

A WARRANTED-ASSERTABILITY DEFENSE OF A MOOREAN RESPONSE TO SKEPTICISM

TIM BLACK

CALIFORNIA STATE UNIVERSITY, NORTHRIDGE

In *Acta Analytica* 23 (2008): 187-205

ABSTRACT: *According to a Moorean response to skepticism, the standards for knowledge are invariantly comparatively low, and we can know across contexts all that we ordinarily take ourselves to know. It is incumbent upon the Moorean to defend his position by explaining how, in contexts in which S seems to lack knowledge, S can nevertheless have knowledge. The explanation proposed here relies on a warranted-assertability maneuver: Because we are warranted in asserting that S doesn't know that p, it can seem that S does in fact lack that piece of knowledge. Moreover, this warranted-assertability maneuver is unique and better than similar maneuvers because it makes use of H. P. Grice's general conversational rule of Quantity—"Do not make your contribution more informative than is required"—in explaining why we are warranted in asserting that S doesn't know that p.*

KEYWORDS: *warranted-assertability, Mooreanism, skepticism, knowledge*

Taken together, the following statements seem to generate a troubling conflict:

- (1) I don't know that I'm not a brain-in-a-vat.
- (2) If I don't know that I'm not a brain-in-a-vat, then I don't know that I have hands.
- (3) I know that I have hands.

These statements generate a conflict because they are independently plausible yet jointly inconsistent. It seems, then, that we must either live with an inconsistency or reject some very plausible claim. Of course, neither of these options is particularly attractive.

In responding to this worry, we can choose from at least two strategies. First, there's the semantic strategy, which maintains that 'know' either is or functions very much like an indexical. On this view, context helps to determine the semantic content of 'know', much like context helps to determine the semantic content of indexicals like 'I' and 'here'.

Moreover, since context affects the semantic content of 'know', it will have an effect on the

semantic content of complex lexical items in which ‘know’ appears, for example, knowledge attributions. Keith DeRose, who is perhaps today’s foremost epistemological contextualist and a proponent of the semantic strategy, puts it this way:

truth-conditions of knowledge ascribing and knowledge denying sentences (sentences of the form ‘S knows that P’ and ‘S doesn’t know that P’ and related variants of such sentences) vary in certain ways according to the contexts in which they are uttered. What so varies is the epistemic standards that S must meet (or, in the case of a denial of knowledge, fail to meet) in order for such a statement to be true. (DeRose 1999a, p. 187)¹

Contextualists seek to resolve the conflict between (1), (2) and (3) in terms of a contextual variability in the truth conditions of knowledge-ascribing and knowledge-denying sentences. They maintain, first of all, that (2) is true across contexts.² (3) is true, they suggest, in contexts in which the epistemic standards are comparatively low, that is, in contexts in which the truth conditions of ‘I know that I have hands’ are comparatively easily satisfied. Finally, (1) is true in contexts in which the epistemic standards are fairly high, that is, contexts in which the truth conditions of ‘I know that I’m not a brain-in-a-vat’ are less easily satisfied.³ According to semantic strategists, then, (1), (2) and (3) do not in fact conflict, even though it seems that they do.⁴

¹ Compare Cohen 1999, p. 57; and Rieber 1998, p. 190. Typically, contextualists are concerned with the truth conditions of knowledge-ascribing and knowledge-denying sentences (see, for example, DeRose 1999a, pp. 187-188). Nevertheless, they often “semantically descend” from talk of the truth conditions of such sentences to talk of the standards for knowledge (see DeRose 1999b). Such a semantic descent is permissible, they suggest, as long as we remember that our primary concern is with the truth conditions of certain sentences.

² See DeRose 1999c, especially pages 29-33.

³ DeRose, who has provided what is perhaps the most influential version of epistemological contextualism, introduces his position in DeRose 1992 and in further detail in DeRose 1999c. Other recent and noteworthy versions of epistemological contextualism are those of Cohen 1988, Lewis 1996, Rieber 1998, Cohen 1999 and Heller 1999. I should also note that for DeRose, as for several other contextualists, *contexts* seem to be *conversations*. So, whatever the criteria are for the individuation of conversations, those same criteria, or at least some very similar ones, will be the criteria for the individuation of contexts.

⁴ Compare the apparent conflict between (1), (2) and (3) to the apparent conflict between

- (1*) I am in Los Angeles,
- (2*) I am not in Los Angeles, and
- (3*) Either I am in Los Angeles, or I am not in Los Angeles (but not both).

We might instead adopt the *pragmatic* strategy in seeking to resolve the apparent conflict between (1), (2) and (3). According to this strategy, ‘know’ is not an indexical, and its semantic content is determined independently of context. It also maintains that there is in fact a conflict between (1), (2) and (3), and it seeks to resolve this conflict in terms of a contextual variability in the *assertability* conditions, rather than the truth conditions, of knowledge-ascribing and knowledge-denying sentences.⁵ Thus, if the pragmatic strategy is to be effective, it must allow us to explain why one of the three conflicting claims seems true even though it’s false.

There are two versions of the pragmatic strategy. According to the *skeptical* version, the semantic content of ‘know’ is such that the standards for knowledge are invariantly quite high. Skeptical invariantists must therefore explain why it seems true that I know that I have hands even though that claim is in fact *false*. I favor a second version of the pragmatic strategy, which I call the *Moorean* version. According to the Moorean, the semantic content of ‘know’ is such that the standards for knowledge are invariantly comparatively low, and we can know across contexts the things we ordinarily take ourselves to know.⁶ If the Moorean pragmatic strategy is to be successful, it must answer two questions:

- (A) How is it that we can know across contexts the things we ordinarily take ourselves to know?
- (B) If I know that I’m not a brain-in-a-vat, for example, why does it seem *true* that I *don’t* know?

The contextualist’s resolution of the apparent conflict between (1), (2) and (3) is similar to a resolution of the apparent conflict between (1*), (2*) and (3*) that notes that it’s *Jane* who utters (1*) but *Matt* who utters (2*).

⁵ Assertability conditions have to do with the conditions under which we may make certain assertions rather than with the conditions under which those assertions are true.

⁶ I do not mean to suggest that *Moore* was (what I am calling) a Moorean invariantist. The main—and perhaps the only—similarity between Moore and the Moorean is this: each maintains that the standards for knowledge are invariant and that we know across contexts the things we ordinarily take ourselves to know. The Moorean also maintains that the standards for knowledge are invariantly low, but Moore himself might not have agreed with this.

An adequate response to (A) requires that we specify the Moorean standards for knowledge, and then explain how our beliefs, even our belief that we're not brains-in-vats, can meet those standards. Since I have responded to questions like (A) in other places,⁷ my concern here will be to answer questions like those asked in (B).

I suggest that the Moorean can employ the notion of *warranted assertability* in explaining why, although it's false that we don't know that we have hands, it can in some contexts seem true.⁸ In particular, the Moorean can maintain that since in some contexts, we are warranted in asserting that we don't know, it can seem true that we don't know. I should also note that in order to determine whether we can provide this sort of explanation, we must set aside any account of warranted assertability which claims that assertions are warranted only if they are true, for the explanation we are concerned to provide crucially involves the idea that certain false assertions are warranted.⁹

Now, DeRose persuasively argues that one important invariantism—Peter Unger's skeptical invariantism—uses warranted assertability improperly, and that we should reject Unger's view on those grounds.¹⁰ DeRose goes on to suggest that Moorean responses must suffer the same fate, for they too will make improper use of the notion of warranted

⁷ Consider in this regard my belief that I'm not a brain-in-a-vat. I argue in Black 2002 and again in Black 2008 that this belief meets the conditions in a methodized version of Robert Nozick's "truth-tracking" account of knowledge (see Nozick 1981, p. 179) and thus that I know that I'm not a brain-in-a-vat. If I am to know that I'm not a brain-in-a-vat, my belief as to whether I'm a brain-in-a-vat must track the truth only through the group of worlds in which perception plays a substantial role in producing my belief. Yet no brain-in-a-vat worlds are among this group. Moreover, my belief as to whether I'm a brain-in-a-vat tracks the truth from the actual world to the farthest world in which it is true both that I'm not a brain-in-a-vat and that perception plays a substantial role in producing my belief – throughout the worlds in which I'm not a brain-in-a-vat and in which perception substantially produces my belief as to whether I'm a brain-in-a-vat, I believe that I'm not a brain-in-a-vat. Thus, I know that I'm not a brain-in-a-vat.

⁸ Warranted-assertability defenses of Moorean (or moderate invariantist) responses are provided by Rysiew 2001, Blaauw 2003, Brown 2005, Pritchard 2005 and Brown 2006; and in outline in Black 2005 and Black 2008. A defense of the Moorean response need not appeal to warranted assertability. See Bach 2005 for a defense that appeals to an error theory.

⁹ Such accounts include knowledge accounts of assertion like those endorsed by Unger 1975, DeRose 1996, Williamson 1996, Williamson 2000 and DeRose 2002.

¹⁰ See DeRose 1999c and especially DeRose 1999a. For Unger's skeptical invariantism, see Unger 1975.

assertability.¹¹ I believe, however, that this need not be the case. In this paper, I defend a Moorean response to skepticism with a legitimate use of the notion of warranted assertability. This supports the Moorean pragmatic strategy, and shows that strategy to be a viable alternative to the semantic strategy. We therefore eliminate one of the chief reasons for preferring the semantic strategy, namely, that there is no viable alternative to it.

What, though, is a legitimate use of the notion of warranted assertability? So far as I know, only DeRose has proposed conditions for a legitimate use of this notion.¹² I hope to show, then, that we can defend a Moorean response to skepticism with a warranted assertability maneuver that meets each of DeRose's conditions. This will do two things. First, even if DeRose has not provided a complete and correct set of conditions, it is clear that his conditions are met by some obviously legitimate uses of the notion of warranted assertability. So, in showing that the Moorean's warranted assertability maneuver meets each of DeRose's conditions, we will have shown that it measures up, at least in this respect, to some obviously legitimate uses. Second, in showing that the Moorean's warranted assertability maneuver meets each of DeRose's conditions, we will have responded to skepticism on our staunchest critic's own terms. This greatly restricts the replies that are available to those who would criticize the Moorean response.

In Section 1, then, I outline DeRose's three conditions for the successful employment of the notion of warranted assertability. In Section 2, I argue that the Moorean response satisfies DeRose's first condition. In Section 3, I defend the Moorean position with a maneuver that employs the notion of warranted assertability, and I show that this maneuver, in addition to satisfying DeRose's first condition, satisfies his last two conditions. Finally, in Section 4, I provide additional arguments for the claim that we should favor the Moorean

¹¹ See DeRose 1999c, p. 210, and DeRose 1999a, p. 203.

¹² In DeRose 1999a.

pragmatic strategy over the contextualist's semantic strategy. In particular, I argue that unlike the Moorean, the contextualist will find it difficult to defend his position with a maneuver that employs the notion of warranted assertability.

1. Warranted Assertability Maneuvers

The Moorean holds that the standards for knowledge are always comparatively low, and thus that they're low both in contexts in which the standards for knowledge are alleged to be high—*high-standards contexts*—and in contexts in which they're alleged to be low—*low-standards contexts*. The Moorean thus adopts Moorean invariantist standards for knowledge:

(MI) S's assertion, 'I know that P' is true if and only if S's belief that P meets certain comparatively low standards for knowledge.¹³

According to the Moorean, then, there are no contexts in which the standards for knowledge are so high as to prevent my belief that I have hands from meeting those standards. Yet if there are no such contexts, why do we judge in some contexts that I *don't* know that I have hands? On behalf of the Moorean, I respond to this question by arguing that there are contexts in which we are *warranted in asserting* that we don't know. And we can mistake the warranted assertability of this claim for its truth, thus explaining why we sometimes judge, albeit mistakenly, that we don't know that we have hands, for example. Let's call maneuvers of this sort *warranted assertability maneuvers*, or WAMs.

DeRose suggests, however, that no successful response to skepticism will include a WAM.¹⁴ His suggestion is that successful WAMs must meet three conditions, and that

¹³ As it is, (MI) is fairly vague, and there are different ways of specifying its standards. Rysiew 2001 prefers to specify them in terms of the relevant alternatives theory of knowledge, according to which S knows that P only if S (or S's evidence) can distinguish P from all relevant alternatives to P. On this view, (MI)'s standards are comparatively low because they require the elimination of a smaller set of alternatives, a set of alternatives that is in fact a proper subset of the set of alternatives the skeptic requires us to eliminate. Brown 2006 would specify (MI)'s standards "in terms of the range of possible worlds across which one's belief matches the facts" (p. 424). The Moorean response I propose here can be modified to work with either of these ways of specifying (MI)'s standards.

¹⁴ See DeRose 1999a, pp. 201-203.

responses that perform a WAM will fail to meet at least one of them. DeRose's conditions are revealed through an examination of an allegedly successful WAM, which is performed in defense of what he calls the 'Don't Know Otherwise' (DKO) account of 'It's possible that P'.¹⁵

(DKO) S's assertion, 'It's possible that P' is true if and only if it's not the case that S knows that P is false.¹⁶

Suppose that Sue knows that my book is on her desk, and hence that it's not the case that she knows that my book isn't on her desk. Suppose, too, that when I ask her about the book, she says, 'It's possible that it's on my desk'. Now, Sue's assertion is true according to (DKO). Yet DeRose claims that her assertion seems "somehow wrong" or even "downright false" given that she knows that the book is on her desk (DeRose 1999a, p. 196). And if she were to say either 'It's *impossible* that your book is on my desk' or 'It's *not* possible that your book is on my desk', she would assert something that would also seem false. This case is therefore one in which we have "[a]n intuition that an assertion is false, where the opposite assertion also seems false, indicating that some intuition here has to be explained away" (DeRose 1999a, p. 200). We may use a WAM, at least in this case, in order to explain away the appropriate intuition: According to DeRose, when Sue says 'It's possible that the book is on my desk', she says something that's weaker than something else she's in a position to assert, namely, 'The book is on my desk'. DeRose goes on to maintain that "[t]o say the weaker thing is to make an inappropriate assertion, and it's this unwarrantedness of the assertion that we're mistaking for falsehood" (DeRose 1999a, p. 197).

DeRose's examination of this case suggests, first of all, that successful WAMs involve cases in which, in one and the same context, both an assertion and its opposite appear

¹⁵ Here, P is to be kept in the indicative mood.

¹⁶ See DeRose 1999a, pp. 196-197, and for a fuller account DeRose 1991. I will not here be concerned with the details of his account of 'It's possible that P'. Instead, I mean to offer this example only as a tool for bringing out DeRose's conditions for successful WAMs.

false, and in which one of these appearances must be explained away. DeRose therefore suggests that

- (I) WAMs are successful only if they need to explain away only an appearance of falsity—in particular, an appearance of falsity that conflicts, in one and the same context, with another appearance of falsity.

But (I) is too restricted, for all that is needed for a successful WAM is two conflicting appearances, one of which must be explained away. The important thing is not that an assertion and its opposite appear *false*, but that an assertion and its opposite appear *to have the same truth value*: Even in cases in which both an assertion and its opposite appear true, we have conflicting appearances that admit of a WAM, for one of those appearances must be explained away. We must revise DeRose's first condition:

- (I') WAMs are successful only if they need to explain away *either* an appearance of falsity that conflicts, in one and the same context, with another appearance of falsity, *or* an appearance of truth that conflicts, in one and the same context, with another appearance of truth.

(I') better captures the truth about the sorts of intracontextual conflicts that facilitate WAMs.

DeRose introduces the second condition after noticing that his WAM appeals “to the generation of a *false implicature* to explain the appearance of falsehood” (DeRose 1999a, p. 198), where an *implicature* is something that is *conveyed* by the making of an assertion, but that's not a part of what the assertion itself *says*.¹⁷ Thus DeRose's second condition:

- (II) WAMs are successful only if they can provide the needed explanation(s) by appealing only to the generation of a false implicature (or false implicatures).¹⁸

¹⁷ See Grice 1989a, p. 24. Suppose, as Grice's famous example goes (see Grice 1989a, p. 33), that A is writing a letter on behalf of a candidate for a job in philosophy. The letter says, “Dear Sir, Mr. X's command of English is excellent, and his attendance at tutorials has been regular. Yours, etc.” A's assertions *convey the idea* that, according to A, Mr. X is no good at philosophy. Nevertheless, A's assertions do not themselves *say* that Mr. X is no good at philosophy.

¹⁸ See DeRose 1999a, p. 200.

DeRose suggests that this condition is in place because there's no way to explain away an appearance of *falsity* by appealing to the generation of a *true* implicature. There must be something false—here, the assertion's false implicature that Sue does not know the whereabouts of the book—that we mistake for the falsity of the assertion itself.

Yet we must revise (II) in light of the revisions we've made to (I), and thus in light of the fact that we might need to explain away an appearance of truth rather than an appearance of falsity. For example, my assertion 'Joe Lieberman is a Republican' is false, but in the right sort of context it can seem true. We can explain the seeming truth of my assertion by appealing to the fact that it has a true implicature, something like this: Lieberman, *qua* legislator, often thinks or acts like a Republican. We should therefore revise (II):

- (II') WAMs are successful only if they can provide the needed explanation(s) [either of apparent falsity or of apparent truth] by appealing only to the generation of the proper sort of implicature(s) [either a false implicature or a true implicature, respectively].

DeRose introduces the third condition after noticing that his WAM appeals to a very general rule of conversation when explaining the unwarrantedness of one of the apparently false assertions.¹⁹ This third condition, formulated so as to account for our revisions to (I) and (II), is:

- (III') WAMs are successful only if they appeal to *general* rules of conversation in explaining why certain apparently false assertions are unwarranted, or why certain apparently true assertions are warranted.

(III') is intended to prevent appeals to special or *ad hoc* rules, for example, rules intended to apply only to 'know' and its cognates.

¹⁹ According to DeRose, his WAM appeals to what he calls the 'Assert the Stronger' conversational rule, according to which "when you're in a position to assert either of two things, then, other things being equal, if you assert either of them, assert the stronger" (DeRose 1999a, p. 197).

Our challenge, then, is to defend the Moorean response to the skeptical puzzle with a WAM that satisfies (I'), (II') and (III'). In what follows, I argue that we can do so. I begin by showing how the Moorean's WAM satisfies DeRose's first condition.

2. Satisfying DeRose's First Condition

The Moorean's WAM satisfies (I') because it begins with an intracontextual conflict of intuitions, that is, with intuitions that conflict within one and the same context. Suppose that we find ourselves in a radical skeptical context, one that is characterized by the fact that it's governed by unusually high standards for knowledge. Suppose, too, that the participants in this context's conversation are cooperative—no one resists the skeptic's machinations, and everyone seriously entertains the skeptic's hypotheses. In such a context, we have the intuition that we *don't* know that we have hands, for example. Yet we also have the intuition that we *do* know: DeRose asks,

[W]hy do we find these claims to know [e.g., that I have hands] plausible even when we're in a context in which the skeptic has raised the standards to such a level that these claims are false? A little caution is in order here. It's controversial just how intuitively correct [I know that I have hands] ... seem[s] to us in such a context. Most of us feel some ambivalence [between knowing and not knowing] ... At any rate, [there is a] very strong pull that [I know that I have hands] continues to exert on (at least most of) us even when the standards are high. (DeRose 1999c, pp. 208-9)²⁰

DeRose himself takes it to be true in radical skeptical contexts not only that it seems that I don't know that I have hands, but also that the claim that I do know continues to exert a "very strong pull ... on (at least most of) us."

²⁰ Compare DeRose 1999c, p. 185.

To make this ambivalence palpable, consider the following example: You're sitting in a classroom surrounded by some of your fellow undergraduates, listening attentively to your philosophy instructor. She's busy explaining Descartes' skeptical arguments and even mentions what she takes to be a contemporary version of those arguments, namely, an argument based on the possibility that you are now a bodiless brain-in-a-vat. The instructor encourages you to entertain this possibility, and you indulge her. When you take this skeptical possibility seriously, you find that you seem not to know some of the things you ordinarily take yourself to know, and this causes it to appear that you don't know that you have hands, for example. After all, if you were a bodiless brain-in-a-vat, you would be having just the perceptual experiences that you're now having.

Still, doesn't it also seem that you *do* know that you have hands? When you compare your present situation to one in which you clearly seemed to know, you find that there is little difference. In fact, the only difference seems to be that you are now seriously contemplating the possibility that you are a bodiless brain-in-a-vat. When you look at your hands on this occasion, you see them as you've seen them on countless occasions when you clearly seemed to know that you have hands. You feel each hand with the other, and they feel just as they did on countless occasions when you clearly seemed to know that you have hands. Your fingers smell faintly of the banana you ate for lunch. You touch the desk in front of you; your hands do not pass through it, and you feel the cool contours of its surface. Given all of these similarities between your present situation and those in which you clearly seemed to know, how is it that it now seems that you *don't* know that you have hands? The perceptual experiences that you are now having are, after all, exactly similar to those you've had in countless situations in which you clearly seemed to know that you have hands. So, even as you're entertaining the brain-in-a-vat skeptical possibility, even as it seems that you don't know that you have hands, it nevertheless also seems that you do know. I submit, then, that

we—Mooreans, skeptics, and contextualists alike—should maintain that both of the following hold in high-standards contexts:

(4) It seems true to say that I *don't* know that I have hands.

(5) It seems true to say that I *do* know that I have hands.

Now, since it adopts (MI), the Moorean response will need to explain away the appearance in (4). But since this is an appearance of truth that conflicts, in one and the same context, with another appearance of truth, the Moorean response satisfies DeRose's first condition.²¹

3. The Moorean's Warranted Assertability Maneuver

Since the Moorean is concerned to defend (MI), she must explain why something that is true—namely, that you know that you have hands—seems false. We can provide this explanation by appealing to the warranted assertability of the claim that you don't know that you have hands. In particular, the Moorean can maintain that the warranted assertability of the claim that you don't know causes it to seem true that you don't know.

Now, if this explanation is to be satisfactory, we must answer two questions. Why are we warranted in asserting that you don't know that you have hands? And how does the warrantedness of this assertion cause it to seem true that you don't know that you have hands?

Let's turn to the first of these questions. Recently, Patrick Rysiew has provided an answer that relies on a distinction between salient and what I'll call relevant* epistemic alternatives. We can define the *relevant* epistemic position* in a certain conversational context as the position that is good enough to rule out all epistemically relevant* alternatives,

²¹ DeRose suggests that we should never want to explain away an appearance of truth. Rather, we'd like to preserve truths and to avoid multiplying falsehoods. But even if we should adopt this attitude in general, we can't always do so. For, in contexts like radical skeptical contexts, which include conflicting appearances of truth, we must maintain that one of the appearances is misleading, and that a proposition that seems true is in fact false.

where “only the *relevant*[*] alternatives affect what is literally expressed by a given knowledge attributing sentence” (Rysiew 2001, p. 490); while the *salient epistemic position* in a certain conversational context is the one that is good enough to rule out all of the salient alternatives, where salience is determined (at least in part) by the “doubts [that] are being entertained [in the context], what alternatives to *p* are ‘in the air’” and the like (Rysiew 2001, p. 490).²² Armed with this distinction, we can account for the fact that I know even in radical skeptical contexts that I have hands. For, according to the Moorean, what’s relevant* is *contextually invariant*. Thus, what’s relevant* in radical skeptical contexts is identical to what’s relevant* in ordinary contexts, namely, an epistemic position that is good enough to rule out all and only the epistemically relevant* alternatives. So, given that I occupy the relevant* epistemic position, I can know even in radical skeptical contexts that I have hands.

Nevertheless, we cannot warrantably assert in a radical skeptical context that I know that I have hands, for doing so implicates, or pragmatically imparts, that I occupy the *salient* epistemic position. Moreover, we *can* warrantably assert that I *don’t* know that I have hands. Why, though, is this the case? Rysiew says that

if, because of the importance of what’s at issue, the range of not-*p* possibilities under consideration—i.e., the *salient* alternatives—expands [so as to include alternatives such as those in which S is a brain-in-a-vat], this just means that were [we then to say, ‘S knows that he has hands’, we] would thereby *impart* that he can rule out *these*. And that might well be false (or might not be known by the speaker to be true). Whereas, by uttering (the falsehood), [‘S *doesn’t* know that he has hands’, we] can impart the important truth ... that his

²² I regret having to deviate from terminology that is both Rysiew’s and well established in the epistemological literature: I call *relevant** alternatives what Rysiew, among many others, calls *relevant* alternatives. I do this in order to avoid confusion later in the paper between epistemically relevant alternatives and relevant assertions (i.e., assertions made in accordance with Grice’s maxim of Relation, which commands, ‘Be relevant!’).

epistemic position isn't so good that he can rule out the sorts of counter-possibilities which are now ... salient. (Rysiew 2001, p. 490)

Still, what accounts for the fact that the assertion 'S doesn't know that he has hands' pragmatically imparts that S's epistemic position is not good enough to rule out the salient alternatives? And what accounts for the fact that the assertion 'S knows that he has hands' pragmatically imparts that his epistemic position *is* good enough to rule out the salient alternatives? Rysiew suggests that

given Grice's Cooperative Principle (CP),²³ not only can we explain, but we should expect, that [these] pragmatic phenomena ... will in general be generated by the makings of knowledge attributions. In particular, I see the maxim of Relation—to wit, "Be relevant!"—as plausibly underwriting [these] phenomena (Rysiew 2001, p. 491)²⁴

Thus, were we to assert 'S doesn't know that he has hands', we would implicate that S is not in the salient epistemic position. In this case, since our conversational partners assume in accordance with the maxim of Relation that our contribution has been made in light of the salient alternatives, they reason that we must intend to communicate that S cannot rule out those alternatives. Moreover, were we to assert '*S knows* that he has hands', we would implicate that S *is* in the salient epistemic position. For, when we make such an assertion, our conversational partners proceed on the assumption, in accordance with the maxim of Relation, that our conversational contribution has been made in light of the salient alternatives. Given this, our conversational partners reason that we must intend to communicate that S can rule out the salient error possibilities. Yet since S cannot rule out all

²³ Grice's Cooperative Principle says, "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 in which you are engaged" (Grice 1989a, p. 26).

²⁴ Jessica Brown, too, has recently suggested that the maxim of Relation underwrites the pragmatic phenomena to which Rysiew—and she—appeals (see Brown 2005 and Brown 2006).

of those possibilities, our conversational contribution violates the maxim of Relation, and so we are not warranted in asserting that S knows that he has hands.²⁵

Yet when we properly understand the maxim of Relation, we can see that our assertion ‘S knows that *p*’ would not violate Relation. This means that Rysiew has not succeeded in his attempt to explain why our true assertion would be unwarranted in radical skeptical contexts. He suggests that our true assertion violates Relation—that is, that it’s not relevant to the on-going conversation—if it would implicate the falsehood that S is in the salient epistemic position with respect to *p*. Properly understood, however, relevance requires only that *the knowledge attribution has a bearing on, or that it is connected with, the matter at hand.*²⁶ And a knowledge attribution can meet this condition even when it pragmatically imparts a falsehood. For example, in a radical skeptical context, I do not occupy the salient epistemic position. Nevertheless, since the point of our conversation is to determine whether I know that I have hands (or that I’m not a brain-in-a-vat), our assertion, ‘He knows that he has hands’, would clearly have a bearing on and a connection with the subject of our conversation. This becomes especially clear when we consider a case in which the relevant* alternatives are, say, *a*₁ and *a*₂, in which the salient alternatives are *a*₂ and *a*₃, and in which S can eliminate *a*₁ and *a*₂ but not *a*₃. Here, even though S doesn’t occupy the salient epistemic position, part of what would be suggested by our true assertion that S knows that *p* is that S can eliminate *a*₂, which is among the salient alternatives. It seems particularly clear in this case that our true assertion would have a bearing on and a connection with the subject of our conversation, in spite of the fact that S fails to occupy the salient epistemic position. Thus, our knowledge attribution in a radical skeptical context would indeed be

²⁵ See Rysiew 2001, p. 491. Brown 2006 maintains that what’s central in this sort of account is not the mentioning of error possibilities, but rather “whether the issue in question is practically important” (p. 422).

²⁶ Compare Grice 1989a’s characterization of relevance—relevance has to do with “subjects of conversation” (p. 27)—and Rysiew 2001’s own characterization—relevant assertions are “to the point” (p. 491). In addition, see Harnish 1976, p. 362; and O’Hair 1969, p. 45.

relevant and would not violate the maxim of Relation. This means that Rysiew hasn't yet explained how our true knowledge attributions would be unwarranted in radical skeptical contexts.

But Rysiew also suggests a different explanation.²⁷ Since the maxim of Relation governs our conversations, my assertion 'S knows that *p*' imparts the following information: S can eliminate all of the epistemically salient alternatives to *p*. But suppose that S *can't* eliminate all of those alternatives. In that case, my assertion 'S knows that *p*' would be unwarranted because it would violate a maxim of *Quality*, either

Do not say what you believe to be false, or

Do not say that for which you lack adequate evidence.²⁸

On this explanation, however, there will be cases in which we would like to provide a WAM, but in which we are prevented from doing so because we won't be able to satisfy condition (III'). Let me explain. In suggesting that the maxim of Relation works in this way with the maxims of Quality, Rysiew breaks the connection between an assertion's having a certain implicature (for which Relation is responsible) and its being either warranted or unwarranted (for which Quality is responsible): On this picture, the fact that an assertion is warranted has nothing to do with its having a true implicature, and the fact that an assertion is unwarranted has nothing to do with its having a false implicature. On Rysiew's suggestion, it is fairly easy to imagine a case in which my assertion complies with the maxim of Relation, has a false implicature and seems false; but in which it also complies with the maxim of Quality because I neither believe it to be false nor lack adequate evidence for it. Suppose, for example, that I am a Moorean who is theoretically informed and who asserts in a radical skeptical context that I know that I have hands. My assertion complies with the maxim of

²⁷ He made this suggestion clear to me in correspondence.

²⁸ See Grice 1989a, p. 27.

Relation (since it bears on the present conversation), has a false implicature (namely, that I can eliminate all the salient alternatives), and seems false. Nevertheless, it complies with the maxims of Quality, for I don't believe it to be false that I know that I have hands, and I don't lack evidence for my assertion. In such a case, we won't be able to perform a WAM that meets condition (III'), for my apparently false assertion is *not* unwarranted. This suggests that we should reject Rysiew's proposal, for it precludes the possibility of providing a successful WAM in some cases in which we'd very much like to do just that.

Our criticisms of Rysiew's proposal make it clear that any adequate proposal must meet the following two conditions:

- (C) It must understand the maxim of Relation as requiring only that assertions have a bearing on or a connection with the matter at hand.
- (D) It must not break the connection between an assertion's having a certain implicature and its being either warranted or unwarranted, for doing so will make it impossible to provide a WAM in some cases.

The resources for a proposal that meets these conditions are to be found neither in Relation nor in Quality, but rather in Grice's maxims of *Quantity* and, in particular, his second maxim of Quantity,

Q2: Do not make your contribution more informative than is required.²⁹

Some have maintained, however, that there's really no distinction between Q2 and Grice's maxim of Relation.³⁰ Granted, there are important similarities between the two. For instance, there are cases in which my conversational contribution is irrelevant precisely because it's more informative than is required. Suppose S asks, 'Is John taller than Fred?'

²⁹ See Grice 1989a, p. 26.

³⁰ Grice himself suggested that there might be no such distinction; see Grice 1989a, p. 27, where he says that "there is perhaps a different reason for doubt about the admission of [Q2], namely, that its effect will be secured by a later maxim, which concerns relevance." Horn 1984, Sperber and Wilson 1986, and Carston 1998 also maintain that the two are not distinct.

All else being equal, I seem to violate Q2 when I respond by saying, ‘Yes. John is taller than Fred, and it’s unseasonably cool in Los Angeles’.³¹ When I make my contribution more informative than it ought to be by saying that it’s unseasonably cool in Los Angeles, I provide irrelevant information. Nevertheless, Q2 is *different* from Relation in significant ways, ways that warrant a distinction between the two. For instance, there are cases—and I provide some examples below—in which my contribution is more informative than is required, but in which it does *not* violate the maxim of Relation. I will say, then, that contributions that violate both Q2 and Relation violate a *maxim of relevance*, and that contributions that violate Q2 without violating Relation violate a *maxim of strength*.

This sort of distinction accords with “our pretheoretical intuitions” (Levinson 2000, p. 380, n. 4), which suggest that a contribution is irrelevant only when it deviates, more or less dramatically, from the subject of the conversation. Stephen C. Levinson suggests that according to “our pretheoretical intuitions,”

relevance is about connectedness and collaborative activity ... [Construing relevance in this way] is also along the lines that Grice intended: the typing of relevance to particular plans and goals is very clear in Grice 1989[b] (p. 222) ... (Levinson 2000, p. 380, n. 4.)

In fact, Rysiew himself suggests that a conversational contribution is relevant when it is “to the point” (Rysiew 2001, p. 491). Thus, our pretheoretical intuitions recommend that we construe relevant assertions as those that are connected with the current subject of conversation, which suggests that we should distinguish maxims of strength from maxims of relevance.

Second, certain cases and prominent examples suggest that there is such a distinction. I’m inclined to think that radical skeptical contexts bring the distinction to the fore—in such

³¹ This example is very similar to one provided by van Rooy 2004, p. 16.

contexts, our knowledge attributions would provide an excess of relevant information, thus violating a strength maxim rather than a relevance maxim. Moreover, prominent examples in which the maxim of Relation is either observed or violated suggest that an assertion is irrelevant when it deviates, more or less dramatically, from the subject of the conversation. Grice himself provides such examples. Suppose that A says, ‘I am out of petrol’, and that B responds, ‘There is a garage round the corner’.³² According to Grice, B’s statement complies with the maxim of Relation and is therefore appropriate. Moreover, it seems to comply with that maxim because, in the first place, it can be taken by A to be topically related to his question.³³ Here’s another example: Suppose that A says, ‘Mrs. X is an old bag’, and that B responds, ‘The weather has been quite delightful this summer, hasn’t it?’ Of B’s response, Grice says that “B has blatantly refused to make what he says relevant to A’s preceding remark” (Grice 1989a, p. 35). And B’s response seems irrelevant, in the first place, because it is not at all topically related to A’s assertion—B’s response deviates quite dramatically from the current subject of the conversation. These examples suggest that relevant assertions are those that are connected with the current topic of conversation, and so they support a distinction between maxims of strength and maxims of relevance.³⁴ Failing to recognize such a distinction compels us to turn our backs on a distinction that is firmly rooted in intuition and in theory, and it compels us to accept the claim that any information that exceeds the threshold of informativeness must count as irrelevant. It seems to me, however, that we should do neither of these things. I suggest, then, that we cling to and utilize the differences between maxims of relevance and maxims of strength.

³² See Grice 1989a, p. 32.

³³ The same sort of thing seems to be happening in the following dialogue: A: ‘Can you tell me the time?’; B: ‘Well, the milkman has come’ (see Levinson 1983, pp. 97-8).

³⁴ Several *theories* also accept and employ a distinction between relevance and strength. See, for example, O’Hair 1969, p. 45; Harnish 1976, p. 362; Hirschberg 1985, p. 65; Welker 1994, p. 31; and Levinson 2000, page 74 and page 380, n. 4.

One way to utilize a maxim of strength is in defense of the Moorean response. This is good news, for, as we have seen, we aren't able to defend that response by appealing to Relation alone or by appealing to Relation working in conjunction with Quality. I now propose the following Moorean WAM, which appeals to the Q2 maxim of strength. In radical skeptical contexts, the conversation between us and our interlocutor concerns possibilities such as those in which I am a brain-in-a-vat. Further, in such contexts, our interlocutor reminds us that if I were a bodiless brain-in-a-vat, I would be having just the perceptual experiences that I'm now having. Given the mentioned error possibility, what is salient to the conversation is an epistemic position that is good enough to rule out even unlikely possibilities, such as the possibility that I am a bodiless brain-in-a-vat. And given that our conversation is governed by Relation, if we were to reply to our interlocutor by asserting 'He knows that he has hands', we would implicate that I occupy the salient epistemic position, a position, moreover, that I do not occupy. Given this, our assertion would be too strong and hence would violate a maxim of strength. Our assertion, then, even though it would be true, would be unwarranted.

Our assertion in a radical skeptical context that I know that I have hands is *stronger* than other assertions in the sense that it is *more informative* than other assertions,³⁵ where all of the information on this scale of informativeness is relevant to the on-going conversation. Regarding informativeness on this scale, we may say that "*p is more informative than q* if the set of states of affairs that *q* rules out is a proper subset of the set that *p* rules out" (Levinson 2000, p. 115). Were we to assert in a radical skeptical context that I know that I have hands, we would convey the information that I can eliminate even the salient alternatives to my having hands. We would not convey this information, however, if we made a different

³⁵ This is Levinson 1983's conception of strength. Compare Leech's narrower conception of strength (see Leech 1983, p. 85).

assertion, for example, the assertion that I don't know that (or whether) I have hands. In this case, we would convey the information that I *cannot* eliminate all of the salient alternatives. The former assertion—that I know that I have hands—is categorical, conveying the information that I can eliminate possibilities that would be left open by the latter, more guarded assertion. Thus, the assertion that I know that I have hands is more informative—and hence stronger—than the other assertion.

Yet why would our assertion in radical skeptical contexts be *too* strong? It would be too strong precisely because it implicates that I am in a position to eliminate even the salient error possibilities, a position that I do not in fact occupy, and because I *am* in a position to make a weaker assertion, one that does not carry such a strong implicature. Although asserting that I know that I have hands would be relevant to the on-going conversation, it would convey information that exceeds the threshold of informativeness, and thus it would violate a maxim of strength.³⁶ This also makes it clear that unlike Rysiew's proposal, our account meets condition (D) above: Since it maintains that an assertion's false implicature is directly responsible for its being unwarranted, our account does not break the connection between an assertion's having a certain implicature and its being either warranted or unwarranted.

But what about our assertion that I *don't* know that I have hands? This assertion, in accordance with the maxim of strength, is warranted in radical skeptical contexts. Once again, in such contexts, what is salient is a position that's good enough to eliminate even unlikely error possibilities. And given that our conversation is governed by Relation, the assertion 'He doesn't know that he has hands' pragmatically imparts the truth that I cannot eliminate all of the salient possibilities. Moreover, since I am in fact in no position to

³⁶ We don't have the option here of saying that our assertion is unwarranted simply because it's false, for we have set aside accounts of warranted assertability according to which an assertion is warranted only if it is true.

eliminate all of those possibilities, our assertion is made in accordance with the appropriate maxim of strength—it does not exceed the threshold of informativeness. This, too, demonstrates that our account meets condition (D), for it maintains that an assertion’s true implicature is directly responsible for its being warranted.³⁷

Let’s now turn to the second question that the Moorean must answer—how does the warrantedness of our assertion that I don’t know that I have hands create the appearance that it’s true that I don’t know? In accordance with the Gricean maxims of strength, we are warranted in asserting in radical skeptical contexts that I don’t know that I have hands. This warranted assertion generates the true implicature that I cannot eliminate all of the salient alternatives to my having hands. In turn, this true implicature, in conjunction with the fact that a more elevated epistemic position is salient in radical skeptical contexts,³⁸ can create the appearance—and lead us mistakenly to suppose—that it’s true that I don’t know that I have hands.

At this point, one might wonder whether the claim that S can eliminate all of the salient alternatives to *p* is in fact an *implicature* of the assertion that S knows that *p*. If it is,

³⁷ Hawthorne 2001 suggests that we can contextualize the concept of freedom—or, at least, a concept of freedom—in the same way that Lewis 1996 contextualizes the concept of knowledge. But we can defend a “Moorean” view of freedom with an analogue of the WAM proposed above and, if nothing else, this should immunize the proposed WAM against the charge of being *ad hoc*. Here, all too briefly, is how a warranted-assertability defense of a “Moorean” view of freedom might go: In contexts of wide philosophical reflection, the conversation between us and our interlocutors concerns considerations having to do with “the neurological springs of [S’s] action and then, indeed, [with the fact] that [S’s] actions are the result of the operations of the laws of nature over which [she] had no control, coupled with the distribution of microparticles in the distant past over which [she] had no control” (Hawthorne 2001, p. 67). Further, in such contexts, our interlocutors remind us that given this causal explainer of S’s actions, she would act just as she does now. What is salient to this conversation is a position that is good enough to rule out the possibility that the mentioned causal explainer provides an adequate causal explanation of S’s actions. And given that our conversation is governed by Relation, were we to reply to our interlocutors by asserting ‘S doesn’t freely ϕ ’, we would truly implicate that S does not occupy the salient position. Moreover, our assertion ‘S doesn’t freely ϕ ’ would be made in accordance with the appropriate maxim of strength: Thinking of the chain of causes as continuous and linear, extending from the efficient causes of S’s actions to connected causes that are ever more distant, our assertion ‘S doesn’t freely ϕ ’ does not implicate that S occupies a position that is good enough to rule out possibilities that are too far out on this chain and hence it does not exceed the threshold of informativeness. Thus, our assertion, even though it would be false, would be warranted and could therefore seem true.

³⁸ We must make reference here to some set of epistemic standards. These standards can be seen as rules that are necessary in order to “work out” the presence of a conversational implicature (and, according to Grice, “[t]he presence of a conversational implicature must be capable of being worked out” (Grice 1989a, p. 31)).

then it ought to be cancelable. Otherwise, the alleged implicature will look like a semantic entailment rather than a pragmatic implicature. Fortunately, the implicature here is cancelable. When we assert that S knows that *p*, our assertion has the implicature that S can eliminate all of the salient alternatives to *p*. We can cancel this implicature by saying, for example, that S knows that *p* but cannot eliminate all of the salient alternatives to *p*. I might say, for example, “I know that the bank will be open on Saturday, but that doesn’t mean that I can rule out all of the salient alternatives to its being open on Saturday. In order to know, I must rule out only the *relevant* alternatives, and it is not a relevant alternative in this case that between today and Saturday, the bank will be audited and then closed by the FDIC.”³⁹ Since the implicatures here are cancelable, the Moorean’s proposal avoids a strong objection.⁴⁰

We’ve now seen that the Moorean’s proposed WAM satisfies all three of DeRose’s conditions. First, it needs to explain away an appearance of truth—the apparent truth of the assertion that I don’t know that I have hands—that would conflict, in one and the same context, with another appearance of truth—the apparent truth of the assertion that I *do* know that I have hands. Second, the Moorean’s WAM explains away this appearance of truth by appealing only to the generation of true implicatures, for example, the true implicature that I cannot eliminate all of the salient alternatives to my having hands. Finally, in explaining why our assertion that I don’t know that I have hands is warranted, the Moorean appeals to general conversational rules, namely, Relation and Quantity. The Moorean response therefore lives up to the challenge: We can defend it with a WAM that meets all three of DeRose’s conditions.⁴¹

³⁹ ‘Relevant’ here is used to mean the same thing as ‘relevant*’.

⁴⁰ Our assertion that S *doesn’t* know that *p* has the implicature that S cannot eliminate all of the salient alternatives to *p*. We can cancel this implicature by maintaining that although S doesn’t know that *p* (because she cannot eliminate all of the relevant* alternatives to *p*), she nevertheless *can* eliminate all of the salient alternatives. See Black and Murphy 2005 for more on the cancelability of alleged semantic entailments.

⁴¹ Leite 2005 objects to WAMs like the one proposed above. He says that in a high-standards context, it must be the case that “[t]he 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

4. *The Moorean versus the Contextualist*

Since the Moorean response can provide all the explanations it must provide, we are well within our rights not only to claim that the Moorean pragmatic strategy is a viable alternative to the semantic strategy as it's employed by, e.g., epistemological contextualists, but also to report that rather than trailing contextualism, Mooreanism is running neck and neck with it. Is there anything that gives Mooreanism the advantage? Note here that since it seems true in high-standards contexts both that I don't know that I have hands and that I do know, the contextualist too must explain away an appearance of truth. And if his response is to be different from the Moorean's, he must explain away the seeming truth of the claim that I know that I have hands. In providing this explanation, the contextualist might perform a WAM of his own. I argue, though, that the contextualist will find it difficult to perform a satisfactory WAM.⁴²

If the contextualist's WAM is to be successful, it must satisfy DeRose's three conditions. To satisfy (III'), the contextualist must specify a general conversational rule

'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" (Leite 2005, pp. 219-220). This is in fact pretty bizarre, and so it's good that we can perform a WAM without committing ourselves to this sort of story. Those who would perform a WAM in a case in which the speaker is moved to announce in a high-standards context that he doesn't know that *p* need not assume that the speaker believes that he knows that *p*, and this is true especially if, as I propose, we perform a WAM without appealing to a maxim of Quality. Indeed, on the WAM proposed above, given that this is a *high-standards* context, it's safe to assume that the speaker takes himself *not* to know that *p*. Thus, since it's a mistake to assume here that the speaker believes that he knows, Leite's criticism does not hold up.

⁴² Of course, this is not to say that the contextualist has no choice but to perform a WAM. The contextualist might appeal to semantic blindness, for example, in explaining away the apparent truth of the assertion that I know that I have hands. Contextualists usually appeal to semantic blindness in order to explain why we have the powerful intuition in some cases that "two speakers [one of whom asserts that she knows, and the other of whom asserts, in the very same context, that she doesn't know] are contradicting one another" (DeRose 2006, p. 318). The contextualist might say here that the extent to which it seems true in radical skeptical contexts that I know that I have hands is just the extent to which I am blind to the contextualist semantics of 'know'; it seems true that I know that I have hands only because we are to some extent semantically blind. Yet this sort of explanation does nothing to push contextualism ahead of Mooreanism, for the Moorean can also appeal to semantic blindness: The extent to which it seems true in radical skeptical contexts that I *don't* know that I have hands is just the extent to which I am blind to the *Moorean invariantist* semantics of 'know'; it seems true that I don't know that I have hands only because we are to some extent semantically blind (see DeRose 2006, Section 2).

according to which we are warranted in asserting in high-standards contexts that I don't know that I have hands, but not warranted in asserting that I do know. I want simply to grant that the contextualist will be able to specify some rule that will satisfy (III'). Furthermore, so as not to place undue restrictions on the contextualist, we should allow him to maintain that either elevated or relaxed epistemic standards are salient in radical skeptical contexts. Given this measure of liberty and the fact that the contextualist has specified a suitable conversational rule, there are four ways in which the contextualist might attempt to explain away the apparent truth of the claim that I know that I have hands. These four ways involve

- the salience of elevated standards and the assertion that I know that I have hands,
- the salience of relaxed standards and the assertion that I don't know that I have hands,
- the salience of elevated standards and the assertion that I don't know that I have hands, and, finally,
- the salience of relaxed standards and the assertion that I know that I have hands.

None of these is successful. Thus, given the plausibility of the Moorean WAM,⁴³ we have good reason to favor the Moorean pragmatic strategy over the contextualist's semantic strategy.

The contextualist's first attempt to provide the needed explanation results in a WAM that fails to meet (II'), which requires that WAMs explain away apparent truths by appealing only to the generation of true implicatures. Suppose that I assert in some radical skeptical context that I know that I have hands. This assertion generates the implicature that I can eliminate all of the salient alternatives to my having hands. But given that *elevated* standards are salient, this implicature is *false*.

⁴³ And given that the Moorean and the contextualist can appeal to semantic blindness with equal effectiveness.

The contextualist's second attempt also fails to satisfy (II'). Suppose that I assert in some radical skeptical context that I don't know that I have hands. This assertion generates the implicature that I can't eliminate all of the salient alternatives to my having hands. But given that relaxed standards are salient, this implicature is *false*.

Third, suppose that I assert in some radical skeptical context that I don't know that I have hands. Given that elevated standards are salient, this assertion produces the *true* implicature that I cannot eliminate all of the salient alternatives to my having hands. Yet this true implicature creates the appearance that it is *true* that I *don't* know, or, perhaps, the appearance that it is *false* that I *do* know. In either case, this WAM cannot be used to explain why it seems *true* that I *do* know that I have hands.

Finally, suppose that I assert in some radical skeptical context that I know that I have hands. Given that relaxed standards are salient, this assertion generates the true implicature that I can eliminate all of the salient alternatives to my having hands. And this is the sort of implicature that might create the appearance that it is true that I know that I have hands.

This final attempt to explain away the apparent truth of the claim that I know that I have hands has the makings of a plausible WAM. Nevertheless, we should reject this WAM, for it is much less realistic than the Moorean's. First of all, it asks us to believe that I would assert in a *radical skeptical* context that I *do* know that I have hands. It is quite unlikely, however, that I would make such an assertion. It seems much more likely that, as the Moorean maintains, I would assert in such a context that I *don't* know that I have hands. Furthermore, the contextualist's WAM suggests that *relaxed* standards are salient in radical skeptical contexts. It seems, however, that this is straightforwardly false. Radical skeptical contexts are characterized by the salience in those contexts of radical skeptical alternatives. And to make salient a radical skeptical alternative is to make salient the idea that such an

alternative must be eliminated if we are to know, and therefore to elevate the epistemic standards.

All of this suggests that the contextualist will find it difficult to perform a WAM in explaining away the apparent truth of the claim that I know that I have hands. Moreover, we have seen that we *can* perform a successful WAM, one that meets all three of DeRose's conditions, in defense of a Moorean response to skepticism, a response that allows us to maintain that we know all that we ordinarily take ourselves to know. This gives us ample reason to prefer the Moorean pragmatic strategy over the contextualist's semantic strategy.

ACKNOWLEDGEMENTS

For their helpful comments and criticisms, I thank Bryan Belknap, Jessica Brown, Albert Casullo, Mark Cullison, Mylan Engel, Jr., Richard Greene, Peter Murphy, Duncan Pritchard, Ray VanArragon, and especially Patrick Rysiew. I also thank audiences at California State University, Northridge; the University of Utah; the October 2002 meetings of the Central States Philosophical Association; the October 2002 Mountain-Plains Philosophy Conference; the March 2003 meetings of the Society for Skeptical Studies (held in conjunction with the American Philosophical Association's Pacific Division meetings); and the April 2003 meetings of the APA's Central Division. Work on this project was supported by a California State University, Northridge Probationary Faculty Development Grant.

REFERENCES

- Bach, K. (2005). The emperor's new 'knows'. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: On epistemology, language and truth*. Oxford: Oxford University Press.
- Black, T. (2002). A Moorean response to brain-in-a-vat scepticism. *Australasian Journal of*

- Philosophy, 80, 148-163.
- Black, T. (2005). Classic invariantism, relevance, and warranted assertability manœuvres. *The Philosophical Quarterly*, 55, 328-336.
- Black, T. (2008). Defending a sensitive neo-Moorean invariantism. In V. F. Hendricks & D. Pritchard (Eds.), *New waves in epistemology* (pp. 8-27). Palgrave Macmillan.
- Black, T. & Murphy, P. (2005). Avoiding the dogmatic commitments of contextualism. *Grazer Philosophische Studien*, 69, 165-182.
- Blaauw, M. (2003). WAMming away at contextualism. *Sats – Nordic Journal of Philosophy*, 4, 88-97.
- Brown, J. (2005). Adapt or die: The death of invariantism? *The Philosophical Quarterly*, 55, 263-285.
- Brown, J. (2006). Contextualism and warranted assertability manoeuvres. *Philosophical Studies*, 130, 407-435.
- Carston, R. (1998). Informativeness, relevance and scalar implicature. In R. Carston & S. Uchida (Eds.), *Relevance theory: Applications and implications*. Amsterdam: John Benjamins.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2, 91-123.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57-89.
- DeRose, K. (1991). Epistemic possibilities. *Philosophical Review*, 100, 581-605.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913-929.
- DeRose, K. (1996). Knowledge, assertion, and lotteries. *Australasian Journal of Philosophy*, 74, 568-580.
- DeRose, K. (1999a). Contextualism: An explanation and defense. In J. Greco & E. Sosa

- (Eds.), *The Blackwell Guide to Epistemology*. Oxford: Blackwell.
- DeRose, K. (1999b). Now you know it, now you don't. In *Proceedings of the Twentieth World Congress of Philosophy: Volume V, Epistemology*. Bowling Green, OH: Philosophy Documentation Center.
- DeRose, K. (1999c). Solving the skeptical problem. In K. DeRose & T. A. Warfield (Eds.), *Skepticism: A contemporary reader*. New York: Oxford University Press.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review*, 111, 167-203.
- DeRose, K. (2006). "Bamboozled by our own words": Semantic blindness and some arguments against contextualism. *Philosophy and Phenomenological Research*, 73, 316-338.
- Grice, H. P. (1989a). Logic and conversation. In H. P. Grice, *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Grice, H. P. (1989b). Meaning. In H. P. Grice, *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Harnish, R. M. (1976). Logical form and implicature. In T. G. Bever, J. J. Katz & D. T. Langendoen (Eds.), *An integrated theory of linguistic ability*. New York: Thomas Y. Crowell.
- Hawthorne, J. (2001). Freedom in context. *Philosophical Studies*, 104, 63-79.
- Hirschberg, J. (1985). A theory of scalar implicature. Dissertation, University of Pennsylvania.
- Heller, M. (1999). Relevant alternatives and closure. *Australasian Journal of Philosophy*, 77, 196-208.
- Horn, L. R. (1984). Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (Ed.), *Meaning, form, and use in context: Linguistic applications*. Washington, D. C.: Georgetown University Press.

- Leech, G. N. (1983). *Principles of pragmatics*. New York: Longman.
- Leite, A. (2005). Some worries for would-be *WAM*mers. *Grazer Philosophische Studien*, 69, 101-125.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge, MA: MIT Press.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549-567.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
- O'Hair, S. G. (1969). Implications and meaning. *Theoria*, 35, 38-54.
- Pritchard, D. (2005). Contextualism, scepticism and warranted assertibility manoeuvres. In J. Keim-Campbell, M. O'Rourke & H. Silverstein (Eds.), *Knowledge and skepticism*, Cambridge, MA: MIT Press.
- Rieber, S. (1998). Skepticism and contrastive explanation. *Noûs*, 32, 189-204.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35, 477-514.
- Sperber, D. & Wilson, D. (1986). *Relevance: Communication and cognition*. Oxford: Basil Blackwell.
- Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford: Oxford University Press.
- Van Rooy, R. (2004). *Relevance implicatures*. Retrieved November 3, 2007, from <http://citeseer.ist.psu.edu/545416.html>.
- Welker, K. (1994). *Plans in the common ground: Toward a generative account of implicature*. Dissertation, The Ohio State University.
- Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, 105, 489-523.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.