Grazer Philosophische Studien
69 (2005), xx–yy.

SOME WORRIES FOR WOULD-BE *WAM*MERS

Adam LEITE
University of Indiana

Summary

DeRose appeals to ordinary English usage to support his contextualist semantics for 'know'-attributions. A common objection holds that though the relevant assertions are both appropriate and seemingly true, their seeming truth arises merely from their appropriateness. This Warranted Assertability Maneuver (WAM) aims to provide a stand-alone objection by providing a reason not to take the ordinary language data at face-value. However, there is no plausible model or mechanism for the pragmatic phenomena WAMmers must postulate. Given what the WAM requires, it is doubtful it could work out in detail.

1. *Introduction*

According to contextualists, the truth-conditions of sentences of the form 'S knows that p'[1] are determined in part by the intentions, purposes, interests and practical circumstances relevant in the conversational context. Two broad lines of evidence have been mustered for this proposal. First, it is claimed to cohere well with the knowledge account of assertion and to provide a unified solution to epistemological puzzles relating to skepticism and lottery-type cases. Second, it is claimed that ordinary English usage supports the contextualist semantics (DeRose 1992; 2002; *forthcoming*). My concern in this paper is with the latter claim; in particular, I will focus on one important example, Keith DeRose's "Bank Cases" (DeRose 1992, 913; cf., Cohen 1999, 58–59). My question is whether a Warranted Assertability Maneuver ("WAM") provides certain noncontextualists (moderate invariantists) with a satisfying response to this appeal to ordinary usage.

1. More precisely: of such sentences on occasions of utterance. I ignore this refinement in this discussion, as nothing turns on it.

DeRose's Bank Cases consist of a pair of cases in which the speaker of a first-person 'know'-attributing or denying sentence is (a) equally confident of the truth of a given proposition p, (b) is in the same epistemic position[2] regarding p, and (c) regards himself as being in the same epistemic position. Yet because of a difference in the costs of error, the speaker asserts, "I know that p" in one case (the low stakes case) while asserting, "I don't know that p" in the other (the high stakes case). In particular, in the low-stakes case, the speaker asserts "I know that the bank will be open on Saturday", while in the high-stakes case—in which the speaker stands to lose his house if his paycheck is not deposited before Monday morning—the speaker asserts "I do not know that the bank will be open on Saturday."

That (a)–(c) are held constant in both cases is essential. If they were not, then the example would no longer provide any evidence for contextualism, since the speaker's behavior would be adequately explained by the shift in non-conversational factors.

DeRose muddies his example by having a conversational participant offer additional reasons for doubt in the high-stakes case. My interest is in a version of the example aimed at showing that *purely practical* factors make a crucial difference. Accordingly, I will focus on a "high stakes" case where there is no variation in the doubts raised. (Assume, again, that requirements (a)–(c) are met).

> Speaker A: Shall we stop at the bank to deposit our paychecks?
> B (Spouse): Nah, we can go tomorrow.
> A: But are they open on Saturdays? Lots of banks aren't.
> B: I was there two weekends ago and checked their hours. They're open until noon. [This much is identical in both cases. The high-stakes case continues as follows.]
> A: Are you sure? If our paychecks aren't deposited by Monday morning, our mortgage check will bounce, and we'll lose our house!
> B: You're right. I don't know they'll be open.

The contextualist makes four claims about this case.

---

2. By "epistemic position" I mean one's condition regarding the epistemic factors, whatever they are, that in addition to truth and belief are required for knowledge.

1. Speaker B's assertion of "I don't know" is conversationally appropriate, in that it is perfectly reasonable, in accordance with the rules of orderly discourse and linguistic usage, and not an odd, forced, or obviously nonliteral use of this sentence. For convenience, I'll use the term "appropriate" to capture all of this.[3]
2. If B were to assert "I know that the bank will be open", the assertion would be inappropriate.
3. What B said in asserting "I don't know" seems to be true.[4]
4. If B were to assert "I know that the bank will be open", what he thereby asserted would seem to be false.

According to the contextualist, this case contrasts with the "low stakes" case on all four counts. The contextualist then argues that in the absence of reasons not to do so, we should take the appearances as a good guide: the truth-conditions for "I know that the bank will be open" and its denial differ in the two cases. The contextualist is guided here by the presumption that if competent and properly informed speakers take an assertion to be both appropriate and true, then all else equal, the assertion is true (DeRose *forthcoming*, sect. 1). Given claims (1)–(4), the bank cases thus *prima facie* favor the contextualist semantics; the burden of proof lies with the objector.

It is crucial to the contextualist's brief that in the high stakes case the assertion, "I don't know …," is appropriate and seemingly true. If the claim were rather that the assertion, "I know the bank will be open", would be inappropriate and seemingly false, it would be less plausible that the *truth-conditions* shift. For if it were inappropriate and seemingly false to assert *both* "I know" *and* "I don't know", then it would be equally plausible that the truth-conditions had remained constant but contextual factors had made it inappropriate to say *anything* about whether one does or doesn't know in the high-stakes case.

Moderate Invariantism[5] is a family of non-contextualist views which

---

3. The relevant notion of appropriateness goes beyond politeness and the like; the speaker is not just being politely concessive or conciliatory.

4. I will talk indifferently of sentences (on occasions of utterance) as being true and of *what is expressed* as being true. Nothing turns on this issue for my purposes here. When I say that a sentence is true, I mean that what is literally expressed by the assertive utterance of it on that occasion is a truth. For convenience, I will likewise use the term 'assertion' sometimes to refer to the asserted sentence, sometimes to what is literally expressed.

5. *Moderate* because non-skeptical; *invariantist* because non-contextualist.

accept the following two claims:

1. There is some "low-stakes" case in which the speaker knows and the knowledge claim is appropriate. I will assume that the low stakes bank case is such a case.
2. Neither the truth-values nor the truth-conditions of 'know'-attributing sentences shift with *anyone's* practical circumstances.

Moderate Invariantists have two primary lines of response to the Bank Cases.[6]

The first denies that the high-stakes assertion of "I don't know …" is really appropriate and seemingly true. It grants that in very similar cases—cases in which the speaker gets worried and loses his confidence or reappraises the strength of his evidence—it would be appropriate and seemingly true for the speaker to say this. But it claims that when specifications (a)–(c) are met, it would be odd for the speaker to declare "I don't know …" (This isn't to claim that it would then be appropriate to assert "I know …"; it might not be appropriate to assert either.) Ordinary cases in which it *looks* appropriate to assert "I don't know p" as a result of high stakes will then get explained away as arising from shifts in the speaker's degree of confidence or appraisal of his own epistemic position (Leite 2004, 346–7; Bach *forthcoming*).

The second, more concessive, response grants that the assertion of "I don't know …" is both appropriate and seemingly true. But it denies that these concessions reveal anything about the sentence's truth-conditions, on the grounds that its seeming truth arises merely from the fact that it is appropriately asserted. This is the Warranted Assertability Maneuver (WAM).[7] It aims to provide a *stand-alone objection* to the contextualist's argument from ordinary language by providing a reason not to take the ordinary language data at face value.

---

6. One could reject claim (2) and still deny contextualism by claiming that it is the practical circumstances of the *subject of the attribution or denial* that matter, not those relevant in the speaker's context. This route is taken by "Subject Sensitive Invariantism" (SSI) (Stanley *forthcoming*; Hawthorne (2004); cf. Kaplan (*forthcoming*), though Kaplan is perhaps best read as making no claim about truth-conditions at all). SSI strikes me as implausible, since it makes incorrect predictions about the appropriateness-conditions of third-person "know"-ascribing sentences. My concern here, however, is the responses available to the Moderate Invariantist.

7. For important criticisms, see DeRose (1999; 2002). For a well-developed attempted WAM, see Rysiew (2001).

My aim in this paper is to evaluate the prospects for this latter response. As I will argue, there is a straightforward reason for pessimism: once one clearly recognizes what it requires, it is extremely hard to see how it could work out in detail.[8]

2. *The* WAM*mer's burdens*

The WAMmer's project has two components, an *explanatory* component and a *justificatory* component. The explanatory component requires a story which, if true, would explain the contextualist's data on the assumption that the moderate invariantist semantics is correct. The justificatory component provides an argument in favor of the explanatory story. I begin with the explanatory component.

2.1 *The* WAM*mer's explanatory tasks*

The WAMmer takes on two main explanatory burdens. The first is to explain why the assertion, "I don't know …," is *appropriate* in the high stakes case even though false. The second is to explain why it seems to ordinary speakers that in asserting this sentence, the speaker has said something true. To meet the first burden, it is not sufficient to explain why it would be inappropriate to assert, "I know …"; in general, to explain why it is inappropriate to assert Ø is not (yet) to explain why it is appropriate to assert not-Ø, since it might instead be inappropriate to assert either. To meet the second burden, it is not sufficient to meet the first, since ordinary speakers can take assertions to be appropriate without taking them to be true. The second burden imposes constraints on how the first is met; the account of why the assertion is appropriate must square with the account of why ordinary speakers are confused.

To meet these burdens, the WAMmer first distinguishes what is *literally expressed* by a given assertive utterance of a given sentence and

8. My main concern is the use of a WAM as a stand-alone objection to the contextualist's argument from ordinary language. If one accepts independent semantical arguments against the contextualist semantics, one might still attempt a WAM to explain the contextualist's examples. My argument below suggests that this won't work unless the WAM appeals to conventional implicatures. Moreover, a WAM isn't the only option here, since there is also the possibility of the less concessive response sketched above. Anticontextualists shouldn't assume that one way or another, a WAM *must* work out.

what is thereby *pragmatically imparted*, that is, what is communicated by the act of utterance even though it is not part of what is literally expressed nor entailed by what is literally expressed. To take a standard example, when asked to evaluate a student, I might reply, "His papers are well formatted and he shows up for class." My utterance literally expresses that his papers are well formatted and that he shows up for class. What is communicated, however, is that I don't think much of his philosophical abilities. There has been hot debate over how this distinction is to be drawn, what role pragmatic factors play in fixing what is literally expressed, and the like. For my purposes here we can prescind from these details.[9]

The WAMmer's next step is to postulate two things: first, an account of what is literally expressed when one asserts "S knows/doesn't know that p" in the circumstances of the contextualist's example; second, an account of what is thereby pragmatically imparted. The details of what is literally expressed will depend upon which particular moderate invariantist view the WAMmer selects, but whichever it is, the assertion of "I don't know …" in the high stakes case will literally express a falsehood. According to the postulated pragmatic theory, however, something conversationally appropriate will thereby be imparted, and the WAMmer postulates, moreover, that according to the norms determining conversational appropriateness, the appropriateness of what is thus communicated outweighs the *prima facie* inappropriateness arising from the falsehood of what is literally expressed. If true, these postulations would thus explain why it would be appropriate in the given conversational circumstances to assert "I don't know …".

Finally, the WAMmer proposes to explain the seeming truth of this assertion by claiming that (we) ordinary speakers confuse what is pragmatically imparted by the utterance with what is literally expressed. If (we) ordinary speakers are confused in this way, and if we recognize what is thus pragmatically imparted to be true, then it is no wonder that the uttered sentence would seem true.[10]

The suggestion is *not* that this confusion explains why the assertion of "I don't know …" is appropriate. Admittedly, an assertion can

9. The phrases "pragmatically impart" and "pragmatically communicate" suggest communicative success. But as I mean to use them, it is possible for something to be pragmatically imparted by a given utterance even if the hearer fails to cotton on.

10. The preceding three paragraphs broadly follow Rysiew (2001).

be appropriate or reasonable when, because of semantic confusion, the speaker mistakenly thinks it true and relevant. In such cases, the assertion is *excusable* and the speaker *blameless*. But this is not what the contextualist had in mind. He meant that someone who is fully aware of the relevant facts and clear about the semantics would judge that the speaker had acted in a fully correct way, in accordance with the relevant norms, not merely done the best he could given his misunderstanding. So the WAMmer rejects the contextualist's data if he claims that the assertion in the high stakes case is appropriate simply because the speaker is confused about what is literally expressed by his utterance. For this reason, the WAMmer's account must allow that the assertion would be appropriate even if the speaker were fully aware of its falsehood. The WAMmer does not discharge his explanatory obligations if he simply says, "The speaker thought he was literally expressing something true and relevant, and that's why his assertion was appropriate."

This point can appear to spell trouble for the WAMmer. His first explanatory burden requires him to explain why the assertion is appropriate—in a way that does not amount to mere excusability—when made by an ordinary speaker. His response to the second explanatory burden claims that ordinary speakers are in error about what is literally asserted in these cases. But if ordinary speakers are confused in this way, then aren't their assertions in these cases at best excusable—and so not appropriate in the way the contextualist requires? The WAMmer's way with the second explanatory task thus appears incompatible with the data which motivate the first.

To avoid this problem, the WAMmer must hold that an assertion can *both* arise from semantic confusion *and* be conversationally appropriate in the richer sense. This requires that semantic confusion be compatible with correctness according to the relevant norms. How could this be? Suppose that a speaker is confused about the truth-conditions of an uttered sentence ('X') but intends to communicate something true and relevant (Ø). Suppose, moreover, that the norms governing appropriate usage involve or generate a general and standard link between sentences like 'X' and contents like Ø, so that the default use of 'X' in these circumstances would be to communicate what he intends to communicate. In such conditions, mistakenly thinking that he is literally asserting Ø is fully compatible with appropriately asserting 'X'—so long as his intention in doing so is to communicate Ø. Hence, the WAMmer sees

the situation like this: an ordinary speaker's assertion of "I don't know …" in the high stakes case is not a misuse of the language, because the speaker abides by the relevant norms in uttering what she does. As long as her utterance aims at communication of what is thereby standardly imparted according to the relevant norms, her assertion is appropriate in the relevant sense; she's just wrong about what she is literally expressing.[11]

The WAMmer therefore postulates a general, default rule governing the use of "I don't know …" such that its utterance in the high stakes case would communicate what, according to the WAMmer, is thereby pragmatically communicated. The relation between the use of such sentences and the communicated contents will then be such that despite the falsity of what is literally expressed: a) The assertion (made by an ordinary speaker) is appropriate in a way that goes beyond excusability resulting from semantic confusion, b) The assertion would be appropriate even if the speaker, being clear about the semantics, recognized it to be false, and c) It is plausible that ordinary speakers could be confused about the truth-conditions of what is literally expressed.

2.2 *The* WAM*mer's justificatory task*

So far the WAMmer has offered a series of postulations which, if true, would explain the contextualist's data while rejecting contextualism. Such postulations count for little, however; they can easily be made whenever one wants to reject a putative truth-condition while accepting the linguistic data offered in its defense. The WAMmer's *justificatory* task, then, is to vindicate postulation of the relevant pragmatic phenomena. As we will see, this is where trouble arises.

2.2a. *Conventional implicature*

The simplest proposal is this. (1) As a matter of convention, asserting "I know/don't know that p" entitles a hearer to make certain inferences which are not licensed by the asserted sentence's semantic content. Thus an assertion of "I don't know that p" would be appropriate, even if false,

11. Ordinary speakers will be unclear about their own minds as a result of this semantic confusion. For instance, they will misunderstand their communicative intentions, stating that in cases like the high stakes bank case, they intend to communicate that they do not know that the bank will be open. Worries about the WAMmer's strategy might be raised on this score, but I leave them to one side.

if what is thereby conventionally imparted is true and conversationally relevant. (2) The conventions function as defaults, so ordinary speakers might well confuse what is imparted by such assertions with their literal semantic content. But a speaker who was fully clear about their semantics would nonetheless regard them as appropriate.

This proposal puts the WAM in the category which Grice termed "conventional implicature". It runs into a fundamental problem. Because the ordinary language data presumptively favor the contextualist's account of the truth-conditions of the relevant sentences, it is question-begging for the WAMmer to claim simply, "This is how we use these expressions: "know"-attributions are conventionally used to convey Ø, even though Ø is not part of the semantic content of what is asserted nor entailed by it." The WAMmer consequently needs a principled argument for regarding the content in question as only *conventionally associated* with the sentences in question, not part of nor entailed by their semantic content. Merely claiming that something is, as a matter of convention, pragmatically imparted does not make it so.

It can be easy to miss this when considering particular proposals. Suppose the WAMmer suggests that "I know that p" is standardly used, as a matter of convention, to indicate (in part) that one is sure enough, or has good enough evidence, to reasonably act on the assumption that p given the risks. If that were so, then it would be reasonable to suggest that an assertion of "I don't know that p" standardly indicates that one is *not* sure enough, or that one does not have good enough evidence, to act on p in one's practical situation. Accordingly, what is literally expressed by the assertion in the high stakes bank case would be false, but what would be conveyed or indicated is that one isn't sure enough, or lacks good enough evidence, to put off the deposit until the next day. The proposal is both initially plausible and meets the WAMmer's explanatory needs. However, because of the way the contextualist takes the truth-conditions of "I know that p" to shift with the practical circumstances, the contextualist holds that if one's utterance of "I know that p" is true, then (as a matter of entailment) one must be sure enough that p, or have good enough evidence that p, to warrant acting on p in one's circumstances. Accordingly, the contextualist holds that what is literally expressed by an utterance of "I don't know that p" is *true* if one fails to meet this condition. The WAMmer has therefore gotten nowhere unless he provides an independent argument against the contextualist's account of the truth-conditions. This is a problem with

the general strategy, not with the particulars of this proposal.[12]

The main test to distinguish conventional implicature from truth-conditional content is *detachability*: conventional implicatures are generated by the use of particular expressions, so an inference is based upon a conventional implicature rather than semantic entailment if it is *not* generated by other ways of formulating the same truth-conditional content. Unfortunately, this test doesn't help the WAMmer. Since the dispute concerns precisely the truth-conditional content of the relevant expressions, there is no neutral way to test for detachability.[13]

2.2b *Conversational implicature and related phenomena*

An independent argument for the WAMmer's postulated pragmatic effect can be provided if the effect is a *conversational implicature* or related phenomenon (such as so-called *explicature* or *conversational impliciture*). Such phenomena are aspects of the speaker's intended meaning, not semantically entailed by what's literally expressed nor purely conventional, but recoverable from the fact that the speaker uttered what s/he did, given the assumption that s/he intended thereby to make an appropriate conversational contribution. They come in two basic varieties: *particularized* and *generalized* (Horn 2004, 5). A particularized conversational implicature or related effect is one which is induced only in a special, quite particular context, while a generalized implicature or related effect is one which is generated in the absence of a special, marked context. Generalized conversational implicatures thus function as *defaults*. This suits them especially well for the WAMmer's purposes, since a default conversational implicature or related effect is most likely to generate just the sort of semantic confusion the WAMmer postulates.[14]

---

12. A similar problem arises for the proposal that in saying "I know …", one is typically attempting to communicate that one's hearers can take one's word on the matter.

13. One could construct a WAM utilizing considerations relating to the *speech-act* conventionally performed (in certain cases) by declaring "I know". In the interests of space, I leave this proposal aside. It works well in first-person cases, but does not generalize to cover contextualist-friendly second and third-person cases.

14. As Levinson comments, "those implicatures that are both derived from observing the maxims and are generalized have a special importance for linguistic theory. For it is these in particular which will be hard to distinguish from the *semantic* content of linguistic expressions, because such implicatures will be routinely associated with the relevant expressions in all ordinary contexts" (1984, 127).

Postulation of a conversational implicature or related effect is justified by means of a conception of conversation and communication as a rational, cooperative enterprise guided by principles such as Grice's general "Conversational Principle":

> (CP) Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange" (Grice [1967] 1989, 26).

To vindicate the claim that an assertive utterance of "U" in given circumstances pragmatically imparts Ø, one shows how a fully rational hearer in those circumstances could "work out"—from (1) what is (according to the semantic theory) literally said, (2) the fact that the speaker uttered those words in those circumstances, and (3) the assumption that the speaker is abiding by the CP—that the speaker meant to communicate Ø. So, for instance, a fully rational conversational participant who heard me respond to a request to evaluate a student by saying, "His formatting is excellent and he shows up for class," would reason: "There is no reason to suppose that he [the speaker] is not observing the principles governing conversation. He means to communicate something that is relevant to the request for an evaluation of student's philosophical abilities. If he thought the student was good he would say so, since there is no reason for him to offer faint praise when he thinks more is warranted. He knew that I would recognize this. So he must mean to communicate that the student isn't very good." According to the Gricean, the claim that the student's lack of ability is pragmatically imparted is vindicated precisely by the fact that this component of the communication can be worked out in this way.

Two points should be emphasized. First, the claim is just that if a plausible line of reasoning of this sort *can be produced by the theorist*, that vindicates postulating the pragmatic phenomenon in question. The claim need not be that ordinary speakers always or ever go through any such reasoning explicitly, nor need it be that such reasoning models a psychologically real process by which actual hearers comprehend speakers' intended meanings.

Second, in cases where the semantics are in dispute, a theorist is not entitled to postulate the pragmatic effect in question unless s/he can provide reasoning (of the relevant sort) by which a rational hearer

could in principle arrive at what is putatively pragmatically conveyed.[15] As Grice put it,

> (…) the final test for the presence of a conversational implicature [has] to be, so far as I [can] see, a derivation of it. One has to produce an account of how it could have arisen and why it is there. And I am very much opposed to any kind of sloppy use of this philosophical tool, in which one does not fulfill this condition (1981, 187; in Neale 1992, 527, Neale's insertions).

This claim is shared by all theorists working within the broadly Gricean framework. It has an important consequence for the anticontextualist WAMmer. The attempted WAM stands or falls with the WAMmer's ability to provide the relevant reasoning in the high-stakes example.

There are four crucial constraints on the WAMmer's attempt to do so.

First, the relevant speech situation is one in which (a) both the speaker and the hearer are neither confused about the semantics nor in error about any relevant facts, and (b) it is common knowledge between them that this is so. Here's why. Suppose that the hearer is semantically confused in just the way that the WAMmer takes ordinary speakers to be. Then, the hearer in the high-stakes case would take the speaker to have literally expressed something true, conversationally relevant, and otherwise in accordance with the principles governing conversation – exactly what, according to the WAMmer, was pragmatically imparted. But then the relevant pragmatic reasoning wouldn't be triggered at all, so the postulated pragmatic effect won't have been vindicated. Likewise, if the hearer takes the speaker to be semantically confused in the way that the WAMmer holds ordinary English speakers to be, then the hearer would reason as follows: "The speaker is confused about the truth-conditions of the sentence he uttered; given what he takes them to be, he thinks he has literally expressed something true and relevant. So that's why he uttered what he did: he took himself to be fully satisfying the conversational principles. So I don't need to look for any further pragmatic communicative intent in order to explain his utterance." On this line of reasoning, the speaker is interpreted as making an *excusable* false utterance. In neither case, then, does the hearer engage in reasoning

15. This point holds even if one has an independent semantical argument against the contextualist semantics, since appealing to conversational implicature isn't the only option here.

that would justify the WAMmer's pragmatic claim: he does not "work out" the pragmatic communication that the WAMmer hypothesizes.

Second, the WAMmer's account of the implicature must at least render it *intelligible* that reflective and thoughtful speakers could mistake what is pragmatically communicated for what is literally expressed. To see the issue here, consider an example. When a speaker evaluates a philosophy student by saying, "His papers are well-formatted and he shows up for class," it is conversationally implicated that the student isn't any good at philosophy. But no halfway intelligent speaker would think that the truth-conditions of the uttered sentence involve the student's being no good at philosophy. In this case, then, it would be exceedingly implausible to postulate such confusion as part of an explanatory theory. This is an example of a *particularized* implicature; the pragmatic effect in this case is generated only via special-case reasoning arising in response to an apparent violation of the conversational maxims. An implicature generated in this way would not justify the WAMmer's postulation of a default, standardized linkage between the relevant expressions and the relevant pragmatically-imparted content, and so it would not justify the postulation of widespread semantic confusion.

Third, given the justificatory purpose of sketching the hearer's reasoning, the reasoning cannot appeal to ancillary knowledge of what speakers usually use "know"-attributions or denials to impart. It might seem that it could. For instance, suppose that speakers generally use "I know …" to impart that they are very confident and use "I don't know …" to impart that they are not very confident, though these things are *not* part of what is literally expressed by the utterance of these sentences. It would then seem that the hearer should be able to reason as follows: 'He said that he didn't know that the bank would be open. This isn't true; he does know it. But speakers generally use sentences of the form "I don't know …." to impart that they are not very confident. So, given that he meant to communicate something true and relevant, and given that he knows that this is how such sentences are standardly used (and knows that I know this too) he must have meant to communicate that he wasn't very confident that the bank would be open.' However, sketching the hearer's reasoning is supposed to *support* or *establish* a claim about what utterances of the relevant sentences are generally used to communicate in addition to what they literally express, so to appeal to such considerations in the sketch of the hearer's reasoning would assume

the very thing that the WAMmer needs to establish. Consequently, the hearer's reasoning must proceed *only* from considerations about the truth-conditions of what was literally expressed, the words the speaker uttered, the circumstances, and the conversational principles; it cannot make use of ancillary information about what is standardly pragmatically imparted by use of the relevant expressions.[16]

The fourth constraint concerns the speaker's attitude to the asserted sentence. By the description of the case, the speaker (1) believes that the bank will be open, (2) is in an evidential position which (according to the moderate invariantist) is adequate for knowledge and knows what his evidence is, and (3) knows what the truth-conditions of "I know the bank will be open" are. So given the moderate invariantist's account of the truth-conditions, the speaker will take himself to be literally expressing a falsehood when he utters, "I don't know the bank will be open tomorrow." Moreover, the hearer will recognize this fact about the speaker, since the hearer is aware of (1)–(3). Consequently, the postulated reasoning must begin with the recognition that the speaker has intentionally asserted a sentence which he takes to literally express a falsehood.

2.2c *The problem*

It is extremely hard to see how the hearer's reasoning could go.

In outline, it looks clear enough. The hearer should reason as follows. "What he just literally expressed—that he doesn't know that the bank will be open tomorrow—is something he takes to be false (and he expects me to realize this). But since he is operating in accordance with the conversational principles, he meant to communicate something true and relevant (and he knows I will recognize this). What he must have meant to communicate is that …," where the ellipsis is filled in with something conversationally appropriate, such as: "he isn't absolutely certain that the bank will be open tomorrow" or "given our circum-

16. Several of Rysiew's proposals fall afoul of this constraint. For instance, his account of the hearer's reasoning from the maxim of Relation has the hearer appeal to ancillary information that "to say, 'I don't know…' is to indicate that one isn't sure" (2001, 491). The speaker in the high-stakes case might be *sure enough for knowledge*, so the "indication" in question will not be a semantic entailment, but rather a pragmatic effect. Cf. his attempt to explain the pragmatic effects of first-person 'know'-ascriptions in terms of the epistemic commitments assumed in assertion (where he has the hearer appeal to the ancillary non-semantic claim that 'I know' adds an element of confidence or certainty to the bare assertion of p) (492).

stances his epistemic position with regard to whether the bank will be open tomorrow isn't good enough (without further checking) to warrant putting off the deposit until tomorrow". But how is the hearer to get from the thought, "He must have meant to communicate something else," to a conclusion about what the speaker must have meant to communicate? As noted previously, he cannot rely upon ancillary knowledge of what such sentences are standardly used to communicate. He must rely only upon the circumstances, the conversational principles, his knowledge of the (moderate invariantist) semantics of the uttered sentence, the fact that the speaker uttered the sentence while believing it to be false, and the assumptions that the speaker isn't confused about what he has literally expressed or mistaken about the relevant facts and aims to abide by the conversational principles. These are meager resources – especially given that what the speaker has done is very strange. It isn't obvious that *anything* would be pragmatically conveyed in these circumstances.

An analogy may help bring out the difficulty. Imagine that a speaker is drinking something. The question of whether it is hot is conversationally salient. The speaker believes it to be at least warm, and the hearer recognizes that he believes this. The speaker announces, "It's not warm. I'd better heat it up." On the assumption that he isn't making a slip, semantically confused or in error about the temperature of his drink, the natural conclusion is that he is not acting in accordance with the conversational principles at all. Whatever he's doing, it's completely bizarre. *But that's exactly what the situation is like, as the WAMmer is conceiving it*. The contextualist takes there to be a scale of epistemic standards, some more stringent than the one which must be met for the truth of a 'know'-attributing sentence according to the moderate invariantist semantics. The WAMmer doesn't dispute that. The contextualist holds that an epistemic standard more stringent than that picked out by 'know' on the moderate invariantist semantics is conversationally salient in the high stakes bank case. The WAMmer agrees. Call the epistemic position which meets that more stringent standard 'absolute certainty' (just to give it a label). The conversationally salient question is whether the speaker has attained the epistemic position *absolute certainty*—a position *better* than that required for knowledge. The speaker believes that he knows, and the hearer recognizes that he believes this. The speaker announces, "I don't know. I'd better check." On the assumptions that he isn't making a slip, semantically confused or in error about his epistemic position and that the moderate invariantist

semantics for 'know' are correct, the natural conclusion is that he is not acting in accordance with the conversational principles at all. Whatever he is doing, it is completely bizarre.

In the analogy, the hearer can only take the speaker to intend to communicate that his drink is not hot. There's no other way to make any sense of his behavior. Likewise, in the bank case we can't make any sense of the speaker's behavior unless we take him to mean to communicate that he isn't *absolutely certain*. In both cases, however, there is something unintelligible in the speaker's behavior; neither speaker is communicating in a fully rational way. This fact importantly limits the explanatory potential of the hearer's ability to divine the speaker's communicative intentions. In the drink case, the hearer's discovery of the speaker's communicative intentions does not ground any sort of *general*, default pragmatic phenomenon. Nor does it make it at all intelligible that competent speakers could generally be confused about the truth-conditions of "It's not warm"—not even when it is uttered in these sorts of conditions. The same conclusion appears to follow in the case of 'knows', if the moderate invariantist semantics is correct. There seems to be no significant difference between the cases.[17]

3. *Some proposals*

I now want to examine the most promising possible responses to this problem. So far as I can see, none fully meets the WAMmer's needs.

3.a *Classical Gricean proposals*

Grice's Conversational Principle encapsulates a set of general maxims, divided into four groups, governing rational conversational interchange (Grice [1967] 1989, 26–7):

> Quality: Try to make your contribution one that is true.
>
> 1. Do not say what you believe to be false.
> 2. Do not say that for which you lack evidence.

17. This consideration presents difficulties for Rysiew's appeal to the maxim of Relation (2001, 491–2).

Quantity:

1. Make your contribution as informative as is required (for the current purposes of the exchange).
2. Do not make your contribution more informative than is required.

Relation: Be relevant.

Manner: Be perspicuous.

1. Avoid obscurity of expression.
2. Avoid ambiguity.
3. Be brief.
4. Be orderly.

I will focus upon possible implicatures arising from the maxims of Quantity, Relation, and Quality, in that order.[18] (I will ignore the maxims of manner, as they seem irrelevant in this context.)

Quantity

The maxims of Quantity generate important scalar effects, as when a speaker reports that someone is "nice-looking", thereby communicating that he isn't drop-dead gorgeous. Since both the contextualist and the WAMmer postulate a scale of epistemic positions, it is natural to expect scalar effects in the high-stakes bank case. Unlike standard scalar cases, however, this case would involve the speaker's *falsely* asserting that a lower point on the scale is not reached in order to communicate that a higher point is not reached. There is no accepted Quantity-based model for scalar effects of this sort. One can say, "He's not cute [stress]" to convey that he is gorgeous. But one can't use "He's not cute" (where this is taken by both parties to be false) to convey that he is *not* gorgeous.

Assertion of a literal falsehood can generate Quantity-based pragmatic effects in certain non-scalar cases. For instance, the assertion "He lost a finger in the accident" ordinarily communicates that he did not lose a thumb. This can be explained via the maxims of Quantity and the lexicalization of 'thumb' (Horn 2004, 16). This pragmatic effect generates phenomena somewhat analogous to what the WAMmer needs. For instance, the assertion "I didn't lose a finger [stress] in the accident", made by someone who is known to have lost a digit, communicates

18. Grice's list has been challenged, but I am aware of no proposed revisions which substantially affect the argument of this section.

that he or she lost a thumb. (Cf., 'rectangle' and 'square'.) Here, what is literally expressed is false, but the assertion is appropriate because what is thereby communicated is true and relevant.

These cases provide a poor model for the high-stakes bank example. In these cases, a term X refers to a general class, only one member of which has a lexicalized term Y referring to it. By saying (in appropriate circumstances) not-X, one implicates that Y. One thus *specifies*, picking out the member of the class referred to by the more specific lexicalized term. But in the high-stakes bank case, the speaker does not (falsely) negate a more inclusive term in order to implicitly apply a more specific lexicalized term.

I am aware of no other ways in which asserting a literal falsehood can generate anything like the relevant sort of pragmatic effects via the maxims of Quantity.

Relation

The conversational issue in the high stakes bank case is whether the deposit can reasonably be put off until the following day. Considerations of relevance suggest that whatever the speaker intends to communicate by uttering, "I don't know …", it is something which will help resolve this issue. So it would seem that the ideal hearer should be able to reconstruct the intended communication as follows. "He intends to make a relevant contribution, and he knows that I recognize his intention to do so. He has said *something* about his epistemic position regarding whether the bank will be open tomorrow. But he takes what he has literally asserted to be false, and he knows that I will recognize this. So he must mean to communicate something else about his epistemic position. It would be relevant to our situation if he weren't in a good enough epistemic position to be comfortable putting off the deposit until tomorrow. So that must be what he intended to convey."

There are several significant problems with this proposal.

First, by themselves considerations of relevance do not enable the hearer to determine what the speaker meant to communicate. If all the speaker said was, "I don't know …" it would be equally possible to interpret him as intending to communicate that he does not *merely* know, but rather is in a *better* position than the one required for knowledge—and hence that it would be best to put off the deposit. The mere utterance of "I don't know …" in the stipulated circumstances is pragmatically underdeterminate, so far as the maxim of Relation goes.

Second, further contextual factors (the speaker's other linguistic behavior—tone, stress, and other utterances) enable the hearer to narrow things down. But this just highlights that the utterance of "I don't know" was completely unnecessary. The speaker could just as well have uttered, "Fitzmallin guzampt fishcake. I guess I'd better check." In fact, the speaker could simply have said "We'd better check," and the hearer, reasoning that the speaker meant to convey something relevant, could have figured out what the speaker meant to convey. Consequently, both the semantic content of the utterance and the fact that it was uttered are completely idle in the account. The point is not that they are redundant; other contextual factors are doing *all* the explanatory work. So vindicating the postulated pragmatic effect via the maxim of Relation doesn't render the speaker's utterance linguistically appropriate in this case. The hearer can figure out what the speaker meant to communicate, but a residual mystery remains: why did he choose *those* words to do so? It looks like a case of linguistic misuse.

Third, this account does not provide any reason to expect a generalized, default implicature. For this reason, it doesn't make it intelligible how ordinary speakers could be confused about the truth-conditions of utterances of the relevance sentences. Nor does it provide any reason to postulate a further pragmatic element of the appropriate sort in ordinary speakers' usage.

Quality

Here's a more promising proposal. For any position on the scale of epistemic positions, if you don't possess it, then you also don't possess any stronger position. (This is a logical relation, a matter of entailment.) According to the moderate invariantist, some position on this scale lower than the highest possible is adequate for knowledge. Suppose, however, that the position required for *being warranted in the high-stakes case in putting off the deposit until Saturday* is stronger than the minimum position required for knowledge. Then not knowing (because one's epistemic position is inadequate for knowledge) will *entail* that one isn't in the stronger position either. Accordingly, "I don't know that the bank will be open" will logically entail (or close enough[19]) something both true and conversationally relevant: that one

19. Calling this an entailment is convenient shorthand, though a bit of a cheat. Given the WAMmer's semantics, "I don't know that p" entails: "P is false or I don't believe that p or

is not in the epistemic position required for being warranted in putting off the deposit. The proposal, then, is that the hearer exploits this fact in order to work out the implicature. The reasoning would go roughly like this. "He literally expressed that he does not know the bank'll be open. But he should take this to be false, given that he believes that the bank will be open and believes that his epistemic position is adequate for knowledge. But given everything else he believes, he would take *I don't know that the bank will be open* to entail *I'm not in the minimum epistemic position required for being warranted in putting off the deposit until Saturday*, and he would expect me to recognize this. Call that entailed proposition "~Ø". ~Ø is both true and directly relevant. There is no natural and quick way for him to literally express it. So ~Ø must be what he meant to communicate, via my recognition that (1) he would take what he literally expressed to be false, (2) ~Ø is directly entailed by what he literally expressed, (3) it is true and relevant, and (4) he has no straightforward and reasonably quick way to express it." On this proposal, the implicature is generated by the interplay of the maxim of Quality, a principle of least effort (given the readily available linguistic resources), and the scalar structure of epistemic positions.[20] The proposal is that the assertion, "I don't know …," in the high-stakes case is a kind of *over-statement*: the speaker is saying something logically stronger (and false) in order to easily communicate something closely related and logically weaker.

This proposal meets the WAMmer's explanatory needs. If general conversational principles dictate that absent a straightforward way to formulate what one intends to communicate, one should make the most easily formulated (even if false) assertion that obviously entails what one intends to communicate, then the relevant implicatures will be generated by utterances of the relevant sentences across a wide range of standard conditions. In particular, assertion of "I don't know that p",

---

my epistemic position regarding p isn't adequate for knowledge." By uttering "I don't know that p", the speaker does not communicate that he does not believe that p is true; if he had wanted to do that, he should have said either "Not-p" or "I don't believe that p" (or something to that effect). So given that the speaker has already indicated that he *does* believe that p is true (by asserting p), the hearer can conclude that the speaker means to focus attention more particularly on the proposition that his epistemic position is not adequate for knowledge. This proposition entails that his epistemic position is not any *better* than the minimum position required for knowledge.

20. The principle of least effort figures centrally in Horn's recent model (2004, 12–17)).

whether true or false, will communicate *that the speaker's epistemic position isn't good enough to warrant acting on p* whenever the required epistemic position is stronger than the minimum position required for knowledge. It would be no wonder, then, if ordinary speakers are semantically confused.

Still, this proposal faces three significant objections.

First, it is not obvious that implicatures are generated in the proposed way; there are no paradigmatic cases of this sort nor any obvious models upon which to draw. What is needed is an ordinary example in which a speaker who is not semantically confused knowingly asserts a falsehood in order to communicate something entailed by it, something which there is no obvious and straightforward way to express. A possible model is provided by "loose" constructions, such as 'He has a square face.' It could be proposed that an utterance of this sentence literally expresses something false but communicates something true and obviously entailed which we cannot easily formulate ("His face is more like a square than like any other easily identified shape", perhaps). However, it is hardly obvious that what is literally expressed here is something false, nor that if it is, this proposal provides the best treatment of loose constructions. I defer here to linguists and philosophers of language.

Second, it is doubtful that the proposal really applies to the case at hand. The proposal requires that there be no *equally simple and straightforward* way to express what the speaker wishes to communicate. But that is not true in the high stakes bank case. The speaker could simply say, "I'm not certain that the bank will be open," or, perhaps, "It's not certain that the bank will be open." So it is not at all obvious that the WAMmer can really justify postulating the necessary pragmatic phenomena.[21]

Finally, any proposal which exploits scalar relations between epistemic positions faces a further explanatory challenge. Quantity-based scalar implicatures are widespread. Assuming scalar relations between epistemic positions, one would expect that asserting "I know that p" could communicate—via the Maxim of Quantity— that one has not

21. This worry is especially significant in light of Horn's conjecture (1972) that because of a redundancy constraint, if use of a given lexical item generates a certain *generalized* conversational implicature, then there will be no lexical item that directly encodes that implicature. (Horn, On the Semantic Properties of the Logical Operators in English. Mimeo. Indiana University Linguistics Club. Cited in Levinson 1984).

attained any *stronger* epistemic position. So on the WAMmer's semantics, it should be appropriate to utter "I know that the bank will be open" in the high stakes case; one would thereby assert something true and communicate something relevant. The WAMmer needs to explain why things don't work this way instead.

3.b *Conversational impliciture*

Consider this example. A parent heartlessly comments to a young child screaming from a scraped knee:

(1) You aren't going to die.

The uttered sentence is false but appropriate, since it successfully communicates that the child will not die *from the scrape*. This as an example of what Kent Bach (1994) calls "conversational *impliciture*." In conversational impliciture, "the speaker is not being fully explicit. Rather, he intends the hearer to read something into the utterance, to regard it as if it contained certain conceptual material that is not in fact there" (126). The conversational impliciture is what *the speaker* literally meant, and is reconstructible through *supplementation* of the literal linguistic content of the utterance.

It can seem that (1) provides the model the WAMmer needs. However, there are two significant reasons why it doesn't. First, conversational implicitures of the relevant type are *detachable* (Bach 1994, 137). However, if the WAM exploits entailments arising from scalar relations between epistemic positions, then it must be modeled on a *non-detachable* pragmatic effect, since these entailment relations hold regardless of the particular words used. Second, implicitures of the relevant type involve "what might be called 'lexical' strengthening, in that what is being communicated could have been made fully explicit by the insertion of additional lexical material" (Bach 1994, 134); the hearer supplies "missing portions of what is otherwise being expressed explicitly" (154). But in the high stakes bank case no additional lexical material would convert the utterance into an explicit expression of only the claim that the speaker's epistemic position isn't good enough to warrant putting off the deposit.

Moreover, in cases of conversational impliciture ordinary speakers easily recognize that the uttered sentence is "strictly speaking" false.

According to the contextualist, this is not so in the high stakes bank case. The WAMmer agrees. Accordingly, the model leaves it unclear how speakers could be confused in the relevant way about the truth-conditions of "know"-attributions.

3.c *Back to ordinary speakers*

At this point, the WAMmer might suggest something like this.[22] "Perhaps the trip through the Gricean apparatus was a mistake: we don't need to explain how an ideal hearer, operating on the assumptions that the speaker is obeying the Conversational Principle and not semantically confused, could "work out" what the speaker intends to communicate; we just need to understand how *actual* hearers do it. And we don't need any special explanation for why ordinary speakers are semantically confused; their grasp of the semantics operates tacitly in the production of their linguistic behavior, so there is no particular reason to think they'd get it right without a lot of work. Here, then, in outline, is how ordinary hearers do it: guided by their tacit understanding of the semantics, they look around for the most relevant proposition "in the ballpark". In the high stakes bank case, the "ballpark" is that the speaker meant to communicate something relevant about his epistemic position. The most relevant thing he might plausibly have meant (given his other behavior) is that his epistemic position isn't adequate to warrant putting off the deposit. So that's what the hearer concludes he meant. And since the hearer isn't clear about the semantics at the conscious level, the hearer will naturally be confused about whether this is what is explicitly expressed by his utterance." This is a last ditch attempt. The trouble with it is simply this: it is completely neutral on the semantics. The contextualist could tell the same story.

4. *Conclusion*

The WAMmer's project has run into a fundamental difficulty. None of the proposals we've considered provide a fully satisfactory fix. So I suggest that we start over. Is the speaker's assertion *really* appropriate in the high stakes case? Consider the conversation again, keeping in

22. Inspired, perhaps, by Relevance theorists.

mind that the second speaker is—and *remains*—as confident that the bank will be open as one ordinarily can be about such things and that he was there just *two weeks* ago and saw no sign that they would be changing their hours (nor does he have any special reason to doubt his memory, etc.). It helps to read it aloud.

> Speaker A: Shall we stop at the bank to deposit our paychecks?
> B (Spouse): Nah, we can go tomorrow.
> A: But are they open on Saturdays? Lots of banks aren't.
> B: I was there two weekends ago and checked their hours. They're open until noon.
> A: Are you sure? If our paychecks aren't deposited by Monday morning, our mortgage check will bounce, and we'll lose our house!
> B: You're right. I don't know they'll be open.

Is it really linguistically appropriate for Speaker B to say "I don't know they'll be open," given that neither his appraisal of his evidence nor his confidence shifts *and he isn't just being conciliatory*? To my ear, it isn't; it is decidedly odd. (This isn't to say that it would be appropriate for him to claim outright that he *does* know, either.) I am therefore inclined to think that moderate invariantists should deny—not explain—the contextualist's putative ordinary language data.

REFERENCES

Bach, K. (1994). 'Conversational Impliciture', *Mind and Language* 9, 124–162.

— (*forthcoming*). 'The Emperor's New 'Knows'', *Contextualism in Philosophy*, (ed.) G. Preyer and G. Peter, Oxford University Press, Oxford.

Cohen, S. (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* 13, 58–89.

DeRose, K. (1992). 'Contextualism and Knowledge Attributions,' *Philosophy and Phenomenological Research* 52, 913–929.

— (1999). 'Contextualism: An Explanation and Defense', *The Blackwell Guide to Epistemology*, (ed.) E. Sosa and J. Greco, 185–203, Blackwell Publishing, Oxford.

— (2002). 'Assertion, Knowledge, and Context', *The Philosophical Review* 114, 167–203.

— (*forthcoming*). 'The Ordinary Language Basis for Contextualism and the New Invariantism', *The Philosophical Quarterly*.
Grice, H.P. (1981). 'Presupposition and Conversational Implicature', *Radical Pragmatics*, (ed.) P. Cole, 183–198, Academic Press, New York.
— (1989) [1967]. 'Logic and Conversation', *Studies in the Way of Words*, Harvard University Press, Cambridge, MA.
Hawthorne, J. (2004). *Knowledge and Lotteries*, Clarendon Press, Oxford.
Horn, L. R. (2004). 'Implicature', *The Handbook of Pragmatics*, (ed.) L. Horn and G. Ward, 3–28, Blackwell Publishing, Oxford.
Kaplan, M. (*forthcoming*). 'Deciding What You Know'.
Leite, A. (2004). 'Is Fallibility an Epistemological Shortcoming?', *The Philosophical Quarterly* 54, 232–51.
Levinson, S. C. (1984). *Pragmatics*, Cambridge University Press, New York, NY.
Neale, S. (1992). 'Paul Grice and the Philosophy of Language', *Linguistics and Philosophy* 15, 509–559.
Rysiew, P. (2001). 'The Context-Sensitivity of Knowledge Attributions,' *Nous* 35, 477–514.
Stanley, J. (*forthcoming*). *Knowledge and Interests*, Oxford University Press, Oxford.