have plenty of doubts, but not *that*.
>
> 341. [T]he questions that we raise and our doubts depend on the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn.
>
> 342. That is to say, it belongs to the logic of our scientific investigations that certain things are in deed not doubted.
>
> 343. But it isn't that the situation is like this: We just can't investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put.[4]

That in a given inquiry some doubts are *hors de combat* has nothing to do with either credulity or limited resources. Rather, it is a matter of the *focus* or *direction* of inquiry. What we are looking into is a function of what we are leaving alone.

In a particular discipline, there will be certain quite general presuppositions that serve to give that discipline its characteristic shape and subject-matter. I like to call them "methodological necessities". Together, they determine the *disciplinary meta-context* for all inquiries of a certain genre. However, they generally are not – and probably could not be – exhaustively catalogued. That is why Wittgenstein insists that certain things are *in deed* not doubted.

The point that what is and is not up for grabs determines the focus of interest applies at all levels. So particular contexts of inquiry, within a given genre, will have their peculiar local presuppositions, as will quite ordinary contexts of epistemic evaluation. But as we saw in the previous section, ordinary contexts typically involve practical concerns.

Fogelin's talk of "no constraints", like Bernard Williams's talk of philosophy as "pure inquiry",[5] suggests that philosophizing is not only free of practical considerations but entirely presuppositionless. However, a form of inquiry that was presuppositionless would be no form of inquiry at all. Fogelin's suggestion that, in studying knowledge qua knowledge, all possible defeaters are salient is not a slip. It points to a fundamental misconception.

The question to ask is this: What presuppositions determine the disciplinary meta-context for investigating knowledge as such. Or rather, since the issue is Cartesian scepticism, what is the disciplinary meta-context for investigating the possibility of knowledge of the external world. We can approach this question by considering a specific instance of AI:

(SAI) I don't know that I am not a brain in a vat.
If I don't know that I am not a brain in a vat,
I don't know that I have hands.
So: I don't know that that I have hands.

In DeRose's schematic presentation of AI, the claim that I have hands is represented by "*O*" . The choice of letter is meant to suggest how the sceptical hypothesis *H* can be used to show that knowledge fails even in the case of the most "ordinary" claim. But the claim that I have hands is very far from ordinary. The oddity of the sceptical defeater is matched by the oddity of the claim it defeats. Even though I have hands and know that I do, this is not something that I would ordinarily have occasion either to claim or to claim to know. Perhaps if I were involved in some grisly accident, I might be relieved to discover and pleased to announce that I have hands. I might even assure you that I know that I have hands: I checked. But outside such outre circumstances, either claim would be distinctly odd.

There is one exception: the claim may not seem so if *we are discussing scepticism*. The claim that I have hands, or that I know that I have hands, is not so much "ordinary" as Moorean. And a Moorean claim is a set-up for the sceptic. Two essential features let it do its work.

One feature is semantic, having to do with content: i.e. with "claim" as what is claimed, what proposition asserted. A Moorean claim – and we could just as well say Cartesian claim – asserts the existence of what Stanley Cavell calls a "generic object": the sort of thing that anyone can recognize, though without being able to say how.[6] Contrast Moorean claims, involving generic objects, with "Austinian" claims, involving things with definite identifying features.[7] To borrow Austin's own example: if I say, looking at a bird, "That's a goldfinch", and you ask me how I know, I can reply "By its red head". This is not the sort of thing that just anyone would know. Being able to spot goldfinches is a matter of (mildly) specialized knowledge. It follows that an Austinian claim might exemplify a particular kind of knowledge (ornithological knowledge, say), but cannot stand for knowledge as such. But this is just what Moorean claims are supposed to do. As claims involving generic objects, they are intended as generic – thus representative – claims. Reference to generic objects is a generalizing device.

The second feature is pragmatic, having to do with "claim" as an act of claiming. As we saw, considered as a speech-act, what the sceptical argument represents as an ordinary claim is really quite extraordinary. Equally clearly, the pragmatic oddity of a Moorean claiming flows from what is claimed. Typically, to assert the existence of a generic object is to give voice to something that anyone can be expected to know; and why would anyone do that? For this reason, Cavell, says that a Moorean claim is defective, by virtue of being entered in a "non-claim" context. However, that is not how the sceptic (or traditional epistemologist) sees things. To be sure, Moorean claims, entered out of the blue, are (seemingly) not tied to any special occasion of utterance, where for one reason or another there might be a question of checking up. But the sceptic sees this as effecting an essential decontextualisation. Detaching a claim from all specific contexts of utterance ensures that its epistemic appropriateness will depend *entirely* on generic epistemic factors. A Moorean claim is intended as a generic claiming. Again, the generic objects play an essential generalizing role, directing our attention to knowledge as such.

From the very outset, then, in the context of philosophical reflection on our knowledge of the world, we are trying to understand knowledge of the world in general. In a sense, the possibility of such knowledge is up for grabs before the sceptical conclusion is reached. This explains why SC contextualists join the sceptic in not backing an

"ordinary" claim against a sceptical hypothesis, as we saw in the case of DeRose.

Now let us turn to how different kinds of claims may be challenged and defeated. Austinian claims invite what we may call *criterial challenges*. The bird is a goldfinch, I claim because it has a red head. You reply: "For all that shows, it could be a goldcrest, for they have red heads too". The challenge invokes a defeater in the form of an alternative factual possibility uneliminated by what I have offered as conclusive evidence (an identifying feature). However, this kind of challenge is as useless to the sceptic as the claim it challenges for it is as specialized as the claim it threatens to defeat. While it might show that I don't know that the bird on the fence is a goldfinch, it has no tendency to show that I don't know that there is a bird on the fence.

In the goldfinch case, I respond to your "How do you know" by citing an *identifying feature*. Clearly, in the case of a generic object, such a response is out of the question. Generic objects thus demand *epistemic* challenges. The suggestion must be that, in some way, I am not well-placed to make my claim: that I am under some kind of *epistemic disability*. For example;

There's a bird on the fence.

How do you know? I can't see anything. (So you can't either.)

But maybe while *you* can't see the bird (from where you are standing, it is hidden by a bush), I can. Local epistemic disabilities are of no use to the sceptic because they are in principle remediable (by an improvement in local circumstances). What the sceptic needs is an epistemic disability that is not tied to specific situations. Any disability that pervasive will of course threaten, to be irremediable. Sceptical hypotheses fill the bill. This is what makes them relevant: they belong essentially to the disciplinary meta-context of the study of knowledge as such. If introducing them raises the standards for knowing, this is an effect of their generic character. Remoteness is neither here nor there. Indeed, as we shall see in the next section, in the context of philosophical reflection on knowledge as such, it is not easy to maintain that sceptical possibilities are all that remote.

To sum up, by prescinding from specific identifying features, Moorean claims invite epistemic defeaters. But by also prescinding from particular circumstances of claiming, they invite the introduction of defeaters citing generic epistemic disabilities. Moorean claims and sceptical hypotheses are made for each other. By involving them, AI offers a paradigm for reflection on knowledge as such. But not so

fast. While I have been speaking of sceptical hypotheses as defeaters for Moorean claims, Moore thought that such claims defeat scepticism. We seem to have a standoff. And if we share Moore's intuition, why not resolve it in our favour by appealing to DeRose's point that we are not playing King of the Mountain? The answer to this is that AI does not reveal the full structure of the sceptic's thought. Sceptical hypotheses are only *indirect* defeaters to Moorean claims. They are direct challenges to a Moorean claim's *implied epistemic commitments*. This is the only way a sceptical hypothesis can suggest an epistemic *disability*.

Suppose, following Moore, I say "Here is a hand" and you ask "How do you know?". I can hardly reply, citing an identifying feature, "By the fingers". This is not because fingers are not identifying of hands: they are. The problem is that fingers are no easier or harder to identify than hands. Here we see another important function of reference to a generic object: to shift the focus of attention from the *character* of an object (which we agree is there) to its *existence*. This shift goes along with the shift from a criterial to an epistemic challenge. So if I agree to entertain your question, I have to give an epistemic response; "There's a hand here because I can see that there is".[8]

Even this will not automatically take the sceptic where he wants to go. For ordinarily taken, this "reply" may be a *dismissal* of the question. Its force may be: "I can see it (you idiot)". Or as DeRose likes to say, "Aw, come on". However, the sceptic has another way of taking this response. "I can see it" becomes "By means of the senses". Thus the sceptic takes "I can see it" to make reference to a *generic source of knowledge*, which must be presumed to be reliable, if the reply that invokes it is to secure epistemic entitlement. Naturally, he does not see himself as putting forward a tendentious re-interpretation of an ordinary response but as making explicit an epistemic presupposition involved in any claim to knowledge of the world: that the senses are a reliable source of worldly information. But am I epistemically entitled to this presupposition, given that I could be a brain in a vat?

We might think that the idea of the senses as a generic source of knowledge is innocent enough. Indeed, who could deny that there are such sources: perception, memory, testimony, and so on. However, while talk of generic sources might be innocent in itself, the sceptic (or traditional epistemologist) interprets it in a tendentious way.

Exploiting the idea of the senses as a "faculty", he encourages us to think of the senses as an information-gathering *module*. This

conception encourages two further thoughts. First, that there is a particular kind of information that the module is designed or adapted gather, so that the sensory module is *informationally bounded*. Second, that the module operates independently of other modules, so that the knowledge it produces is epistemically and semantically independent of all collateral commitments. Accordingly, we can be in possession of the evidence of the senses without being in possession of any knowledge lying outside their informational bounds. The evidence of the senses constitutes an *autonomous stratum* of knowledge. (We are by now well on the way to foundationalism, in its classical form.)

In the context of this conception of the senses, sceptical hypotheses really come into their own. If I am a brain in a vat, my senses are wildly unreliable as a source of information about the world. But surely, they would still tell me something. The thought is by now almost irresistible that they tell me how things look or appear. The evidence of the senses is *phenomenal* evidence. But since these senses are autonomous, phenomenal knowledge must be independent of knowledge of the world. Unfortunately, as sceptical hypotheses also show, phenomenal evidence radically underdetermines worldly facts. It must, since it is knowledge that we would possess even if an appropriate sceptical hypothesis were true. In the very way that they serve to set the sensory module's informational bounds, sceptical hypotheses suggest that those bounds cannot be crossed.

The disciplinary meta-context for investigating knowledge (of the world) as such is rich in presuppositions. It is therefore by no means evident that the sceptic is studying ordinary knowledge. I do not believe that he is.[9] Rather, his very way of framing his questions involves the creation of a special subject-matter: knowledge as such. And to suppose that knowledge of the world, as such, is even a potential object of theory or reflection, we have to conceive of our epistemic capacities in a special way. The sceptic changes the subject in more ways than one.

6. A DILEMMA FOR SCC

From the standpoint of issue contextualism, SCC's chief failing is that it is not contextualist enough. Behind its contextualism is a form of invariantism: the idea that there is a simple scale by which epistemic standards can be judged relaxed or demanding, no matter what the subject at issue. But this view is difficult to maintain, even for advocates of SCC themselves.

SC contextualists want to insulate ordinary knowledge-claims from the results of sceptical reflections. But in practice, this turns out to be easier said than done.

Let us consider two further rules of conversational presupposition suggested by Lewis:

> *Rule of Actuality*. The possibility that actually obtains is never properly ignored.
> *Rule of Resemblance*. "[If] one possibility saliently resembles another [then] if one of them may not properly be ignored, neither may the other."

As Lewis notes, actuality is not eliminated by the subject's evidence. Unfortunately, any possibility uneliminated by the subject's evidence – including sceptical possibilities – resembles actuality in this salient respect. So sceptical possibilities are never properly ignored. Result: scepticism. Lewis suggests making an *ad hoc* exception for resemblances of this type. It would, he says, be better to avoid *ad hocery*. But with admirable candour, he admits that he does not know how to do this.

It should be clear what to say about this. The supposed salient resemblance between our "normal" world and sceptical alternatives is that, in all worlds, the same possibilities remain uneliminated by our "evidence". But this supposed resemblance is an artifact of the special meta-context created by the presuppositions of "traditional" (sceptical) epistemology. Outside of that meta-context, this alleged resemblance has no salience (no relevance) whatsoever. Of course, in saying this, we go beyond SCC to issue contextualism, not modifying Lewis's approach but abandoning it.

To see that this is not just a problem for Lewis, let us turn again to DeRose. An apparent advantage of DeRose's position, as compared with Lewis's, is that it is more purely externalist. DeRose explains the plausibility of the sceptic's claim that we not know not-H by appeal to his modified subjunctive conditionals analysis. Since his approach makes no mention of evidence's failing to eliminate possibilities, he may seem set fair to avoid Lewis's problem. But on closer examination, this apparent advantage proves illusory.

A well-known problem for the subjunctive conditionals analysis is given by Nozick himself in his grandmother case.

> *Grandmother*. A grandmother learns that her grandson is alive and well when he visits her. But if he were dead or gravely ill, the family would find a way to shield her from this upsetting news. Here, SCA is violated: if the grandson were not alive and well, the

> grandmother would still believe that he was. But we are reluctant to deny that she knows her grandson is alive, when she can see that he is.

DeRose agrees with Nozick: examples like this show the need to link the analysis with *methods* of belief-formation. But where Nozick favours a more complex statement of the sensitivity requirement, involving an explicit reference to methods, DeRose suggests that, in determining the range of worlds across which *S*'s belief needs to be sensitive, we should place "heavy emphasis...upon similarity with respect to the method of belief-formation utilised by *S*" (DeRose 1995, p. 196).

This move, plausible enough in its own way, threatens to saddle DeRose with Lewis's problem. In his initial presentation of his response to scepticism, DeRose's language strongly suggests that his idea of "distant" possible worlds involves a *content-based* measure: sceptical possibilities are remote in that they invoke worlds in which things happen that are wildly at variance with our ordinary view, indeed in which most of our ordinary beliefs are false. But if, in judging which worlds are relevantly close, we are to weigh similarity of methods "very heavily", it is up for grabs whether vat-worlds are remote. The sceptic claims they are not, since they resemble the actual world in respect of the role of experience in belief-formation. This is just what Lewis's Rules of Actuality and Resemblance also suggest.

DeRose is under pressure to go along with Lewis here. For not merely does he concede that the subjunctive conditionals analysis of knowledge needs to be linked with methods of belief formation, he sees, that to account for the intuitive appeal of scepticism, methods need to be individuated along the lines suggested by the sceptic.

A belief that a sceptical hypothesis is false will be insensitive. However, mere insensitivity doesn't seem to capture the appeal of scepticism. DeRose himself brings this out in an ingenious way by noticing the problem created by "naked" sceptical hypotheses: e.g. the hypothesis that I falsely believe that I have hands. While insensitive, such a hypothesis would be a very poor candidate for "*H*" in an instance of AI. Clearly, what is wrong with naked sceptical hypotheses is that they stipulate that I am wrong about some ordinary belief without saying how I came to be so mistaken. But if we ask ourselves what makes the brain-in-a-vat hypothesis so disquieting – compared with an ineffective, naked sceptical hypothesis – there is

only one answer: the envatted brain is provided with the *same experience* that an ordinary, embodied person might enjoy. But once we build this reference to common experience into our idea of method of belief formation, and once we commit ourselves to weighing similarity of methods very heavily in deciding which worlds count as "close", sceptical worlds are brought close to the actual world. Sceptical possibilities become relevant alternatives in the most mundane situations. Result: scepticism. At the end of the day, DeRose and Lewis are in the same boat.

Again, the source of the problem is the unidimensional character of SCC, with its emphasis on raising the standards for knowledge-ascriptions. As we noticed at the very outset, this amounts to a significant invariantist element in SCC. In DeRose's version, context determines how far out from the actual world we need to go in determining whether we are in an epistemically strong position with respect to believing that *P*. But the distance measure remains fixed. This is a bad idea. For as we just saw, in order to explain why scepticism is *ever* appealing, we have to allow that the alleged resemblances between our situation and that of the brain in a vat are sometimes salient; and then the problem is to explain in some *non-ad hoc* way why they aren't always so. Issue contextualism provides the answer, for it allows us to argue that a shift of disciplinary meta-context doesn't simply raise epistemic standards (on a fixed scale): it changes the subject. In DeRose's terms, which possible worlds count as close and which as distant *also* changes with context. This is a point that DeRose himself is under pressure to concede, though to concede it is to give up on SCC altogether.

7. CONCLUSION

Issue contextualism links the apparent plausibility with the presuppositions of epistemology, conceived as the study of knowledge as such. But merely to call attention to these presuppositions is not to show that epistemology, so conceived, is in any way objectionable. Even so, the very idea of such a subject is more peculiar than is generally recognized.

Why are we inclined to suppose that it is possible to theorize about knowledge (or knowledge of the world) as such? Why do we think that we ought to be able to "understand human knowledge in general"? Barry Stroud insists that there is nothing obviously absurd about the quest for such understanding. According to Stroud, "we

can and do reflect in very general terms on human beings and their place in the world''. When we do, we find both similarities and differences between human beings and ''other parts of nature''. On the difference side of the ledger, humans, unlike rocks and tree branches,

> do not just move, they do things. Other animals also do things, but humans differ from them in the extent to which they think about what they do, and then act as a result of that deliberation. And, not co-incidentally, also know things about the world around them. Knowledge is essential to deliberation and to informed action (Stroud, 1996, p. 122).[10]

Such reflections lead naturally to questions about knowledge itself:

> What exactly is knowledge, and how do human beings know the sorts of things they have to know to live the kind of lives they lead. These are very general questions, but they are not for that reason alone illegitimate... (Stroud, 1996, p. 123)

There is something odd about these remarks. If we want to know how human beings come to know what they need to know, to live the lives they lead today, we might have thought that, say, the history of science and technology would be the place to go for enlightenment. How could the study of ''knowledge as such'', which abstracts from everything having to do with human life and human interests have anything to tell us in response to the sorts of questions that Stroud raises? In truth, the crucial move is made before Stroud gets to knowledge. It is made when Stroud refers to ''human beings and other parts of nature''. For issue contextualists, practices of argument and inquiry do not belong to *nature* but to *culture*. Accordingly, they have particular histories rather than a common essence. Stroud's apparently commonsense reflections depend on a kind of scientism that he does not examine, or even really acknowledge.

I cannot now further contest the presuppositions of sceptical epistemology's disciplinary meta-context, though it will be obvious to anyone acquainted with the philosophy of the last century that they are eminently contestable. My point is a narrower one. The interest of sceptical hypotheses is tied to the curious meta-context we have been excavating. Thus, to the extent that we find reason to reject any or all of its presuppositions, sceptical hypotheses lose their significance. Merely mentioning them gives them no call on our attention and never did. By supposing the contrary, SCC offers only a shallow and misleading diagnosis of scepticism. If anything, it obscures the features of sceptical argumentation that we need to focus on. By contrast, issue contextualism helps us ask exactly the questions that we need to ask.

NOTES

[1] In Sosa/Kim (eds.) 2000, selections from my own work on issue contextualism are set in a section on relativism. I offer some thoughts about the relation of issue contextualism to relativism in Williams 2001, ch. 19.
[2] Stroud 1984, pp.39. Quoted by DeRose in DeRose and Warfield (eds.) 1999, p. 3.
[3] Ibid.
[4] Wittgenstein 1969, paragraph numbers in original.
[5] Williams 1978. See especially ch. 1. For some extended criticism of Williams, see my *Unnatural Doubts*, pp. 211f.
[6] Cavell 1979. See especially ch. 6. My discussion here is deeply indebted to Cavell. But I have some criticisms too. See *Unnatural Doubts*, ch. 4.
[7] Austin's views, as discussed here, are found in his paper "Other Minds" in Austin 1961. For some excellent recent discussion of Austin, see Kaplan 2000.
[8] A point well made by Cavell. See Cavell 1979, p. 161.
[9] Neither does Kaplan. But Kaplan takes me to task for not sufficiently appreciating Austin's response to the sceptic. However, my criticism of Austin is not that he is wrong, but that he is so unsympathetic to traditional epistemology that he is not interested in excavating its disciplinary meta-context, and so not interested in understanding why doing epistemology that way might even seem to be an appealing project.
[10] Stroud's essay is a response to my *Unnatural Doubts*.

REFERENCES

Austin, J. L.: 1961, *Philosophical Papers*, Oxford University Press, Oxford.
Cavell, S.: 1979, *The Claim of Reason*, Oxford University Press, Oxford.
Cohen, S.: 1988, 'How to be a Fallibilist', *Philosophical Perspectives* **2**, 91–123.
DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* **104**, 1–52. Reprinted in K. DeRose, and T. Warfield: 1999, *Skepticism: A Contemporary Reader*, Oxford University Press, Oxford, 183–219. References to this reprinting.
DeRose, K.: 1999, 'Contextualism: An Explanation and Defence', in J. Greco, and E. Sosa (eds.), *The Blackwell Guide to Epistemology*, Blackwell, Oxford, 187–205.
Fogelin, R.: 2003, *Walking the Tightrope of Reason*, Oxford University Press, Oxford.
Kaplan, M.: 2000, 'To What Must an Epistemology Be True?', *Philosophy and Phenomenological Research* **61**, 279–304.
Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**, 549–567. Reprinted in K. DeRose and T. Warfield: 1999, *Skepticism: A Contemporary Reader*, Oxford University Press, Oxford, 220–239. References to this reprinting.
Sosa, E. and J. Kim (eds.): 2000, *Epistemology: An Anthology*, Blackwell, Oxford.
Stroud, B.: 1984, *The Significance of Philosophical Scepticism*, Oxford University Press, Oxford.
Stroud, B.: 1996, 'Epistemological Reflection on Knowledge of the External World', *Philosophy and Phenomenological Research* **56**, 345–358. Reprinted in B. Stroud: 2000, *Understanding Human Knowledge*, Oxford University Press, Oxford, 122–138. References to this reprinting.

Williams, B.: 1978, *Descartes: The Project of Pure Inquiry*, Pelican, Harmondsworth.
Williams, M.: 1992, *Unnatural Doubts*, Blackwell, Oxford (1998, Princeton University Press, Princeton).
Williams, M.: 2001, *Problems of Knowledge*, Oxford University Press, Oxford.
Wittgenstein, L.: 1969, *On Certainty*, Harper, New York.

Department of Philosophy
The Johns Hopkins University
347 Gilman Hall
3400 North Charles Street
Baltimore, MD 21218
USA
E-mail: mwilliams@jhu.edu