

Chapter Three

Is Knowledge Contextual?

Consider this triad: (1) I know that those black-and-white equine animals are zebras; (2) I don't know that those animals aren't cleverly painted mules; (3) I know that those animals are zebras only if I know that they aren't cleverly painted mules. Assuming a univocal reading of all terms throughout, the triad is inconsistent. And yet all three of its members seem plausible. Denying (1) seems objectionably skeptical; denying (2) seems naively optimistic; denying (3) strikes many as ridiculous. Something must give. Epistemic contextualists propose that not all terms have a univocal reading throughout. In particular, they propose that knowledge ascriptions (sentences of the form "S knows that P" or "S does not know that P") have contextually variable truth conditions. In normal contexts (2) is false; in skeptical contexts (1) is false; but in no context is (3) false. Stewart Cohen defends contextualism on the grounds that it provides the best solution to the puzzling triad of claims. Earl Conee rejects contextualism. Conee argues that contextualists have failed to sufficiently motivate the semantic hypothesis that their dissolution of the puzzle presupposes. Conee also argues that we don't need the contextualist hypothesis about knowledge ascriptions in order to explain the puzzle; rather, we can equally well explain it based on an invariantist semantics for knowledge ascriptions.

Contextualism Contested

Earl Conee

It should be helpful to situate epistemic contextualism (EC) in a broader context. Epistemology will creep gradually into the discussion.

Most broadly, semantic contextualism is the thesis that something semantic about a symbol varies with some differences in contexts that involve tokens of the symbol. The kind of semantic contextualism of most direct interest to epistemologists asserts

that different sentence tokens of the same type have different truth conditions. That is, some differences in the context where a sentence is used affect which conditions must obtain in the world for the use of the sentence to state a truth.

For instance, the sentence “Smith’s thimble is large enough” uncontroversially displays this variation. Let’s consider a context in which Smith’s friends are discussing whether the thimble fits her. A friend says, “Smith’s thimble is large enough.” Assuming that Smith’s thimble is an ordinary one that Smith has used comfortably for years, the friend states something true by uttering this sentence token. Suppose that the same sentence is uttered simultaneously by a household mover. He has been talking about items that are sizable enough to require two movers to carry, and he mistakenly thinks that Smith has a gigantic thimble. The claim made by that token of the sentence does not state a truth.

In this example the context of the subject of the sentence, the thimble, does not vary. What does vary is the context in which the sentence is used. This is called variation in the “attributer” context, or “attributer relativity.” Variations that affect truth conditions in the context of the subject of the sentence are called “subject” contextual variations, or “subject relativity.” We shall be focusing on claims of attributer relativity.

Indexical expressions provide other illustrations of uncontroversial attributer relativity. Words like “I,” “you,” and “it” in a sentence contribute to the sentence’s having different truth conditions in different contexts. For instance, the first-person pronoun, “I,” is the subject of a sentence that is used to state a truth only if the predicate of the sentence is true of the one who is asserting something with the sentence.

The clear attributer relativity of truth conditions for sentences with a pronoun might seem at first glance to show that attributer relativity is of negligible philosophical interest. No doubt “You are happy” makes a true claim just in case the individual addressed by the attributer is happy. The philosophical questions about persons seem entirely unaffected by the innocuous contextualist claim about how the pronoun “you” refers to them. Is a person a material object, an immortal soul, or something else? Does a person exist as a fetus in early pregnancy? Are the mental properties of a person identical to physical properties? The contextualism is neutral about all of this.

On second thought, though, the general conclusion is too hasty. The contextualism seems to take a side in some philosophical controversies. It claims that sentences of the form “I am F” are true just when the one using the sentence is F. If this is correct, then apparently there has to exist some such entity as the one who is using the sentence for any of these sentences to be true. Yet this runs contrary to no-self theories. They say that there are no such beings as ones who use sentences. According to one no-self view, experiences exist. When these experiences bear certain relations to one another, such as later conscious episodes bearing a “recollective” relation to earlier experiencings, we say that the experiences have “the same subject.” That is not strictly true. There is no one entity that undergoes the experiences. There are just particular events of experiencing and relations among them.

Thus, our supposedly bland semantic claim – that sentences with first-person pronoun subjects say something true just when the predicate is true of one using the sentence – appears to conflict with this no-self view. So the semantics now seems to take a philosophical side. Perhaps the claim that epistemic terms exhibit attributer relativity will likewise turn out to exclude some existing philosophical options.

But on third thought, the assertion of first-person pronoun attributer relativity is no particular threat to the no-self view. The contextualist semantic claim employing the phrase “the one using the sentence” is among the claims that the no-self theory addresses. Our sample version of the theory declares that all such statements are false. To be plausible, the theory must also give some explanation of how such statements are often reasonable to make, in spite of their falsity. Whatever the explanation is, we should expect it to carry over to the assertion of first-person pronoun attributer relativity. The fact that this is a semantic assertion does not make it less likely that the theory can explain its plausibility.

The final merits of no-self views lie in the details of their defensibility. Generally, we see that semantic assertions are neither automatically irrelevant nor immediately decisive.

This illustration should be borne in mind as we consider EC. Many epistemological issues may be unaffected by the existence of attributer relativity in the truth conditions of sentences that include epistemic terms. By itself this is neither a strength nor a weakness in either EC or the unaffected theories. An EC thesis may have some substantial epistemic implication. The implication in conjunction with some epistemological theories may appear to induce surprising truth values. This result need not lead sensible proponents of the affected theories to give them up. They may be able to interpret the sentences to explain away any difficulties.

Another non-epistemic example of purported attributer relativity illustrates how contextualism characteristically copes with certain sorts of conflict. Suppose that Robinson is Chair of the Boring Committee (“BC”). At the time for a meeting of the BC, Robinson sees that all of the members of the BC are in the room. He declares, “Everyone is present.” The members present would be led by this to think that the membership of the BC is present. Yet “everyone,” the subject term of Robinson’s sentence, seems to be about every person in the world. Of course none of the members of the BC is led to believe that every person in the world is present.

One reasonable semantic hypothesis is that quantifiers like “everyone” have a contextually narrowed range. A use of a sentence with the subject “everyone” states a truth just when the predicate of the sentence is true of all of the people who are made relevant by the context of use. In Robinson’s context, his “everyone” would be restricted to all who are members of the BC.

Though this attributer relativity is a reasonable hypothesis, such examples can also be accommodated while allowing that “everyone” always ranges over, well, everyone. One alternative view is that sentences with “everyone” typically are overly general and therefore false. But people who hear such claims immediately ignore those who are irrelevant to the purposes at hand. This practice of ignoring is widely taken for granted. That makes such false claims efficient ways to communicate something that is about only those who are in the group of contextually determined interest.

An asset of this contextually unvarying (“invariantist”) semantic view of such sentences emerges in the following continuation of the example. Suppose that someone on the BC, say Stickler, opposes Robinson’s claim. Stickler says, “No, what you said is not true, because not everyone is here. The Pope is not here.” Probably Stickler’s comment would come off as unamusing, at best. But it is not easy to see that Stickler is saying something false.

Quantifier contextualists can find truth in Stickler’s comment. They can say that Stickler has shifted the conversational context, perhaps by mentioning the Pope,

someone who is not in the group previously assumed to be relevant. So in Stickler's mouth "everyone" ranges over a wider group, perhaps all people worldwide. Thus, in that context it is correct for Stickler to assert "not everyone is here." The contextualist can note that this is compatible with the truth of the claim made by Robinson with "everyone," on the grounds that Robinson's claim had a crucially narrower contextually determined range.

This contextualist response does not wholly avoid the difficulty, though. The response finds a truth in the claim made by Robinson and it finds a truth in a claim made by Stickler. The remaining difficulty is this. It seems that Stickler was entirely accurate (though unamusing, unhelpful, and maybe obtuse). Yet the contextualist view does not find only truth in what Sticker says. Stickler said that Robinson's original claim was untrue. The contextualist view that we are considering must disagree with this. The contextualist must say that Stickler himself is wrong on that point, because the truth conditions of Robinson's original claim were met. The contextualist view implies that the original claim was true if and only if the members of the group in question were all there, and they all were. So our impression that Robinson was entirely accurate must be erroneous.

It is not obvious how best to interpret quantifiers in cases like this. The contextualist position may be overall best. In any event, contextualism applied to quantifiers displays a characteristic feature. Contextualism finds truth at the expense of contradiction. That has a nice constructive ring to it. But it runs a risk of interpretive failure. Concerning some cases in which one sentence seems to be used to deny something that another sentence says, contextualists characteristically allege that they are not really flat denials. This happens whenever there is a difference in what the contextualism counts as a semantically relevant context. Any such contextual difference implies that the two claims do not just contradict one another, even if one sentence used seems to negate the other.¹ Most commonly, contextualist proposals yield compatible truth conditions for the sentences. For instance, they say that Robinson and Stickler can both be right about those whom they respectively call "everyone."

Moral relativists are the contextualists of ethical theory. They hold roughly that something about the context of a moral evaluation, perhaps something in the culture or the personal values of the evaluator, contributes to the truth conditions of the resulting evaluation. Moral relativists have always been a relatively small and besieged group, though not lacking in able philosophers. The reason seems to be this. The prevalence of fundamental disagreements about moral evaluations is a virtually ineluctable element of the ethical situation. Among non-philosophers there are vigorous disputes about capital punishment, the treatment of animals, gun control, and so forth. Or, at least, their exchanges give every appearance of being disputes, and the participants clearly regard them as such. Among philosophers there are vigorous disputes among consequentialists, Kantians, virtue theorists, and divine command theorists. Or, at least, they certainly seem to disagree. In both cases the disputants often speak from more or less different cultural and personal backgrounds. When these differences are ones that a moral relativist claims to be crucial, the relativist cannot find contradictory claims, and may find only compatible claims. Yet it is deeply dubious that any such conciliatory interpretation is correct. Moral relativism has some resources to try to explain away the appearance of disagreement. But often it seems clear both that the disputants are fully competent with the language they are using and that those

on each side fully intend to use it to deny what the other side affirms. These facts are hard for a relativist to explain away. These cases are trouble for moral relativism.

To the extent that apparent disagreement over epistemic claims is similarly witting and heartfelt, it is similarly difficult to find credible a reconciling EC interpretation. Some arguments about skepticism appear to have that character. Like moral relativists, EC theorists can work to explain away the appearance of contradiction or disagreement. But also as in the moral case, often it seems clear that competent speakers use words with the intent to be discussing the same thing and to be respectively affirming and denying something about it. These cases are trouble for EC.

Finally by way of preliminaries, we should be alert to a misleading temptation. It is tempting to think that semantic contextualism is obviously correct for all evaluative expressions. It is a plain fact that people routinely apply differing standards of evaluation in different contexts. For instance, suppose that Jones works on a Sabbath. Smith counts Jones as “evil” because Jones is thereby disobeying a strict edict of Smith’s religious faith. Robinson denies that Jones is “evil” because Jones’s work is intended to promote a social cause that Robinson admires. Numerous similar examples exist for evaluations of beauty, humor, morality, and many other things. In particular, the application of differing standards of judgment plainly occurs as well in epistemic evaluations, such as attributions of knowledge, justification, and good reasoning. Are these contextually varying truth conditions in action?

Not necessarily. The fact that different standards are routinely applied in making an evaluative judgment does not imply the correctness of semantic contextualism about the contents of the judgment. To see this, note that differing standards are routinely applied in making judgments that are uncontroversially context invariant. For instance, usually the genuineness of a purported piece of US currency is judged by a quick look, sometimes it is judged by a more careful inspection, and occasionally it is rigorously judged using a high-tech device. In spite of these differing standards, the content of the uses of “genuine US currency” is clearly the same. Other examples abound.

Thus, we cannot validly infer that contextual differences yield differing truth conditions simply because differing standards are often applied. For semantic contextualism to be correct, there must be something else about an expression, something beyond its susceptibility to differing standards of judgment, that makes the truth conditions of tokens of the expression differ with context.

Epistemic Contextualism

The most widely discussed form of EC holds that the truth conditions for tokens of sentences that include “knows” (and cognate expressions) vary with the attributer’s context. Broadly speaking, what varies with some differences in the context in which a “knows” sentence is used is the strength of the epistemic position that the subject of the sentence must be in, in order for the sentence to assert a truth.

Many versions of EC share this core. Typically, the strength of epistemic position required is said to vary in a range that allows, at its low end, many true attributions in everyday contexts concerning ordinary judgments based on perception, memory, testimony, and perhaps also inductive generalization and high probability. At the high

end of the range of variation, the typical EC truth conditions are demanding enough to make true many skeptical denials of “knowledge” of the external world.

This range corresponds to a genuine phenomenon of “knowledge” attributions. Fluent English speakers regularly make confident “knowledge” attributions about numerous routine matters of daily life: the color of a shirt when it is seen in good light, the occurrence of a death when it is read about in a newspaper, the identity of the first US President when it is recalled from some source or other, the sum of two large numbers when they are totaled on a calculator, and so on. Fluent speakers typically become increasingly hesitant about “knowledge” attributions as the practical significance of the right answer increases. Fluent speakers typically doubt or deny “knowledge” attributions when coming to grips with impressive skeptical arguments.

The heart of the contextualist explanation of this phenomenon is that people are responding in an approximately accurate way to shifting standards for correct “knowledge” attributions. In the context of everyday life, “knowledge” attributions have truth conditions that are weak enough to be met by seeing a color in good light, reading a report in a reliable publication, recalling a famous historical fact, adding on a calculator, and so forth. As the truth value of a proposition matters more, the context shifts so that the truth conditions for “knowledge” of the proposition require somehow excluding ever more potential sources of error. In a context where skeptical considerations are prominent, the skepticism somehow gives rise to truth conditions for “knowledge” attributions that are so demanding that external world beliefs do not meet them.

There are two noteworthy ways in which the contextualist explanation of this phenomenon involving “knowledge” attribution could be misleading.

Loose Talk

The first way that EC could be misleading is by being fundamentally mistaken. It may be that all “knowledge” attributions have the same truth conditions, but people apply contextually varying standards for making the attributions. The most plausible unvarying standard for truth is very high, but not unreachably high. In ordinary contexts, when nothing much turns on it, people will claim knowledge, and attribute knowledge to themselves and others, in belief and in speech, on a basis that is significantly weaker than what is actually required. This is an efficient way to communicate assurance about the proposition and to facilitate taking it for granted.

According to the view now under consideration, fluent speakers actually realize, at least tacitly, that this knowledge ascription is just loose talk. The realization is shown by the fact that, if asked whether some proposition to which knowledge is ascribed on some such basis is *really* known, or *truly* known, or *really and truly* known, fluent speakers have a strong inclination to doubt or deny that it is.² Only the most conspicuous facts of current perception, the clearest memories, triple-checked calculations, and the like will often pass some such “really and truly” test. The present view has it that the answer to this “really and truly” question reveals what a speaker judges to be knowledge when she is trying her best to apply her best thought as to what is the actual standard for the truth of a “knowledge” attribution.

A contextualist has a ready explanation for the results of this “really and truly” test. The explanation is that raising a question about “knowledge” using these words serves

to impose more stringent truth conditions. So it is no wonder that propositions that pass previous, more lenient standards are no longer said to be “known.”

There is an important liability of this contextualist response. There are many big topics that receive similar results on the “really and truly” test. For instance, when fluent speakers consider whether they are “really and truly” happy, or they consider who are “really and truly” their friends, or they consider what is “really and truly” worth striving for, they tend to have doubts about their previous casual attributions of happiness, friendship, and worthwhile goals. It is possible that the English terms for these categories of happiness, friendship, and so forth all have contextually varying truth conditions, and that the standards are raised by these “really and truly” questions. But that is not what seems to be happening as we ponder these questions. It seems that our aim is to set aside loose talk about an important topic and tell the truth about it, the very same topic that was previously more casually discussed. The contextualist who holds that the “really and truly” test changes the truth conditions is bound to count this view of what we are doing as a mistake.³ That is a liability.

Strict Truth

The second way in which contextualism about “knows” may be misleading would be by being correct. The reason why this might mislead is that it might suggest that the correctness of EC has a significant implication concerning some philosophical issue about knowledge. The fact is that the truth of EC in general would have no such implication.

For instance, suppose that “really and truly” does constitute a context in which “know” has truth conditions with different and higher standards than those of ordinary uses of “know.” It may be that all philosophical discussion of knowledge is carried on in that one context. That is, it may be that every issue about knowledge, from the Gettier problem to the extent of scientific knowledge, from the nature of justification to the merits of external world skepticism, has been discussed solely in this single “really and truly” context. Granting epistemic contextualism, this claim of a single philosophical context has some plausibility. Philosophy generally gets going only when we get serious about a topic. So it may be that “really and truly” is one colloquial way to direct attention to the truth conditions that arise when we take this philosophical sort of attitude toward our investigation.

If all of this is so, then “knows” is context sensitive. But appreciating that fact would tell us nothing new about our philosophical issues. All sides to the philosophical disputes would have been discussing the same thing all along. At most, we would have reason to avoid importing into philosophical discussions of knowledge examples in which “knows” is less scrupulously applied. Given the correctness of EC, these may well not be cases of the philosophical topic of knowledge. Not relying on such cases is sensible. But we have reason not to rely on them in any event. Whether or not EC is correct, relatively casual uses of a term are more likely to be mistakes.

More detailed contextual views have more philosophically consequential implications. For instance, some EC views say enough about what differentiates contexts, and what the differing truth conditions are, to imply that skeptics and non-skeptics typically operate in contexts that make each position correct about what it typically

calls “knowledge.” The present point is only that contextualism in general says nothing about when there are context shifts. So even if they exist, they may not affect philosophical questions about knowledge.

Again, suppose that EC is correct in broad outline, but now suppose that philosophy does not set a single context. In skeptical contexts, the requirements for “knowledge” are severe; in various non-skeptical contexts, including some in which philosophical issues are under discussion, the requirements are variously less severe. There are still no immediate epistemological consequences, even for the merits of external world skepticism. The severe requirements of skeptical contexts might be satisfiable by felicitous external world judgments, the standards actually imposed by skeptics to the contrary notwithstanding. The outline of one such possibility is this. It might be that in skeptical contexts, only 100 percent of the right sort of probability is enough. But it is enough (along with true belief). There need not also be something that skeptics often require: internally possessed evidence that excludes all possibility of false belief. Skeptics see various ways in which our judgments often both lack this probability and are possibly false. Skeptics draw the extreme conclusion that the external world judgments must have the support of entailing evidence. But perhaps external world judgments are in the right way 100 percent probable when they are correctly causally connected to the fact known.⁴ This might occur only in optimal cases of perception and memory. The skeptical context would set the 100 percent probability requirement, while the skeptic would mistakenly impose a more stringent standard. EC does not imply that skeptics have infallible theoretical insight into the truth conditions for “knows” that hold when they are being skeptical.

Again, EC leaves the main epistemological issues unresolved. Only versions that are specific about the contextual truth conditions do more. Those versions must be defended on their own particular merits. The general thesis of EC is neutral about the epistemically consequential details.

A Delicate Balance

EC has never been obviously true. In this way it differs from the contextual variation in truth conditions exhibited by indexicals and comparatives like “large” and “near.” EC is defended not by direct semantic reflection among fluent speakers, but by its explanatory strengths. The classic explanatory role is to explain how both ordinary “knowledge” attributions and skeptical arguments denying “knowledge” can seem so plausible. The classic contextualist view is that each is correct. The reason is that ordinary “knowledge” ascriptions often meet the truth conditions of “knows” in everyday contexts, while skeptical denials of “knowledge” are right according to the truth conditions for “knows” operative in their contexts.

This reason does not quite yield the data to be explained. The data are the *plausibility*, or the *reasonableness*, of both ordinary “knowledge” attributions and skeptical “knowledge” denials. The classic contextualist proposal offers us a way to have those claims come out *true*. In order to use this truth to account for the plausibility in question, the explanation must add something about how the truth engenders the plausibility.

The most straightforward way to do this is to add something along the following lines. When we are thinking about ordinary speakers, and skeptics, are we in some way

cognizant of the truth conditions of their respective “knowledge” ascriptions? We are in some way cognizant that the respective truth conditions are met. And so we find the respective claims plausible.

This explanation attributes to us some sort of grip on the contextual variation in standards for “knowledge” ascriptions. We are not said to think explicitly that there is a contextual variation. But we are supposed to be guided by our understanding of “knows” to count some differences in context as decisive to truth when we make our plausibility judgments.

The problem is to understand how the correctness of EC can be so difficult to recognize that it has been missed entirely until lately, and many philosophers who consider it still sincerely deny it. How can it be that fluent English speakers are well enough attuned to this variation for it to guide some of their plausibility judgments, while none noticed that “knows” is context sensitive until recently, and some are unable to recognize this by semantic reflection, even under the tutelage of a contextualist?

There is no contradiction in this. It is not a hopeless situation for EC. There may be some adequate psychological explanation of how one can have a concept that one commonly applies in a context-sensitive way, while being unable to spot this feature of the concept when one tries. It is difficult to understand this sort of grasp, though. For a start, it would be nice to be shown other examples of concept possession with *almost* these same features, say, examples where the context sensitivity is sufficiently hidden to have been universally overlooked but is eventually recognizable by virtually all proficient speakers. The oddity of this sort of grasp of a concept poses a challenge to EC. It is another liability of the view.⁵

Notes

- 1 Contextualists find differences in truth conditions. This does not automatically make the claims compatible. However, if there is any contextual variation in truth conditions, then the two claims are not the direct contradictories of one another that they seem to be. Also, denials of contradiction are what a straightforward application of contextualism engenders. It is possible to add a special provision for disputes. A special provision might claim that in a dispute something overrides each speaker’s separate context in determining truth conditions. It might be claimed that this yields one context for both of the pair of syntactically contradictory utterances. This sort of view will have a hard time locating one interpretation that is credible for both utterances, while otherwise interpreting in the spirit of the contextualism. Why does the basis for the unifying interpretation override the particular contexts of use only in disputes, and not everywhere? Furthermore, justifying this sort of special provision is especially difficult when uses of syntactically contradictory sentences are not public, but rather formulate successive thoughts by someone weighing the merits of opposing positions. There is no intuitive third context to appeal to in assigning unifying truth conditions, while the successive contexts seem tied for suitability as context determiners.
- 2 “Really” by itself, and “truly” by itself, sometimes seem to function as intensifiers, as in “a really hot day,” and “a truly long drive.” But “really and truly” seems rather to function as a way to get serious about the truth of the modified claim. Similarly functioning terms include “actually” and “genuinely.” “Really and truly” will be the phrase of choice here.
- 3 The contextualist who denies that the “really and truly” test changes the truth conditions seems to have no reasonable principled way to say why the truth conditions do change in skeptical contexts.

- 4 This is very roughly the position of Fred Dretske about knowledge in general, in his *Knowledge and the Flow of Information* (1981).
- 5 The negative perspective on contextualism that I express in this exchange has been influenced by the work of several philosophers. The following essays have been most influential: Schiffer (1996), Sosa (2000), and Feldman (2001).

References

- Dretske, F. (1981) *Knowledge and the Flow of Information*. Oxford: Blackwell.
- Feldman, R. (2001) Skeptical problems, contextualist solutions. *Philosophical Studies* 103: 61–85.
- Schiffer, S. (1996) Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society* 96: 317–333.
- Sosa, E. (2000) Scepticism and contextualism. *Philosophical Issues* 10: 1–18.

Contextualism Defended

Stewart Cohen

1 Contextualism and Skeptical Paradoxes

Contextualism has been proposed as a way to resolve stubborn epistemological paradoxes. Where P is some commonsense proposition about the external world (e.g., I see a zebra) and H is some skeptical hypothesis (e.g., I see a cleverly disguised mule), the paradox takes the following form:

- 1 I know P.
- 2 I do not know not-H.
- 3 I know P only if I know not-H.

These propositions constitute a paradox because each is independently very plausible, yet jointly they are inconsistent. Because our intuitions about knowledge lead to paradox, skepticism threatens.

Most philosophers who consider this paradox are unwilling to accept skepticism. Instead, they attempt to provide an anti-skeptical resolution of the paradox. But one cannot satisfactorily respond to the paradox by simply denying one member of the inconsistent set. The paradox arises because each proposition seems true. To resolve the paradox, one must explain the intuitive appeal of the denied proposition. Otherwise, we have no explanation for how the paradox arises. So an anti-skeptical resolution must resolve the paradox in a way that preserves the truth of (1) while explaining the appeal of the propositions that threaten a skeptical result, i.e. (2) and (3).

One robust feature of our intuitions regarding (1) to (3) is that they tend to vacillate. We begin by confidently holding (1), that we know many things about the world. Under pressure from skeptical arguments, we see the appeal of (2). Given that (3) is

compelling, and that (3) and (2) entail the falsity of (1), we begin to waver in our assent to (1). But we find this unsatisfactory given our commonsense rejection of skepticism. But again (1) together with (3) entails that (2) is false. But when we consider the skeptical arguments for (2), we tend to fall back into a skeptical frame of mind.

The contextualist resolution of the paradox proposes that we take these intuitions at face value. Rather than rejecting one of these intuitions as mistaken, contextualism attempts to explain away the apparent inconsistency of our intuitions by arguing that they reflect the contextually varying truth conditions for knowledge ascriptions. Contextualism holds that ascriptions of knowledge are context sensitive – that the truth value of sentences of the form “S knows P” depend on contextually varying standards for how strong one’s epistemic position with respect to P must be in order for one to know P.

There are various ways of cashing out this notion of the strength of one’s epistemic position. But contextualist resolutions of the paradox are alike in holding that the context of ascription determines which proposition of the skeptical paradox gets denied. According to contextualism, in everyday contexts (2) is false, in stricter “skeptical contexts (1) is false, but in no context is (3) false. This resolution of the paradox explains the appeal of skeptical arguments by allowing that the claims of the skeptic are true, relative to the very strict context in which they are made. But the resolution is anti-skeptical in that it preserves the truth of our everyday knowledge ascriptions relative to the everyday contexts in which they are made. In essence, contextualism concedes that there is some truth to skepticism, but contains the damage by holding that the skeptical claims are true only relative to atypically strict contexts.

2 Alternative Accounts

We can think of contextualism as involving two theses:

- (a) Ascriptions of knowledge are context sensitive.
- (b) The context sensitivity of knowledge ascriptions provides the basis for resolving the skeptical paradox.

Criticisms of contextualism can challenge (a), or grant (a) while denying (b). Earl Conee raises criticisms of both kinds.

Conee challenges (a) by arguing that there are alternative ways of accounting for our intuitions that saddle us with the paradox. He suggests the possibility that the truth conditions for knowledge ascriptions are invariant – most plausibly “very high, but not unreachably high” – but that the standards for making appropriate, though not strictly speaking true, attributions vary with context. In everyday contexts, “when nothing much turns on it” people will ascribe knowledge to themselves and others on a basis that is “significantly weaker than what is actually required.” But competent speakers realize that this is loose talk, as evidenced in their willingness to withdraw their claims when challenged as to whether the proposition is “really and truly” known. Conee notes that “only the most conspicuous facts of current perception, the clearest memories, triple-checked calculations and the like” will pass the “really and truly” test.

Conee's point here is that there is nothing in the intuitive data to support contextualism over this alternative account. But although the account Conee suggests is consistent with the intuitive data, I don't see that this account does explain the intuitive data as well as contextualism. The problem for Conee's alternative is that competent speakers, under skeptical pressure, tend to deny that we know even the most conspicuous facts of perception, the clearest memories, etc. For the contextualist, this is because the strictness of the standards in skepticism make our skeptical denials true, in those contexts. But Conee's alternative incorrectly predicts that we should hang on to such knowledge ascriptions, even in the face of skeptical "really and truly" challenges.

As Conee notes, the contextualist should say that when one applies the "really and truly" test, one is in fact raising the standards, and in this new context, the truth conditions for sentences of the form "S knows P" are so strict as to falsify the ascriptions. This would seem to commit the contextualist to taking a similar line for a whole host of predicates that behave similarly when subjected to the "really and truly" test – for example, happy. But according to Conee, that is not what seems to be happening when we consider questions like "Am I really and truly happy?" It seems to us that we are setting aside loose talk and trying to get to the truth of the matter. According to Conee, this is a liability for contextualism.

There is no doubt that when one is in what the contextualist wants to view as a high standards context, one has a feeling of enlightenment regarding the correct application of the predicate in question. We feel as if we are seeing the truth of the matter that has, up until that point, eluded us. So we have the "witting and heartfelt" sense that, contrary to what contextualism holds, our skeptical judgments conflict with our previous everyday judgments. So the contextualist has to argue that we are in all of these cases mistaken in our meta-judgment that our skeptical judgments conflict with our earlier judgments. Conee holds that this is a liability for contextualism. But the extent to which this is a liability will depend on the plausibility of attributing to us this kind of mistake. I will return to this issue below.

Conee goes on to raise objections to (b). He suggests two ways in which contextualism might be misleading, even if it turns out that ascriptions of knowledge are context sensitive. First, it might be the case that the skeptic (or the skeptic in us) is mistaken in thinking that the stricter skeptical standards are unsatisfiable. Second, it may be that all philosophical discussions are carried out in one strict, "really and truly" context. If either of these possibilities obtains then the fact that ascriptions of knowledge are context sensitive will not, by itself, settle the main *philosophical* questions about knowledge.

Conee is right that even granting the context sensitivity of knowledge ascriptions, nothing strictly follows about what the standards are in particular contexts. That is, (b) is logically stronger than (a). For the contextualist solution to work, it must be that the standards in everyday contexts are low and attainable, and the standards in skeptical contexts are high and unattainable. The contextualist proposes that if we take our shifty intuitive judgments, both skeptical and commonsensical, at face value, then we get just those results.

Of course, the contextualist has no *proof* that our intuitive judgments are correct. What contextualism offers is a non-skeptical way to resolve the skeptical paradox. It proposes an explanation of our inconsistent inclinations that accounts for the appeal of skeptical arguments, while still preserving the truth of our everyday knowledge ascriptions. Now for all that, contextualism may be false. But the fact that it can do

this much shows that there is much to recommend it. In the end, whether contextualism should be adopted depends on how it compares with other proposals for resolving the paradox.

This brings us back to Conee's proposals. The first is that the standards are context sensitive, but we are mistaken in thinking that the strict skeptical standards are not met. Perhaps this is so. But absent some defense of a particular proposal as to how we are mistaken, as to how we can satisfy what appear, in some contexts, to be the very strict standards for knowledge, the proposal does not count for much. (Conee mentions Dretske's information-theoretic account, but I don't think he is seriously endorsing it.) At best, it shows how the falsity of the contextualist solution to the epistemological paradox is consistent with the thesis that ascriptions of knowledge are context sensitive. And this is something the contextualist can concede.

Moreover, even if we had a plausible account of how we could meet the stricter standards, such an account would have to explain why we were inclined to make the mistake of thinking that we cannot meet those standards, when in fact we can. Recall that these are intuitive judgments. It is not clear that our intuitive skeptical judgments are made on the basis of some false theory we hold about what is required for knowledge. If the correct theory entails that our skeptical intuitions are mistaken, the question still remains as to why we have those intuitions. Only with such an explanation would we have a resolution of the paradox.

Conee's second proposal is that although the standards for knowledge vary across contexts, there is a single standard that is always in effect in philosophical discussions of knowledge, "from the Gettier problem, to the extent of scientific knowledge, from the nature of justification to the merits of skepticism all have been discussed solely in the 'really and truly' context." This is another possibility wherein "knowledge" is context sensitive, but that fact has no implications for important philosophical issues.

Again, Conee is correct that nothing philosophically significant strictly follows from the thesis that ascriptions of knowledge are context sensitive. The point of contextualism is to provide an explanation for the intuitive data in a way that can explain the appeal of skepticism while still preserving the truth of our everyday knowledge ascriptions. The fact that contextualism can provide such an explanation does not entail that it is the only possible account, even on the supposition that ascriptions of knowledge are context sensitive. So we need to compare the relative merits of competing accounts.

An important fact about our epistemological intuitions is that even when we are discussing philosophy, we vacillate between thinking that we know and thinking that we do not. When we discuss the Gettier problem, we readily allow that we know all sorts of things. We use these intuitions about when we know as data in constructing our theory of knowledge. But when we discuss skepticism, we have intuitions that we don't know the very things we thought we knew when discussing the Gettier problem.

How do we explain this shiftiness in a way that avoids skepticism while still explaining its appeal? Contextualism proposes a way to do this in terms of contextually shifting truth conditions. The view that Conee mentions provides no such explanation. Since our intuitions vacillate even in philosophical contexts, if, as Conee suggests, there is a single standard for all such contexts, then some of those intuitions are mistaken. But which ones are mistaken, and why, if they are mistaken, do we have them? Conee's suggestion is merely the bare statement of a possibility that leaves

entirely open whether skepticism is true, and if not, why we find it appealing. Thus the contextualist explanation of our intuitions, while not entailed by the thesis that ascriptions of knowledge are context sensitive, has much more to recommend it than the possibility that Conee raises.

3 Is the Contextualist Model Coherent?

Of course, a theory's explanatory power does not count for much if the theory is independently implausible. And Conee argues that contextualism taken as a whole entails an implausible story about competent speakers. He notes that contextualism explains the appeal of the premises of the skeptical paradox by arguing that each is in fact true, relative to a particular context. Thus contextualism must hold that the appeal of the premises of the paradox reflects our grasp of their truth conditions. So it must be that judgments of competent speakers, to the extent that their judgments are so guided, are sensitive to the contextual shifts in the truth conditions.

But at the same time, contextualism involves a kind of error theory with respect to certain meta-judgments about knowledge. According to the view, when we are in skeptical contexts where (1) is false, we mistakenly think our denial of (1) conflicts with our everyday assent to sentences of that form. That is, we fail to recognize the context sensitivity of our own judgments even though we grasp, in some sense, their shifting truth conditions. Conee claims that this account of competent speakers is "difficult to understand" and challenges the contextualist to provide other examples of our use of context-sensitive predicates that fits the model the contextualist proposes in the case of "knows."

As I have argued, there are precedents for the contextualist model (see Cohen, 1999, 2001). Consider ascriptions of flatness. You can lead competent speakers to question their everyday ascriptions of flatness by making salient "bumps" that ordinarily we do not pay attention to. As Peter Unger demonstrated, taking this strategy to the extreme, for example, by calling attention to microscopic surface irregularities, one can lead competent speakers to worry whether anything is really flat (Unger, 1975). But Unger's case for flatness skepticism is interesting precisely because many who feel the pull of flatness skepticism look back on their previous flatness ascriptions and think they may have been wrong.

Should we worry that all along we have been speaking falsely when we have called things "flat"? Surely not. Philosophical reflection will convince most that ascriptions of flatness are relative to context-sensitive standards. Roads that count as flat in a conversation among Coloradans do not generally count as flat in a conversation among Kansans. And while one can truly ascribe flatness to a table in everyday conversations, one might not be able to truly ascribe flatness to that same table when setting up a sensitive scientific experiment. If we implicitly raise the standards high enough (by making salient microscopic bumps), then perhaps, relative to that context, no physical surface really is flat. But of course, that does not impugn our ascriptions of flatness in everyday contexts where the standards are more lenient.

So the controversy over whether anything is flat can be resolved by noting that ascriptions of flatness are context sensitive. And to the extent that competent speakers' ascriptions of flatness are guided by these contextually shifting truth conditions, they

in some sense grasp those truth conditions. But then why can we get competent speakers to question their everyday flatness ascriptions by implicitly raising the standards? It must be that although competent speakers are guided in their flatness ascriptions by contextually shifting truth conditions, such speakers can fail to realize that their flatness ascriptions are context sensitive. And because they can fail to realize this, they can mistakenly think that their reluctance to ascribe flatness, in a context where the standards are at the extreme, conflicts with their ascriptions of flatness in everyday contexts. That is to say, they conflate the proposition expressed by “X is flat” at a strict context, with the proposition expressed by that sentence at a more lenient context. And this is precisely what the contextualist argues occurs in the case of knowledge ascriptions.¹

Having said that, I should note that there is an important difference between the two cases. Contextualist theories of flatness ascriptions, once proposed, gain easy and widespread acceptance among most people. But, as Conee points out, contextualist theories of justification/knowledge do not. This is something a contextualist – one like me anyway who relies on the analogy – needs to explain.

One thing the contextualist can say here is that there are varying degrees to which competent speakers are blind to context sensitivity in the language. Thus, while everyone can readily see the context sensitivity of indexicals like “I” and “now,” it can take some amount of reflection to convince most that “flat” is context sensitive. And for a term like “knows,” it may be very difficult even after some amount of reflection for competent speakers to accept context sensitivity. It may take subtle philosophical considerations concerning the best way to resolve a paradox in order to “see” the context sensitivity of “knows.”

Here is a somewhat speculative attempt to explain why it is difficult to accept that “knows” (and “justified”) is context sensitive. Justification and knowledge are normative concepts. To say that a belief is justified or constitutes knowledge is to say something good about the belief. We value justification and knowledge. But contextualist theories are deflationary. Contextualism about knowledge says that most of our everyday utterances of the form “S knows P” are true, even though the strength of S’s epistemic position in those instances does not meet our highest standards. In the same way, contextualism about flatness says that most of our everyday utterances of sentences of the form “X is flat” are true, even though X’s surface may fall short of perfect flatness.

In other words, contextualism is a “good news, bad news” theory. The good news is that we have lots of knowledge and many surfaces are flat; the bad news is that knowledge and flatness are not all they were cracked up to be. We find this much easier to accept in the case of flatness than knowledge, because ascriptions of flatness do not have the normative force that ascriptions of knowledge/justification do.

Note

- 1 A perhaps less clear case involves ascriptions of solidity. With the rise of our understanding of atomic physics, it became clear that many of the objects we think of as solid actually contain quite a bit of empty space. This led some people to conclude that no physical objects are solid. Others found this conclusion to contradict common sense. No doubt some

vacillated between agreeing with the conclusion and rejecting it. Again, I assume that philosophical reflection will convince most that the correct thing to say here is that ascriptions of solidity are context sensitive. Here it seems most plausible to view ascriptions of solidity as involving a kind of implicit quantification and the context as governing a shifting domain. So for x to be solid, it must contain no spaces between its parts. But what counts as a part depends on the context. If this is correct, then we have another case where competent speakers' judgments are guided by contextually shifting truth conditions, but the speakers are unaware of this fact.

References

- Cohen, S. (1999) Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives* 13: 57–89.
- Cohen, S. (2001) Contextualism defended: comments on Richard Feldman's "Skeptical problems, contextualist solutions." *Philosophical Studies* 103: 87–89.
- Unger, P. (1975) *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.

Contextualism Contested Some More

Earl Conee

Stewart Cohen offers EC as a way to solve an epistemological problem that he counts as a paradox. Here is the standard illustration of the problem. While I see an animal in a zoo enclosure:

- (1) I know that I see a zebra.
- (2) I do not know that I do not see a cleverly disguised mule.
- (3) I know that I see a zebra only if I know that I do not see a cleverly disguised mule.

Stew reports that our intuitions about the truth of (1) to (3) robustly tend to vacillate. We are drawn to (1) by common sense. Skeptical arguments lead us to (2). (3) is compelling all along, and inferences from (3) and each of (1) and (2) lead us to the denial of the other.

Stew's EC theorist has it that each of (1) and (2) is true at a context from which it is found intuitive. The intuitive appeal of (1) is to be explained by the proposal that there are relatively lenient truth conditions for (1) in everyday contexts. The plausibility of a skeptical conclusion is to be explained by the proposal that skeptical claims that imply (2) are true at contexts with relatively strict truth conditions for sentences using "know."

As I write, a painting is in the news. Thoughts about the reports provide a model of an intellectual conflict similar to (1) to (3) that does not allow a contextualist resolution.

A woman bought a painting for \$5. It is a large canvas covered with bright squiggles and blotches of paint. Eventually the woman happened to contact an art professor about it. She was led to believe that the artist might well be the renowned twentieth-century American painter, Jackson Pollock. An expert found a Pollock fingerprint in the paint. He declared it a genuine Pollock. An art institute examined the painting. After careful study, the institute announced that the painting is not by Pollock. There the matter stands.¹

Considering the case, I vacillate between two conclusions.

C1: The painting is by Jackson Pollock.

C2: The painting is not by Jackson Pollock.

On different bases, each seems true. I am tempted to accept C1 when I think:

How could Pollock's fingerprint get in a painting that looks like a Pollock to an expert, if Pollock didn't paint it? There is no other plausible way.

I am tempted to accept C2 when I think:

The art institute has its reputation at stake. They are aware of the fingerprint evidence. Yet they say that the work is not by Pollock.

I incline toward C1 again when I note that the art institute is relying heavily on the lack of an established history for the work, while I doubt that it is peculiar to lose track of one painting. I incline toward C2 again when I note that I may know little of the reputable art institute's full grounds for their negative conclusion. I alternate between C1 and C2 on the basis of these reasons. I become perplexed and reluctant to judge.

This is no paradox. There is a straightforward empirical fact of the matter (leaving vagueness aside). I happen to have pretty strong evidence for each conclusion with no overriding consideration. I alternate between C1 and C2 by employing different standards of judgment, each of which is normally sufficient for reasonable belief. Some of my standards make a mistake here. This is not mysterious. We have no good grounds to think that such standards are universally correct.

Finally about this case, a reconciling contextualist interpretation of my perplexity is out of the question.² The thoughts that I formulate with C1 and C2 are incompatible. Now let's return to Stew's (1) to (3), and its EC resolution.

Intuitive responses do incline us toward accepting (1) and (2), though we can become perplexed and reluctant to judge. It may be misleading to classify the responses as intuitions. Neither (1) nor (2) is just credible on its own. Each seems true, when it does, on the basis of something else. The bases differ. (1) seems true when we naively consider the facts of an ordinary zoo visit. (2) seems true partly as a consequence of plausible skeptical reasoning. For instance, a typical key premise alleges that a visitor to the zoo has no evidence that distinguishes the animal seen from a cleverly disguised mule.

The appeal of (1) and (2) derives in one way or another from other considerations. I shall call any such basis a "standard for judgment." It is of some importance that the plausibility of (1) and (2) derives from standards for judgment, that this joint plausibility is perplexing, and that we have no reason beyond the limited credibility of our standards to regard them as unerring. This all makes (1) to (3) significantly like my C1

and C2. There the truth conditions do not vary. I am drawn to C1 and C2 by credible but fallible standards, while I have no overriding evidence about which standards support a falsehood. The counterpart invariantist view of (1) to (3) beckons.

Does EC have an explanatory advantage? The EC resolution is crucially incomplete. Where (1) to (3) seem true, an EC theorist locates satisfied truth conditions. But satisfied truth conditions do not explain plausibility. What directly explains the plausibility to a person of a sentence is whatever leads the person to regard the sentence as true. An explanation of plausibility must cite something that influences the person's thinking.

Presumably EC theorists will supplement their explanation with claims about our use of commonsense standards for judging (1) and skeptically encouraged standards for judging (2). To involve varying truth conditions in the account, EC theorists will also claim that our standards are sufficient for the truth of (1) and (2) in their respective contexts. They might say that we regard each context's truth condition itself as our standard there.

Imputing a standard for judgment yields the same immediate explanation of plausibility, whether or not we add that it establishes truth. So the distinctive EC claim that our standards establish the truth of (1) and (2) is idle for this explanatory purpose.

Some of the standards have a notable vulnerability that EC overlooks. Our inclination toward (2) results from some more or less philosophical reasoning. It is alluring enough to draw us toward (2). But like most alluring philosophical arguments, it is subject to reasonable objection. At best, its success remains in doubt. EC sweeps the doubts aside and declares that the reasoning in support of (2) yields genuine truth conditions for knowledge. No philosophical reasoning deserves such deference.

Is the EC claim that the relevant standards establish truth an asset on other grounds? A complete account of our responses to (1) through (3) also says why we reasonably and confidently employ the pertinent standards. EC theorists can hold that our standards for judging (1) and (2) are somehow tacitly acquired in learning the meaning of "know." They might say that our learning partly consists in acquiring a rule for accessing a context's truth condition, and the rule yields the standards that make (1) and (2) plausible to us. EC theorists can then hold that the analytic link of these standards to the meaning of "know" explains why the standards seem so trustworthy to us.

Invariantists can offer a better explanation of our reasonable confidence in the standards. Invariantists can observe that although the standards that we apply to (1) and (2) may include ones we are taught in learning "know," they are not therefore semantically guaranteed to be necessary or sufficient conditions for the application of the term. They can be stereotypes offered to guide usage, well supported rules of thumb, or apparent implications of convincing theories. Nothing readily available shows us which appealing standards fail here. An account along these lines explains the attractiveness of (1) and (2) in a way that better accommodates other features of the situation. Allowing our standards to be fallible lets us be right about the conflict that we see among the propositions that attract our acceptance as we consider (1) to (3).

Brief Replies

R1. Stew's EC is the conjunction of

- (a) Ascriptions of knowledge are context sensitive,

and

- (b) The context sensitivity of knowledge ascriptions provides the basis for resolving the skeptical paradox.

Stew writes that I opposed (a) by offering a particular invariantist view of knowledge to account for “our intuitions that saddle us with the paradox.”

I did not offer the invariantist view to explain what Stew regards as our intuitions about knowledge. As may now be clear, I believe that those intuitive responses have no simple connection to the actual conditions on knowledge. I offered the view as a possibility about the nature of knowledge that accords with our having a specific intelligent basis for a variety of the knowledge attributions that the view counts as false. The basis is that those attributions are efficient loose talk. The “really and truly” test is supposed to show that this status is how we ourselves think of the attributions in our best moments.

R2. Stew suggests that there is a telling vacillation in intuitive knowledge judgments within philosophy. During Gettier problem discussions we readily allow that we know many things, while discussions of skepticism give rise to intuitions that we do not know those things. EC can offer shifting truth conditions within philosophy to explain the vacillation. I discussed the possibility of a context-sensitive view of “knows” that finds within philosophy no variation of truth conditions. This view leaves the vacillation unexplained.

I was thinking of EC as Stew’s (a) only. I presented the particular context-sensitive view of “knows” as an example of how (a) could be true while leaving philosophy unaffected. In support of this sort of context sensitivity, I offered just the observation that philosophy seems to occupy a “serious” context.

What does account for the vacillation within philosophy? When we consider the Gettier problem, most of us (some philosophers are spoilsports about this) apply commonsense standards for knowledge. Any of them will do. When we think hard about skepticism, we tend to apply more severe standards. (Here too there have been spoilsports, notably G.E. Moore.) These are reasonable proclivities, if only because the standards are initially plausible. Philosophers’ use of such standards explains the vacillation. Any claimed connection between this variation in our standards and the actual truth conditions requires argument.

R3. Stew offers a speculative account of why some find it difficult to accept EC (where EC is understood as the conjunction of (a) and (b)). We regard knowledge as a good thing. If EC is true, our external world “knowledge” is not so hot. We only meet disappointing truth conditions. Some hold out against EC to avoid this disappointment.

Probably this accounts for some resistance to EC. It does not hit home with me. My resistance derives primarily from two sources. First, it seems to me that skeptical views deny the very knowledge that is otherwise widely affirmed. EC implies otherwise. On further reflection, the thoughts continue to seem incompatible.

Explanatory success could conceivably outweigh this reflective liability. That brings us back to points about explaining our knowledge attribution inclinations. Invariantists can credibly explain our reasonably confident use of the varying standards for judgment that EC theorists invoke.³ The distinctive EC thesis is that some of our

relatively high standards, and some of our less high standards, are contextual truth conditions, or derive from them. This gives an EC explanation no advantage. It incurs the costs of caving in to dubious skeptical reasoning and requiring some durable impressions of incompatibility to be errors.

Notes

- 1 Some of these claims about the case may turn out to be inaccurate. By the time you read this, the question may well be settled. In any event, the example as described serves to make the point.
- 2 The sentence types display some contextual variation. “The painting” has different referents in different contexts and probably the name “Jackson Pollock” does too. But to formulate my thoughts with C1 and C2, we must hold fixed the referents of those two expressions. The resulting thoughts are irreconcilable.
- 3 Toward providing an invariantist explanation of our attitudes to the standards, I have just mentioned stereotypes, rules of thumb, and inference from theories. An adequate explanation would elaborate. The point is only that some invariantist explanatory possibilities are at least as initially credible as the envisaged EC account.

Contextualism Defended Some More

Stewart Cohen

1 When Should We Appeal to Contextualist Resolutions of Philosophical Paradoxes?

Earl Conee raises important issues about when we should appeal to a contextualist resolution of a philosophical problem. He argues, in effect, that contextualists are too quick to appeal to a contextualist resolution of our conflicting intuitions regarding knowledge. His strategy is to present another case that he argues is structurally identical to the knowledge case where a contextualist resolution looks to be uncalled for. This, he argues, should make us question the validity of appealing to a contextualist resolution of the knowledge case.

Conee’s case involves a conflict over whether a particular painting is by Jackson Pollock. On the one hand, an expert testifies that the painting contains Pollock’s fingerprint. On the other hand, an art institute has testified that the painting is not by Pollock. This leads us to vacillate over whether or not the painting is by Pollock.

About this case, Conee makes the following points:

- (a) Different considerations pull us in different directions regarding a particular matter of fact.
- (b) The considerations in both directions are fallible but we have no overriding basis that tells in favor of one set of considerations over the other.
- (c) We find this situation perplexing.

All the same, according to Conee, “This [case] is no paradox. There is a straightforward empirical fact of the matter (leaving vagueness aside).”

Conee then notes that the skeptical paradox is like his Pollock case in respect of (a) to (c). But, Conee argues, there is no temptation in the Pollock case to appeal to a contextualist semantics for the predicate “... is by Pollock.” Clearly here, the truth conditions are invariant. And this, he argues, should lead us to question whether a contextualist resolution of the skeptical case is correct.

Conee is surely correct that we should not necessarily be attracted to a contextualist solution anytime we have a case that satisfies (a) to (c). And I agree that a contextualist resolution of the Pollock case would be very implausible. But I will argue that there are important disanalogies between the two cases that make a contextualist resolution of the skeptical case plausible.

Perhaps the most important difference is that, as Conee himself points out, there is nothing paradoxical about the question of whether the painting is by Pollock. It is a straightforward empirical matter. Our only problem is that we lack empirical information – information that presumably would settle the matter.

But the knowledge case is not a straightforwardly empirical matter. The question about whether we know is arguably not an empirical question. In fact, there need not be any disagreement between skepticism and common sense about any empirical fact. In the skeptical case, the issue is conceptual. Different considerations pull us in different directions concerning whether or not a particular set of empirical facts is properly classifiable as an instance of knowledge.

Conee questions whether it is correct to call our conflicting responses about the knowledge case “intuitions.” For our each response is not “credible on its own,” but rather based on “something else.” I agree that our responses are based on various considerations. Our response that we know we see a zebra is based on the facts of the case as described. Our response that we don’t know we don’t see a cleverly disguised mule is based on the appeal of various skeptical principles. But I don’t see that they are not intuitions. Our response that we know we see a zebra is an intuitive judgment about whether the concept of knowledge applies in a particular case. And our response that we don’t know we don’t see a cleverly disguised mule is an intuitive judgment about the cogency of a particular argument. The fact that nothing empirical is at stake makes it appropriate to think of these responses as intuitive judgments.¹ And the fact that the Pollock case is clearly empirical makes it clear that our conflicting judgments in that case are not intuitive. We cannot decide by intuition whether or not a particular painting is a Pollock.

So there is an important difference between the Pollock case and the skepticism case. The former involves empirical judgments about an empirical matter of fact, whereas the latter involves intuitive judgments about whether a conceptual matter, or to ascend semantically, about whether a particular predicate applies. Of course, even if we are dealing with intuitive judgments concerning the application of a predicate, it does not follow that a contextualist treatment is called for. There must be a semantical model of the predicate that makes a contextualist treatment plausible. Different contextualists have proposed different models.

Lewis argues that ascriptions of knowledge involve an implicit quantification (Lewis, 1996). S knows P just in case S’s evidence eliminates every alternative. But what counts as an alternative that S’s evidence must eliminate depends on context. So for Lewis, the context sensitivity of “knows” is a function of contextual restrictions on the domain of quantification.

On my own view, the context sensitivity of knowledge is inherited from one of its components, i.e. justification (Cohen, 1999). In general, when a predicate has both a comparative form and a *simpliciter* form, context will determine the degree to which the predicate has to be satisfied in order to be satisfied *simpliciter*. So, for example, context determines how flat a surface must be in order to be flat *simpliciter*. Analogously, context will determine how justified a belief must be in order to be justified *simpliciter*. So if justification is a constituent of knowledge, and “justified” is context sensitive, that provides some reason to think “knows” will be context sensitive.

So this makes for another important difference between the skeptical paradox and Conee’s Pollock case. Contextualists about knowledge have provided semantic models for “knows” that explain its context sensitivity in a way that allows the dispute about knowledge to be resolved. But I can see no way to model the predicate “... is by Pollock” that would explain the sort of context sensitivity required to resolve the dispute over whether the painting is a Pollock.

2 Explaining the Skeptical Paradox: Contextualism versus Invariantism

Conee also argues that the contextualist does not really explain the appeal of the sentences that constitute the skeptical paradox. On the contextualist view, those sentences come out as true, in those cases where they strike us as true. But Conee notes that “satisfied truth conditions do not explain plausibility.” According to Conee, the contextualist needs an explanation of the plausibility of those sentences involving “... something that influences the person’s thinking.”

But he then goes on to explain quite well how a contextualist can explain the appeal of the sentences that constitute the paradox:

EC theorists can hold that our standards for judging (1) and (2) are somehow tacitly acquired in learning the meaning of “know.” They might say that our learning partly consists in acquiring a rule for accessing a context’s truth condition, and the rule yields the standards that make (1) and (2) plausible to us. EC theorists can then hold that the analytic link of these standards to the meaning of “know” explains why the standards seem so trustworthy to us.

I could not have said it better myself. But Conee goes on to argue that invariantists can offer a better explanation of our intuitions regarding the sentences constituting the paradox.

Invariantists can observe that although the standards that we apply to (1) and (2) may include ones we are taught in learning “know,” they are not therefore semantically guaranteed to be necessary or sufficient conditions for the application of the term. They can be stereotypes offered to guide usage, well supported rules of thumb, or apparent implications of convincing theories. Nothing readily available shows us which appealing standards fail here. An account along these lines explains the attractiveness of (1) and (2) in a way that better accommodates other features of the situation. Allowing our standards to be fallible lets us be right about the conflict that we see among the propositions that attract our acceptance as we consider (1) to (3).

So on Conee's view, it is better to explain our intuitions about the sentences (1) to (3) in a way that allows us to be wrong about them. By so doing we can vindicate our judgment that our skeptical endorsement of (2) conflicts with our commonsense endorsement of (1). Recall that on the contextualist view, that judgment is a mistake that results from conflating contexts.

So here we have a clear contrast between the contextualist response to the paradox, and Conee's invariantist account. On the contextualist view, our first order judgments (1) to (3) are correct. More specifically, our judgment that (1) is true is correct when evaluated at everyday contexts, our judgment that (2) is true is correct when evaluated at stricter skeptical contexts, and our judgment that (3) is true is correct at every context. But our meta-judgment that (1) to (3) are inconsistent is a mistake. On the invariantist view, we are correct in our meta-judgment that (1) to (3) are inconsistent, but we are mistaken in at least one of our first-order intuitive judgments concerning (1) to (3).

Each view provides an explanation for the sort of mistake it attributes to us. On the invariantist view, our mistaken endorsement of both (1) and (2) results from the fallibility of stereotypes, rules of thumb, and our ability to see logical implications.

On the contextualist view, our mistaken judgment that (1) conflicts with (2) results from our conflating contexts. And the contextualist can appeal to the very same kind of mistake in the case of the puzzle concerning flatness to avoid the charge of special pleading.²

So which explanation is better, the invariantist or the contextualist? It is important to see that the contextualist explanation of our intuitions in the paradox, if correct, constitutes a resolution of the paradox. It resolves the apparent inconsistency of (1) to (3) in a way that accounts for the cogency of skeptical arguments without compromising our strong sense that we know much about the world.

But the invariantist explanations suggested by Conee do not, in themselves, resolve the paradox. According to those explanations, there is a genuine conflict between our commonsense judgment that (1) is true and our skeptical judgment that (2) is true. Owing to the use of stereotypes, or rules of thumb, and so on, we are mistaken in at least one of those judgments. The problem is that we are unable to determine which is the mistaken judgment. As Conee notes, "Nothing readily available shows us which appealing standards fail here." In the end, contextualism will have to be weighed against worked-out versions of invariantism which explain which standards fail and why.³

3 Objections to Contextualism

Conee mentions two reasons in particular to worry about the feasibility of a contextualist resolution of the paradox: "It incurs the costs of caving in to dubious skeptical reasoning and requires some durable impressions of incompatibility to be errors." Let's take these in reverse order.

As Conee says, contextualism holds that our impression that (1) to (3) are inconsistent is an error. But (1) to (3) constitute a paradox, precisely because we have the "durable impression" that each is true and that as a set they are inconsistent. But any response to the paradox will require that at least one of these impressions is erroneous. While the contextualist holds that our impression that (1) to (3) are jointly inconsistent is in error, the invariantist holds that our impression that they are all true is in error. Moreover, as I argued earlier, the contextualist can bolster the plausibility of attributing

such a mistake to us by appealing to an analogous mistake in our sometimes paradoxical thinking about flatness.⁴ So I don't see that contextualism here is at a comparative disadvantage.

In what way is the skeptical reasoning dubious? Surely there is no obvious error in the reasoning. Otherwise we would not be stuck with a stubborn paradox. We might say the reasoning is dubious in that we are confident that the conclusion is false. Surely this is true. Again, the reason we are stuck with the paradox is that we find the conclusion of the skeptical argument to be intuitively unacceptable.

Does contextualism cave in to skeptical reasoning? Yes and no – the former because contextualists allow that the sentences expressing skeptical conclusions are true, evaluated at skeptical contexts, the latter because the contextualists hold that those very same sentences are false when evaluated at everyday contexts.

Does this concede too much to the skeptic? Of course, this is something about which reasonable people can disagree. But by my lights, what is truly surprising and intuitively repugnant about skepticism is the thought that all along in our everyday lives, when we have been uttering sentences of the form “S knows P,” we have been speaking falsely. Contextualism denies that this is so.

Contextualism does explain the stubborn appeal of skeptical arguments by allowing that there is some truth in skepticism. But the strategy of the contextualist is to limit the damage. So although contextualism does concede there is some truth in skepticism, it limits the damage by showing how this concession still allows for the truth of our everyday knowledge ascriptions. To me this is no worse than conceding that in certain limited and unusual contexts, we can truly say “Nothing is flat,” while at the same time holding that many of our everyday utterances of the form “X is flat” are true.

Notes

- 1 Of course there is a general problem about what an intuition is, but to discuss it would go beyond the scope of this essay, not to mention my understanding.
- 2 I discuss the “flat” case in my first contribution to this exchange. The contextualist need not hold that our first-order knowledge ascriptions are infallible. The contextualist need only hold that the explanation of what is occurring in the skeptical paradox is that our judgments reflect our grasp of the correct standards for the context.
- 3 I am indebted to Earl Conee here for helping me to see this issue more clearly.
- 4 Again, see the first response to Conee in this exchange.

References

- Cohen, S. (1999) Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives* 13: 57–89.
- Lewis, D. (1996) Elusive knowledge. *Australasian Journal of Philosophy* 74: 549–567.