Contextualism, Skepticism, and the Gettier Problem

Berit Brogaard

Synthese **139/3 (2004), 367-86**

Abstract: The contextualist epistemological theories proposed by David Lewis and others offer a view of knowledge which awards a central role to the *contexts* of knowledge attributions. Such contexts are held to determine how strong an epistemic position must be in order to count as knowledge. Lewis has suggested that contextualism so construed can be used both to ward off the skeptic and to solve the Gettier problem. A person knows P, he says, just in case her evidence eliminates every possibility that not-P, where the domain of 'every' is determined by the context. Lewis provides a list of rules that can tell us, for a given context, which not-P possibilities must be eliminated and which can properly be ignored. But his account entails, counterintuitively, that knowledge can truly be attributed even to a person in a Gettier situation provided only that the attributor is ignorant of the fact that the person is gettiered. It has been criticized on those grounds by S. Cohen. In this paper I shall argue that most other forms of contextualism suffer the same fate as Lewis's. The allies of contextualism haven't yet shown us whether contextualism can succeed in maintaining a notion of ordinary knowledge while resisting the absurdity that knowledge can be a matter of sheer good luck. At the end of the paper I shall suggest a possible solution to the problem by showing how Cohen's line of criticism leads to a modified conception of what sort of justification a belief must have to count as knowledge in ordinary contexts.

Introduction

Imagine that Hank sees what appears to be a sheep on a hill. What Hank actually sees is a rock. It happens, though, that behind this rock, and out of Hank's view, is a sheep. This is a typical Gettier counterexample to the thesis that justified true belief is knowledge. In this case, Hank does indeed enjoy

a justified true belief that there is a sheep on the hill, but he does not actually know that there is a sheep on the hill, since it is merely accidental that one is present.

In his paper, 'Elusive Knowledge' Lewis has defended a contextualist solution to the Gettier problem.[1] On a traditional, maximally strict account of what knowledge is, a person can know that P only if she has conclusive justification for her true belief that P. In other words, she can know that P only if her evidence eliminates every possibility that not-P. The idea here is that a person who *knows* that P should not be able to have *exactly the same evidence* in a possible situation in which P is not the case. If a person has evidence that P, then she has perceptual experiences or memories that indicate that P is the case. If she could have the very same P-indicating perceptual experiences or memories even in a possible situation in which P is *not* the case, then her actual evidence does not eliminate every possibility that not-P.

On Lewis's contextualist account of knowledge, in contrast, maximal justification is not always needed. There are contexts in which we are permitted to ignore certain possibilities that not-P. Lewis's definition of knowledge goes as follows.

> *Definition of knowledge*: a person knows that P just in case her evidence eliminates every possibility that not-P except for those possibilities that we are 'properly ignoring'.[2]

What is properly ignored varies from context to context. Thus it is the context which determines how strong a person's epistemic position[3] must be with respect to P in order for her to know that P.[4]

The main virtue of contextualism is, familiarly, that it can do away with skepticism. The problem with skepticism is this: if maximal justification of beliefs is required in order that something be known, then it is impossible to know anything.[5] Contextualism permits us to ignore the not-P possibilities raised by the skeptic in ordinary contexts, and yet still attend to them in philosophical contexts.

But contextualism, Lewis argues, works for the Gettier case, too: Hank fails to know that there is a sheep on the hill (P) because Hank's evidence does not eliminate the not-P possibility that there is only a sheep-shaped rock on a sheepless hill. We, the attributors of knowledge to Hank, are attending to this not-P possibility. Thus this possibility may not—in this context—be properly ignored. Accordingly, Hank does not know that there is a sheep on the hill. Lewis hereby solves the Gettier problem by securing that the possibility cannot properly be ignored that there is no Gettier situation—but rather only a situation that gives rise to a false belief.

Suppose Hank, now, tells Dick that there is a sheep on the hill. Dick himself knows that this is so, and he is wondering whether he can attribute to Hank the corresponding knowledge. Dick sees that Hank is looking toward the putative sheep on the hill, but he does not realize that Hank is in fact looking at a sheep-shaped rock. In his paper 'Contextualist Solutions to Epistemic Problems' Cohen argues that Lewis's solution to the Gettier-problem is flawed.[6] Cohen's argument turns on an appeal to our ordinary intuitions about what counts as knowledge. His main point of criticism is that, on Lewis's account, the subject is counted as knowing P even in a Gettier situation, provided only that the context is such that the attributor is ignorant of the fact that the subject is in a Gettier situation. Thus Dick can truly attribute knowledge to Hank. Cohen, in contrast, argues that it is counterintuitive to assert that a person in a Gettier situation can ever have knowledge, and that it is counterintuitive to suppose that, for purposes of attributing knowledge, one person can be in a better position than another in virtue of his greater ignorance.

This is Cohen's critique of Lewis. But the sort of problem he raises is not unique to Lewis's contextualism. In this paper I shall argue that most other forms of contextualism suffer the same fate as Lewis's. The advocates of contextualism haven't yet shown us whether contextualism can succeed in maintaining a notion of ordinary knowledge while resisting the absurdity that knowledge can be a matter of sheer good luck.[7] At the end of the paper I shall suggest a possible solution to the problem by showing

how Cohen's line of criticism leads to a modified conception of what sort of justification a belief must have to count as knowledge in ordinary contexts.

Four Rules of Relevance

When Lewis says that S knows that P just in case all not-P possibilities have been eliminated except for those we are properly ignoring, he takes the class of possibilities properly ignored to vary with context. But which facts about the context determine which possibilities may properly be ignored?

According to Lewis, eight rules of relevance serve to determine these possibilities. I shall here discuss only four of these, the rest being irrelevant to my thesis.[8] The first of the rules I consider corresponds to the truth condition on the traditional analysis of knowledge as true justified belief. Lewis does not have special rules for justification and belief. The belief and justification conditions are built into his initial definition. The rest of the rules, then, tell us when the justification condition must be obeyed and when it can be weakened.

(1) The Rule of Actuality says that if a possibility that not-P is actual, then it can never properly be ignored by the attributor. This rule (the 'truth condition') captures our intuition that you cannot know P if P is false.[9]

(2) The Rule of Belief says that "a possibility that the subject believes to obtain is not properly ignored, whether or not he is right so to 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."[10] This rule captures our intuition that you cannot know P if your belief that P is contrary to the evidence you have.

(3) The Rule of Resemblance says that if one not-P possibility saliently resembles another, then if one of them may not properly be ignored by the attributor, neither may the other. A possible situation saliently resembles an actual situation if, for example, it is such that the subject can have the same evidence in each. This rule corresponds to an additional 'Gettier' criterion that a true justified belief must satisfy in order to be knowledge.[11] In the Hank case described above, for example, the possibility that

Hank sees merely a rock on a sheepless hill resembles actuality perfectly in respect of Hank's evidence. The actual case of Hank seeing merely a rock on a sheepless hill may not be ignored. Hence, the actual Gettier case may not properly be ignored either.

(4) The Rule of Attention says that the attributor may not properly ignore a not-P possibility to which she is attending. If I call to your attention the possibility that Hank is perhaps not looking at a sheep, but rather at a sheep-shaped rock, then you may not properly ignore this possibility.

When knowledge gets attributed, then two parties are involved: the knowing subject and the attributor. The application of the Rules of Resemblance, and Attention in a given context depends on what is salient to, or paid attention to by, the attributor in that context. These rules select out what *counts as knowledge* in different contexts. They tell us whenever a possibility that not-P can properly be ignored, and the answer they give depends, not on facts about the subject to whom knowledge is ascribed, but rather on whether or not the attributor in that context is aware of the given possibility. Since attributors in different contexts may ignore different possibilities, the very same knowledge attribution may be true in one context of attribution and false in another.

The Rules of Actuality and Belief, in contrast, are applied to the output of this counts-as-knowledge function, so that everything which satisfies these rules *is* knowledge.[12] That is: the Rules of Actuality and Belief are applied independently of who the attributor is: no attributor can properly ignore an actual not-P possibility or not-P possibilities which the subject believes or ought to believe to obtain. If an attributor ignores an actual situation or alternatives thereto which the subject believes or ought to believe to obtain, then the sentence used to make the knowledge attribution is false. The Rules of Actuality and Belief are 'rules' in only a very derivative sense: they are rules which only a god could follow. We might better call them 'conditions' in the sense that they may be satisfied even independently of anyone having any awareness of this fact.

Cohen's Objection

Consider again Hank, who thinks he sees a sheep on the hill. It is possible that there is no sheep on the hill but only a rock that looks like a sheep. This not-P possibility might be actual and yet Hank's evidence would not be changed thereby. The not-P possibility resembles actuality perfectly in respect of Hank's evidence. Suppose Tom is standing next to Hank. Hank tells Tom that there is a sheep on the hill. Tom can see that Hank is looking in the direction of the sheep on the hill, but he sees that in front of the sheep there is a sheep-shaped rock and suspects that Hank might be looking at the latter. In Lewis's terms, the possibility that there is no sheep on the hill but only a rock that looks like a sheep is *salient* to Tom, and so Tom cannot properly ignore it. Hence, Tom is right, in this context, to hold that Hank does not know that there is a sheep on the hill. But consider another person, Dick, who stands some distance from Tom and who has no notion that Hank may be seeing only a rock. The resemblance between actuality and the not-P possibility that Hank sees merely a rock that looks like a sheep is not salient to Dick. Thus, Dick is not aware that Hank is in a Gettier situation. Since the not-P possibility that Hank sees merely a sheep-shaped rock is not salient to Dick, the Rule of Resemblance allows that this not-P possibility may properly be ignored. Hence, Dick can truly attribute to Hank the knowledge that there is a sheep on the hill.

The example is due to Cohen, for whom, given that Hank sees only a rock,

> it is very strange to suppose that there is any context of attribution in which one can truly say of [Hank] that he knows there is a sheep on the hill. The sentence, '[Hank] knows there is a sheep on the hill' looks false (at the world and time), regardless of who happens to be uttering it.[13]

Lewis might explain away this problem by saying that Cohen runs together what can be properly ignored in two different contexts of attribution. What can properly be ignored by Tom in his context of attribution, Lewis might say, should be kept apart from what can properly be ignored by Dick in his context of attribution. Cohen is (Lewis might say) tacitly encouraging his readers to take the perspective of Tom, whose epistemic situation is after all superior to that of Dick. It is from Tom's perspective that it

is strongly counterintuitive to hold that there is a context in which Hank can truthfully be *attributed* the knowledge that there is a sheep on the hill.

But this won't work. Cohen's objection feeds upon Lewis's requirement that a false belief can never be knowledge, or, in Lewis's terms, that actuality may never be properly ignored. For example, if Hank sees *merely* a sheep-shaped rock on a sheepless hill, then no one could in any context truly attribute to Hank knowledge that there is a sheep on the hill. But the attributor who does not know that Hank is in a Gettier situation (wherein the hill is not sheepless) can, in Lewis's view, truly attribute to him such knowledge.

Cohen spells out his argument along the following lines. Compare three cases. $Hank_1$ is looking at a sheep-shaped rock that happens to have a sheep behind it. $Hank_2$ is actually looking at a sheep, but has evidence that should have made him believe that he is looking at a sheep-shaped rock on a sheepless hill. $Hank_3$ is looking at a sheep-shaped rock on a sheepless hill. All three believe the same thing, namely that there is a sheep on the hill. On traditional analyses of knowledge, $Hank_1$ fails to know because he is in a Gettier situation. $Hank_2$ fails to know because, while what he believes is true, this belief is contrary to his evidence. $Hank_3$ fails to know because his belief is false.

On Lewis's account, all three persons fail to know in the context of attribution in which the attributor is apprised of all the information conveyed in our previous paragraph. For in each case the attributor may not then properly ignore the not-P possibility that there is only a sheep-shaped rock on a sheepless hill. $Hank_1$ is in a situation where that not-P possibility saliently resembles actuality, and so by the Rules of Resemblance and Actuality the attributor cannot properly ignore it. $Hank_2$ ought to believe that the not-P possibility that there is only a sheep-shaped rock on the hill is actual, because he has evidence to that effect. By the Rule of Belief, the attributor cannot properly ignore that not-P possibility. In $Hank_3$'s case, the not-P possibility that there is only a sheep-shaped rock on the hill *is* actuality. Hence, by the Rule of Actuality, this not-P possibility, too, cannot properly be ignored.

The attributor is accordingly right to rule out that $Hank_1$ knows, but only if she sees that he is in a Gettier situation. If she doesn't see that, then she can truly attribute to $Hank_1$ knowledge that there is a sheep on the hill. False and unjustified beliefs fail to count as knowledge in any context, but this is not so of Gettier-beliefs, and this asymmetry between what may be properly ignored in the case of false and unjustified beliefs and what may be properly ignored in the case of Gettier-beliefs provides one ground for Cohen's criticism.

Cohen suggests that, one way to amend Lewis's account in such a way that it can solve the Gettier problem might be to allow that the application of the Rules of Actuality and Belief, too, should vary with the attributor. This, however, would permit an attributor to take false beliefs to be knowledge if it is not salient to the attributor that they are false. Cohen, therefore, concludes by suggesting that we must endorse the thesis that the Rule of Resemblance should be construed as being like the Rules of Actuality and Belief, in that its application is independent of the attributor.

One way for the Rule of Resemblance to apply to a context of attribution independently of who the attributor is, is to give up the salience qualification. The salience qualification says that if a not-P possibility is *salient* to the attributor, then it may not properly be ignored. A not-P possibility that resembles actuality in such a way that the subject would have the same evidence in that not-P possibility as in the actual world may, then, not be ignored by any attributor in any context. It would not matter whether the attributor is aware of this possibility or not.

Then, however, we are faced with skepticism; for the world of the skeptic (which includes brains in vats, evil demons and sundry other monsters) is a not-P possibility that resembles actuality perfectly with respect to a knowing subject's evidence. With the salience qualification in the Rule of Resemblance we can say that the world of the skeptic is not salient to us in ordinary contexts, at least not before the skeptic brings it to our attention. Without the salience qualification we are at the skeptic's mercy. In fact, if we drop the salience qualification, we can then never know ordinary things. For in ordinary contexts our evidence is never strong enough to rule out every not-P possibility.

We are thus still left with the problem that an attributor who attributes knowledge to a subject in a Gettier situation can truly do so if only she does not realize that the subject is in a Gettier situation. But it surely seems unsatisfactory that one's ignorance can enable one to be right about a knowledge attribution.[14] The sentence in our opening paragraph to the effect that Hank knows that there is a sheep on the hill seems to be false in all contexts and for all attributors; for Hank is after all not looking at a sheep, but merely at a sheep-shaped rock. We could save Lewis's account by sacrificing the salience criterion in the Rule of Resemblance. But a criterion of this sort is required in order to secure that we can have ordinary knowledge.

DeRose's and Cohen's Proposal

Cohen's criticism of Lewis, I shall now argue, pertains to most forms of contextualism. DeRose combines a tracking condition with contextualism.[15] The problem, if we were to assign knowledge status to an accidentally true belief, is that we would have had the belief even if it were false. To the extent that a belief's falsehood would have no influence whatsoever on our certainty concerning it, we have the best kind of reason we could have for rejecting it as an instance of knowledge. Beliefs are simply too feebly supported when believing them has little to do with whether they are true or not. DeRose credits Nozick for this insight, and makes it central to his own epistemology. He suggests, following Nozick, that it be a necessary condition for knowledge that beliefs track the truth:[16]

Truth Tracking: if P were false, S wouldn't believe that P.[17]

Truth Tracking involves a subjunctive conditional, and so is true when, and only when, its consequent is true in the closest possible worlds in which its antecedent is true. Suppose I truly believe there is a sheep on the hill because I am looking at one. Whether I know it depends on whether I believe it in the closest worlds in which my belief is false. Since the hill contains no sheep look-alikes in the closest worlds in

which my belief is false, then I do indeed know there is a sheep on the hill. For I would not have believed it had my belief been false.

But suppose now that while I am looking at the sheep on the hill, I ask myself whether I am justified in believing there is a sheep on the hill. I wonder whether I can exclude the possibility that the alleged sheep is a dog camouflaged as a sheep. I don't believe it is. But I don't know it is not. For in the closest worlds in which my belief is false, the alleged sheep is a dog camouflaged as a sheep, and in those worlds I believe there is a sheep on the hill. So by Truth Tracking I don't know the alleged sheep is not a dog camouflaged as a sheep. But if knowledge is closed under (known) implication, then if I don't know there is not a sheep look-alike on the hill, I don't know there is a sheep on the hill.

This sort of reasoning is the crux of the skeptic's argument. In keeping with Truth Tracking I don't know I'm not a brain in a vat. For in the closest world in which I am one I would believe I am not one. But if I don't know I'm not a brain in a vat, then I don't know I have hands. The only way to maintain Truth Tracking, unrestrictedly, as a necessary condition on knowledge, and also hold that I know I have hands, even though I don't know I'm not a brain in a vat is to reject the Closure Principle. This is Nozick's strategy. But this scheme is unavailable to the contextualist. For she wants to grant that in those contexts in which it is true that I don't know I'm not a brain in a vat, it is true too that I don't know I have hands. And it is unclear why Closure should obtain only in stringent contexts.

The contextualist in favor of Truth Tracking adopts a different strategy. DeRose argues that we should integrate Truth Tracking into a contextualist solution to the puzzle.[18] DeRose calls the method by which the skeptic raises the standards for knowledge the 'Rule of Sensitivity':

> *Rule of Sensiivity*
>
> when it is asserted that some subject knows that P, the standards for knowledge tend to be raised to such a level as to require S's belief in P to track the truth for it to count as knowledge.[19]

Under ordinary low standards I know I have hands, because this belief tracks the truth under lower standards. So the Rule of Sensitivity is satisfied. By Closure I know I'm not a brain in a vat. But that knowledge isn't claimable. For if it were claimed, then by the Rule of Sensitivity we would (tend to) raise the standards to such a level as to require my belief that I'm not a brain in a vat to track the truth. As DeRose puts it,

> on our solution, we do know, for instance, that we're not BIVs, according to ordinary low standards for knowledge. But, though ['we don't know that we are not brains in vats'] is false when evaluated according to those ordinary low standards, we're able to explain its plausibility … by means of the fact that the high standards at which [it] is true are precisely the standards that an assertion or denial of it put into play. Since attempts to assert [it] are bound to result in truth, and attempts to deny it are destined to produce falsehood, it's no surprise that we find it so plausible.[20]

In arguing thus, DeRose is admitting that it may be true that we know that P when P doesn't track the truth under ordinary low standards. Whether the Rule of Sensitivity applies depends on whether the attributor attributes to the subject the knowledge in question. Once the knowledge is claimed, the standards tend to be raised to such a level as to require the subject's belief to track the truth.[21]

One of the main problems with DeRose's theory is that unclaimable knowledge is a peculiar form of knowledge. As Jim Stone argues, my knowledge that I'm not a brain in a vat "can have no epistemic bearing on the question of whether or not I am one, should it arise"—for if it arises, I no longer know I'm not a brain in a vat. Stone finds such contextualism unacceptable because it dignifies "insensitive, unclaimable knowledge" with the title of knowledge. The insensitive, unclaimable knowledge that I'm not a brain in a vat which "must vanish whenever I wonder if I am one" isn't worthy of the title.[22]

Just how damaging is that line of objection against contextualism? Doesn't it beg the question against DeRose's proposal? I think not. The heart of DeRose's proposal, recall, is that the knowledge

that I'm not a brain in a vat *is* unclaimable, and so is not supposed to have bearing on the question of whether it is true that I am one. Yet the belief that I'm not a brain in a vat doesn't even seem justified. Certainly, if justification is a matter of whether one can give, or have, good reasons for holding the belief, then I cannot be justified in believing that I'm not a brain in a vat. For giving, or having, those reasons should be enough to raise the standards to such a level as to require the belief to track the truth, as it would make me consider whether or not I am a brain in a vat.

There is another way in which DeRose's theory cannot do the job as advertised. The Rule of Sensitivity is supposed to capture the idea that only beliefs that are relevantly related to the facts the beliefs are about can be knowledge. For that reason DeRose's proposal can apparently avoid the above Gettier counterexample, and thus seems to be a good candidate for a rule that can substitute Lewis's Rule of Resemblance. If it were false that there is a sheep on the hill, Hank would still believe there is a sheep on the hill. For he believes there is a sheep on the hill, because he is looking at a sheep-shaped rock. Thus, Hank's belief doesn't tracks the truth, and just as long as the Rule of Sensitivity applies in this case, the claim that Hank knows there is a sheep on the hill is false.

Unfortunately, DeRose's theory suffers from the same problems as Lewis's, for recall that in ordinary contexts where our standards for justification are low, beliefs that fail to track the truth may count as unclaimable knowledge. It may be true that I know I'm not a brain in a vat under ordinary low standards; for if no one claims the knowledge, the Rule of Sensitivity doesn't apply. Similarly, if no one attributes to Hank the knowledge that there is a sheep on the hill, then Hank silently knows that there is a sheep on the hill; for whether the Rule of Sensitivity applies depends on whether the attributor attributes to the subject the knowledge in question. So if no one claims the knowledge, the rule does not apply. Hence, if there is such a thing as unclaimable knowledge, as is required by DeRose's solution to the skeptical problem, Hank knows there is a sheep on the hill. That, as far as I can see, presents a problem for DeRose if he wishes to sustain the idea that knowledge cannot be a matter of sheer good luck.

Cohen's main criticism, I take it, is that it is counterintuitive to *assert* that a person in a Gettier

situation can ever have knowledge, and that it is counterintuitive to suppose that, for purposes of attributing knowledge, one person can be in a better position than another in virtue of his greater ignorance. But even this criticism appears to hold up here. When it is asserted that some subject knows P, the standards for knowledge only *tend* to be raised to such a level as to require S's belief in P to be track the truth. If the attributor were unaware that Hank is looking at a sheep-shaped rock, then the attributor wouldn't raise the standards. Consequently, the attributor's claim that Hanks knows there is a sheep on the hill is true. DeRose might insist that the Hank case exemplifies a case where the standards are raised. But if it isn't obvious to the attributor that Hank's belief doesn't track the truth, then what would make him raise the standards in this case, but not in the case where Hank is actually looking at a sheep?

Of course, truth-tracking contextualism doesn't exhaust the contextualist theories on the market. The contextualist may, for example, rely on a relevant alternatives theory of sorts, and attempt to escape the trouble that way. But the relevant alternatives theorist is put in difficulties, too. Cohen argues that S knows P just in case S's reasons rule out all relevant alternatives.[23] There are two ways in which alternatives become relevant, namely by the subject possessing an intersubjectively evident defeater or by the alternative being likely in the actual world. But which defeaters are evident, and which alternatives are likely depend on the way the world is together with the standards of evaluation. If our world is one in which we rarely hallucinate, are rarely deceived by demons etc., skeptical alternatives are not relevant relative to normal standards. Nonetheless, Cohen argues, "if frequently we hallucinate or are deceived by a demon, these alternative can be relevant. If it frequently happens that members of my society are kidnapped and turned into brains-in-a-vat, then that alternative can be relevant".[24] But it is not at all clear that Cohen's relevant alternatives agenda avoids the difficulties he raised against Lewis. I believe that Jones owns a canary in virtue of having seen him buy one the same morning. I am ignorant of the fact, however, that a divorce court deeded Jones's canary over to his wife, and that thousands of miles away Jones's aunt dies and leaves Jones her canary. I don't posses a defeater, and it is rare, suppose, that

divorce courts award canaries to wives. So, that alternative is not relevant relative to normal standards. Yet it is incontestable that I do not know Jones owns a canary.[25] And to insist that the relevance of an alternative doesn't vary with the standards of evaluation but only with the way the world is, or that the availability of a defeater varies neither with obviousness nor with the standards of evaluation would undermine contextualism. For then there would not ever be a way of taking the threat of skepticism seriously.

Inferential Knowledge

But the contextualist need not claim defeat yet. The nub of the Rule of Sensitivity is that in strict contexts, one's belief must track the truth if it is to stand the test in the epistemic court. The main problem with DeRose's proposal is that the Rule of Sensitivity applies whenever we make a knowledge claim. Consequently, any assertion to the effect that a subject knows that she is not a brain in a vat is guaranteed to be false. The problem is that DeRose treats ordinary knowledge as being so elusive that the mere mentioning of a skeptical possibility will trigger a switch of context. But this treatment does not do justice to our ordinary use of the term 'knowledge'. Ordinarily, it will take some convincing to make ordinary people doubt that the world is as they experience it. That is, the possibility that we are deceived is typically taken to be remote. Therefore, if the Rule of Sensitivity is to leave room for ordinary knowledge, it must apply only when the attributor seriously questions the evidence the subject has for believing P. A New Rule of Sensitivity may be formulated as follows:

> *New Rule of Sensitivity*
>
> When the evidence a subject has for believing P is called into doubt, the standards for knowledge tend to be raised to such a level as to require S's belief in P to track the truth for it to count as knowledge

The New Rule of Sensitivity is compatible with the claim that skeptical hypotheses are not seriously entertained in ordinary contexts. Thus, the rule allows for the possibility that we know that the skeptical hypotheses are false in ordinary contexts. And it allows for the possibility that such knowledge is claimable. The latter contention is controversial. It is widely held that we can be sufficiently justified in believing such ordinary claims as the claim that I have hands. The controversy lies with the contention that we can be sufficiently justified in denying skeptical hypotheses. Surely, my sheep-indicating experience gives me a good reason to believe that there is a sheep on the hill but—it is argued—it doesn't give me a good reason to believe that the sheep-shaped entity on the hill is not a dog camouflaged as a sheep.[26]

What is needed to resist the conclusion that we cannot be justified in denying skeptical hypotheses is an argument to the effect that if P is justified, then the consequences of P are justified as well. That is, to resist the conclusion it is not necessary that we present an argument to the effect that the reasons for believing P are also reasons for believing the consequences of P. The latter claim runs into difficulties. We need only argue that justification is closed under known entailment. Peter Klein has presented the following argument for closure of justification:[27] my reasons for believing that there is a sheep on the hill need not be the same as my reasons for believing that I am not looking at a dog camouflaged as a sheep. If my belief that there is a sheep on the hill is sufficiently justified, then that belief ought to be sufficient to justify its consequences. For any evidence against the consequences of a belief P would be evidence also against P itself. Another way to look at it is this. Any consequence of P will be true at least in those worlds in which P is true. So there is a sense in which a consequence of P is justified to the extent P is.

However, the idea that the consequences of a belief P are as justified as P is correct only if we apply some version of the indicator analysis to justification. The indicator analysis, as formulated by Luper and more recently Sosa, says that S knows that P only if S's belief that P is based on the fact that a reason R holds, and P is true in the closest worlds in which R holds (Or: R would obtain only if P were

true).[28] Luper suggests that we understand the subjunctive conditional in a way that is slightly stronger than usual: the subjunctive conditional is true so long as P holds throughout the worlds close to the actual in which R. The closure principle for justification (and knowledge) is sustained by this account. However, the indicator analysis allows Gettier-style counterexamples. This is because Gettier-style counterexamples involve inferential beliefs that may satisfy the indicator requirement even if they are inferred from beliefs that do not. I believe that Mary, a student in my logic class, is born on a weekday based on the fact that every student in my logic class is born on a weekday. My belief satisfies the indicator requirement. But my belief need not be sufficiently justified. It is not sufficiently justified if the reason I believe that every student in my logic class is born on weekday is that my crystal ball (accidentally) gave me the (right) answer.

It may be suggested that we can overcome this problem if we combine the indicator analysis with a reliability condition.[29] But deduction from a true belief is a (hyper)reliable method. So one would need to argue that the method employed above is that of deducing one belief from another, where the latter is acquired by an unreliable method (e.g. a crystal ball). In any event, I believe there is a more natural way to avoid the above counterexample and the obvious obstacles for reliability conditions (such as the problem of generality). The motivation for the indicator requirement was that the consequences of a belief are as justified as the belief itself, as evidence against a consequence of a belief is evidence against the belief itself. Thus, for chains of entailed beliefs, the first in the chain must meet the indicator requirement:

> *Oblique Indicator Condition*:
>
> For chains of entailed beliefs, the first in the chain must meet the indicator requirement if it and subsequent beliefs are to count as known.[30]

With the revised conditions in place of Lewis's Rule of Resemblance we get a contextualism that is

contextualist with respect to *the sort of justification* one must have to be adequately justified in holding a belief. Which sort of justification one must have to be adequately justified varies with the demands the attributor imposes on justification. However the mere mentioning of a not-P possibility isn't sufficient to raise the epistemic standards. For even if it is mentioned that we are brains in a vat, there remains the possibility that the attributor gives the claim no credit, because she has no doubts concerning a knowledge claim that entails its falsehood. But as soon as the attributor *genuinely doubts* that S knows P, she raises the standards to such a level as to require P to track the truth. She thereby changes the meaning of 'know', as she now seeks a sort of justification that requires that P track the truth. One might go so far as to say that knowledge under ordinary low standards and knowledge under raised standards are different kinds of knowledge.[31]

The present proposal defuses the problems with DeRose's theory. The application of the oblique condition is not a function of context. So whether or not gettiered subjects can have knowledge does not vary with the effective epistemic standards. The oblique requirement eliminates that possibility once and for all. Hence Cohen's criticism is averted.

Under ordinary low standards beliefs that fail to track the truth achieve knowledge status when they are correctly inferred from beliefs that satisfy the indicator requirement. So the knowledge that I'm not a brain in a vat is claimable. The proposed contextualist theory hence appears to survive the sort of objection urged by Stone. At least it allows room for the claimability of the knowledge that I'm not a brain in a vat. But does it adequately explain how it is that my knowledge that I'm not a brain in a vat bears on the question of whether or not I am one—should it arise?

I believe it does. Actually considering whether or not P doesn't automatically raise the standards to a level at which I don't know I am not a brain in a vat. For under ordinary low standards, I can settle the issue if I reason as follows: 'I have good reasons for believing that I have hands. And if I have hands, I cannot be a brain in a vat'. This Moorean response to the skeptic, of course, is elusive. Once I seriously entertain the possibility that my brain is disembodied, I must agree with the skeptic. But this should not

come as a surprise, as the elusiveness of knowledge is essential to the contextualist enterprise. In ordinary contexts—contextualism tells us—it is perfectly legitimate to use sense-perception to justify our beliefs about the external world and our relation to it. That is, ordinary contexts demand only that knowledge, properly so-called, must satisfy the oblique requirement. The skeptic wants more. She wants knowledge that satisfies the tracking requirement.

So my response to Stone is this: under ordinary low epistemic standards, I can address the question of whether or not I'm a brain in a vat without beginning to doubt that I know I'm not one. The reason is this: I do not doubt that I know I have hands under ordinary low standards.[32] But then it follows that I do not doubt that I know I'm not a brain in a vat; for if I did, I would also doubt that I know I have hands.

Under raised epistemic standards, by contrast, beliefs must be independently justified against an empirical background, not logically guaranteed on the basis of derivations from an empirically justified belief. I am required to work from a position in which the belief that I'm not a brain in a vat is unjustified to one in which it is, without appealing to empirically justified beliefs that entail it. But obviously I cannot do so. So I don't know I'm not a brain in a vat under the raised standards.

The difference between the constraints imposed on justification under ordinary low standards versus raised standards is reflected in the following closure arguments:

A:

I know that I have hands

If I know I have hands, I know I'm not a brain in a vat

So I know that I'm not a brain in a vat

B:

I don't know I'm not a brain in a vat

If I don't know I am not one, then I don't know I have hands

I don't know that I have hands

Given A I know I'm not a brain in a vat, because I believe it by recognizing that my belief that I have hands entails that I'm not a brain in a vat. Since the soundness of A requires that I never seriously doubt that I know I'm not a brain in a vat, it can be sound only under ordinary standards. If B is sound, I'm already considering the question of whether or not I'm a brain in a vat on *independent grounds*; otherwise I wouldn't have come to the conclusion that I don't know I'm not a brain in a vat. So B can be sound only under raised standards.

Conclusion

Cohen's argument against Lewis's solution to the Gettier problem is that it seems counterintuitive that an attributor who is ignorant of the fact that a subject is gettiered can for this very reason truly attribute knowledge to that subject, while another attributor, who is aware of this fact, cannot. The problem arises because the Rule of Resemblance has a salience qualification that makes it applicable only if the attributor is aware that the rule applies. We cannot forego the salience qualification without being faced with skepticism. For if we do, the not-P possibility that the skeptic's world is actual resembles what is ordinarily believed to be actuality with respect to every subject's evidence. Thus, it may never properly be ignored. But this would mean that we can know (nearly) nothing. The skeptic's world would be lurking around every corner.

I have argued that Cohen's criticism pertains to most forms of contextualism. DeRose, for example, argues that we should integrate Nozick's tracking condition into a contextualist solution to the skeptical problem. Whether the tracking condition applies in a given situation depends on whether the attributor attributes to the subject the knowledge in question. For once the knowledge is attributed, the standards tend to be raised to such a level as to require the subject's belief to track the truth. This

proposal admits that we can have the knowledge that we aren't brains in a vat in ordinary contexts, as long as the knowledge isn't claimed. So the skeptical inference from 'I don't know I'm not a brain in a vat' to 'I don't know I have hands' is unavailable. Unfortunately, allowing the application of the tracking condition to vary with context has the undesirable consequence that a gettiered subject can have knowledge so long as the attributor does not attribute the knowledge to her.

Contextualism is defensible, however, if we adopt a modified conception of what sort of justification a belief must have to count as knowledge under ordinary low standards. On the proposed contextualist theory, the consequences of P may be as justified as P, even if the evidence for P is not transmitted to the consequences. Just as long as the subject's evidence is an indicator of P, the belief P can sustain its consequences under ordinary low standards. This liberty with respect to adequacy of reason clarifies how we can know under ordinary low standards that some incredible yet alternative explanation of our empirical evidence is false without opening the door to knowledge that is a matter of sheer good luck.[33]

References

Cohen, S.: 1987, 'Knowledge, Context, and Social Standards', *Synthese* 73, 3-26.

Cohen, S.: 1997, 'Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery', *Australasian Journal of Philosophy* 76, 289-306.

Cohen, S.: 2002, 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research* 65, 309-329.

DeRose, K.: 1995, 'Solving the Skeptical Problem', *The Philosophical Review* 104, 1-52.

Dretske, F.: 1967, 'Epistemic Operators', *Journal of Philosophy* 67, 1007-1023.

Hawthorne, J.: 2002, 'Deeply Contingent A Priori Knowledge', *Philosophy and Phenomenological Research* 65, 247-269.

Klein, P.: 1995, 'Skepticism and Closure: Why the Evil Genius Argument Fails', *Philosophical Topics*

23, 213-236.

Luper, S.: 1984, 'The Epistemic Predicament: Knowledge, Nozickian Tracking, and Skepticism', *Australasian Journal of Philosophy* 62, 26-60.

Luper, S.: 1987, 'The Causal Indicator Analysis of Knowledge', *Philosophy and Phenomenological Research* 47, 563-587.

Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-567.

Lewis, D.: 1998, 'Index, context, and content', in *Papers in Philosophical Logic*, Cambridge University Press, Cambridge, 21-44.

Nozick, R.: 1981, *Philosophical Explanations*, Clarendon Press, Oxford, 172-196.

Sosa, E.: 1999, 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-152.

Stone, J.: 2000, 'Skepticism as a Theory of Knowledge', *Philosophy and Phenomenological Research* 60, 527-545.

Wedgwood, R.: 2002, 'Internalism Explained', *Philosophy and Phenomenological Research* 65, 349-369.

Wright, C.: 2002, '(Anti-)Skeptics Simple and Subtle: G. E. Moore and John McDowell', *Philosophy and Phenomenological Research* 65, 330-348.

NOTES

[1]Lewis (1996).

[2]*Lewis, ibid.*, p. 554.

[3]A person's epistemic position is stronger the more not-P possibilities her evidence (for P) rules out.

[4]*Context* is here to be understood as *context of attribution.* The context of attribution includes the parts of physical reality upon which the attributor is focusing upon or attending to in any given case. The attributor may also focus on mere possibilities. For example, she may focus upon the possibility that she and the knowing subject are brains in vats. The possibilities focused upon need not be maximally specific; they need only be specific enough for the

purpose in hand (that is: possibilities need not be entire possible worlds). This way of understanding context seems to be in keeping with what Lewis has to say in his (1996). For a fuller discussion of context, see his (1998).

[5]On Lewis's view, knowledge is closed under strict implication rather than known implication (1996, p. 564). Thus, the closure principle for knowledge says that if you know that P, and P entails Q, then you must also know that Q. 'You have two hands' entails the denial of, say, the brain in a vat hypothesis. If you fail to know that you are not a brain in a vat, then you also fail to know that you have two hands.

[6] Cohen (1997).

[7]I here rely on a (widely) shared intuitive understanding of *luck* and *accident*. Arguably, knowledge is compatible with some degree of luck. See e.g. Sosa, 'Beyond Internal Foundations to External Virtues', manuscript, R. Wedgwood (2002, p. 355) and Hawthorne (2002, pp. 254 and 265-6). As Sosa points out, all knowledge depends on the luck of "proper teaching, of proper upbringing, of proper constitution" (p. 14). If the luck of our constitution is of no concern, Sosa asks, then why is the luck of our emplacement so detrimental to knowledge? To provide an analysis of when luck is detrimental to knowledge would be beyond the scope of this paper. The arguments presented in this paper can be seen to rest on the (perhaps disputable) assumption that the sort of luck exemplified in the standard Gettier cases is incompatible with knowledge. Thanks here to an anonymous referee.

[8] The four rules mentioned here are restrictive rules. They tell us what possibilities may not be properly ignored. In addition to these rules, Lewis mentions four permissive rules: the Rule of Reliability, which says that the possibility that a reliable method fails may properly be ignored, two Rules of Method, which say that the possibility that a sample is not representative or that the best explanation of our evidence is not the true explanation may be properly ignored, and the Rule of Conservatism, which says that we may ignore that we know that certain possibilities are ignored in ordinary contexts.

[9] Notice that the Rule of Actuality does not require that the attributor pay attention to the grounds of the subject's belief (e.g. if the subject believes, *truly*, that it is raining on the *false* grounds that the streets are wet, then the Rule of Actuality does not require us to pay attention to the fact that the streets are *not* wet); it requires only that the attributor does not ignore the actual possibility that P is false, where P is the propositional content of the knowledge claim contained in the aforementioned definition of knowledge.

[10]Lewis (1996, p. 555). Suppose it's raining. I have evidence that ought to make me believe that it's *not* raining (my friend told me so). But I am superstitious, and I think that it always rains on the first day of February. Hence, I pick up in my own weird fashion the true belief that it's raining. By the Rule of Belief, I fail to know that it's raining, since I am ignoring the possibility that it's not raining—a possibility that I ought, in that context, to believe to obtain.

[11]This is the fourth 'Gettier' condition that some philosophers have been searching for.

[12] Lewis has one disclaimer to the thesis that context-sensitive knowledge is knowledge *per se*. He says that knowledge of this sort (what he calls 'elusive knowledge') is unstable, since it does not survive a shift of attention. A better form of knowledge is more stable, stands more chance of surviving a shift of attention. But still: elusive knowledge *is* knowledge, he says; only it is not claimable knowledge. See pp. 562ff.

[13]Cohen (1996, p. 298).

[14] It might seem that Lewis is not quite of the same opinion. He thinks that it might make sense to say that a subject's cleverness can limit his knowledge. If a subject is clever enough to think up skeptical possibilities, then we cannot attribute ordinary knowledge to him. (1996, pp. 561ff.) Lewis also accepts that the cleverness or stupidity of *the attributor* determines what the attributee can or cannot know, but surely he would not accept this context-dependence when the attributee is gettiered. It seems, however, that Cohen is accusing him of being forced to accept exactly that.

[15]Lewis's unrestricted Rule of Resemblance is a pure internalist condition in that the epistemic agent can ascertain *a priori* whether the rule is satisfied. But no internalist condition can be at the service of the contextualist as a substitute for the Rule of Resemblance (this was basically the teaching of Gettier). DeRose's Rule of Sensitivity is more promising, as it is an externalist condition.

[16] See Nozick (1981), and DeRose (1995). DeRose argues that S's belief is *insensitive* if S would believe that P if P were false.

[17]Strictly speaking, the condition should say: S believes P via method M, and if it were the case that not-P while S applies M, then it would not be the case that S believes P via M. Otherwise Truth Tracking would admit obvious counterexamples, such as Nozick's Grandmother case (*ibid.*).

[18] DeRose, (*ibid.*)

[19] DeRose (*ibid.*, p. 36).

[20] DeRose (*ibid.*, pp. 39-40)

[21] A similar notion of unclaimable knowledge appears to be entailed by Lewis's theory. See note 12 above. On Lewis's account the standards are raised to such a level as to require that our evidence (for P) eliminates the possibility that not P only when we attend to the not-P possibility. But arguably an attributor pays attention to the not-P possibility that we're brains in a vat simply by asserting that we know we're not brains in a vat.

[22] Stone (2000, pp. 536-537).

[23] Cohen (1987). Arguably, both Lewis and De Rose are relevant alternative theorists, too. However, the difference between Cohen's theory and the two aforementioned is that, on Cohen's theory, whether an alternative is relevant depends, not only on whether it is salient to the attributor in her context of attribution, but also on what the world is like. As a result, some alternatives can be too remote to be relevant under ordinary low standards. Since objective probabilities do play a role in determining what counts as a relevant alternative Cohen's position is more like that of a non-contextualist relevant alternatives theorist, like Dretske, than is Lewis's or DeRose's.

[24] Cohen (1987, fn 31).

[25] The world in which Jones's canary is awarded to Jones, but in which Jones's aunt doesn't die is, arguably, close enough to the actual world, and so is a relevant alternative. But that triggers a new problem. If the alternative is always relevant, I can never know that someone owns a canary. On the other hand, if the alternative is relevant only when the attributor attends to it and the attributor does not attend to the facts about the divorce, I know Jones owns a canary (contrary to our intuitions). To say that this alternative is relevant only if it is actually obtains, or if S is gettiered seems *ad hoc*. The same *ad hoc* strategy would have been available to Lewis who could have supplemented his rules with: 'if P is a Gettier belief, not-P may not properly be ignored'.

[26] For a defense of this claim, see Dretske (1967, pp. 1015-1016). For a more recent argument for transmission-failure, see C. Wright (2002). Tracking is the most obvious reason for transmission failure. For those who deny tracking, however, there are other reasons for transmission failure. One might argue, for example, that the sheep-indicating experience can be an adequate reason for the belief '*that* is not a sheep look-alike' only together with the reason 'if *that* is a sheep, then it is not a sheep look-alike'. (I.e. the belief in question is inferential; compare note

29). But this is not the place to explore what it takes for a reason to be adequate. However we are going to explain adequacy, the inference from the reasons to the belief in question will probably be required to be cogent.

[27] See Klein (1995).

[28]See e.g. S. Luper (1984), (1987); and Sosa (1999). Sosa calls a belief that satisfies the indicator condition 'safe'. According to Sosa, a belief is safe just in case it could not easily have been false given the reasons on which it is based.

[29]Sosa, for example, requires that beliefs are virtuous (i.e. "one's belief derives from a way of forming beliefs that is an intellectual virtue, one that in our normal situation for forming beliefs would tend strongly enough to give us beliefs that are safe"). See E. Sosa, 'Beyond Internal Foundations to External Virtues', manuscript (p. 7). Or one might simply require that R-type events be widely accompanied by P-type events. However, I find it more intuitive to resist the Gettier-style counterexamples by formulating the indicator requirement as suggested below. One difficulty for this view is that of explaining why poorly justified beliefs in the truth of necessary statements fail to be knowledge. My belief that 2+2=4 satisfies the indicator requirement. Whatever my reasons are for holding it, my belief is true, as my belief is true in every world. But my belief need not be sufficiently justified. It is not sufficiently justified if it is based on my (unreliable) teacher's testimony. However, one might argue that beliefs acquired via testimony and other indirect methods (such as dreaming and guessing) are inferential. For example, one might argue that my 'real' reasons for believing that 2 + 2 = 4 are that my teacher told me so and that my teacher is reliable. If my teacher is unreliable, the indicator condition fails (since the reasons on which my belief is based do not obtain).

[30] I owe this formulation of the oblique requirement to an anonymous referee. This approach fends off also those Gettier counterexamples that are not obviously associated with inferential belief. S, who is looking at a box that seems to contain a vase, forms the belief that there is a vase in the box. Unbeknownst to her, however, she is looking at a vase-hologram. The box contains a machine that is rigged to display a vase-hologram only when a vase presses on a lever. She thus meets the indicator requirement. Yet she fails to meet the oblique requirement if we take her belief to derive from 'that thing in the box is a vase' (where 'that thing' refers to the hologram).

[31]Compare this to Cohen's recent suggestion (inspired by Sosa's distinction between animal and reflective knowledge) that basic (perceptual) knowledge is a different kind of knowledge, which "represents a relatively

minimal cognitive achievement—something that relatively unsophisticated beings routinely obtain". See Cohen (2002), pp. 326-7. Holistic (reflective) knowledge, on the other hand, requires inductively formed belief regarding our reliability. The present proposal differs in various respects from Cohen's. For example, Cohen's basic knowledge requires tracking and so is not closed under known entailment, and his holistic knowledge requires that we can know independently that we are not systematically deceived. On the present proposal, on the other hand, low-standard knowledge requires that the subject's evidence is an (oblique) indicator of P, whereas high-standard knowledge requires that the subject's belief tracks the truth. Nonetheless, both theories are attempts to accommodate conflicting demands by supposing that there are two kinds of knowledge.

[32] Genuine doubt on the part of the subject or the attributor would eliminate the possibility of S knowing P. If the subject doubts that she knows P on the basis of evidence or argument, then by the Rule of Belief that possibility may not be ignored. If the attributor doubts that S knows P, then she pays attention to a not-P possibility, which for that reason may not be ignored.

[33]I would like to thank Duncan Pritchard, Matt McGrath, Joe Salerno, Barry Smith and two anonymous referees for helpful comments on an earlier draft of this paper.