

CONTESTING CONTEXTUALISM

Patrick RYSIEW
University of British Columbia, Canada

Summary

According to Keith DeRose, the invariantist's attempt to account for the data which inspire contextualism fares no better, in the end, than the "desperate and lame" maneuvers of "the crazed theory of 'bachelor'", whereby *S*'s being unmarried is not among the truth conditions of '*S* is a bachelor', but merely an implicature generated by an assertion thereof. Here, I outline the invariantist account I have previously proposed. I then argue that the prospects for sophisticated invariantism—either as a general approach, or in the specific form I have recommended—are not nearly as dim as DeRose suggests.

1.

No one denies that whether you know that *p* depends on many features of you and your situation. Of course, people disagree over which such features are important—whether knowing depends on what else you believe, on objective features of your situation over and above whether *p* is true, on certain counterfactuals holding of you, on your evidence for (/against) *p*, on the etiology of your belief, and so on. The natural thought, though, is that once such factors are fixed, it's either true that one knows or it isn't—changes in such things as what the speaker has in mind in attributing knowledge to you don't affect the truth conditions of what he says, and so don't affect whether he speaks truly in saying that you know. But contextualists deny precisely this. In Keith DeRose's words, "the truth conditions of sentences of the form '*S* knows that *p*' or '*S* does not know that *p*' vary in interesting ways depending on the context in which they are uttered" (DeRose 1992, 914), and 'context' here means none other than such things as the interests, expectations, and so forth of knowledge attributors (e.g., DeRose 1999, 189–190, Cohen 1999, 57).

Now, various theorists have argued that contextualism faces certain theoretical difficulties;¹ others have claimed that the data that contextualists cite in supporting their view—including, and especially, the fact that there is a manifest flexibility in our willingness to attribute knowledge—are in fact compatible with thinking that the truth conditions of knowledge sentences don't shift with changes in context. As to the latter claim, however, the 'invariantist' needs to do more than simply “‘chalk it all up to pragmatics’—to claim that the varying epistemic standards that govern our use of ‘knows’ are only varying standards for when it would be appropriate to say we know” (DeRose 2002, 194), and just leave it at that. Granted: since our willingness to utter a sentence, and our intuitions as to whether it's true, may be due as much to pragmatic as to semantic factors, any direct inference from our verbal behavior to semantic claims would be a *non sequitur*. But by the same token, simply to say that the context-sensitivity of knowledge attributions is explained by pragmatics is not to show how that's so.

But while (nonsceptical²) invariantists have tended to rest content with “a ‘bare warranted assertability maneuver’” (DeRose 1999, 199; cf. 2002, 176), this is no longer the case. Elsewhere (Rysiew 2001), for example, I have presented an account of how pragmatic phenomena of a familiar Gricean sort, together with an invariantistic conception of the truth conditions of knowledge sentences, can be used to account for the flexibility in our willingness to attribute knowledge; and others have recently offered such accounts as well.³ Here, in Section 2, I briefly outline the account I have proposed. In Section 3 I argue, following Brown (*forthcoming*), that the reasons DeRose offers for being sceptical about the prospects of sophisticated invariantism as a general approach are open to question. In Section 4 I argue that DeRose's reasons for finding the specific form of sophisticated invariantism I

1. For example, many complain that contextualism does not really help to address scepticism (see Feldman 1999 and 2001, Klein 2000, Kornblith 2000, and Sosa 2000); Hofweber (1999), Rysiew (2001), and Schiffer (1996) argue that the contextualist's error theory—our getting mixed up about contexts—is in one way or another problematic; Stanley (2004) argues that ‘knows’ doesn't behave like context-sensitive terms with which we're familiar; and Cappelen and Lepore (2003) argue that contextualists rely on illegitimate context shifting arguments.

2. In the view of many, nonsceptical invariantism is the only variety of invariantism worth wanting. (Cf. Cohen 1999, 83 and DeRose 1999, 202.)

3. These include Bach (*forthcoming*), Brown (*forthcoming*), and Pritchard (*forthcoming*). (Importantly, Bach's proposal does not rely on any form of ‘WAM’.)

have offered unsatisfactory are not decisive. If these arguments are on the right track, then sophisticated invariantism—whether as a general approach, or along the lines of the account I have proposed—remains a viable alternative to contextualism.

2.

The key idea of sophisticated invariantism, and what separates it from invariantism *simpliciter*, is that it attempts to explain the data which inspire contextualism by appealing to a distinction between the literal contents of knowledge sentences and the information generated by (and inferable from) specific utterances thereof. How one states the particular sophisticated invariantist view I favor depends on how neutral one wishes to remain about the truth conditions of knowledge sentences. In neutral terms, the basic idea is this: the proposition literally expressed by the sentence ‘*S* knows that *p*’ is (unsurprisingly) that *S* knows that *p*; and, minimally, that *S* knows that *p* entails that *S* believes that *p*, *p*, and that *S* is in a good epistemic position with respect to *p*. It is a familiar idea, though, that what a given sentence means need not be what the speaker means in uttering it—i.e., what he intends to communicate. To capture this, I say that an utterance of ‘*S* knows that *p*’ pragmatically imparts (or conveys) a proposition of the form, ‘*S*’s epistemic position with respect to *p* is good enough ...’, where the ellipsis is completed according to the interests, purposes, assumptions, etc., of the speaker. So, for instance, if I say, “I know who’s leading the NBA in scoring”, what I say—i.e., the proposition literally expressed—is true just in case the requirements for knowing are fulfilled. But, in a given case, I might say such a thing, not in order to remark upon my knowing (/not knowing) *per se*, but to get across the information that, for instance, my epistemic position is good enough that you can take my word on the matter. (I am able to do so partly in virtue of the fact that ‘*S* knows that *p*’ entails that *S* is in a good epistemic position with respect to *p*, and partly in virtue of the fact that language-users count on others to make their conversational contributions conform to the Cooperative Principle⁴—see below.)

4. Or Bach & Harnish’s (1979) ‘communicative presumption’, or Sperber & Wilson’s (1986) ‘principle of relevance’. I tend to favor a Gricean framework, but a ‘relevance theoretic’ one, say, would serve just as well.

Obviously, given that whether my epistemic position with respect to p is good enough in the *operative* sense is going to depend upon our shifting interests, purposes, etc.—on whether we’re just chatting casually, whether we’re about to place a large bet, and so on —, there may be contextual variation in the information pragmatically imparted by the relevant utterance. Add to this the fact that our ‘semantic intuitions’—our intuitions about the truth conditions of sentences—are generally insensitive to the semantic/pragmatic distinction (we often read things into what is literally said, even by ourselves, without realizing we’re doing so), and we have all we need to explain why it can seem that the truth conditions of knowledge-attributing sentences depend on context when in fact they do not.

Now, in my paper, I go beyond this neutral framework and adopt, for the sake of illustration, a particular elaboration of it. Specifically, I adopt a relevant alternatives (RA) view, whereby the goodness of epistemic position required for knowing that p is cashed out in terms of the subject’s ability to rule out the relevant not- p alternatives. Further, I suggest that the relevant alternatives are fixed by what we, *qua* normal humans, *take to be* the likely counter-possibilities to what the subject is said to know. As the standard of relevance is not unreasonably high, we preserve our judgments as to subjects’ (/our) possessing knowledge in a wide variety of situations. And as the standard is fairly invariantistic, the present view gives us a fairly invariantistic RA semantics for knowledge sentences: the mere mentioning of a certain not- p alternative, or the mere fact that the speaker has certain not- p possibilities in mind when he says, “ S knows that p ,” doesn’t affect what that sentence means (what it literally expresses). What such things do or can affect, however, is just what information a given utterance of a knowledge sentence is liable to pragmatically impart. So, that I’m just a (bodiless) brain in a vat (a BIV) isn’t likely to be a relevant alternative to my having hands (since the BIV possibility is something normal people hardly ever consider); but if that possibility has been raised, and I counter with, “No, I know that I have hands,” I’m liable to be taken to be implying (i.e., to mean) that I can somehow rule that possibility out. And so it usually goes, on this view: while the truth of the sentence requires only that the relevant alternatives be ruled out, a knowledge attributor typically means to get across the more specific information that the subject’s epistemic position is good enough that he’s able to rule out the salient alternatives—i.e., those alternatives that are actually ‘in

the air'. And since, patently, there will be variation in what the salient alternatives (to p) are, we can expect to find a corresponding flexibility in our willingness to attribute knowledge.

As I say in the paper (2001, 487, 501), dissatisfaction with either the whole RA approach to knowledge or with my specific version thereof shouldn't be taken to cast doubt on the general framework being employed. (The reader is free to use his/her own preferred account of knowledge in fleshing-out the framework sketched above.) The idea is just to give the reader a sense of how that framework might be deployed, given a specific understanding of what's required for knowledge. It's not required that that specific understanding be the one I've chosen to operate with.

Nor is this ecumenicalism unmotivated. For all the general framework requires are some quite uncontroversial features of rational communication operating on only a partial, but in itself irreproachable, account of the truth condition of the relevant sentences. That is, given only a minimal view of the truth conditions of knowledge sentences—viz., that ' S knows that p ' implies that S believes that p , p , and that S is in a good epistemic position with respect to p —, on the assumption that our linguistic exchanges are governed by something like Grice's Cooperative Principle (CP), we have every reason to expect that there will be a flexibility in our willingness to attribute knowledge along the lines I have suggested. Specifically, it's because speakers strive to conform, and are known to so strive, to the maxim of relation ("Be relevant") that, in uttering a sentence of the form, ' S knows that p ', the speaker is naturally taken to intend/mean that S 's epistemic position with respect to p is 'good enough' *given the epistemic standards that are operative in the context in question*. For it is only if speakers are understood to be intending to communicate information about how the subject fares *vis-à-vis* the contextually operative standard(s)—hence, i.a., about whether the subject can or cannot rule out any contextually salient alternatives—that they can be seen as striving to be maximally relevantly informative (hence, as conforming to CP) (2001, 491–2).⁵

Now, it is being assumed here that information can be pragmatically generated by a given utterance without requiring for its generation the complete semantics of the sentence uttered. After all, ordinary speak-

5. The injunction to make one's conversational contribution 'maximally relevantly informative' is Harnish's (1976) suggested encapsulation of CP.

ers seem rather adept at communicating information by extra-semantic means; but even those whose job it is to provide truth conditional analyses of certain familiar sentence types can seldom agree on what they are. Further, there seem to be clear cases in which the pragmatic generation of certain information depends upon semantic features of the relevant sentences while leaving the exact truth conditional analysis thereof underdetermined. Thus, you don't need to have successfully Chisholmed sentences of the form '*O* looks *F* [to *S*]' in order to infer, from the fact that someone says, "The table looks black," that they have some doubt about whether it really is black—all you need to know is that, since *x*'s looking *F* doesn't entail that *x* is *F*, if the speaker does think that the table is black, that's what he should say, since that would be more informative. Similarly, you don't have to have anything worth calling a theory of belief in order to find it natural to say "I believe that *p*" in order to indicate that you're not absolutely certain as to *p*—since your believing that *p* would already be obviously inferable from you're saying "*p*", and since '*S* believes that *p*' doesn't imply that *p*, you can make explicit use of the term in order to emphasize the contrast with the categorical, "*p*", thereby flagging your lack of certainty.⁶

And so too in the present case: you don't have to have a complete and satisfactory theory of knowledge in order to think that speakers might utter knowledge sentences in order to get across quite specific information about how people (themselves or others) measure up to the *operative* epistemic standards. All you need is to see is that if I didn't think that *S* did so measure up, it would be odd, indeed uncooperative, of me to say, "*S* knows that *p*"; for whatever exactly knowledge is, we know that *S*'s knowing that *p* entails that *S* is in a good epistemic position with respect to *p*. "And why", the hearer may well ask, "would he say something that means [semantically implies] *that* if he didn't in fact think that *S* measured up to the epistemic standards that are in play, and so was able to put to rest any doubts or rule out any not-*p* possibilities that had just been raised?"⁷

6. Even if some of our pretheoretic intuitions surrounding the semantics of the expressions in question end up being explained away, some such inputs to truth-conditional analyses of the relevant expressions are necessary (otherwise, we won't know where to begin); and the semantic facts being assumed in the examples just given—that neither '*x* looks *F*' nor '*S* believes that *x* is *F*' implies that *x* is *F*—seem as widespread and uncontroversial as any.

7. Cf. Grice's suggestion as to a non-conversational analogue of the maxim of relation:

Notice a second quite general point that can be gleaned from the examples cited just above. Those examples might be seen as relying on DeRose's 'assert the stronger' rule:

[T]here's a very general conversational rule to the effect that when you're in a position to assert either of two things, then, other things being equal, if you assert either of them, you should assert the stronger. (2002, 175; 1999, 197)

However, it is less clear whether the examples illustrate this rule if we take the 'strength' in question to be such that '*p*' is stronger than '*q*' just in case it "impl[ies] but [is] not implied by" '*q*' (*ibid.*). For instance, if what's at issue is whether *x* is *F*, intuitively, "*x* is *F*" is 'stronger' than either "*S* believes that *x* is *F*" or "*x* looks *F*".—Certainly, someone's saying either of the latter would be less informative than the simple, "*x* is *F*".⁸ But it seems false that *x*'s being *F* implies(-but-isn't-implied-by) the proposition that *x* appears *F* [to *S*], or that *S* believes that *x* is *F*. So, if it is 'assert the stronger' which is the conversational rule at work in examples like those just given, the 'strength' in question is something like the overall informativeness of one's conversational contribution, given what's at issue.⁹ Better, then, to see the examples just given as illustrating Grice's maxim of quantity—"Make your [conversational] contribution as informative as is required (for the current purposes of the exchange)" (1989, 26)—, and say that uttering a sentence which expresses a proposition that is weaker (in DeRose's sense) than something else one is in a position to say is not the only way of being less informative than one might have been.¹⁰ Either way, the important point

"I expect a partner's contribution to be appropriate to the immediate needs at each stage of the transaction" (1989, 28).

8. Once again, assuming that what the conversants are interested in is *x* and its features, and not facts about *S*. If the latter were what they were interested in, being told that *S* believes that *x* is *F*, e.g., might well be more informative than being told that *x* is *F*.

9. It was because I was thinking of 'strength' in this way that it seemed to me that both "I know that *p*" and "I know that not-*p*" were stronger than "It's possible that *p*" (assuming that what's at issue is whether *p*): either of the former evinces a commitment to something that entails a univocal answer as to whether *p*, whereas, whatever its exact truth-conditional content, 'It's possible that *p*' does not; hence, from a speaker's saying the latter, one can infer that the speaker doesn't (take himself to) know either that *p* or that not-*p*.—Such was my thinking. Even so, it is of course a mistake to suggest that one might execute a 'WAM' of both 'I don't know that *p*' and 'I don't know that not-*p*' using DeRose's 'assert the stronger' rule and his stated notion of strength (see my 2001, 510, n. 30, and DeRose 2002, 198–199, n. 40).

10. Thanks to Kent Bach for discussion on this point.

is that, just as you needn't have a full and satisfactory account of the truth conditions of the relevant sentence in order to be able to use it to get across, or to glean from its use, information that's plausibly not part of its truth conditional content, you don't have to restrict yourself to strictly logical relations among the propositional contents of the sentences people utter, or might have uttered, in order to explain why certain information is pragmatically generated by their saying what they do—just as relevant are such things as which questions the parties are concerned to resolve. These are the sorts of facts about linguistic communication which it is important to keep in mind in assessing attempts to explain certain linguistic behavior by appeal to pragmatic factors, as well as in evaluating DeRose's reasons for being sceptical about the viability of sophisticated invariantism; it's to those reasons that we now turn.

3.

With a view to developing a general account of why some 'warranted assertability manoeuvres' (WAMs) succeed and others fail, DeRose gives the following examples of each:

- (i) 'lame' WAM (1999, 197; 2002, 174): A "crazed philosopher of language" gives an account whereby *S*'s being unmarried is not among the truth conditions of the sentence, '*S* is a bachelor'; rather *S*'s being unmarried is simply among the warranted assertability conditions of such a sentence.
- (ii) credible WAM (1999, 196–7; 2002, 174–5): If someone knows that *p*, it seems wrong, even false, for him to say (merely), "It's possible that *p*"; this might tempt us to think that 'It's possible that *p*' is true iff *S* 'doesn't know either way' (DKEW)—i.e., iff (1) *S* doesn't know that *p* and (2) doesn't know that not-*p*. However, a defender of the 'doesn't know otherwise' (DKO) account, who thinks that (2) alone is a genuine truth condition of the sentence, could argue that the appearance of (1)'s being a truth-condition derives from the fact that 'It's possible that *p*' is weaker than 'I know that *p*', and therefore that, from the fact that someone asserts the former, we can infer that he doesn't know that *p*.

What makes (ii) a credible WAM and (i) a miserable failure? According to DeRose, the reasons for according (ii), but not (i), some legitimacy, are as follows. First, (ii) is motivated by genuinely conflicting intuitions: where someone knows that p , it would seem wrong for him to say *either* “It’s possible that p ” *or* “It’s not possible that p ”; so we have good reason to think that *some* apparent falsehood is going to have to be explained away here (*ibid.*; 2002, 192). But there is a conspicuous lack of genuinely conflicting intuitions in the case of (i): “in the problem cases for the crazed theory of bachelor ... there’s no such pressure to have to explain away any appearances: It seems false to say of married men that they are bachelors, and it seems true to say of them that they are not bachelors” (1999, 198; 2002, 192). Further, as it stands, (i) is just a *bare* warranted assertability maneuver, whereas the second WAM, utilizing “a very general conversational rule” (‘assert the stronger’—see above), explains the apparent falsity in terms of the generation of a false implicature. Here is DeRose’s illustration, wherein the speaker knows that p , but has said only, “It’s possible that p ”:

To assert that weak possibility statement is unwarranted and generates the false implicature that the speaker doesn’t know that P by the following Gricean reasoning, which is based on the assumption that the speaker is following the “assert the stronger” rule: “If he knew that P , he would have been in a position to assert something stronger than ‘It’s possible that P_{ind} ’, and thus would have asserted some stronger thing instead. But he did assert ‘It’s possible that P_{ind} ’, not anything stronger. So he must not know that P .” We can easily mistake the falsehood of the implicature generated by such an assertion of “It’s possible that P_{ind} ” and the resulting unwarrantedness of the assertion for the falsehood of the assertion itself. (2002, 175; cf. 1999, 197)

DeRose is of course right that WAMing shouldn’t be allowed to be carried too far, so that, “[w]henever your theory seems to be wrong because it is omitting a certain truth-condition ... you can simply claim that assertions of the sentences in question generate implicatures to the effect that the condition in question holds” (2002, 173). But, on the face of it, the sophisticated invariantist view outlined above satisfies DeRose’s criteria for a successful WAM. Thus, while it can seem wrong to claim that I know that I’m not a BIV, for example, it can seem no less wrong to say that I *don’t* know that I’m not a BIV. After all, I know that I have hands, and BIVs are handless (cf. DeRose

1995, 2).¹¹ In terms of this same example, I explain the apparent falsity of “I know that I’m not a BIV” by appeal to the falsity of what would be conveyed by an utterance thereof—viz., that I can rule out the possibility that I’m a BIV; and the conversational rules appealed to—Grice’s CP and, specifically, the maxim of relation—are as general as one could like.

But matters are not so straightforward. For, if DeRose is right, when we look closer we find reason to doubt whether an invariantist WAM such as the one just described can really succeed. Thus, it’s a central idea of the present account that speakers can utter knowledge sentences in order to convey information that plausibly is not part of those sentences’ truth conditional contents. And since it is the information actually (pragmatically) conveyed that guides speakers in their knowledge-attributing (/denying) practices, speakers will sometimes resist uttering a (literal) truth because they rightly sense that in doing so they’d be committing themselves to something false.¹² So, once again, I suggest that ‘I know I’m not a BIV’ is one such sentence; it is, I claim, true (given that my being a BIV is not a relevant alternative), but an utterance thereof will standardly convey that the speaker can rule out the possibility that he’s a BIV, and speakers (understandably) might not want to commit themselves to that. In fact, speakers might go so far as to say something that’s literally false—e.g., “I *don’t* know that I’m not a BIV”—in order to communicate a salient truth (viz., that they can’t rule out the possibility that they are BIVs). Nor is this surprising if speakers are, as seems plausible, concerned with communicating (though not necessarily uttering only literal) truths. DeRose, however, finds this sort of claim incredible:

[I]nvariantists do not begin with a good candidate for WAMing, and they have to explain away as misleading intuitions of truth as well as intuitions of falsehood. For in the “low standards” contexts, it seems appropriate and

11. See the discussion of the (apparent) reality of conflicting intuitions in the cases in question, below.

12. As discussed below, the claim is not that speakers would necessarily have transparent knowledge of the matter, in just these terms. Like many others, I think that ordinary speakers (including us, most of the time) simply don’t distinguish between information that’s semantically encoded in sentences and that which is pragmatically generated by various utterances thereof. And why should they (we)? What they (we) care about, first and foremost, is the overall communicative content of utterances, not the relative contributions thereto of semantic and pragmatic factors.

it seems true to say that certain subjects know and it would seem wrong and false to deny that they know, while in the “high standards” context, it seems appropriate and true to say that similarly situated subjects don’t know and it seems inappropriate and false to say they do know. Thus, whichever set of appearances the invariantist seeks to discredit—whether she says we’re mistaken about the “high” or the “low” contexts—she’ll have to explain away both an appearance of falsity and (much more problematically) an appearance of truth. (2002, 193)

[But,] except where we engage in special practices of misdirection, like irony or hyperbole, don’t we want to avoid falsehood both in what we implicate and (especially!) in what we actually say? (2002, 192–3; 1999, 199)

Thus, as DeRose sees it, the data in question are simply unfriendly to an invariantistic handling: when we restrict our judgments (as we should) to their proper contexts, we see that our intuitions in the cases in question do not in fact conflict; and once those judgments *are* properly understood, we see that any invariantist account of the data is going to face the unenviable task of having to explain why we would utter a literal falsehood in order to communicate some truth.

As Brown (*forthcoming*) argues, however, on neither of these counts is DeRose’s argument persuasive. First, if we consider the sorts of cases the proper handling of which is at issue, it seems that every party to the current debate is going to have to explain away an appearance of *some* truth. As DeRose says of the independently plausible but (apparently) mutually inconsistent propositions which together constitute the sceptical “argument from ignorance”, for instance, “something plausible has to go” (1995, 2). Of course, the contextualist denies that conflicts of intuitions such as the one generated by the argument from ignorance, or more generally by comparing what it seems “appropriate and true” to say in high standard cases with what it seems “appropriate and true” to say in low standards cases, are genuine.—According to him, the appearance of conflict stems from our conflating contexts; so no apparent truths need actually be denied. But such cases draw our interest and seem to us to be of theoretical import precisely because it seems to us *to be true* that we have conflicting intuitions; and, as Brown says (*forthcoming*: Section 2), that *that* isn’t in fact the case—that our intuitions are, in fact, mutually consistent—is among the contextualist’s central claims. So if having to explain away an apparent truth is in itself a sign that one is headed down the wrong theoretical path, that’s bad news for everyone.

Granted, the contextualist doesn't have us *deliberately* uttering falsehoods so as to communicate truths; rather, on his view we simply fail to see that certain things—e.g., that certain of our judgments really are in conflict—*are* false. But, as will be clear from our discussion to this point, this response rests on the misunderstanding of the sort of sophisticated invariantist view being defended here. For it is essential to that view that speakers often conflate what is semantically expressed by a sentence with what an utterance thereof pragmatically imparts: it is among the sophisticated invariantist's central claims that our pretheoretic intuitions as to what we're 'saying' are insensitive to the semantic/pragmatic distinction (Rysiew 2001, Sections 3–4, Bach 1994 & *forthcoming*, Brown *forthcoming*, Section 3).¹³ To the extent that this is so, whether speakers see themselves as 'saying' something true/false will depend on the perceived truth value of the information that is merely pragmatically conveyed—if they think it's true, they won't see themselves as expressing falsehoods, deliberately or otherwise. Further, as Brown points out (*ibid.*), there are cases in which certain theorists have found it plausible to think that a literally false sentence can seem true because of the information an utterance thereof would communicate—those theorists may be wrong about this, of course; but the idea is “part of a standard and well-entrenched approach in the philosophy of language.” Given that this is so, intuitions of the sort DeRose is invoking here cannot, in themselves, be used to arbitrate between sophisticated invariantism and contextualism. To suppose that they can begs the question.¹⁴

13. It is this because of this that I have been hesitant (2001, 510 n. 32; cf. nn. 26, 29) to use 'implicature' in stating my view, preferring to speak more neutrally of information that is merely '(pragmatically) imparted' or 'conveyed'. In implicatures, properly so-called, one means what one (literally) says but also something else. But in less than fully literal speech—e.g., in Bach's (1994) implicatures, one of the cases Brown discusses—what one means is not what one (literally) says, but the more specific information that is generated by one's utterance.

14. And not just against sophisticated invariantists. It begs the question against the many philosophers and linguists—for instance, Bach (1994), Sperber & Wilson (1986), Recanati (1989), Carston (1988), and Gibbs & Moise (1997)—who have argued that in very many cases, the literal truth conditions of the uttered sentence fails to correspond to what the speaker intends to communicate. In their view, irony, figurative speech, and so forth, are best viewed as limiting cases on a continuum of less than fully literal speech.

4.

In addition to being pessimistic about the prospects for sophisticated invariantism as a general approach, DeRose finds my particular brand thereof unsatisfactory. He writes:

... While, as Rysiew points out, his account is “underwritten” by the very general conversational rule, “be relevant,” his is not the type of account I have urged is to be desired. On Rysiew’s relevant-alternativist theory of knowledge attributions, an assertion of “S knows that P” carries two separate but related meanings ... Where the contextually salient alternatives diverge significantly from the relevant alternatives, the injunction to “be relevant” gets the listener to fasten on the second, pragmatic meaning, since the semantic content of “S knows / doesn’t know that P” is clearly irrelevant to the purposes at hand. Thus, on Rysiew’s treatment (490) of my “high-standards” Bank Case, I utter the falsehood that I “don’t know” that the bank is open on Saturday, but all is well, for my wife can be counted on to ignore the falsehood I strictly said, since mere “knowing” / “not knowing” is not what is relevant to our there unusually heightened concerns, and focus instead on the second, pragmatic meaning that my assertion carries—that it’s not the case that I have a true belief and can rule out all of the *salient* alternatives to P. I am highly suspicious of accounts that help themselves to two such meanings in the way Rysiew’s does, disliking not only the loss of economy in explanation, but also worrying that we will not be able to combat all manner of absurd theories about the truth-conditions of various sentence if defenders of those theories are able to posit separate pragmatic meanings that do the work of accounting for usage in the troublesome cases, allowing their account of the truth-conditions to sit safely off in the corner, above the fray ... (2002, 198, n. 17)

As should be clear from the summary given above, however, that assertions of knowledge sentences “carry two meanings” is not a bare stipulation of the present account. The leading idea of that account, rather, is simply that we need to distinguish between the content of a given sentence and what one might convey in uttering it on a given occasion; and that, with some help from general conversational rules (as well as parties’ knowledge of the context, their general background knowledge, and so forth), uttering a given sentence (in context) can, and often does, serve to communicate information over and above, or even quite other than, what the sentence itself expresses.—Not, of course, that this means that the latter information (in the present case, that the

subject can/can't rule out the salient alternatives) is something that the sentence 'carries around', as it were. On the contrary, the whole point of appreciating pragmatic phenomena for what they are is to avoid such things as theorists' introducing extra meanings by hand.

Then again, if the role of the meaning of knowledge sentences on the present account was to trigger, thanks *solely* to the irrelevance of that semantic content in the cases in question, the hearer's inference to the information that is said to be (merely) pragmatically generated, that would be a problem. For then that content could be just about anything, and its connection to the information I allege to be pragmatically conveyed would be mysterious—the latter would be in no interesting way "generated by the semantic meaning together with general principles" (*ibid.*). However, the present account does assign knowledge sentences' semantic content an essential role in leading the hearer to the information I claim is pragmatically conveyed, and not just in virtue of its (sometimes) irrelevance. Let me clarify how this is so.

In outlining the general framework above (Section 2), I indicated how a minimal but still substantive account of knowledge sentences' truth conditions, interacting with familiar Gricean principles (especially, the maxim of relation), might explain the generation of the information which I claim is imparted by actual attributions/denials of knowledge; and in my paper I suggested that, supposing a particular invariantistic elaboration of knowledge sentences' truth conditions to be correct, we can predict the pragmatic generation of the sort of data that have seemed to some to favor contextualism. Granted, neither the discussion of Section 2, nor that of the hearers' inferences in the cases in question that I've offered—e.g., in the high standards bank case (2001, 490–2)—requires the specific invariantistic RA semantics I described above.¹⁵ But this doesn't show that the generation of information that is said to be merely pragmatically conveyed does not essentially rely on the meaning of knowledge sentences. For recall (Section 2) that such information can depend for its generation on (presumed) semantic features of the sentences involved while leaving the correct truth conditional analysis of those sentences underdetermined. And, in the

15. Much less does it require that hearers have a complete and explicit grasp of the distinction between relevant and merely salient alternatives, as opposed to some (admittedly fallible, though on the whole reliable) judgments as to the sorts of cognitive-epistemic abilities that are and aren't required for knowing.

present case, if ‘*S* knows that *p*’ didn’t semantically imply (at least) that *S* is in a good epistemic position with respect to *p*, the inference I allege—from a speaker’s saying “I [/don’t] know that *p*” to his [/not] being (or, his representing himself as believing that he is [/not]) in a good enough epistemic position to be able to rule out the salient not-*p* alternatives—would not go through.

For example, in the reconstruction just mentioned, I portray the hearer as relying on the fact that “‘to say “I don’t know ...” is to indicate that one isn’t sure”’ (2001, 491); and the account of why first-person knowledge claims might be used to express one’s confidence (492)—hence, why their denials might be used to indicate that one isn’t sure—requires that the subject’s fulfilling the conditions for knowing is seen to be both implied by ‘*S* knows that *p*’ and presumed by a person’s asserting, simply, that *p*. Further, though I do not say this in the paper, in interpreting her husband’s utterance the hearer in this example is plausibly seen as relying as well, even if only implicitly, on the literal content of ‘I don’t know the bank will be open’, inasmuch as she has no reason to think either that the speaker does not take himself to have a true belief (there is no indication that he *doubts* the bank will be open on Saturday), or that he takes himself not to be in a good epistemic position with respect to the bank’s being open, period (he’s been to the bank on a very recent Saturday, for example, and under other circumstances would be credited with knowing). Together with the assumption that he is striving to conform to the maxim of relation, these points make it natural to suppose that when the husband says, “I don’t know ...”, he means, and is taken to mean, that his epistemic position isn’t so good that he can rule out the possibility, which his wife has just explicitly raised, that the bank has very recently changed its hours.—Or so it seems to me; and I could just be wrong about this. But I don’t wish to say that my discussion of these matters is decisive, or that it couldn’t be improved upon.¹⁶ The present point, rather, is that the specific semantic content of knowledge sentences plays an essential role on the account being defended here, even if its playing that role does not require that the RA semantics described above be correct. What the eliminability of the latter shows is just that a particular invariantistic account of what

16. For example, Brown (*forthcoming*, Section 4) argues that I under-rate the role of practical importance (as against salience) in our knowledge-attributing behavior, and offers her own reconstruction of (e.g.) the high standards bank case.

constitutes the goodness of epistemic position required for knowing is not singled out by the phenomena in question. That is why, as I have emphasized, the general sophisticated invariantist framework I've described is largely neutral as between differing accounts of the exact truth conditions of knowledge sentences.

The foregoing would, of course, constitute a problem for the present view if it was intended to show that an invariantistic RA semantics is the correct account of the truth conditions of knowledge sentences. But it is not; the task, rather, is to show that that account, indeed any (non-sceptical) invariantistic semantics, *could* be correct: just as DeRose's WAM is intended to show that the DKEW theorist's proposed extra truth condition on 'It's possible that *p*' can be explained away (assuming that DKO is correct, and using CP), the present account is intended to show that *if* an invariantist semantics such as RA is true of knowledge sentences, that semantics, together with CP, can account for the data which have been taken to favor contextualism.

That said, though, if we do distinguish between the truth conditions of a given type of sentence and the information generated by particular tokenings thereof, how are we going to go about evaluating this or that account of what the former are? Indeed, how are we going to decide whether to adopt a contextualist or a non-contextualist stance towards those truth conditions? Well, of course, there are no fast and easy answers to these questions. But this is hardly news. We start off with some intuitions as to the truth conditions of a certain class of sentences; in the face of one or another reason—some conflicting intuitions, e.g.—for thinking that some other account of those truth conditions, or some other view of their sensitivity to context, might be correct, we have to decide which account which makes best overall sense of all of the relevant data—hence, which intuitions should in the end be directly reflected in our final account of the truth conditions, and which should be explained away. The process is seldom quickly decided; but it is, unfortunately, all too familiar, and there is no reason to think that any view which issues from it is going to inevitably be *ad hoc*.

Finally, it's worth pointing out that the view of the context-sensitivity of knowledge attributions I have been defending here may actually be supported by some of DeRose's own observations (cf. Rysiew 2001, 492):¹⁷ As DeRose and others have suggested, in asserting that

17. This is not to say that what follows is its sole support. What follows requires first-

p one represents oneself as knowing that p (see DeRose 2002, 179ff.). DeRose (2002) uses this fact as a basis for a new kind of argument for contextualism. Arguably, though, it points away from contextualism and towards sophisticated invariantism. For if asserting that p already involves representing oneself as knowing that p , this gives us reason to suppose that, when speakers who are striving to conform to the CP do actually say, “I know that p ”, they might do so in order to get across information that wouldn’t already be obviously inferable from their saying, simply, “ p ” (including, that they take themselves to know that p). Whereas, the flexibility in our knowledge-attributing practices is most naturally seen as supporting contextualism when we take speakers to be guided solely by what the sentences they utter mean—no more, and no less.

5.

As DeRose says, “[i]t is clear and uncontroversial that knowledge attributions are *in some way* governed by varying epistemic standards” (2002, 187). What is controversial is whether those varying standards bear upon the truth conditions of knowledge attributions. Here, I have argued that we need not be as pessimistic as is DeRose about the viability of sophisticated invariantism, either as a general approach, or in the form I have recommended. How, then, are we to decide between contextualism and sophisticated invariantism? The question is one of deciding which account of the context-sensitivity of knowledge attributions, the contextualist’s or the sophisticated invariantist’s, is favored by the balance of considerations. The goal here has been to suggest that the balance doesn’t obviously tip the contextualist’s way.¹⁸

person attributions; the general rationale for the approach (Section 2) does not.

18. Thanks to Kent Bach (who also suggested the title), Martijn Blaauw, Jessica Brown, Adam Leite, and Mohan Matthen for their thoughtful and helpful comments on earlier drafts of this paper.

REFERENCES

- Bach, K. (1994). 'Conversational Implicature', *Mind and Language* 9, 124–162.
- (forthcoming). 'The Emperor's New "Knows"', *Contextualism in Philosophy: On Epistemology, Language and Truth*, (ed.) Gerhard Preyer and Georg Peter, Oxford University Press, Oxford.
- Bach, K. and R. M. Harnish (1979). *Linguistic Communication and Speech Acts*, MIT Press, Cambridge, Massachusetts.
- Brown, J. (forthcoming). 'Contextualism and Warranted Assertibility Manoeuvres', *Philosophical Studies*.
- Cappelen, H. and E. Lepore (2003). 'Context Shifting Arguments', *Philosophical Perspectives 17: Language and Philosophical Linguistics*, 25–50.
- Carston, R. (1988). 'Implicature, Explicature, and Truth-theoretic Semantics', *Mental Representations: The Interface Between Language and Reality*, (ed.) R. M. Kempson, 155–181, Cambridge University Press, Cambridge.
- Cohen, S. (1999). 'Contextualism, Skepticism, and The Structure of Reasons', *Philosophical Perspectives 13: Epistemology*, (ed.) James E. Tomberlin, 57–89, Ridgeview, Atascadero CA.
- DeRose, K. (1995). 'Solving the Sceptical Problem', *The Philosophical Review* 104.1, 1–52.
- (1999). 'Contextualism: An Explanation and Defense', *The Blackwell Guide to Epistemology*, (eds.) J. Greco and E. Sosa, 185–203, Blackwell, Cambridge, MA.
- (2002). 'Assertion, Knowledge and Context', *The Philosophical Review* 111.2, 167–203.
- Feldman, R. (1999). 'Contextualism and Skepticism', *Philosophical Perspectives 1: Epistemology*, 91–114.
- (2001). 'Skeptical Problems, Contextualist Solutions', *Philosophical Studies* 103, 61–85.
- Gibbs, R. W., and J. F. Moise (1997). 'Pragmatics in Understanding What is Said', *Cognition* 62, 51–74.
- Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press, Cambridge, MA.
- Harnish, R. (1976). 'Logical Form and Implicature', *An Integrated Theory of Linguistic Abilities*, (eds.) T. Bever, J. Katz, T. Langendoen, Crowell. Reprinted in S. Davis (ed.), *Pragmatics: A Reader*, Oxford University Press, 1991.
- Hofweber, T. (1999). 'Contextualism and the Meaning-Intention Problem',

- Cognition, Agency and Rationality*, (eds.) K. Korta, E. Sosa, and X. Arrazola, 93–104, Kluwer, Dordrecht, Boston, and London.
- Klein, P. (2000). 'Contextualism and the Real Nature of Academic Skepticism', *Philosophical Issues* 10, 108–16.
- Kornblith, H. (2000). 'The Contextualist Evasion of Epistemology', *Philosophical Issues* 10, 24–32.
- Lewis, D. (1979). 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* 8, 339–359.
- Pritchard, D. (*forthcoming*). 'Contextualism, Scepticism, and Warranted Assertibility Manoeuvres', *Knowledge and Skepticism*, (ed.) J. Keim-Campbell, M. O'Rourke & H. Silverstein, MIT Press, Cambridge, MA.
- Recanati, F. (1989). 'The Pragmatics of What is Said', *Mind & Language* 4.4, 295–329.
- Rysiew, P. (2001). 'The Context-Sensitivity of Knowledge Attributions', *Noûs* 35.4, 477–514.
- Schiffer, S. (1996). 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society* 96, 317–333.
- Sosa, E. (2000). 'Skepticism and Contextualism', *Philosophical Issues* 10, 1–18.
- Sperber, D., and D. Wilson (1986). *Relevance: From Communication to Cognition*, Blackwell, Cambridge, MA.
- Stanley, J. (2004). 'On the Linguistic Basis for Contextualism', *Philosophical Studies* 119, 119–146.