

## A WARRANTED-ASSERTABILITY DEFENSE OF A MOOREAN RESPONSE TO SKEPTICISM

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 (other) indexicals like 'I' and 'here'. Moreover, since context affects the semantic content of 'know', context will have an effect on the semantic content of complex lexical items in which 'know' appears, for example, on the semantic content of knowledge attributions. Epistemological contextualists, who employ the semantic strategy in addressing conflicts like those between (1), (2) and (3), have put the point 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)<sup>1</sup>

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.<sup>2</sup> Next, (3) is true 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. Yet, finally, (1) is true as well 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. Since I fail to meet the epistemic standards that are in place in such contexts, I truly impute to myself a lack of knowledge.<sup>3</sup> According to semantic strategists, then, (1), (2) and (3) do not in fact conflict, even though it seems that they do.<sup>4</sup>

---

<sup>1</sup> Compare Cohen, 1999, p. 57; and Rieber, 1998, p. 190. DeRose, along with other epistemological contextualists, is primarily concerned with the truth conditions of knowledge-ascribing and knowledge-denying sentences, sentences such as ‘Keith knows that he has hands’ and ‘Stew doesn’t know that he’s not a brain-in-a-vat’ (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). DeRose seems to believe that such a semantic descent is permissible as long as we remember that our primary concern is with the truth conditions of certain sentences. So, like DeRose and others, I often talk of the standards for knowledge rather than of the truth conditions of knowledge-ascribing and knowledge-denying sentences.

<sup>2</sup> See DeRose, 1999c, especially pages 29-33.

<sup>3</sup> 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. Each of these contextualisms seeks to establish that we need not support the outright denial of either (1), (2) or (3). 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.

<sup>4</sup> 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.

The contextualist’s resolution of the apparent conflict between (1), (2) and (3) is quite similar indeed 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\*). Once we know this, it is not at all difficult to see that there’s no conflict between (1\*), (2\*) and (3\*).

Yet we can also 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.<sup>5</sup> 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.

According to the *skeptical* pragmatic strategy, the semantic content of ‘know’ is such that the standards for knowledge are invariantly quite high. On this view, it is very difficult, if not impossible, to meet those standards. Skeptical invariantists accordingly maintain that we never (or only very rarely) know the things we ordinarily take ourselves to know. They must therefore explain why, even though it’s false, it seems true that I know that I have hands.

I favor a second pragmatic strategy, which I call the *Moorean* pragmatic strategy. According to this strategy, the semantic content of ‘know’ is such that the standards for knowledge are invariantly comparatively low. The Moorean invariantist accordingly maintains that we can know across contexts the things we ordinarily take ourselves to know.<sup>6</sup> If the Moorean pragmatic strategy is to be successful, then, 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

---

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

<sup>6</sup> I do not mean to suggest that *Moore* was (what I am calling) a Moorean invariantist, for there are differences as well as similarities between Moore’s claims and those of the Moorean invariantist. The main—and perhaps the only—similarity between the two is this: both maintain that the standards for knowledge are invariant and that we know across contexts the things we ordinarily take ourselves to know. The Moorean invariantist also maintains that the standards for knowledge are invariantly *low*, but Moore himself might not have agreed with this.

*don't know?*

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,<sup>7</sup> 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, it can in some contexts seem true that we don't know that we have hands.<sup>8</sup> In particular, the Moorean can maintain that since we are in some contexts warranted in asserting that we don't know, it can seem true that we don't know. As Patrick Rysiew has recently noted, however,<sup>9</sup> those who wish to respond to skeptical puzzles in this way have not done as much as they should have done. He says that

sophisticated (nonsceptical) invariantists [i.e. Mooreans] have heretofore rested content with “a ‘bare warranted assertability maneuver’” [DeRose, 1999a, p. 199]; and insofar as such would-be sophisticated invariantists have denied that context affects the truth conditions (versus the ‘warranted assertability’ conditions) of knowledge-attributing sentences and simply left it at that, their sophistication has been mere pretense. To date, no one has offered a compelling argument showing that a nonsceptical invariantism really is compatible with our manifestly flexible knowledge-attributing practices. (2001, p. 479)

---

<sup>7</sup> See <Author's Paper A> and <Author's Paper C>.

<sup>8</sup> The Moorean need not employ the notion of warranted assertability in providing this explanation. See Bach, Forthcoming for an explanation that appeals to an error theory.

<sup>9</sup> In Rysiew, 2001.

So far, as Rysiew here suggests, only skeptical invariantists have employed the notion of warranted assertability.<sup>10</sup> Moreover, Keith DeRose, who is perhaps today's foremost epistemological contextualist, persuasively argues that one very important skeptical invariantism—Peter Unger's—uses warranted assertability improperly, and that we should reject Unger's view on those grounds.<sup>11</sup> DeRose goes on to suggest that Moorean invariantisms must suffer the same fate, for they too will make improper use of the notion of warranted assertability.<sup>12</sup> 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—if not the chief reason—for preferring the semantic strategy, namely, that there is no viable alternative to it.

What, however, 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.<sup>13</sup> I hope to show, then, that we can defend a Moorean invariantist response to skepticism with a warranted assertability maneuver that meets each of DeRose's conditions. This will do two things for the Moorean. First, even if DeRose has not provided a complete and correct set of conditions, it seems clear that those conditions are met by some obviously legitimate uses of the notion of warranted assertability. So, in showing that her warranted assertability maneuver meets each of DeRose's conditions, the Moorean will have shown that her use of that notion measures up, at least in this respect, to some obviously legitimate uses. Second, in showing that her warranted

---

<sup>10</sup> Among the most influential warranted-assertability defenses of skepticism are those proposed by Unger, 1975 and by Stroud, 1984. For a recent defense of this sort, see Stone, 2000.

<sup>11</sup> See DeRose, 1999c and, especially, DeRose, 1999a.

<sup>12</sup> See DeRose, 1999c, p. 210; and DeRose, 1999a, p. 203.

<sup>13</sup> In DeRose, 1999a.

assertability maneuver meets each of DeRose's conditions, the Moorean will have responded to skepticism on her staunchest critic's own terms. This greatly restricts the replies that are available to those who would criticize the Moorean's response.

Here, then, is the plan: In Section 1, 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 the last two of his conditions. We can therefore defend the Moorean response with a legitimate use of the notion of warranted assertability. This supports the claim that the Moorean pragmatic strategy is a viable alternative to the semantic strategy, thus eliminating one of the chief reasons for preferring the semantic strategy. In Section 4, I demonstrate the value of the Moorean position: it can explain our intuitions, both epistemic and otherwise, in a variety of contexts. Finally, in Section 5, I provide additional arguments for the claim that we should favor the Moorean 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 invariantist 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*—as well as 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.<sup>14</sup>

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, and therefore no contexts in which the standards are so high as to keep me from knowing that I have hands. Yet if there are no such contexts, why do we judge in some contexts that it's true that I don't know that I have hands? On behalf of the Moorean, I respond to this question by arguing that even though there are no contexts in which the standards keep us from knowing that we have hands, there are nevertheless some 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.<sup>15</sup> 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.<sup>16</sup> His suggestion is that successful WAMs must meet three conditions, and that 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':<sup>17</sup>

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

---

<sup>14</sup> As it is, (MI) is fairly vague, and there are different ways of specifying its comparatively low standards for knowledge. I prefer 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. Thus, the Moorean's standards for knowledge are comparatively low because she requires 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.

<sup>15</sup> In the same way, we can explain why we judge that I don't know that I'm not a brain-in-a-vat.

<sup>16</sup> See DeRose, 1999a, pp. 201-203.

<sup>17</sup> Here, P is to be kept in the indicative mood.

knows that P is false.<sup>18</sup>

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). And 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 false. Thus, successful WAMs begin with two conflicting appearances, both of which are appearances of falsity. One of these appearances must be explained away, and no matter which one it turns out to be, it's sure to be an appearance of falsity. DeRose therefore suggests that

---

<sup>18</sup> 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. I leave the details of his account, as well as potential criticisms of it, for another time.

- (I) WAMs are successful only if they need to explain away *only* an appearance of falsity—and, 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, then, 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.<sup>19</sup>

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*, according to H. P. Grice, is something that is *conveyed* by the

---

<sup>19</sup> DeRose’s criticism of Unger’s WAM rests primarily on the claim that it fails to meet this first condition. Unger’s WAM must explain away not only “an appearance of falsity” but also “an appearance of truth” (DeRose, 1999a, p. 201). DeRose claims, quite rightly, that it seems appropriate and true in ordinary low-standards contexts to say that I know that I have hands, and wrong and false to deny that I know (see DeRose, 1999a, p. 201). Unger cannot escape these appearances. However, since Unger maintains that the standards for knowledge are both high and invariant, he is committed to the claim that I do *not* know in ordinary contexts that I have hands. Unger is therefore committed to each of the following claims:

- (i) I do *not* know in ordinary contexts that I have hands.
- (ii) It seems *false* in ordinary contexts that I do *not* know that I have hands.
- (iii) It seems *true* in ordinary contexts that I *do* know that I have hands.

Since Unger maintains that I don’t know that I have hands, he must explain away both the apparent falsity of the assertion that I don’t know and the apparent truth of the assertion that I do know. But these two appearances do *not* conflict; they are perfectly compatible. Thus, Unger’s WAM fails to meet DeRose’s first condition.

making of an assertion, but that's not a part of what the assertion itself *says*.<sup>20</sup> 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).<sup>21</sup>

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, when I notice a particularly long swimming pool, I might say, "That swimming pool is a mile long!" My assertion here is literally false, but it seems (or it can in the proper context seem) true. In cases like this, it seems that there must be a true implicature—here, the assertion's true implicature that the pool is remarkably long—that is responsible for the apparent truth of the assertion itself. 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].

---

<sup>20</sup> 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.

<sup>21</sup> See DeRose, 1999a, p. 200.

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.<sup>22</sup> This suggests a third condition, formulated so as to account for our revisions to (I) and (II):

- (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.

This third condition is intended to prevent appeals to special or *ad hoc* rules, for example, rules intended to apply only to ‘know’ and its cognates.

The challenge for the Moorean, then, is to defend her response to the skeptical puzzle with a WAM that meets all three of DeRose’s conditions. In what follows, I argue that she can do so. This, of course, speaks in favor of the Moorean response. I begin by showing how the Moorean’s WAM satisfies DeRose’s first condition.

## ***2. Satisfying DeRose’s First Condition***

Does the Moorean’s WAM begin with an intracontextual conflict of intuitions, i.e., intuitions that conflict within one and the same context? Indeed it does. Suppose that we find ourselves in a radical skeptical context, one that is characterized by the fact that it’s governed by unusually high (i.e., skeptical) 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

---

<sup>22</sup> According to DeRose, one of the apparently false assertions is unwarranted because it violates 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).

that we have hands (for example), but we also have the intuition that we *do* know. Speaking of such a context, DeRose himself says,

(Not-C) [the claim that I know that I have hands] seems plausible because it's true when evaluated at the standards most normally applied to it. But, it will be asked, why do we find these claims to know 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 (not-C) does seem to us in such a context. Most of us feel some ambivalence [between (not-C) and (C)] ... At any rate, [there is a] very strong pull that (not-C) continues to exert on (at least most of) us even when the standards are high ... (DeRose, 1999c, pp. 208-9)<sup>23</sup>

In spite of DeRose's plea for caution here, it seems evident that he 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" (DeRose, 1999c, pp. 209). Following DeRose, I'll call this our *ambivalence* in radical skeptical contexts.

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

---

<sup>23</sup> Compare DeRose, 1999c, p. 185.

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 it seemed clear that you know, you find that there is little difference. In fact, the only difference seems to be that you are now entertaining 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 it seemed clear that you know that you have hands. You feel each hand with the other, and they feel just as they did on countless occasions when it seemed clear that you 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 it seemed clear that you 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 it seemed clear that you 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 the Moorean adopts (MI), she'll 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, she satisfies DeRose's first condition.<sup>24</sup>

### **3. *The Moorean's Warranted Assertability Maneuver***

Since the Moorean is concerned to defend (MI), she must explain why something that is true on her account—namely, that you know that you have hands—seems false. She can provide this explanation by appealing to the warranted assertability of the claim that you don't know that you have hands.<sup>25</sup> 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, the Moorean must answer two questions. First, why are we warranted in asserting that you don't know that you have hands? Second, how does the warrantedness of the assertion that you don't know that you have hands create the appearance that it's true that you don't know?

Let's turn to the first of these questions. Recently, Rysiew has provided an answer that relies on a distinction between salient epistemic alternatives and what I'll call germane epistemic

---

<sup>24</sup> DeRose suggests that we should (or that we might) 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. Given this, the contextualist too must explain away an appearance of truth here, namely, the appearance in (5). I argue in *<Author's Paper D>* that the contextualist cannot provide this explanation. His explanation must appeal either to features of this particular context itself or to features of some other context. In either case, however, his explanation violates principles to which he himself subscribes. This dilemma, it seems to me, poses a very serious problem for contextualism.

<sup>25</sup> This suggests that we can sometimes be warranted in making false assertions. This should not be a problem, however, since there *are* false assertions that are also warranted. Suppose, for example, that you invoke a controversial theorem while giving a presentation at a conference. In addition, suppose that even though you've told no one so, you believe that you have *proved* the theorem. In this case, your assertion seems warranted even if it is false and your would-be proof turns out to be unsuccessful. (For arguments against the claim that S's assertion that P can be warranted even if S doesn't know that P, and even if P is false, see Unger, 1975, Chapter VI; and Williamson, 2000, Chapter 11.)

alternatives. We can define the *germane epistemic position* in a certain conversational context as the position that is strong enough to rule out all epistemically germane alternatives, where “only the [*germane*] 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 strong enough to rule out all of the salient alternatives, where salient alternatives are 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).<sup>26</sup> 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 germane is *contextually invariant*. Thus, what’s germane in radical skeptical contexts is identical to what’s germane in ordinary contexts, namely, an epistemic position that is strong enough to rule out all and only the epistemically germane alternatives. So, given that I occupy the germane 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

---

<sup>26</sup> I regret having to deviate from terminology that is both Rysiew’s and well established in the epistemological literature: I call *germane* alternatives what Rysiew (and many others) call *relevant* alternatives. I do this in order to avoid confusion between epistemically relevant alternatives and relevant assertions (i.e., assertions made in accordance with Grice’s maxim of Relation, which commands, ‘Be relevant!’), which we’ll discuss later in the paper.

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

Questions remain, however. 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 so strong as 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* strong enough to rule out the salient alternatives? Rysiew suggests that

given Grice’s Cooperative Principle (CP),<sup>27</sup> 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)<sup>28</sup>

Thus, 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 is to the point and that it has been made in light of the salient alternatives. Given

---

<sup>27</sup> 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).

<sup>28</sup> 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, Forthcoming). Grice’s maxim of Relation concerns relevance. He says, ‘Under the category of Relation I place a single maxim, namely, “Be relevant”’ (Grice, 1989a, p. 27).

this, our conversational partners reason that we must intend to communicate that S can rule out the salient error possibilities. Moreover, 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.<sup>29</sup>

Rysiew may now maintain the following: (i) It can be true even in radical skeptical contexts that I know that I have hands, for even in those contexts I can rule out all of the germane alternatives to my having hands; (ii) in radical skeptical contexts, we are warranted in asserting that I don’t know that I have hands, for doing so accords with the maxim of Relation—my assertion is relevant since it pragmatically imparts the truth that I cannot eliminate all of the salient alternatives; and (iii) in radical skeptical contexts, we are not warranted in asserting that I know that I have hands, for doing so violates the maxim of Relation—my assertion isn’t relevant since it pragmatically imparts the falsehood that I can eliminate all of the salient alternatives.

Yet when we properly understand the maxim of Relation, we will see that my assertion ‘S knows that *p*’ would *not* violate that maxim. This means that Rysiew has not succeeded in his attempt to explain why my true assertion would be unwarranted in radical skeptical contexts. He suggests that my true assertion violates the maxim of 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*. This suggestion is mistaken, however, for we shouldn’t go so far as to claim that the relevance of a knowledge attribution requires that S occupy the salient epistemic position. Properly understood, relevance requires only that *the knowledge attribution has a*

---

<sup>29</sup> See Rysiew, 2001, p. 491.

*bearing on, or that it is connected with, the matter at hand.*<sup>30</sup> And a knowledge attribution can meet this condition even when S fails to occupy the salient epistemic position. 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 germane 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 a salient alternative. 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.<sup>31</sup> Thus, our knowledge attribution in a radical skeptical context would indeed be 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 has also suggested a different explanation.<sup>32</sup> 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

---

<sup>30</sup> 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; O'Hair, 1969, p. 45; and van Rooy, 2004, pp. 17-8.

<sup>31</sup> Thanks to ... for this suggestion.

<sup>32</sup> He suggested this in correspondence.

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.<sup>33</sup>

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 Relation (since it is relevant to the on-going conversation), has a false implicature (namely, that I can eliminate all of the salient alternatives), and seems false. Nevertheless, it complies with the maxims of Quality, for I don’t believe that it to be false that I know that I have hands, and I don’t lack evidence for my assertion. In such a case, no assertion is *both* apparently false *and* unwarranted, which means that we won’t be able to perform a WAM that meets condition (III), for that condition requires some assertion that *is* both

---

<sup>33</sup> See Grice, 1989a, p. 27.

unwarranted and apparently false. 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 like very much 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 that they are connected 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 makes it impossible to provide a WAM in cases in which we'd like to be able to do so.

But what can the Moorean do now? Must she concede that she cannot defend her proposal with an adequate WAM, thus conceding that Mooreanism still trails contextualism? I don't think so. The resources for an adequate WAM are to be found in the neighborhood neither of Grice's maxim of Relation nor of his maxims of Quality, but rather in the neighborhood of his maxims of *Quantity* and, in particular, his second maxim of Quantity,

Q2: Do not make your contribution more informative than is required.<sup>34</sup>

Some have maintained, however, that there's really no distinction between Q2 and Grice's maxim of Relation.<sup>35</sup> 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?' All else being equal, I

---

<sup>34</sup> See Grice, 1989a, p. 26.

<sup>35</sup> 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 that the two are not distinct.

seem to violate Q2 when I respond by saying, ‘Yes. John is taller than Fred, and it’s unseasonably cool in Los Angeles’.<sup>36</sup> 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. There are, for instance, cases 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*.

I want now briefly to point out some of the advantages of making this distinction. First, such a 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 a proper construal of relevance will capture

our pretheoretical intuitions that 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.)<sup>37</sup>

P. F. Strawson presents a Principle of Relevance, which he describes in the following way:

... stating is not a gratuitous and random human activity. We do not, except in social desperation, direct isolated and unconnected pieces of information at each

---

<sup>36</sup> The example is very similar to one provided by van Rooy, 2004, p. 18.

<sup>37</sup> See, too, Levinson, 1987. In fact, there are times at which Sperber and Wilson themselves, in spite of the fact that they famously reject this Gricean construal of relevance (see Sperber and Wilson, 1986), suggest that we *should* construe relevance in this (Gricean) way. Suppose, for example, that Peter asks Mary, ‘Do you want some coffee?’ Mary responds, ‘Coffee would keep me awake’. Of a case like this, Sperber and Wilson say, “The explicit content of [Mary’s] utterance does not directly answer Peter’s question; it is therefore not relevant as it stands” (Sperber and Wilson, 1986, p. 35).

other, but on the contrary intend in general to give or add information about what is a matter of standing or current interest or concern. (Strawson, 1971, p. 92)

Indeed, Rysiew himself suggests that a conversational contribution is relevant when it is “to the point” (Rysiew, 2001, p. 491).<sup>38</sup> Thus, our pretheoretical intuitions recommend that we construe relevant assertions as those that are connected with the current subject of conversation. To stay true to our pretheoretical intuitions, then, 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 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’.<sup>39</sup> 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.<sup>40</sup> 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

---

<sup>38</sup> Others make similar suggestions. See van Rooy, 2004, in which he writes, “If I ask you, for instance, whether John is sick, I would be very puzzled by your answer *Yes, John is sick, and it is warm in Africa*. The second conjunct to the answer seems to be completely *irrelevant* to the issue, and thus should not be mentioned” (p. 18). See, too, van Rooy and Schulz, 2004, in which they write, The “kind of information [that] is really of *relevance* in the specific situations” is “the information the questioner is interested in” (pp. 6-7).

<sup>39</sup> See Grice, 1989a, p. 32.

<sup>40</sup> 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).

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. In providing examples like this one, Grice himself seems to suggest that relevant assertions are those that are connected with the current topic of conversation.<sup>41</sup> This supports a distinction between maxims of strength and maxims of relevance. So, until someone demonstrates that there is no such distinction, I suggest that we cling to and utilize the differences between maxims of relevance and maxims of strength. 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.

A Gricean maxim of strength is now available to us for use in providing a Moorean WAM. This is good news, for, as we have seen, we won't be able to provide such a WAM by appealing to the maxim of Relation alone or by appealing to Relation working in conjunction with a maxim of Quality. I now propose the following Moorean WAM, which appeals to a maxim of strength, namely, to Grice's Q2 maxim of Quantity. 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 a strong epistemic position, one strong enough to rule out even unlikely possibilities, such as the possibility that I am a bodiless

---

<sup>41</sup> Several theories 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.

brain-in-a-vat. And given that our conversation is governed by the maxim of 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, if we were to assert that I know that I have hands, 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,<sup>42</sup> 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 assertion, for example, the assertion that I’m not sure whether I have hands, or the assertion that I don’t know that (or whether) I have hands. In either of these cases, 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 either of the latter assertions. Thus, the assertion that I know that I have hands is more informative—and hence stronger—than those other assertions.

Yet why would our assertion in radical skeptical contexts be *too* strong? It would be too strong precisely because it implicates that I can eliminate even the salient error possibilities,

---

<sup>42</sup> This is Levinson, 1983’s conception of strength. Compare Leech’s narrower conception of strength (see Leech, 1983, p. 85).

which I cannot in fact eliminate. Although our assertion 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. 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 and 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 strong 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 eliminate all of those possibilities, my 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.

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,<sup>43</sup> can create the appearance that it's true that I don't know that I have hands. In this way, then, the warrantedness of our assertion that I don't know that I have hands can lead us mistakenly to suppose that it's true that I don't know.

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 our assertion that S knows that *p*. If it is, then it ought to be cancelable. Otherwise, the alleged implicature will look very much like a semantic entailment rather than a pragmatic implicature, and we will have a very strong objection to the Moorean's proposed explanation. Fortunately, the implicature here is indeed cancelable. When we assert that S doesn't know that *p*, our assertion has the implicature that S cannot eliminate all of the salient alternatives to *p*. Yet we can cancel this implicature by maintaining, for example, that although S doesn't know that *p*, she nevertheless *can* eliminate all of the salient alternatives to *p*. The plausibility of this cancellation is highlighted by cases in which S cannot eliminate all of the germane alternatives to *p*, but in which she can eliminate each member of the set of salient alternatives, {*a*<sub>18</sub>, *a*<sub>19</sub>, *a*<sub>20</sub>}, which is completely distinct from the set of germane alternatives, {*a*<sub>1</sub>, *a*<sub>2</sub>, *a*<sub>3</sub>, ..., *a*<sub>17</sub>}.

In the case in which we assert that S *does* know that *p*, our assertion has the implicature that S *can* eliminate all of the salient alternatives to *p*. Here again, we can cancel this implicature by saying, for example, that S knows that *p* but *cannot* eliminate all of the salient alternatives to *p*. The plausibility of this cancellation is highlighted by cases such as those in which radical skeptical alternatives are salient, that is, by cases in which S can eliminate all of the germane

---

<sup>43</sup> 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)).

alternatives to  $p$ , but in which she can't eliminate every member of the set of salient alternatives,  $\{a_{18}, a_{19}, a_{20}\}$ , which is completely distinct from the set of germane alternatives,  $\{a_1, a_2, a_3, \dots, a_{17}\}$ . So, since the implicatures here are cancelable, the Moorean's proposal avoids a strong objection.<sup>44</sup>

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 that conflicts, in one and the same context, with another appearance of truth. The Moorean must explain why, in radical skeptical contexts, a false assertion appears true. And in those contexts, it seems true to say both that I don't know that I have hands and that I do know. So, no matter which of these appearances the Moorean must explain away, she must explain away an appearance of truth that conflicts, in one and the same context, with another appearance of truth. Second, the Moorean's WAM explains away this appearance of truth by appealing only to the generation of true implicatures. The warrantedness of our assertion that I don't know that I have hands generates the true implicature that I cannot eliminate all of the salient alternatives to my having hands, which in turn causes us mistakenly to suppose that I don't know that I have hands. Finally, in explaining why our assertion that I don't know that I have hands is warranted, the Moorean appeals to very general conversational rules, namely, to Grice's maxim of Relation and to his maxims of Quantity. The Moorean therefore lives up to the challenge: as we've now seen, she can adequately defend her response to the skeptical puzzle with a WAM that meets all three of DeRose's conditions.

---

<sup>44</sup> See <Author's Paper B> for more on the cancelability of alleged semantic entailments.

#### ***4. Applying the Moorean Response in Other Cases***

As it turns out, maxims of strength are valuable tools for the Moorean. Besides employing them as we already have, the Moorean can enlist their services in order to explain the warrantedness and unwarrantedness of other assertions. She can use them, for example, in providing an explanation of the unwarrantedness in a radical skeptical context of the assertion ‘I have hands’. In this section, I demonstrate how the Moorean can explain our intuitions, both epistemic and otherwise, in a variety of contexts. This will support the claim that maxims of strength are general conversational rules, and that they apply not just to ‘know’ and its cognates, but in other cases as well.

##### *4a. ‘He has hands’ as asserted in a radical skeptical context*

The Moorean can explain the unwarrantedness in a radical skeptical context of the assertion ‘He has hands’. 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. Furthermore, in such contexts, our interlocutor reminds us that if I were a bodiless brain-in-a-vat, *I would have no hands*. Given the mentioned error possibility, what is salient to the conversation is a strong position, one strong enough to eliminate even unlikely error possibilities, such as the possibility that I am a bodiless brain-in-a-vat. If we were to reply to our interlocutor by asserting ‘He has hands’, we would implicate that I occupy the salient position, a position, moreover, that I do not occupy. Thus, the assertion that *I have hands*, even though it is true, is unwarranted, for it is too strong and hence violates a maxim of strength. It implicates that I can eliminate even the salient error possibilities, which I cannot in fact eliminate.

*4b. 'He doesn't have hands' as asserted in a radical skeptical context*

Moreover, we can provide an explanation of this sort of the fact that in radical skeptical contexts we may *not* assert 'He doesn't have hands'. Were we to assert 'He doesn't have hands', our assertion would convey the information that I'm in a strong enough epistemic position to rule out the possibility that I *do* have hands. This is far too informative, for I am in fact in no position to eliminate that possibility. So far, we've done two things in our conversational context: we have made salient the possibility that I am a bodiless brain-in-a-vat, and we have asserted that I don't have hands. Yet neither of these things puts me in position to rule out the possibility that I have hands. (Note that each is consistent with my having hands.) Thus, our assertion 'He does not have hands' is too strong—it conveys the information that I can eliminate the possibility that I have hands when I am in fact in no position to eliminate that possibility.

*4c. 'He knows that these animals are zebras' as asserted in a non-radical skeptical context*

We can also provide an explanation of this very sort of the unwarrantedness of certain assertions in non-radical skeptical contexts. Suppose, for example, that in a context in which the salient alternatives include the alternative that the animals over there in the zoo's zebra cage are cleverly disguised mules, Tom asserts that Sam knows that those animals are zebras. The Moorean can explain the unwarrantedness of Tom's assertion, which suggests that our WAM applies in a wide range of cases, cases involving radical skeptical contexts as well as cases involving non-radical skeptical contexts. Non-radical skeptical contexts are characterized by the salience in those contexts of less-than-radical skeptical possibilities, possibilities that do not call into question all our knowledge of the external world. Non-radical skeptical possibilities include, for example,

the possibility that these animals are cleverly disguised mules. Why, then, in non-radical skeptical contexts, would it be unwarranted for Tom to assert that Sam knows that these animals are zebras? This question is especially pressing for the Moorean, who maintains that it is *true* that Sam knows that these animals are zebras.

Once again, the Moorean can maintain that Tom's assertion would be unwarranted because it would be too strong and hence violate a maxim of strength. It would be too strong because it would implicate that Sam is in a strong enough position to eliminate even the salient error possibilities, such as the possibility that these animals are cleverly disguised mules. Yet, in a context like Tom's, in which the cleverly-disguised-mule possibility is salient, Sam does *not* occupy the salient epistemic position (even though he *does* occupy the germane epistemic position). Tom's assertion would therefore be too strong.

##### ***5. The Moorean versus the Contextualist***

We have now seen that the Moorean can provide all the explanations she must provide. This supplies ample support for the claim that the Moorean pragmatic strategy is a viable alternative to the semantic strategy (as it's employed by, e.g., epistemological contextualists). Thus, at this point, we are well within our rights to report not only that Mooreanism does not trail contextualism, but that the two are running neck and neck. Is there anything that gives Mooreanism the advantage? I think there is. Notice 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.<sup>45</sup>

If the contextualist's WAM is to be successful, it must satisfy DeRose's three conditions. To satisfy the third condition, the contextualist must specify a general conversational rule 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. A maxim of strength might fit the bill here, but I want simply to grant that the contextualist will be able to specify *some* rule that will meet DeRose's third condition. Furthermore, so as not to place undue restrictions on the contextualist, we should allow him to maintain either that *elevated* epistemic standards are salient in radical skeptical contexts or that *relaxed* standards are salient in those 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,

---

<sup>45</sup> Of course, this is not to say that the contextualist has no choice but to perform a WAM. It has been suggested to me, for example, that the contextualist might appeal to semantic blindness 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, Forthcoming, p. 3). To provide the explanation we seek, the contextualist might say 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: She can say that 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. Perhaps the only way, then, for the contextualist to pull ahead of the Moorean is to perform a superior WAM. As I argue, however, the contextualist cannot do this.

- the salience of relaxed standards and the assertion that I know that I have hands.

I argue that none of these is successful. Thus, given the plausibility of the Moorean WAM,<sup>46</sup> 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 DeRose's second condition, 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*. This first attempt fails to meet DeRose's second condition.

The contextualist's second attempt also fails to meet that condition. 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. (In fact, this is just the WAM that the Moorean proposes, and since the

---

<sup>46</sup> And given the fact that the Moorean and the contextualist can appeal to semantic blindness with equal effectiveness.

contextualist wants to distance himself from the Moorean, he shouldn't want in the first place to pursue this option.) The third option is therefore of no use to the contextualist.

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. It seems to meet each of DeRose's three conditions—it needs to explain away an appearance of truth that conflicts, in one and the same context, with another appearance of truth; it explains away this appearance of truth by appealing only to the generation of true implicatures; and, as we have granted, it can appeal to a general conversational rule. Nevertheless, we should reject this WAM, for it is much less realistic than the Moorean's WAM. 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 that I *don't* know that I have hands if I were to find myself in a radical skeptical context. 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. This is simply an expression of elevated epistemic standards. It seems quite unrealistic, then, to suggest that relaxed standards are salient in radical skeptical contexts.

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 invariantist response to skepticism, a response that allows us to maintain that we know all that we ordinarily take ourselves to know. We have also seen that the Moorean response can explain our intuitions, epistemic and otherwise, in a variety of contexts. We therefore have reasons in abundance for preferring the Moorean pragmatic strategy to the contextualist's semantic strategy.<sup>47</sup>

---

<sup>47</sup> Acknowledgments...

## REFERENCES

Bach, Kent. (Forthcoming) “The Emperor’s New ‘Knows’,” in *Contextualism in Philosophy: On Epistemology, Language and Truth*, edited by Gerhard Preyer and Georg Peter (Oxford: Oxford University Press).

<Author’s Paper A>

<Author’s Paper B>

<Author’s Paper C>

<Author’s Paper D>

Brown, Jessica. (2005) “Adapt or Die: The Death of Invariantism?,” *The Philosophical Quarterly* 55: 263-285.

\_\_\_\_\_. (Forthcoming) “Contextualism and Warranted Assertability Manoeuvres” *Philosophical Studies*.

Carston, Robyn. (1998) “Informativeness, Relevance and Scalar Implicature,” in *Relevance Theory: Applications and Implications*, edited by Robyn Carston and Seiji Uchida (Amsterdam: John Benjamins).

Cohen, Stewart. (1988) "How to be a Fallibilist," *Philosophical Perspectives 2: Epistemology*, 91-123.

\_\_\_\_\_. (1999) "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives 13: Epistemology*, 57-89.

DeRose, Keith. (1991) "Epistemic Possibilities," *Philosophical Review* 100: 581-605.

\_\_\_\_\_. (1992) "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52: 913-929.

\_\_\_\_\_. (1999a) "Contextualism: An Explanation and Defense," in *The Blackwell Guide to Epistemology*, edited by John Greco and Ernest Sosa (Oxford: Blackwell).

\_\_\_\_\_. (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).

\_\_\_\_\_. (1999c) "Solving the Skeptical Problem," in *Skepticism: A Contemporary Reader*, edited by Keith DeRose and Ted A. Warfield (New York: Oxford University Press).

\_\_\_\_\_. (Forthcoming) "'Bamboozled by Our Own Words': Semantic Blindness

and Some Arguments Against Contextualism,” *Philosophy and Phenomenological Research* [online] Available from World Wide Web:  
(<http://pantheon.yale.edu/~kd47/Bamboozled.pdf>).

Grice, H. P. (1989a) “Logic and Conversation,” in H. P. Grice, *Studies in the Way of Words* (Cambridge, MA: Harvard University Press).

\_\_\_\_\_. (1989b) “Meaning,” in H. P. Grice, *Studies in the Way of Words* (Cambridge, MA: Harvard University Press).

Harnish, Robert M. (1976) “Logical Form and Implicature,” in *An Integrated Theory of Linguistic Ability*, edited by Thomas G. Bever, Jerrold J. Katz and D. Terence Langendoen (New York: Thomas Y. Crowell).

Hirschberg, Julia. (1985) *A Theory of Scalar Implicature*, PhD dissertation, Computer and Information Science, University of Pennsylvania.

Heller, Mark. (1999) “Relevant Alternatives and Closure,” *Australasian Journal of Philosophy* 77: 196-208.

Horn, Laurence R. (1984) “Toward a New Taxonomy for Pragmatic Inference: Q-based and R-based Implicature,” in *Meaning, Form, and Use in Context: Linguistic Applications*, edited by Deborah Schiffrin (Washington, D. C.: Georgetown University Press).

- Leech, Geoffrey N. (1983) *Principles of Pragmatics* (New York: Longman).
- Levinson, Stephen C. (1983) *Pragmatics* (Cambridge: Cambridge University Press).
- \_\_\_\_\_. (1987) "Minimization and Conversational Inference," in *The Pragmatic Perspective*, edited by Jef Verschueren and Marcella Bertuccelli-Papi (Amsterdam: John Benjamins).
- \_\_\_\_\_. (2000) *Presumptive Meanings: The Theory of Generalized Conversational Implicature* (Cambridge, MA: MIT Press).
- Lewis, David. (1996) "Elusive Knowledge," *Australasian Journal of Philosophy* 74: 549-567.
- O'Hair, S. G. (1969) "Implications and Meaning," *Theoria* 35: 38-54.
- Rieber, Steven. (1998) "Skepticism and Contrastive Explanation," *Noûs* 32: 189-204.
- Rysiew, Patrick. (2001) "The Context-Sensitivity of Knowledge Attributions," *Noûs* 35: 477-514.
- Sperber, Dan and Deidre Wilson. (1986) *Relevance: Communication and Cognition* (Oxford: Basil Blackwell).

Stone, Jim. (2000) "Skepticism as a Theory of Knowledge," *Philosophy and Phenomenological Research* 60: 527-545.

Strawson, P. F. (1971) "Identifying Reference and Truth-Values," in P. F. Strawson, *Logico-Linguistic Papers* (London: Methuen & Co Ltd.).

Stroud, Barry. (1984) *The Significance of Philosophical Scepticism* (New York: Oxford University Press).

Unger, Peter. (1975) *Ignorance: A Case for Scepticism* (Oxford: Oxford University Press).

Van Rooy, Robert. (2004) "Relevance Implicatures" [online] Available from World Wide Web: (<http://citeseer.ist.psu.edu/545416.html>).

Van Rooy, Robert and Katrin Schulz. (2004) "Pragmatic Meaning and Non-monotonic Reasoning: The Case for Exhaustive Interpretation" [online] Available from World Wide Web: (<http://turing.wins.uva.nl/~vanrooy/exhLP-r.pdf>).

Welker, K. (1994) *Plans in the Common Ground: Toward a Generative Account of Implicature*, PhD thesis, Linguistics, The Ohio State University.

Williamson, Timothy. (2000) *Knowledge and Its Limits* (Oxford: Oxford University Press).