# The Semantics of Knowledge Attributions*

*Nikola Kompa*

Universität Münster, Germany

The basic idea of conversational contextualism is that knowledge attributions are context sensitive in that a given knowledge attribution may be true if made in one context but false if made in another, owing to differences in the attributors' conversational contexts. Moreover, the context sensitivity involved is traced back to the context sensitivity of the word "know," which, in turn, is commonly modelled on the case either of genuine indexicals such as "I" or "here" or of comparative adjectives such as "tall" or "rich." But contextualism faces various problems. I argue that in order to solve these problems we need to look for another account of the context sensitivity involved in knowledge attributions and I sketch an alternative proposal.

*Keywords:* conversational contextualism, context sensitivity, indexicality, knowledge attributions, semantics

## I

## The Basic Idea

Different things go by the name "contextualism." The form of contextualism I want to focus on has come to be called conversational or attributor contextualism. The conversational contextualist's basic claim is that a given knowledge attribution may be true in one context but false in another, owing to differences in the attributors' conversational contexts. More fully, the claim is this:

> [**EC**]: An attributor, A, in conversational context $C_A$, might say something true in saying: "X knows that P," whereas another attributor, B, in conversational context $C_B$, might say something true in saying: "X doesn't know that P." So A's knowledge attribution and B's knowledge denial can both be true; and this can be explained by pointing out differences between $C_A$ and $C_B$.

Call this the Epistemic Claim. Note that A and B are talking about the same person X, the same P, and the same time t. Consequently, X's epistemic position is exactly the same in both cases. Accordingly, the contextualist claims that knowledge attributions are *context sensitive*: their truth-value is sensitive to contextual variation. But the context in question is not the putative knower's context but rather the attributors' conversational context.

Now exactly what kind of context sensitivity is at issue here? Commonly, the context sensitivity of knowledge attributions is traced back to an alleged context sensitivity of the word "know." More specifically, contextualists who take a stance on the question at all try to compare the word "know" either to genuine indexicals such as "I" and "now," or to comparative adjectives such as "tall," "flat" or "rich." So the epistemic claim is commonly amended by a Semantic Claim that could, roughly, be put thus:

> [SC] Indexicals and predicative uses of comparative adjectives are obviously context sensitive. And the context sensitivity of the word "know" can be modelled on one of these two kinds of context sensitivity.

An example due to Keith DeRose will help to illustrate the contextualist point (DeRose 1992, 913).[1] He wants us to consider the following two cases:

> *Bank Case A:* My wife and I are driving home on Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." I reply, "No, I know it'll be open. I was just there two weeks ago. It's open until noon."
>
> *Bank Case B*: My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, "Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as I was before that the bank will be open then, still, I reply, "Well, no. I'd better go in and make sure."

According to the conversational contextualist, the truth of the knowledge attribution in case A and the knowledge denial in case B is due to changes in the attributors' conversational contexts. So what exactly has changed from one context to the other? First, in case B it is much more important to be right—much more is at stake. Secondly, a certain error possibility is mentioned in case B that has been given no attention in case A. And thirdly, the error possibility is taken seriously in case B.

All the changes are supposed to be changes in the attributors' conversational contexts. So the next question is: what is a conversational context? Since a full answer to this question would take us into such difficult areas as the dynamics of discourse and the interplay of conversational mechanisms, I will not attempt to fully answer the question here but only outline what I take to be a promising starting point. To a first approximation, think of a conversational context as the set of shared presuppositions of the participants in a conversation.

The idea goes back to Robert Stalnaker who proposes to identify a context with what is presumed to be common to the participants in the discourse (cf. Stalnaker 1998, 98). But a conversational context is not just any old set of presuppositions. It is the set of presuppositions the participants take for granted in pursuing their conversational aims. More specifically, the participants' shared interests and intentions,

the purpose or point of the conversation, and the shared background assumptions determine what is presupposed in the conversation. And a conversational context is characterized by the set of those shared presuppositions.[2] Now when the question is whether one could truly say of X that he knows that P, then the participants' shared background assumptions, their shared interests, intentions, and purposes are relevant to answering that question.

They help setting a certain standard: the standard X must live up to in order to know that P. So I take it that the contextualist should subscribe to the following **P**ragmatic **C**laim:

> **[PC]** A Conversational Context is characterized by what is presupposed by the participants in a conversation. And what is presupposed in a conversation is, in turn, determined by the participants' shared assumptions, interests, intentions, and purposes.

So let us agree to call someone a contextualist only if he subscribes to the epistemic claim [EC], the semantic claim [SC], and the pragmatic claim [PC]. But now the contextualist seems to face several problems. So in what follows, I will first discuss some of the problems. I will then, in a second step, discuss a suggested solution. The solution which is put forth by Keith DeRose captures an important insight. But DeRose tries to underpin the important point with the help of the semantic claim [SC] and thereby drains the solution of much of its explanatory power. So I will claim that, in the end, the contextualist is well-advised to give up the semantic claim for another account of the context sensitivity involved in knowledge attributions. In a last step, I will sketch an alternative account of the context sensitivity involved and outline how it can help to solve the problems.

## Problems

(1) The contextualist claims that one and the same knowledge attribution can be true when made in one context and false when made in another context. And this may be so even if the would-be knower, the thing known, and the time of the attribution are held fixed in both cases. Now suppose that attributor A in context $C_A$ says something true in uttering "X doesn't know that P," while attributor B in context $C_B$ says something true in uttering "X knows that P." Suppose further that attributor A considers B's knowledge attribution, being well aware of the low standards employed in B's *context*. But then it looks as if A would say something true were he to say:

> **[1]** Attributor B says something true in uttering "X knows that P" but X doesn't know that P (Kompa 2002, 5).

(2) Timothy Williamson raises a somewhat similar problem. He points out that the contextualist seems to be "committed to the assertion

> **[2]** 'Everyday propositions are true and I don't know it' " (Williamson 2001, 26).

For the contextualist has to concede that at least some of our everyday knowledge attributions are true. Yet the contextualist, in exposing his view, is doing epistemol-

ogy. He is thereby considering far-fetched possibilities. And according to the standard contextualist account, the possibilities thereby become relevant (cf., e.g., David Lewis' *rule of attention*, Lewis 1996). So the contextualist will also agree that he doesn't know anything unless he can rule out these possibilities—something he cannot do, presumably.

(3) Thirdly, doesn't the contextualist have to concede that the last line of the following dialogue (due to Palle Yourgrau 1983, 295) is what B should say according to contextualism, although it sounds rather odd?

A: Is that a zebra?

B: Yes, it is a zebra.

A: But can you rule out its being merely a cleverly painted mule?

B: No, I can't.

A: So, you admit you didn't know it was a zebra?

B: No, I *did* know *then* that it was a zebra. But after your question, I no longer know.

In other words, it seems as if a contextualist is committed to accept sentences such as:

**[3]** I knew that P a moment ago, but now I no longer know.

Now similar mechanisms seem to be at work in these cases. (i) In all three cases something like the following **P**rinciple of **C**ross-**C**ontext **A**ttribution **[PCCA]** is employed: It is possible to attribute knowledge across contexts. That is to say, it is possible to attribute knowledge even when the attributor and the subject of the attribution do not share a context. (Actually, that is what the contextualist has been doing all along, so he'd better accept the principle). (ii) And in Williamson's and Yourgrau's examples something like the following **P**rinciple of **C**ontext **C**hange **[PCC]** seems to be invoked: a context change can be brought about by mentioning or at least considering a hitherto ignored error possibility.

So far, we have collected the following conflicting data: on the one hand, we have the examples the contextualist provides in order to support his claim. The examples seem to show something important about our ordinary usage of the word "know," namely that there are cases in which a speaker—given his interests and intentions—truly says that a person X knows that P, while another speaker—with different interests and intentions—truly says that X doesn't know that P. Attributing knowledge seems to be a context sensitive matter. On the other hand, all of the above sentences [1] to [3] sound odd. But if knowledge attributions were context sensitive, then these sentences shouldn't sound odd. Moreover, the sentences seem to follow from the contextualists' basic claim together with two seemingly innocuous principles of context change and cross-context knowledge attribution.

## II

Now the question I want to address in what follows is this: is it possible to explicate the context sensitivity at issue in knowledge attributions in such a way as to account

for the conflicting data? In particular, is it possible to explicate it in such a way as to account for the problematic sentences?

## Indexicality

Keith DeRose makes a proposal how to characterize the context sensitivity so as to solve the above problems. He suggests that the real commitments of contextualism have to be couched in metalinguistic terms. And the analogy here is with indexical expressions.

> One hour ago I was in my office. Now I am in the word processing room. How can I truly say where I was an hour ago? I cannot truly say, "I was here," because I wasn't here; I was there (DeRose 2000).

And similarly with knowledge attributions, for speaker B in Yourgrau's dialogue cannot simply say, "I did know then that it was a zebra," just as DeRose cannot say "I was here" when he wants to tell us where he was an hour ago. The best DeRose can say is: "My previous location claim was true." And the best B can say is: "My previous knowledge claim was true"—according to DeRose. And we could, following DeRose, handle the other problematic cases more or less analogously.

For example, a claim like "In context $C_1$ she knew that P while in context $C_2$ she doesn't know" could be rephrased thus: "In context $C_1$ she counted as knowing while in context $C_2$ she doesn't count as knowing," or alternatively: "She meets the standards set by context $C_1$ but fails to meet the standards set by context $C_2$." (cf. DeRose ibid.).[3] So what it comes down to is this: disquotation is not always as straightforward as one might think. But that is hardly news anyway. What might be news is that knowledge attributions are a case in point.

But even if we accept DeRose's diagnosis in order to get the contextualist's commitments straight, we are not yet out of trouble. For if we accept DeRose's diagnosis we thereby seem to buy into another serious problem for contextualism. In his diagnosis, DeRose is relying on an analogy between the word "know" and genuine indexicals such as "I" or "here."[4] He endorses the semantic claim [SC]. But as various philosophers have convincingly argued, it is highly questionable whether the word "know" is an indexical. Consider the following sceptical argument [SA], where H might be any skeptical hypothesis, e.g., that I am a Brain in a Vat, and O might be any ordinary hypothesis, e.g., that I have hands:

> **[SA]** I don't know that not-H.
>
> If I don't know that not-H, then I don't know that O.
>
> Therefore: I don't know that O (DeRose 1995, 1).

Now as Stephen Schiffer has pointed out, the contextualist who takes "know" to be an indexical will give the following diagnosis of the paradoxical nature of [SA]:

> We instinctively know that the conclusion-asserting sentences of [SA] would express a false proposition in a quotidian context in which sceptical hypotheses weren't at issue, and we mistakenly suppose that it's asserting the same false proposition in [SA]. In other words, [SA] strikes us as presenting a profound paradox merely because we're ignorant of what it's really saying, and this because we don't appreciate the indexical nature of knowledge sentences (Schiffer 1996, 325).

But the contextualist solution involves, as Schiffer argues, a very implausible error theory: We are taken in by [SA] simply because we are mistaken about which propositions our utterances express. Moreover, the error theory is in tension with the semantics. For if "know" were an indexical, we shouldn't be confused about which propositions our knowledge attributions express since we are usually not confused about which propositions our indexical utterances express (cf. Schiffer ibid.).

So one of the main problems for the proponent of an indexical semantics for "know" is this: competent speakers are well aware of the context sensitivity of genuine indexicals but they are not (at least not fully) aware of any alleged context sensitivity of the word "know." Indexicals wear their context sensitivity on their sleeves. The word "know" obviously doesn't (cf. Kompa 2002, 10/11).[5]

This connects up with another problem. If "know" is an indexical, then speaker A who is employing demanding standards and therefore *denies* that X knows that P and speaker B who is employing relaxed standards and *claims* that X knows that P do not really disagree. For A doesn't deny what B asserts. A says something like the following: X doesn't know that P relative to high standards. While B says something like the following: X knows that P relative to low standards. As Richard Feldman points out: "It is *always* possible to explain the conflicting inclinations we have in these cases by appeal to context dependence or ambiguity: [. . .] Similarly, whenever people seem to disagree, it is possible to say that there is no real disagreement, but that they are asserting and denying different propositions" (Feldman 2001, 72).[6] So if "know" were an indexical, then the disagreement between A and B would be only apparent. But that is not what we want to say on an intuitive basis, at least not in all cases. There are cases where we seem to have a genuine conflict.

## Comparative Adjectives

So the indexical account of "know" is fraught with problems. But if "know" isn't an indexical, then we need another account of the context sensitivity involved. Here is an alternative proposal. Maybe "know" is more aptly be grouped together with context sensitive expressions such as ". . . is tall" or ". . . is rich" or ". . . is flat."

Unfortunately, to try to assimilate the context sensitivity of the word "know" to the context sensitivity of predicative uses of comparative adjectives such as "tall" or "flat" is not a very promising move either. For although the proposal fares better—if only slightly—with respect to the problems the indexical approach faces, it has problems of its own.[7] The most pressing problem is that there are lots of linguistic data speaking against the proposal—as, e.g., Jason Stanley has made clear (cf. Stanley [ms], 11 ff). Stanley points out, for instance, that predicative uses of comparative adjectives allow for modification, as in "He is very tall." But "know" doesn't allow for any such modification. And while we can say something like "five feet tall" or "20 years old," there is no—as he calls it—"natural measure phrase" with "know."

Moreover, there is a comparative form of all the comparative adjectives (hence their name). We have "taller than," "richer than," "flatter than," etc. But there is no comparative form of "know," and so on.[8] Now DeRose emphasizes that

> the best kind of case for contextualism is not an indirect argument that takes as its premise that some other term, like "tall," is context-sensitive, and then argues that because "knows" is so similar to "tall," "knows" too is context-sensitive (DeRose [ms], 22).

Since he is not relying on such an argument, the defectiveness of the analogy between "know" and comparative adjectives is no objection to his contextualist proposal. But then he owes us an account of the context sensitivity of "know," because now the context sensitivity of "know" seems to be a context sensitivity *sui generis*. So pending an account of this "new" kind of context sensitivity, his proposal is bound to be somewhat *ad hoc*. Also, he has to admit that the kind of context sensitivity involved has the following interesting feature: competent speakers are not fully aware of it. It takes some courses in epistemology to make them recognize it. And if that is so, then there are semantic facts that are not accessible even to competent speakers. Consequently, the semantics of "know" comes out rich but partly inaccessible.

## III

Let us take stock. On the one hand, there are examples such as DeRose's bank case that seem to show that whether a given knowledge attribution is true or false is an interest-relative matter. Knowledge attributions seem to be context sensitive, given a notion of a conversational context as a set of shared presuppositions, determined by the participants' interests, intentions, etc. On the other hand, we have sentences like "I knew that P a moment ago, but now I no longer know" that we can hardly make sense of, but which seem to follow from the contextualist's basic claim together with two simple principles of context change and cross-context knowledge attribution. Moreover, to try and explicate the context sensitivity involved in terms of indexicality or similar phenomena has been shown unsatisfactory. The semantic claim [SC] has to be given up. So the contextualist has to come to terms with the problematic sentences and he has also to provide another account of the context sensitivity involved in knowledge attributions.

The problematic sentences seem to follow from the contextualist's basic claim together with the principle of cross-context attribution **[PCCA]** and the principle of context change **[PCC]**. Accordingly, one could either give up the basic claim[9] or try to find fault with the principles. I will opt for the latter. Of course, that doesn't obviate the need for another account of the context sensitivity involved—given that the semantic proposals discussed above have been considered unsatisfactory. So I will also very briefly sketch an alternative account.

### A Sketch

Here is the basic but still very sketchy idea. The word "know" is context sensitive. Whether someone can truly be said to know that P is an interest-relative matter. But the context sensitivity at issue is much more subtle than in the case of indexicals. And it is not only that the word "know" exhibits a sensitivity to contextual factors that is less obvious than in the case of indexicals. The relevant contextual factors themselves are more complex too. They are harder to pin down than in the case of indexicals. More specifically, the word "know" is, as Michael Williams puts it, a honorific term. Whether someone deserves to be described as someone who knows that P depends, among other things, on whether we want to rely on his claim that P or not. This in turn depends on our interests and intentions, on what is at stake, on

which alternatives we have reasons to take seriously, etc. So an utterance of "X knows that P" is true in context C only if X might reasonably count as knowing that P, given the interests and intentions (manifest in the participants' shared presuppositions) operative in C. (That is not to deny that P must be the case if X is to know that P. That is a necessary condition for knowledge if ever there was one.)

But the context sensitivity is not only very subtle. It is also very pervasive because it affects not only "know" but predicative uses of most other adjectives—and not just comparative adjectives—as well. For example, whether a given object O might truly be said to be red, say, depends on whether O is as close to being paradigmatically red as the participants care in the context in question. And this in turn depends on the participants' interest and intentions.[10] But now someone might object that the participants' interests and intentions alone cannot make any of their claims (e.g., their claim that O is red) true. That would be to cheap. The rules of language use impose certain constraints on how the terms in question are to be employed. So the participants have to comply with the rules. They have to be competent, that is. Still, there is always room for error and misinformation. So the participants better be well-informed. But, most importantly, the participants' interest and intentions have to be reasonable too—or so the objection goes. I agree. The participants have to be competent. Also, error is always possible. (For example, the participants might be mistaken about whether X might reasonably count as knowing that P—even given their own interests and intentions). And there is a distinction between "true by the standard that in fact governs the context" and "true by the standard that should reasonably govern the context." The question of whether the standard that is in fact operative is the one that should be operative is a sensible question. But given that the participants in the conversation are in fact reasonable and competent, the two things may well coincide. And if that is so, then for most predicative uses of adjectives the following seems to hold: if a competent speaker says of a particular object X that it is such and such, then what she says is true if X might reasonably count as being such and such in the context at hand, given the interests and intentions operative in the context and given that the interests and intentions be themselves reasonable. And analogously in the case of knowledge attributions.[11]

## Semantics vs. Pragmatics

But now one might wonder whether the contextual mechanism operative in determining the truth-value of knowledge attributions aren't too complex to be part of the meaning of "know" on any plausible analysis of its meaning. One might therefore think that the context sensitivity involved has to be a pragmatic phenomenon. So what is it, a pragmatic or a semantic phenomenon? At first glance, both answers seem to enjoy a certain plausibility. The proponent of the semantic approach (call him the semanticist) will say something like the following. If speaker A says "X knows that P," then what he said (or the proposition he thereby expressed) can be rendered thus:

[S] "X knows that P relative to standard S."

The semanticist has to concede that meaning facts might not be transparent, for we are not aware of any such contextual relativity to standards of knowledge.

Moreover, he has to admit that intuitions about what has been said in a given utterance are not reliable since we wouldn't take A as having said any such thing as [S]. The semanticist also faces the second problem mentioned above, the problem that there seems to be a genuine disagreement between speaker A who claims that X knows that P and speaker B who denies that X knows that P. But if the semanticist is right, there is no disagreement here, at least not at the level of what is said. Speaker A simply said that X knows that P relative to standard S1, while speaker B said that X doesn't know that P relative to standard S2. (Of course the semanticist could try to argue that there is disagreement nonetheless—it just doesn't come in at the level of what is said. It comes in over the question of whether we should, in a given case, adopt standard S1 or standard S2.)

The proponent of the pragmatic approach (call him the pragmatist), on the other hand, will hold that the context sensitivity involved is a pragmatic phenomenon. It doesn't show up at the semantic level—the level of what is said or the proposition expressed. (Of course one could argue, as some philosophers do, that the semantic level is not the level of what is said. One might, for example, posit a semantical deep structure that is by definition not transparent. But let us put that aside for the moment.) But now the resulting contextualism is in danger of collapsing into invariantism. For if the context sensitivity is pragmatic, how can it be relevant to the truth-value of a knowledge attribution? Accordingly, the pragmatist who wants to stick with the epistemic claim [EC] has to hold that pragmatic phenomena such as speakers interests and purposes affect the truth-value of a knowledge attribution. And the idea is not that pragmatic factors such as speakers' interests might help determine the semantic value of one of the terms uttered. That happens all the time. Rather, these factors do not manifest themselves at the semantic level—the level of what is said or expressed—at all. But they affect the truth-value of knowledge attributions nonetheless. Consequently, the link between semantics and truth-value has to be severed. And that is a fairly radical departure from tradition. But there is also a conservative touch to the pragmatist picture. According to the pragmatist, A said that X knows that P and B said that X doesn't know that P. They really disagree over whether X knows that P. And they are not ignorant of what they say. Rather, they say exactly what they think they say.

So it all depends on what the correct adequacy conditions for a theory of meaning are. If a theory of meaning is adequate only of it construes the meaning of a term as fully transparent to those competent with the term, then the context sensitivity of "know" has to be a pragmatic phenomenon: standards are invoked for explanatory purposes. They help to explain why a given knowledge attribution has the truth-value it has. But they needn't be part of the semantics to serve that purpose. If, on the other hand, a theory of meaning is inadequate if it allows pragmatic factors to be relevant to determining the truth-value of a given knowledge attribution (in the way discussed above), then we will have to take the context sensitivity to be a semantic phenomenon—conceding that there are semantic facts that are not accessible even to competent speakers. So in building up our semantic theory, we would "end up positing parameters that are not built into the deep psychology of ordinary competence"—as John Hawthorne puts it (Hawthorne 2004, 109). We would thereby buy into some variant of what Stephen Schiffer has called the *meaning-intention problem*, though (cf. Schiffer 1992 & Hofweber 1999).

## The Two Principles

I don't want to take a stance on the question of whether the context sensitivity at issue is to be construed semantically or pragmatically. My sketchy account of the context sensitivity involved in knowledge attributions is neutral on the question of how these further issues are to be resolved. All I want to claim here is that if the account outlined above is at least roughly on the right track, then we can solve the above problems by dismissing the two principles. So let us take a closer look at the two principles.

Take the principle of cross-context knowledge attributions first. I think that DeRose is right as to the contextualist commitments. They are best couched in metalinguistic terms. But DeRose fails to explain why that should be so. Part of his explanation is that "know" is an indexical. But that doesn't explain what needs to be explained. If I want to say where I was an hour ago, I could simply say: "I was *there*, in my office." There is no need to go metalinguistic in the case of indexicals. An indexical expression can easily be substituted in cross-context reports. So if "know" were an indexical, then a similar substitution should be available to make cross-context knowledge attributions. But, obviously, no such substitution is available.[12]

And given that knowledge attributions are context sensitive in the way indicated above, one shouldn't expect a cross-context knowledge attribution to be a simple matter. I suspect that [PCCA] seemed innocuous at first only because the context sensitivity involved in knowledge attributions was modelled on the case of indexicals. But indexicals allow us to think of a context in much more simple terms than does the more subtle form of context sensitivity at issue here.

Moreover, in order to successfully make a cross-context attribution such as "X knew that P," it has to be clear which standard governs the attribution. In the case of indexicals, the context of utterance governs. In the case of expressions such as "local" or nearby" both the subject's context and the attributor's context might govern. As Hawthorne rightly points out, "local is flexible with regards to whose location is relevant" (Hawthorne 2004, 103). But in this respect at least, "know" seems to work more like an indexical. That is to say that in order to properly attribute knowledge across contexts, it might be necessary to adopt the standard that is supposed to govern the attribution. But, presumably, one cannot simply adopt a certain standard at will.

Now using a metalinguistic paraphrase such as "A said something true when he said 'X knows that P'" might help. Firstly, it makes the reference to the standard/context in question much more obvious. It is therefore less misleading than a disquotational attribution such as "X knew that P." Secondly, it seems to allow us to talk about the standard which governs a certain knowledge attribution without forcing us to adopt that standard. But maybe that is still not good enough. Take again sentence [1]. A says:

> **[1]** Attributor B says something true in uttering "X knows that P" but X doesn't know that P.

Now sentence [1] simply sounds odd. The more one comes to think about contextualism, the less odd it might sound. But still, something is amiss with [1]. Now what A does is to approve of B's standard and at the same time adopt a different standard. But maybe one cannot approve of a different standard than the standard one is adopting in making a knowledge attribution. After all, what is at issue is not just warranted assertibility, but truth. There might be something like a pragmatic contra-

diction in approving of one standard and at the same time adopting a different standard when it comes to the question of whether a given knowledge attribution is true or not. And sentence such as [1] and [2] might be defective for exactly that reason. So the principle of cross-context attribution needs to be modified in order to take these difficulties into account.

Now take the second principle, the principle of context change. It should be clear by now that contexts don't change as easily as is suggested by the problematic sentences discussed above. Since a context is characterized by the shared assumptions, interests, purposes, and intentions of the participants, the mere mentioning or even the considering of an alternative will hardly affect a change in context—i.e., a change in the point of the conversation, the shared assumptions, and the interests and intentions of the participants. Of course, it might do so. But it needn't do so. We usually don't switch back and forth between contexts as easily as the examples suggest. Consequently, Yourgrau's dialogue would—in real life—take a quite different course. Either B thinks that there is some reason to take the alternative mentioned by A seriously. Then he will say something like: "You are right, I don't know and I didn't know a moment ago either." But no context change has been affected by the mere mentioning of the possibility by A. Or B will refuse to take the alternative seriously and say, with DeRose, something like: "Painted mules [. . .]! C'mon! That's absurd. Get outta here with that crazy idea" (DeRose [ms], 4).

## Notes

* I want to thank the organizers and the participants of the *Bled Conference on Contextualism* 2004 for making it such an interesting and stimulating conference. I would also like to thank the participants of the colloquium on theoretical philosophy at the University of Muenster in summer 2004. And I am especially grateful for conversation with and comments from Oliver Scholz, Sebastian Schmoranzer and Stephen Schiffer.

1. Here is another example due to Stewart Cohen (Cohen 2000, 95): "Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask if anyone knows whether the flight makes any stops. A passenger Smith replies, "I do. I just looked at my flight itinerary and there is a stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary, anyway. It could contain a misprint. They could have changed the schedule since it was printed, etc." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago on the basis of the itinerary. They decide to check with the airline agent."
2. As I said, this is only a starting point. The contextual factors that have to enter into a characterization of a conversational context may well turn out to be too complex to be amenable to any such treatment in terms of the participants' shared presuppositions at all.
3. If, on the other hand, you tried to disquote a metalinguistic utterance such as "My previous knowledge claim was true" uncritically into an object language statement such as "I knew that P," you would thereby commit what DeRose calls the *fallacy of semantic descent.* (DeRose 2000). Here is an example of his to illustrate the point. The phrase "it is possible that P" has, arguably, an epistemic reading, something like: "for all I know, P." Now let speaker A say "It is possible that P" and thereby say something true (on the epistemic reading). And let speaker B say "It is not possible that P" and thereby also say something true (on the epistemic reading). Now, of course we are not forced to conclude that the same thing is both possible and not possible.
4. Whether he has to rely on the analogy is a different matter. As I try to make clear in the last part of the paper, one could combine the metalinguistic part of the diagnosis with another account of the context sensitivity involved.
5. Wayne Davis speaks, very aptly, of semantic blindness in these cases.

6. And he goes on to claim that there are many case, for example in cases of a moral controversy, in which "a contextualist solution to our wavering intuitions" (ibid., 73) is not plausible at all. So the question is whether the explanation for our "wavering intuitions" regarding knowledge attributions is to be modelled on cases of moral controversy or rather on those case where we aptly appeal to context sensitivity.
7. It fares better for the following reasons. (i) Although speakers are, to some extent, aware of the context sensitivity of comparative adjectives (and predicative uses thereof), they are not as fully aware of it as they are in the case of indexicals. (ii) If A says "X is rich" and B says "X is not rich," then they seem to disagree—just as in the case of knowledge attributions.
8. A couple of qualifications are in order here. Stanley himself rightly observes that there are negative comparative adjectives like "flat" or "small" or "young" which do not co-occur with measure phrases either (Stanley 2002, 13). And as Rosemarie Rheinwald and Sebastian Schmoranzer have pointed out to me in conversation, to say that there are no comparative forms of "know" is too strong. The following, e.g., seems to be a perfectly good utterance: "He knew better than all the others how dangerous the chemicals were." Also, as soon as we talk about justification instead of knowledge, all our problems disappear. I am indebted to Sebastian Schmoranzer for helpful discussion of these issues.
9. One could simply deny the context sensitivity of knowledge attributions. One might try an invariantist explanation instead—to use a term coined by Peter Unger—and claim that standards for knowledge are invariant across contexts. If the invariantist wants to avoid scepticism he will take the standards to be invariant but low. (He might claim, though, that what varies from context to context are standards for warranted assertibility. Also, he could hold that a change in truth-value is brought about only by a change in the inferential context but not by a change in the conversational context, cf. Pritchard 2002). Consequently, speakers would usually speak truly if they claimed knowledge in ordinary cases. But if they denied knowledge in high standard cases, they would say something strictly speaking false (but they might nevertheless implicate something true, in the Gricean sense). Of course the invariantist is also free to adopt high standards—if he is sympathetic to the skeptic. But either way, invariantism doesn't square well with our intuitions concerning the truth-value of knowledge attributions. If I deny knowing that P because stakes are very high or because I am attending to error possibilities that I take to be relevant but cannot rule out, I do say something strictly speaking true—at least that is what the above examples seem to show. The invariantist has to explain away these intuitions.
10. Here is an example due to Charles Travis that might help to illustrate the point: "Consider the sentence 'The ball is round', and two cases of its use. *Case A*: What shape do squash balls assume on rebound? Pia hits a decent stroke; Jones watches. 'The ball is round', she says at the crucial moment. Wrong. It has deformed into an ovoid. Jones did not say the ball to be as it was, so spoke falsely. *Case B*: Fiona has never seen squash played. From her present vantage point the balls seems a constant blur. 'What shape is that ball?', she asks. 'The ball is round', Alf replies; truly, since that it is the sort of ball a squash ball (and this one) is. It is not, e.g., like a very small rugby ball. So there are both true things and false things—thus a variety of different things—to be said of a given ball, and of the way it is at a given time, in the words 'The ball is round', used so as to have meant (as used) what they mean (in English)." (Travis 1996, 454). And here is a second example: "Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports: 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leave chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood." (Travis 1997, 89)
11. Or maybe there are simply different kinds of predicates: those that can truly be applied if applied in accordance with the standard that is in fact operative in the context (given that the participants are competent), and those that can truly be applied only if applied in accordance with the standard that should reasonably govern the context. The predicates ". . . is red" or ". . . is tall" might be of the first kind, the predicates ". . . is morally wrong" or ". . . is a good argument" might be of the second kind. What about "know"? That depends on whether there is a fact of the matter which standard is to reasonably govern a given knowledge attribution.
12. This might be taken to be a point in favor of a pragmatic construal of the context sensitivity involved. For one might argue that a substitution would be available only if a speaker would (implicitly or explicitly) say something about a standard or any such thing when making a knowledge attribution. Suppose speaker A says: "X knows that P." If he would thereby say something like "X knows that P relative to standard S," then something like the following should be a perfectly good report of his utterance: "A said that X knows that P relative to standard S." But since no such report is available, A didn't say anything about a standard.

## References

Cohen, S. 2000: "Contextualism and Skepticism", *Philosophical Issues* 10, 94–107.

DeRose, K. 1992: "Contextualism and Knowledge Attributions", *Philosophy and Phenomenological Research* 52, 913–929.

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

DeRose, K. 2000: "Now you know it, now you don't", *Proceedings of the Twentieth World Congress of Philosophy*, Bowling Green, Ohio, Volume 5.

DeRose, K. ms: "The Ordinary Language Basis for Contextualism and the New Invariantism", available online: http://pantheon.yale.edu/~kd47/OLP.htm.

Feldman, R. 2001: "Skeptical Problems, Contextualist Solutions", *Philosophical Studies* 103, 61–85.

Hawthorne, J. 2004: *Knowledge and Lotteries*, Oxford University Press, Oxford & New York.

Hofweber, T. 1999: "Contextualism and the Meaning-Intention Problem", in K. Korta, E. Sosa and J. Arrozola (eds.) *Cognition, Agency and Rationality*, Dorderecht, Kluwer.

Kompa, N. 2002: "The Context Sensitivity of Knowledge Attributions", *Grazer Philosophische Studien* 64, 1–18.

Lewis, D. 1996: "Elusive Knowledge", *Australasian Journal of Philosophy* 74, 549–67.

Pritchard, D. 2002: "Two Forms of Epistemological Contextualism", *Grazer Philosophische Studien* 64, 19–55.

Schiffer, S. 1992: "Belief Ascription", *The Journal of Philosophy* 89, 499–521.

Schiffer, S. 1996: "Contextualist Solutions to Scepticism", *Proceedings of the Aristotelian Society* 96, 317–333.

Stalnaker, R. 1998: "On the Representation of Context", *Journal of Logic, Language, and Information* 7, reprinted in Stalnaker 1999: *Context and Content*, New York & Oxford, 96–113, quotations are from Stalnaker 1999.

Stanley, J. ms: "On the Case for Contextualism", www.umass.edu.

Travis, C. 1996: "Meaning's Role in Truth", *Mind* 105, 451–466.

Travis, C. 1997: "Pragmatics", in (eds.) Bob Hale & Crispin Wright: *A Companion to the Philosophy of Language*, Blackwell, Oxford: 87–107.

Williamson, T. 2001: "*Comments on Michael Williams'* Contextualism, Externalism and Epistemic Standards", *Philosophical Studies* 103, 25–33.

Yourgrau, P. 1983: "Knowledge and Relevant Alternatives", *Synthese* 55, 175–190; reprinted in M. Williams (ed.) 1993: *Scepticism*, Dartmouth, Aldershot, 251–266, quotations are from Williams 1993.

Received: October 2004

*Nikola Kompa*
*Westfälische Wilhelms-Universität*
*Philosophisches Seminar*
*Domplatz 23*
*D-48143 Münster*
*Germany*
*e-mail: nkompa@uni-muenster.de*