

Keith DeRose, “Contextualism and Knowledge Attributions”

I. Contextualism: initial exposition

Consider the following cases.

Bank Case A. My wife and I are driving home on a 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 in this case 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 on Saturday. 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.”

Assume that in both cases the bank *will* be open on Saturday and that there is nothing unusual about either case that has not been included in my description of it. It seems to me that (1) when I claim to know that the bank will be open on Saturday in case A, I am saying something true. But it also seems that (2) I am saying something true in Case B when I concede that I *don’t* know that the bank

Keith DeRose, “Contextualism and Knowledge Attributions,” *Philosophy and Phenomenological Research* 52 (1992): 913–29.

will be open on Saturday. Yet I seem to be in no better position to know in Case A than in Case B. It is quite natural to say that (3) If I know that the bank will be open on Saturday in Case A, then I also know that it will be in Case B.

Is there any conflict here among (1), (2), and (3)? I hope not, because I want to investigate and defend a view according to which all three of them are true. Of course, it would be inconsistent to claim that (1) and (2) are true, and also hold that (4) If what I say in Case A in claiming to know that the bank will be open on Saturday is true, then what I say in Case B in conceding that I don't know that the bank will be open on Saturday is false. But there is a big difference between (3) and (4), and this difference is crucial to the view I want to investigate and defend.

We may, following Peter Unger, call the view I want to investigate a "contextual"¹ theory of knowledge attributions: it is a theory according to which the truth conditions of sentences of the form "S knows that p" or "S does not know that p" vary in certain ways according to the context in which the sentences are uttered.² The contextualist can deny (4) even while admitting that I am in no better position to know in Case A than in Case B. The contexts of my utterances in the two cases make it easier for a knowledge attribution to be true in Case A than in Case B.

There are important contextual differences between Case A and Case B which one might think are relevant. First, there is the importance of being right. In Case B, a lot hinges on whether or not the bank will be open on Saturday, while in Case A it is not nearly as important that I be right. One might think that requirements for making a knowledge attribution true go up as the stakes go up.³

Second, there is the *mentioning* of a possibility. In Case B my wife raises the possibility that the bank may have changed its hours in the last two weeks. One might think that if this possibility has been mentioned, I cannot truly claim to know that the bank will be open on Saturday on the ground that two weeks ago it was open on Saturday unless I can rule out the possibility that the bank's hours have changed since then. On the other hand, perhaps I don't have to be able to rule out this possibility in order to truly say I know if, as in Case A, no such possibility has not been suggested.⁴

Third, there is the *consideration* of a possibility. Since my wife raised the possibility of the bank changing its hours in Case B, I have that possibility in mind when I utter my sentence. Perhaps, since I am considering this possibility, I must be able to rule it out in order to truthfully claim to know that the bank will be open on Saturday. On the other hand, in Case A I am not considering the possibility, so perhaps I do not have to be able to rule it out in order to truthfully say that I know that the bank will be open on Saturday.⁵ (Of course, it must still be *true* that the bank will be open on Saturday in order for me to know that it will be.)

Again following Unger, we may call someone who denies that the types of contextual factors we have just looked at affect the truth conditions of knowledge attributions an "invariantist." According to the invariantist, such features of an utterance of a knowledge attribution do not affect how good an epistemic position the putative knower must be in for the attribution to be true. In considering the Bank Cases, for instance, the invariantist will assert (4), which seems very plausible, and will therefore deny either (1) or (2). Typically, the invariantist will deny (1). In fact, Unger uses the term "invariantism" to denote the position that

the standards for true knowledge attributions remain constant *and very high*—as high as they can possibly be. This position I will call “sceptical invariantism,” leaving the more general term “invariantism” to denote any position according to which the truth conditions for knowledge attribution do not vary in the way the contextualist claims they do, whether or not the standards are said to be very high. I will then use “non-sceptical invariantism” to refer to a position according to which the standards are held to be constant but relatively low.⁶ The *sceptical* invariantist will deny (1). She may admit that I am *warranted in asserting* that I know in Case A or that it is *useful for me to say* that I know, but will insist that what I say in claiming that I know is, strictly speaking, false. On the other hand, similar maneuvers can be used by the *non-sceptical* invariantist to deny (2). A non-sceptical invariantist may admit that I *should not say* that I know in Case B, because my wife mistakenly thinks that I must be able to rule out the possibility that the bank has changed its hours in order to know that the bank will be open on Saturday, and saying that I know will lead her to believe that I can rule out that possibility. Still, my wife *is* mistaken about this requirement, and if I were to say that I knew, I would be saying something that is, though misleading, true. Thus, it is *useful for me to assert* that I *don’t* know. But for all its usefulness, my assertion is, strictly speaking, false.

Contextualists, of course, can disagree about what types of features of the context of utterance really do affect the truth conditions of knowledge attributions and to what extent they do so. I will not here enter into this thorny issue, although I have a preference for the more “objective” features—like the importance of being right and what has been said in the conversation—and tend to discount as relevant to truth conditions such “subjective” features as what possibilities the speaker is considering.⁷ In this paper I address some *general* issues that confront any contextualist. In Part II, I distinguish between contextualism and a very prominent theory of knowledge which has been called the “relevant alternatives” theory (RA), and in Part III, I respond to an important objection to which *any* form of contextualism seems vulnerable.

By thus isolating and defending contextualism, I will do much to clear the way for contextualist resolutions to sceptical arguments. Contextual theories of knowledge attributions have almost invariably been developed with an eye towards providing some kind of answer to philosophical scepticism. For some sceptical arguments threaten to show, not only that we fail to meet very high requirements for knowledge of interest to philosophers seeking absolute certainty, but also that we don’t meet the truth conditions of ordinary, out-on-the-street claims to know. They thus threaten to establish the startling result that we never, or almost never, truly ascribe knowledge to ourselves or to other human beings. According to contextual analysis, when the sceptic presents her arguments, she manipulates various conversational mechanisms that raise the semantic standards for knowledge, and thereby creates a context in which she can truly say that we know nothing or very little. But the fact that the sceptic can thus install very high standards which we don’t live up to has no tendency to show that we don’t satisfy the more relaxed standards that are in place in ordinary conversations. Thus, it is hoped, our ordinary claims to know will be safeguarded from the apparently powerful attacks of the sceptic, while, at the same time, the persuasiveness of the sceptical arguments is explained.⁸

Many find such contextualist resolutions of sceptical arguments very attractive, especially since their main competition is the sceptical invariantist resolutions according to which the persuasiveness of various sceptical arguments is explained in a way as alarming as it is simple: They seem persuasive because they are indeed sound and successfully establish the startling conclusion that we never or almost never truly ascribe knowledge.⁹ But many, while finding the contextualist resolutions a preferable alternative to an unacceptably radical form of scepticism, at the same time feel an initial resistance, closely related to the appeal of (4), to the thought that contextual factors of the types I've mentioned can really affect whether or not a subject knows.¹⁰ While many are willing to accept this thought in order to avoid the sceptical conclusion, there remains a feeling that the contextualist is asking them to swallow pretty hard—although perhaps not quite so hard as the sceptical invariantist would have them swallow. As contextualists have rushed to apply their theories to the problem of scepticism, this initial resistance has not yet been adequately addressed. I will address this resistance, as well as some explicit objections to contextualism that have been raised in the philosophical literature and which are based on the source of this resistance, in Part III below. But first, in Part II, we must carefully distinguish contextualism from RA.

II. Contextualism and ‘relevant alternatives’

The most popular form of contextualism, I think it is fair to say, is what has been called the “relevant alternatives” view of knowledge (RA). But we must be careful here. As we shall see, it is a bit tricky to say just in what sense RA is a contextualist view. According to RA, a claim to know that *p* is made within a certain framework of relevant alternatives which are incompatible with *p*. To know that *p* is to be able to distinguish *p* from these relevant alternatives, to be able to rule out these relevant alternatives to *p*. But not every contrary of or alternative to *p* is a *relevant* alternative.¹¹ In an ordinary case of, say, claiming to know that some animals in a zoo are zebras, to borrow an example introduced by Fred Dretske,¹² the alternative that they are cleverly painted mules is *not* a relevant alternative, and one need not be able to rule it out in order truly to claim to know that the animals are zebras. But in an extraordinary case, that alternative might be relevant. How can it become relevant?

In one of the standard presentations of RA, Alvin Goldman (1976) presents various factors which can affect the range of relevant alternatives. These factors may be divided into two groups. First, there are features of the putative knower's situation; these I will call “subject factors.”¹³ A subject in an ordinary situation can be truly said to know that what he sees up ahead is a barn even if he cannot rule out the possibility that it is just a barn facade. But, Goldman points out, if there are a lot of such facades in the putative knower's vicinity, then the possibility that what the person is seeing is just a facade *is* a relevant alternative, and the person does not know that he is seeing a barn, even if what he sees happens to be an actual barn (pp. 772–73).

Second, there are features of the speaker's situation, which I will call “attributor factors.” Goldman writes, “It is not only the circumstances of the putative knower's situation, however, that influence the choice of alternatives. The

speaker’s own linguistic and psychological context are also important.” Goldman suggests that “if the speaker is in a class where Descartes’s evil demon has just been discussed,” then certain alternatives may be relevant which ordinarily are not (p. 776).

Insofar as a relevant alternatives theorist allows attributor factors to influence which alternatives are relevant, he is a contextualist. An invariantist can be a relevant alternatives theorist if he allows only subject factors to influence which alternatives are relevant.¹⁴ Consider two situations in which Henry has a good, clear look at what he takes to be—and what, in fact, is—a barn. In Case C there are no barn facades around, but in Case D the area Henry finds himself in is (unbeknownst to him) teeming with barn facades, although Henry is luckily looking at the only actual barn in the area. This does not seem to be a pair of cases in which Henry is in equally good positions to know that what he is seeing is a barn; the conditional, *If Henry knows in Case C, then he knows in Case D* does not seem to be true, so the invariantist can agree that a sentence attributing knowledge to Henry in Case C can be true, while one attributing knowledge to him in Case D is false. And he can use the idea of “relevant alternatives” to explain the difference. Thus, although most versions of RA allow attributor factors to be relevant and are therefore contextualist views, an RA theorist need not be a contextualist.

Of course, in first-person present tense knowledge claims, the attributor of knowledge and the putative subject of knowledge are in the same situation (they are the same person at the same time). If Henry says, “I know that that’s a barn,” there is no difference between the speaker and the putative knower. In this situation the invariantist RA theorist will allow only factors that attach to Henry qua putative knower (e.g. the presence or lack of facades in his vicinity) to matter in evaluating his claim for truth, while the contextualist will also allow factors that attach to Henry qua attributor of knowledge (such as whether or not the issue of facades has been raised in the conversation) to matter.¹⁵

Although Goldman draws the distinction between what I am calling subject factors and attributor factors, he does not explain the importance of this distinction. I am stressing it because it is crucial to some of the important claims RA theorists have wanted to make about the *meanings* of knowledge attributions.¹⁶ Gail Stine, for example, writes:

In Dretske’s zoo example, the animal’s being a mule painted to look like a zebra is not a relevant alternative. So what one means when one says that John knows the animal is a zebra, is that he knows it is a zebra, as opposed to a gazelle, an antelope, or other animals one would normally expect to find in a zoo. If, however, being a mule painted to look like a zebra became a relevant alternative, then one would literally mean something different in saying that John knows that the animal is a zebra from what one meant originally and that something else may well be false. (Stine (1976), p. 255)

But here we must be very careful. Much depends on *how* the animal’s being a painted mule has become a relevant alternative. Suppose that it has become a relevant alternative due to a change in subject factors: There has been a zebra shortage and many zoos (even reputable zoos) *have* been using painted mules in

an attempt to fool the zoo-going public. This could come about without the speaker's knowing it. Would one then *mean* something different by saying that John knows that the animal is a zebra? I think not.

The meaning of "meaning," of course, is difficult to get hold of. But there seems to be a fairly straightforward and important sense in which one *does* mean something different if the range of relevant alternatives has been changed by attributor factors but does *not* mean something different if the range of relevant alternatives has been changed only by subject factors. Stewart Cohen, whose version of RA clearly is a contextualist one, writes that he

construes "knowledge" as an indexical. As such, one speaker may attribute knowledge to a subject while another speaker denies knowledge to that same subject, without contradiction. (Cohen (1988), p. 97)

This lack of contradiction is the key to the sense in which the knowledge attributor and the knowledge denier mean something different by "know." It is similar to the sense in which two people who think they are in the same room but are in fact in different rooms and are talking to each other over an intercom mean something different by "this room" when one claims, "Frank is not in this room" and the other insists, "Frank is in this room—I can see him!" There is an important sense in which both do mean the same thing by "this room," in which they are using the phrase in the same sense. But there is also an important sense in which they do not mean the same thing by the phrase; this is the sense by which we can explain the lack of contradiction between what the two people are saying. To use David Kaplan's terminology, the phrase is being used with the same *character*, but with different *content*.¹⁷ Similarly, in Bank Case B from Part I of this paper, when, in the face of my wife's doubt, I admit that I don't know that the bank will be open on Saturday, I don't contradict an earlier claim to know that I might have made before the doubt was raised and before the issue was so important because, in an important sense, I don't mean the same thing by "know" as I meant in the earlier claim: While "know" is being used with the same *character*, it is *not* being used with the same *content*. Or so the contextualist will claim.

But if the range of relevant alternatives is changed by subject factors, the meaning of "know" is not in the same way changed. If very many nearby banks *have* discontinued their Saturday hours in the last two weeks, then it seems that my original claim to know may well have been false, and if I admit that I did not know after this surprising fact about local banks is called to my attention, I will be taking back and contradicting my earlier claim to have known.

Recall the two cases in which Henry has a good, clear look at what he takes to be a barn. (In Case C, there are no barn facades around, but in Case D, the fields are filled with barn facades, but Henry is luckily looking at the only actual barn in the area.) In each case, insert two people in the back seat of the car Henry is driving, and have the first say to the second, "Henry knows that that is a barn." It seems that, in the sense under discussion, what the first person *means* by "knows" in each of the two cases is the same. In Case C what she is saying is true, while in Case D it is false. The presence of the barn facades has changed the *truth value*, but not the *truth conditions* or the meaning (content), of the first person's knowledge attribution.

“CONTEXTUALISM AND KNOWLEDGE ATTRIBUTIONS”

So attributor factors affect the truth values of knowledge attributions *in a different way* than do subject factors: attributor factors working in such a way that they affect the content of the attribution, but subject factors working in a different way that does not affect its content. These different ways can be explained as follows. Attributor factors set a certain standard the putative subject of knowledge must live up to in order to make the knowledge attribution true: They affect *how good an epistemic position the putative knower must be in to count as knowing*. They thereby affect the truth conditions and the content or meaning of the attribution. Subject factors, on the other hand, determine whether or not the putative subject lives up to the standards that have been set, and thereby can affect the truth value of the attribution *without* affecting its content: They affect *how good an epistemic position the putative knower actually is in*.¹⁸

To make use of the character/content distinction, the “character” of “S knows that p” is, roughly, that S has a true belief that p and is in a *good enough* epistemic position with respect to p; this remains constant from attribution to attribution. But how good is good enough? This is what varies with context. What the context fixes in determining the “content” of a knowledge attribution is how good an epistemic position S must be in to count as knowing that p. The mentioning of alternatives like painted mules, or barn facades, or changes in banking hours, when there is no special reason for thinking such possibilities likely, can be seen as raising the strength and changing the content of “know” because the ability to rule out such alternatives would only be relevant if one were after a strong form of knowledge (if one were requiring the putative knower to be in a very good position in order to count as knowing).

Subject factors, then, are best construed, not as affecting the truth conditions of knowledge attributions, but rather as affecting whether those truth conditions are satisfied. This fact severely limits RA’s prospects for explaining variations in the content of knowledge attributions. RA, for all I’ve said, may be a helpful tool for determining or explaining why certain attributions of knowledge have the *truth values* they have.¹⁹ Note, however, that for RA to be successful in this capacity, it *must* allow subject factors to affect the range of relevant alternatives, for, as Goldman’s barn cases (cases C and D) clearly show and as is evident in any case, subject factors can affect these truth values.

But RA theorists have wanted to make claims about the *meaning* of knowledge attributions²⁰. Many of them have thought that the meaning of knowledge attributions changes from case to case depending upon various factors, and they have thought that this change in meaning *amounts to* a change in the range of alternatives that are relevant.²¹ But we can now see that the content of a given knowledge attribution cannot be specified by citing what the range of relevant alternatives is, because that range is a function of subject factors (which do not affect the content of the attribution) as well as attributor factors (which do). There can be a drastic change in the range of relevant alternatives from one attribution to another without there being any change in meaning between the two attributions, then, because the change in the range of relevant alternatives can, and often will, be the result of differences in subject factors, which will not have any affect on the meaning of the attribution.²²

III. The objection to contextualism

Having distinguished contextualism from RA, I will now seek to defend contextualism from a certain type of important objection. The obvious attraction of contextualism, besides (and closely related to) the resolution of sceptical arguments it purportedly provides, is that it seems to have the result that very many of the knowledge attributions and denials uttered by speakers of English are true—more than any form of invariantism can allow for, and certainly more than sceptical invariantism can allow for. Thus, recalling the Bank Cases, contextualism allows us to assert both (1) and (2) and many of us will find both (1) and (2) compelling. Unfortunately, contextualism seems to be vulnerable to a certain type of powerful objection which is closely related to the appeal of (4). Suppose, to recall an example we've already considered, that two people see some zebras in a zoo. Palle Yourgrau constructs the following conversation, and claims that “something is amiss” in it:

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.²³

This absurd dialogue is aimed at contextualists who think that the mentioning of a possibility incompatible with what one claims to know is enough to require that one rule the possibility out before one can truly claim to know. But this type of attack can work against other contextualists, also. Dialogues much like the above dialogue but with the following last lines seem equally absurd:

B': No, I did know then that it was a zebra. But now that it has become so important that it be a zebra, I no longer know.

B'': No, I did know then that it was a zebra. But now that the possibility of its being a painted mule has occurred to me, I no longer know.

The general point of the objection is that whether we know something or not cannot depend on, to use Peter Unger's words, “the contextual interests of those happening to use the terms on a particular occasion” (Unger (1984), p. 37).

How shall the contextualist respond? The objection as I have put it forward, though it explains much of the initial resistance many feel toward contextualism, is based on a mistake. The contextualist believes that certain aspects of the context of an attribution or denial of knowledge attribution affect its content. Knowledge claims, then, can be compared to other sentences containing other context-sensitive words, like “here.” One hour ago, I was in my office. Suppose I truly said, “I am here.” 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. The meaning of “here” is fixed by the relevant contextual

“CONTEXTUALISM AND KNOWLEDGE ATTRIBUTIONS”

factors (in this case, my location) *of the utterance*, not by my location at the time being talked about.

Similarly, the contextualist may admit that the mentioning of the painted mules possibility affects the conditions under which one can truthfully say that one knows an animal to be a zebra: one now must be able to rule out that possibility, perhaps. But the contextualist need not, and should not, countenance the above dialogue. If in the context of the conversation the possibility of painted mules has been mentioned, and *if* the mere mention of this possibility has an effect on the conditions under which someone can be truly said to “know,” then any use of “know” (or its past tense) is so affected, even a use in which one describes one’s past condition. B cannot truly say, “I did know then that it was a zebra”; that would be like my saying, “I was here.” B *can* say, “My previous knowledge claim was true,” just as I can say, “My previous location claim was true.” Or so I believe. But saying these things would have a point only if one were interested in the truth-value of the earlier claim, rather than in the question of whether in the *present* contextually determined sense one knew and knows, or didn’t and doesn’t.

Yourgrau writes of the zebra case, “Typically, when someone poses a question regarding whether we really know that P obtains rather than some alternative to P, if we cannot satisfactorily answer the question, we conclude that our earlier claim to know was faulty” (p. 183). But do we? We do not stubbornly repeat ourselves, to be sure: “Still, I know that it is a zebra!” We *might* even say, “I don’t know” or “I didn’t know.” All of this the contextualist can handle. But do we (or should we) admit that our *earlier* claim was *false*? I am on the witness stand being questioned.

Lawyer: Were there any zebras in the zoo on April 23?

Me: Yes.

L: Do you know that?

M: Yes.

L: How do you know?

M: I saw some there.

L: So, you knew that they were zebras?

M: Yes.

L: Could you rule out the possibility that they were only cleverly painted mules?

M: No, I suppose not.

L: So, did you really know that they were zebras?

M: Is there any reason to think that they were painted mules, of all things?

L: Just answer the question!

Well, how should I answer the question? If there is no special reason to think they were painted mules then *I* certainly wouldn’t want to admit that I didn’t know they were zebras, but maybe I’m just being stubborn. Suppose I do admit it:

M: I guess I didn’t *know* that they were zebras.

L: Aha! The witness has contradicted his earlier claim. First he says

that he knew; now he says he didn't. Now which is it, Mr. DeRose?

Surely something is amiss in *this* dialogue; my lawyer should object. I haven't contradicted my earlier claim, as much as it looks as if I have. It would be as if the following had occurred. While standing in a bright yellow room, I said, "This room is yellow." The lawyer then dragged me by the ear into a room in which all was grey and got me to say, "This room is grey," and now he is jumping all over me: "First he says. 'This room is yellow,' then he says, 'This room is grey.' Which is it?" The contextualist maintains that something very much like this has happened in my original dialogue with the lawyer. Of course, there is room for the invariantist to deny this contextualist claim. But it is *far* from clear that in cases like the one Yourgrau brings to our attention, we should admit that our earlier claim was false or that our later claim contradicts it.

So, the objection that whether we know something or not does not depend on contextual factors of the type we have been considering is based on a mistake. But Unger does not make this mistake when he raises an objection similar to the one we have been considering.²⁴ He writes of "our belief that the semantics of these expressions ["know" is one of the expressions being considered] is appropriately *independent*, that the conditions do not depend on the contextual interests of these happening to use the terms on a particular occasion" (Unger (1984), p. 37). Insofar as we do have *this* belief, that the conditions *for truly saying* that someone knows do not depend on the sorts of contextual factors we have been discussing, then contextualism goes against at least one of our beliefs. But it seems that much of the appeal of this belief derives from the plausibility of the thesis (with which the contextualist can agree) that whether *we know* something or not does not depend on such factors. The answer to the question, "Does she know?", in whatever context it is asked, including a philosophy paper, is determined by facts independent of contextual factors (or what I have been calling attributor factors). These contextual or attributor factors affect the content of the question, but once the question is asked with a specific content, its answer is determined by subject factors, which are precisely the kinds of factors which *can* very plausibly be thought to affect whether or not the subject knows. Going back to our opening examples, the contextualist can affirm (3) *in any context in which it is uttered*: If I know in Case A, then I know in Case B. Of course, the contextualist must deny (4), and (4) sounds very plausible, but much of the appeal of (4) comes from the plausibility of (3). And since we *must* give up either (1), (2), or (4), those who, like me, find (1) and (2) *very* plausible will be well-motivated to give up (4), especially since (3) can still be affirmed.

In general, then, when it looks as if the contextualist has to say something strongly counter-intuitive, what he must say turns out to be, on the contrary, something fairly theoretical concerning the truth conditions of certain sentences. Do we really have strong intuitions about such things? At any rate, the contextualist can go along with the simple facts that we all recognize: that if I know in Case A, then I know in Case B, and that whether we know something or not does not typically depend on our current interests or on other such contextual factors.²⁵

Notes

- 1 I take the terms “contextualism” and “invariantism” from Unger (1984).
- 2 The importance of this theory will not be confined to knowledge attributions. For instance, in DeRose (1991) I argue that S’s assertion, “It is possible that P,” where the embedded P is in the indicative mood, is true if and only if (1) no member of the relevant community *knows* that P is false and (2) there is no relevant way by which members of the relevant community can come to *know* that P is false. As I there argue, there is a great deal of flexibility in the matter of who is and is not to be counted as a member of the relevant community and what is and is not to be counted as a relevant way of coming to know: That these matters are determined by aspects of the contexts in which the statement is made. If, as I am here defending, there is a contextually-determined variation on how good an epistemic position one must be in to count as knowing, then—since epistemic possibilities have entirely to do with what is and is not known and what can and cannot come to be known in certain ways—this variation will affect the content of epistemic modal statements as well: As the standards for knowledge go up, and it becomes harder and harder for a knowledge attribution to be true, it will become easier and easier for an assertion of epistemic possibility to be true.
- 3 That the importance of being right is an important contextual factor is suggested in Austin (1961), p. 76, fn. 1. Dretske denies the importance of this factor in (1981a), pp. 375–76.
- 4 David Lewis (1979) stresses this contextual factor, presenting an interesting account of how the mentioning of sceptical possibilities can affect the range of relevant alternatives by means of what he calls a “rule of accommodation.” In Chapter 3 (see especially section I) of DeRose (1990). I argue that Lewis’s account is not complete, and I locate an independent mechanism of standard changing which, I now believe, is *at least* as important (and probably considerably more important) to the application of contextualism to the problem of scepticism as is the mechanism Lewis has located.
- 5 Alvin I. Goldman (1976) stresses the importance of what possibilities the speaker is considering.
- 6 While Unger does not even consider the view that the standards for true knowledge attributions don’t change but are held constant at a fairly low level, non-sceptical invariantism is defended (at least conditionally) by Robert Hambourger (1987). Hambourger argues that *if* the standards are constant (Hambourger does not believe that this antecedent is true), then they must be fairly low (pp. 256–57). In the terminology I have introduced, Hambourger is arguing that if some form of invariantism is correct, it must be a form of *non-sceptical* invariantism.
- 7 My main reason for discounting as relevant to truth conditions the matter of what the speaker is thinking, at least with respect to spoken interactions between people, is that I don’t think that one should be able, merely by a private act of one’s own thought to drastically “strengthen” the content of “know” in such a way that one can truly say to someone who is quite certain that he is wearing pants, “You don’t know you’re wearing pants,” without there having been anything in the conversation to indicate that the strength of “know” has been raised.

There *might* yet be a fairly tight connection between what raises the truth condition standards and what speakers *tend to* think or perhaps what they *should* think of the standards as being. Perhaps the truth condition standards are what a typical speaker would take them to be or should take them to be, given what has gone on in the conversation. But it seems unfair to one's interlocutor for the truth condition standards of a *public, spoken* knowledge attribution to be changed by an idiosyncratic, private decision. It is far more plausible to suppose that when one is *thinking to one's self* about what is or is not "known," the content of "know" is directly tied to the strength the thinker intends.

- 8 While, as I've said, contextualist theories (including contextualist versions of RA) are almost invariably developed with an eye towards philosophical scepticism, the most thoroughly worked out contextualist attempts to resolve the problem of scepticism that I am aware of are to be found in Unger (1986), Cohen (1988) (see also Cohen (1987)), and DeRose (1990), especially Chapter 3. Fred Dretske has also applied this type of theory of knowledge to the problem of scepticism in several places. See Dretske (1970), (1971), (1981a), and (1981b).
- 9 See Unger (1975).
- 10 A typical objection one meets in presenting contextualism, as I know from personal experience, is: "How can *our* context have anything to do with whether or not *Henry* knows?", where Henry is a character in an example and so is not present in the room.
- 11 See Goldman (1976), p. 772 Stine (1976), p. 249; and Dretske (1970), p. 1022.
- 12 See Dretske (1970), pp. 1015–1016.
- 13 Please note that by "subject factors" I do *not* mean *subjective* (as opposed to objective) factors. I rather mean factors having to do with the putative *subject* of knowledge and her surroundings (as opposed to the attributor of knowledge).
- 14 Thus, what Goldman calls the "first view" of RA, according to which "a complete specification" of the putative knower's situation determines "a unique set of relevant alternatives" (pp. 775–76), is an invariantist version of RA. Goldman does not endorse this view; he says he is "attracted by the second view" (p. 777), which clearly is a contextualist version of RA.
- 15 Some factors, I believe, will both affect how good an epistemic position the speaker/putative knower is in *and* (at least according to the contextualist) how good a position he must be in to make his knowledge claims true. Thus, they will be both subject and attributor factors.
- 16 I further discuss the importance of this distinction between subject factors and attributor factors and the resulting contextualist view according to which content varies in response to attributor factors in Chapter 1 of DeRose (1990). In particular, I there discuss, in addition to the issues treated in the present paper, the advantages such a view according to which content varies over a *range* has over theories like that put forward in Malcolm (1952) according to which there are two distinct senses of "know", a strong sense and a weak one.
- 17 See Kaplan (1989), esp. pp. 500–507.
- 18 Unger makes a similar division in (1986), where he distinguishes between the "profile of the context," which corresponds roughly to how good a position the putative knower must be in to count as knowing, and the "profile of the facts,"

“CONTEXTUALISM AND KNOWLEDGE ATTRIBUTIONS”

which corresponds roughly to how good a position the putative knower actually is in (see esp. pp. 139–40). Unger does not there discuss RA, and so does not use the distinction to distinguish contextualism from RA. He does, however, introduce an important complication which I have ignored in this paper, since it has little effect on the points I’m making here. Unger points out that there are many different aspects of knowledge and that in different contexts, we may have different demands regarding various of these aspects. Thus, for example, in one context, we may demand a very high degree of confidence on the subject’s part before we will count him as knowing while demanding *relatively* little in the way of his belief being non-accidentally true. In a different context, on the other hand, we may have very stringent standards for non-accidentality but relatively lax standards for subject confidence. As Unger points out, then, things are not quite as simple as I make them out to be: Our standards are not just a matter of how good an epistemic position the subject must be in, but rather of how good in which respects. Stewart Cohen also suggests a related division, his more closely aligned with the spirit of RA. See note 22 below.

- 19 Thus what I take to be RA’s *basic idea*—that to know that P, one must be able to rule out all of the relevant alternatives to P—may be sound.
- 20 RA’s *basic idea* (see note 19, above) is not about contextual *variations in meanings*. In deed, as I’ve pointed out, an RA theorist can be an invariantist. It is, then, in going *beyond* this basic idea that RA theorists have, by my lights, gone wrong by tying the meaning of a given attribution too closely to what the range of relevant alternatives is.
- 21 In addition to the Stine passage we have looked at, see, for example, Goldman (1976), pp. 775–77 (esp. p. 777), where Goldman seems to think that *what proposition is expressed* by a given knowledge attribution is specified by what the range of relevant alternatives is. Something similar seems to be suggested in Lewis (1979), esp. pp. 354–55. Lewis seems to think of the “conversational score” of a given context, with respect to knowledge attributions and epistemic modal statements, to be something that can be specified by giving the range of possibilities that are relevant in that context.
- 22 A different view which escapes this problem but is still well within the spirit of RA is that the character of “S knows that p” is that S has a true belief that p and can rule out all alternatives to p that are *sufficiently probable*. The context of utterance can then be seen as fixing the content by determining just how probable an alternative must be to count as being sufficiently probable. Something like this alternative view is suggested by Cohen (1988), according to whom context determines “how probable an alternative must be in order to be relevant” (p. 96). (This view is only *suggested* by Cohen because he never says that this probability level for alternative relevance is *all* that context fixes in determining the content of an attribution.) Expanding this idea, we might then take aspects of the putative knower’s situation to affect how *probable a given alternative is*. Instead of the meaning being specified by the range of alternatives that are relevant, this view, more plausibly, has it specified by the standards (in terms of probability) alternatives must meet to count as relevant. This still seems more precise than my admittedly vague talk of *how good an epistemic position* one must be in to count as knowing. I fear, however, that this precisification will not work. Among other reasons for doubting that the notion of probability can do all

the work assigned to it here is this: The complication Unger raises about the many different aspects of knowledge (see note 18 above) shows that no single measure like the probability an alternative must have to be relevant can capture all that context does in fixing the content of a knowledge attribution. This probability standard of alternative relevance can be, at best, one among several aspects of knowledge the standards for which are fixed by context.

- 23 Yourgrau (1983), p. 183. The absurdity of such a conversation, along with the worry that it causes problems for theories of knowledge attributions like the one I am investigating, was originally suggested to me by Rogers Albritton, who has been making such suggestions since well before Yourgrau's article came out.
- 24 Actually, Unger does make this mistake at one point, not about knowledge but about flatness. Throughout his epistemological writings, Unger compares knowledge attributions with claims about the flatness of objects. In (1984). Unger describes an invariantist semantics for "flat" according to which an object must be as flat as possible in order for a sentence like "That is flat" to be true of it, and a contextualist semantics for "flat" according to which how flat something must be in order for a sentence like. "That is flat" to be true of it varies with context, and he claims that there is no determinate fact as to which semantics is correct. In attacking the contextualist semantics for "flat," Unger writes: "How can the matter of whether a given surface is *flat*, in contradistinction to, say, whether it is suitable for our croquet game, depend upon the interests in that surface taken by those who happen to converse about it? This appears to go against our better judgement" ((1984), p. 39). But the contextualist need not and should not claim that "the matter of whether or not a given surface is flat" depends "upon the interests in that surface taken by those who happen to converse about it," although the contextualist *will* say that the truth conditions *for the sentence* "That is flat" do depend upon such contextual interests. I believe that the above passage is just a slip on Unger's part; he is usually more careful in making his attack on contextualism. But it is revealing that Unger makes this slip: It shows how easy it is to confuse the claim (a) that whether or not something is flat or is known does not depend on contextual interests with the claim (b) that the truth conditions for a sentence about flatness or about knowledge do not depend on contextual interests, which does not follow from (a).
- 25 I am indebted to Robert M. Adams, Rogers Albritton, Peter Unger, and an anonymous referee for *Philosophy and Phenomenological Research* for comments on previous drafts of this paper.

References

- Austin, J. L., (1961) "Other Minds," in J. L. Austin, *Philosophical Papers* (Oxford University Press): 44–84.
- Cohen, Stewart, (1987) "Knowledge, Context, and Social Standards," *Synthese* 73: 3–26.
- , (1988) "How to be a Fallibilist," *Philosophical Perspectives* 2: 91–123.
- DeRose, Keith, (1990) "Knowledge, Epistemic Possibility, and Scepticism," UCLA doctoral dissertation.
- , (1991) "Epistemic Possibilities," *The Philosophical Review* 100: 581–605.
- Dretske, Fred, (1970) "Epistemic Operators," *Journal of Philosophy* 67: 1007–1023.