CLASSIC INVARIANTISM, RELEVANCE, AND WARRANTED ASSERTABILITY MANŒUVERS

TIM BLACK

The Philosophical Quarterly 55 (2005): 328-336

Jessica Brown effectively contends that Keith DeRose's latest argument for contextualism fails to rule out contextualism's chief rival, namely, classic invariantism. Still, even if her position has not been ruled out, the classic invariantist must offer considerations in favor of her position if she is to convince us that it is superior to contextualism. Brown defends classic invariantism with a warranted assertability maneuver that utilizes a linguistic pragmatic principle of relevance. I argue, however, that this maneuver is not as effective as it might be. I propose a different warranted assertability maneuver—one that utilizes a pragmatic principle of strength—that affords a more successful defense of classic invariantism, and that helps to establish that classic invariantism is superior to contextualism.

Recently, Keith DeRose has provided yet another influential argument for epistemological contextualism.[1] His argument involves the Knowledge Account of Assertion, or KAA, and the Context-Sensitivity of Assertion, or CSA.

KAA: The standards for when one is in an epistemic position warrantedly to assert that *p* are identical to those that comprise a truth condition for 'I know that *p*'.[2]

CSA: How well positioned one must be with respect to *p* in order warrantedly to assert that *p* depends on context.[3]

These premises are meant to establish epistemological contextualism, that is, the claim that the truth conditions of knowledge attributions depend on certain features of the attributor's context. Moreover, contextualists contend that since these features can shift from context to context, the

[1] See K. DeRose, 'Assertion, Knowledge, and Context', *Philosophical Review*, 111 (2002), pp. 167-203.

[2] See DeRose, 'Assertion, Knowledge, and Context', p. 179. Such an account is also defended in P. Unger, *Ignorance* (Oxford UP, 1975), Chapter VI; and in T. Williamson, *Knowledge and its Limits* (Oxford UP, 2000), Chapter 11.

[3] See DeRose, 'Assertion, Knowledge, and Context', p. 181.

truth conditions of knowledge attributions can shift from context to context. In some contexts, then, the truth conditions of knowledge attributions can be quite demanding, so demanding, in fact, that one cannot truthfully attribute knowledge in those contexts. Yet the truth conditions of knowledge attributions are comparatively relaxed in other contexts, in which, as a consequence, one can truthfully attribute knowledge.

Jessica Brown argues, however, that DeRose's argument for contextualism fails to rule out contextualism's chief rival, namely, classic invariantism, or CI:

CI: The truth conditions of knowledge attributions depend neither on the subject's context nor on the attributor's context.[4]

DeRose's argument is effective against CI, of course, only if its premises are true. He defends KAA by appeal to the Knowledge Rule, or KR, and the Uniqueness Claim, or UC.[5]

KR: One must know that *p* in order to be well-enough positioned with respect to *p* to be warranted in asserting that *p*.

UC: KR is the only rule governing how well positioned one must be in order warrantedly to assert that *p*.

Brown argues that UC must concern a notion of warrant that is narrower than the one involved in CSA.[6] This is so, Brown argues, because UC is defensible only when KR concerns what she calls warrant*, which is a narrower notion of warrant than the one involved in CSA. Thus, given that KR and UC together establish KAA, KAA and CSA involve different notions of warrant. It therefore seems that from KAA and CSA, we cannot legitimately infer any claim that is

[4] See J. Brown, 'Adapt or Die: The Death of Invariantism?', *The Philosophical Quarterly*, 55 (2005), pp. 263-285. Brown also contends, as have others, that DeRose's arguments do not rule out subject-sensitive invariantism, according to which the truth conditions of knowledge attributions depend on the subject's context but not on the attributor's context. (See Brown, 'Adapt or Die', pp. 267-272; and T. A. Blackson, 'An Invalid Argument for Contextualism', *Philosophy and Phenomenological Research*, 68 (2004), pp. 344-5.) I will not in this paper be concerned with this sort of invariantism.

[5] See DeRose, 'Assertion, Knowledge, and Context', p. 180.

[6] See Brown, 'Adapt or Die', pp. 275-278.

incompatible with CI. Unless another defense of KAA is in the offing, then, we have no reason to think that DeRose's argument rules out CI.

But where does this leave the non-skeptical classic invariantist (hereafter, *the invariantist*)? Her position has not been ruled out by DeRose's argument, but what can she say in its defense? In particular, asks Brown, how will the invariantist explain CSA? Brown opts for an explanation that involves a warranted assertability maneuver, or WAM, which

> consists in the claim that our intuitions about when it's correct to attribute knowledge reflect the warranted assertability conditions rather than the truth conditions of knowledge attributions. … Using [a] Gricean model …, the invariantist may suggest that, in [a] high standards context, the attribution of knowledge is literally true but not warrantedly assertible since it pragmatically conveys a falsehood.[7]

Even though other explanations of CSA are available to the invariantist,[8] a defense that utilizes a WAM is important. For, if such a defense is successful, it helps non-skeptical classic invariantism—hereafter, *invariantism*—to regain its footing, leaving it well defended and leaving contextualism with no discernible advantage over invariantism.

Given that Brown has effectively maintained that DeRose's argument fails to rule out CI, the invariantist's fate rests on her WAM. If her WAM does not succeed, then, so far as her arguments go, contextualism still has the upper hand. I argue in this paper that her WAM is not as effective as it might be. Lurking nearby, however, is a different WAM, which we can utilize in providing a more successful defense of invariantism.

---

[7] Brown, 'Adapt or Die', p. 265.

[8] See K. Bach, 'The Emperor's New 'Knows'', in Gerhard Preyer and Georg Peter (eds.), *Contextualism in Philosophy: On Epistemology, Language and Truth* (Oxford UP, Forthcoming); and J. Hawthorne, *Knowledge and Lotteries* (Oxford UP, 2004).

1. The Invariantist's Relevance-Based Warranted-Assertability Maneuver

Let's begin by examining Brown's WAM. She provides her WAM in defense of an invariantist response to puzzles that stem from DeRose's bank cases.[9] In one case—call it *easy*—DeRose and his wife are deciding whether to deposit their pay cheques on Friday or to wait until Saturday; DeRose "was at the bank just two weeks ago on a Saturday, and found that it was open until noon";[10] and "no disaster will ensue if we waste a trip to the bank on Saturday only to find it closed."[11] In easy, DeRose asserts that he knows that the bank will be open on Saturday, and almost all of us would say that his assertion is true.[12] Yet in the other case—call it *tough*—"disaster, not just disappointment, would ensue if we waited until Saturday only to find we were too late: We have just written a very large and very important check, and will be left in a *very* bad situation if the check bounces, as it will if we do not deposit our paychecks before Monday."[13] In tough, DeRose asserts that he *doesn't* know that the bank will be open on Saturday. Here again, DeRose claims, almost all of us would say that his assertion is true.[14]

Nevertheless, the invariantist maintains that DeRose knows both in easy *and in tough* that the bank will be open, and thus that if he were to assert in tough that he knows that the bank will be open, his assertion would be true. Yet if DeRose knows in tough that the bank will be open, why may he not *say* that he knows? Brown seeks to answer this question with a WAM. She writes,

> A natural way for a C[lassic] invariantist to explain how, in tough, the true knowledge ascription is not warranted is to appeal to the rule of relevance … .

[9] See DeRose, 'Assertion, Knowledge, and Context', pp. 169-70; and K. DeRose, 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52 (1992), pp. 913-29.
[10] DeRose, 'Assertion, Knowledge, and Context', pp. 169-70.
[11] DeRose, 'Assertion, Knowledge, and Context', p. 170.
[12] See DeRose, 'Assertion, Knowledge, and Context', p. 170.
[13] DeRose, 'Assertion, Knowledge, and Context', p. 170.
[14] See DeRose, 'Assertion, Knowledge, and Context', p. 170.

> The conversation between DeRose and his wife in tough concerns the practical question of whether they should wait in the queue to deposit their pay cheques or wait until the next day, a Saturday. Further, in tough, DeRose's wife reminds DeRose how important it is that the cheques are paid in before Monday, and raises the possibility that the bank has changed its hours since the last time DeRose was there. Given the stakes and the mentioned error possibility, it seems that what is relevant to the conversation in tough is a strong epistemic position, one strong enough to rule out even unlikely errors, such as the possibility that the bank has changed its hours. Thus, were DeRose to reply to his wife by asserting, 'I know that the bank is open' he would implicate that he is [in] such a strong epistemic position. For, if he were not in such a strong epistemic position, his assertion would not be relevant to the on-going conversation. Since, by hypothesis, DeRose is not in such a position, the assertion that he knows would be unwarranted even though it is true.[15]

This explanation, however, seems to undermine the invariantist's primary aim, namely, to allow for the fact that DeRose *knows*, even in tough, that the bank will be open. Brown maintains that what is relevant to the conversation in tough is a strong epistemic position, one that is strong enough to rule out even unlikely errors, such as the possibility that the bank has changed its hours. On a plausible and fairly widely accepted account of knowledge, however, S knows that *p* only if S is in the relevant epistemic position with respect to *p*.[16] Given this conception of knowledge, Brown's explanation seems committed to the claim that DeRose knows in tough only if he is in an epistemic position that is strong enough to rule out even

[15] Brown, 'Adapt or Die', pp. 280-281.

[16] That is, S knows that *p* only if S is in an epistemic position that is strong enough to rule out all relevant error possibilities.

unlikely errors. Yet, by hypothesis, he is not in such a strong epistemic position in tough. Brown's explanation therefore seems committed to the claim that DeRose does *not* know in tough that the bank will be open on Saturday.

We can deal with this problem by introducing a distinction between relevant epistemic positions and salient epistemic positions:[17] The *relevant* epistemic position is the position that is strong enough to rule out all epistemically relevant error possibilities (where the relevant error possibilities are those that one must eliminate if one is to know); while the *salient* epistemic position is the one that is strong enough to rule out all the salient error possibilities (where the salient error possibilities are those that the conversational participants have in mind, or that they've mentioned, no matter how unlikely). Given this distinction, we can account for the fact that DeRose knows in tough that the bank will be open. For, according to the invariantist, what's relevant in tough is identical to what's relevant in easy, namely, an epistemic position that is strong enough to rule out all epistemically relevant error possibilities, but *only* those possibilities. Thus, given the above conception of knowledge, and given that DeRose occupies the relevant epistemic position, he can know in tough that the bank will be open on Saturday.

Nevertheless, DeRose's assertion in tough that he knows that the bank will be open is not warranted, for it would implicate that he is in the *salient* epistemic position. We should therefore understand Brown as maintaining that S's assertion is irrelevant to the on-going conversation (and hence unwarranted) if (a) his assertion would implicate that he is in the salient epistemic position, but (b) he is not in that position. We should prefer this to the claim that Brown seems to make, namely, that S's assertion is irrelevant to the on-going conversation if (a) his assertion would implicate that he is in the relevant epistemic position, but (b) he is not in *that* position.

[17] See P. Rysiew, 'The Context-Sensitivity of Knowledge Attributions', *Noûs*, 35 (2001), pp. 477-514, especially p. 488.

For, if S's assertion would implicate that he is in the relevant epistemic position, then it certainly seems that it would not be irrelevant to the on-going conversation, whether or not S is in fact in the relevant epistemic position.

We've now come to a better understanding of Brown's invariantist WAM, an understanding that relies on a distinction between relevant and salient epistemic positions. Still, we should at this point examine the claim that S's assertion is irrelevant to the on-going conversation if (a) his assertion would implicate that he is in the salient epistemic position, but (b) he is not in that position. There is reason, I think, to question this claim. For we shouldn't go so far as to claim that the relevance of an assertion requires that the asserter occupy the implicated epistemic position. In fact, it seems sufficient for relevance that *the assertion has a bearing on, or that it is connected with, the matter at hand*.[18] And an assertion can meet this condition even when the asserter fails to occupy the implicated epistemic position. DeRose's assertion in tough, for example, *does* have a bearing on and *is* connected with the matter at hand. In tough, DeRose and his wife are concerned with whether the bank will be open on Saturday, and his assertion 'I know that the bank will be open' would indeed be connected with this matter. Thus, his assertion would be relevant in tough. This gives us reason to think that he—or that his assertion—violates no rule of relevance. If this is correct, then we have reason to think that Brown has not yet explained how, in tough, DeRose's true assertion is unwarranted. For the success of her explanation depends on the claim that his assertion would violate a rule of relevance, and we now have reason to think that it would violate no such rule.

[18] Compare Grice's own characterization of relevance, which includes the claim that relevance has to do with "subjects of conversation" (H. P. Grice, 'Logic and Conversation', in H. P. Grice, *Studies in the Way of Words* (Harvard UP, 1989), p. 27). Also, compare the characterizations of relevance in R. M. Harnish, 'Logical Form and Implicature', in T. G. Bever, J. J. Katz, and D. T. Langendoen (eds.), *An Integrated Theory of Linguistic Ability* (Thomas Y. Crowell, 1976), p. 362; in S. G. O'Hair, 'Implications and Meaning', *Theoria*, 35 (1969), pp. 38-54, especially p. 45; and in R. van Rooy, 'Relevance Implicatures', ms. [online at http://citeseer.ist.psu.edu/545416.html], pp. 17-8.

2. A Strength-Based Warranted-Assertability Maneuver

What must the invariantist do? Must she concede at this point that even though DeRose's arguments do not rule out invariantism, she cannot defend her proposal with an adequate WAM, thus conceding that invariantism still trails contextualism? I don't think so. Brown is most decidedly on the right track here, for an adequate WAM is to be found in the neighborhood of Grice's second maxim of Quantity,

Q2: Do not make your contribution more informative than is required,[19]

and we can see the rule of relevance to which Brown appeals as a submaxim of Q2: I can violate Q2 by providing irrelevant information. Suppose S asks, "Is John taller than Fred?" All else being equal, I violate Q2 when I respond by saying, "Yes. John is taller than Fred, and it's unseasonably cool in Los Angeles."[20] Yet this is not the only way to violate Q2: I can violate Q2 by providing information that, even though it's relevant, exceeds the threshold of informativeness. Let's call submaxims of the former sort *submaxims of relevance*, and of the latter sort *submaxims of strength*.

We should make this distinction, at least for the time being, because doing so has three advantages over not doing so. First, such a distinction fits with "our pretheoretical intuitions,"[21] which suggest that a contribution is irrelevant only when it deviates, more or less dramatically, from the subject of the conversation. In fact, several theorists construct their theories on and around such intuitions. For example, P. F. Strawson presents a Principle of Relevance, which he describes in the following way:

---

[19] Grice, 'Logic and Conversation', p. 26.

[20] This example is similar to one provided by van Rooy, 'Relevance Implicatures', p. 18.

[21] S. C. Levinson, *Presumptive Meanings: The Theory of Generalized Conversational Implicature* (MIT Press, 2000), p. 380, n. 4.

> … 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 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.[22]

Thus, our pretheoretical intuitions—as well as theories built around those intuitions—seem to suggest that relevant assertions are those that are connected with the current subject of conversation. This makes room for a distinction between submaxims of strength and submaxims of relevance. To stay true to our pretheoretical intuitions, then, we should accept such a distinction.

Second, certain cases and prominent examples suggest that there is such a distinction. I'm inclined to think that tough brings the distinction to the fore—in tough, our knowledge attributions would provide an excess of relevant information, thus violating a strength submaxim rather than a relevance submaxim. Moreover, prominent examples in which Grice's maxim of Relation is either observed or violated suggest that an assertion is irrelevant only when it deviates from the subject of the conversation.[23] 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'.[24] According to Grice, B's statement accords with the maxim of Relation and is therefore appropriate. Moreover, it seems to accord with that maxim because, in the first place, it can be

---

[22] P. F. Strawson, 'Identifying Reference and Truth-Values', in P. F. Strawson, *Logico-Linguistic Papers* (Methuen & Co Ltd, 1971), p. 92. See, too, Stephen C. Levinson, who 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 …" (Levinson, *Presumptive Meanings*, p. 380, n. 4). Others make similar suggestions. See van Rooy, 'Relevance Implicatures', p. 18, at 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."

[23] Grice's maxim of Relation concerns relevance. He says, 'Under the category of Relation I place a single maxim, namely, "Be relevant"' (Grice, 'Logic and Conversation', p. 27).

[24] See Grice, 'Logic and Conversation', p. 32.

taken by A to be topically related to his question.[25] 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. This makes room, then, for a distinction between submaxims of strength and submaxims of relevance.

Third, distinguishing relevance submaxims from strength submaxims allows us to remain faithful to Gricean and neo-Gricean pragmatic theories. Several such theories make a distinction such as ours. Note that Grice himself, in tying "relevance to particular plans and goals,"[26] seems to allow for a distinction such as the one we make between strength submaxims and relevance submaxims.[27] In addition, Levinson maintains that Grice's maxims of Quantity, which are concerned with the amount of information provided by conversational contributions, can generate GCIs, or generalized conversational implicatures.[28] He says, however, that the "maxim of Relation or Relevance … has pertinence only to the immediate, ever variable, conversational goals: it generates PCIs [particularized conversational implicatures], not GCIs."[29] This clearly indicates that, for Levinson, we should distinguish the maxim of Relevance from certain (other) maxims concerning the amount of information provided by conversational contributions.[30] This, of course, is just the sort of distinction I'm concerned to make when I distinguish between relevance submaxims and strength submaxims. So, until someone demonstrates that there is no such distinction, I suggest that we cling to and utilize the differences between submaxims of

---

[25] 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 S. C. Levinson, *Pragmatics* (Cambridge UP, 1983), pp. 97-8).

[26] Levinson, *Presumptive Meanings*, p. 380, n. 4.

[27] See Grice, 'Logic and Conversation'.

[28] See Levinson, *Presumptive Meanings*, p. 20. Regarding GCIs, Levinson says, "An implicature i [from utterance U] is *generalized* iff U implicates i *unless* there are unusual specific contextual assumptions that defeat it" (Levinson, *Presumptive Meanings*, p. 16).

[29] Levinson, *Presumptive Meanings*, p. 74. Regarding PCIs, Levinson says, "An implicature i from utterance U is *particularized* iff U implicates i only by virtue of specific contextual assumptions that would not invariably or even normally obtain" (Levinson, *Presumptive Meanings*, p. 16).

[30] Compare Harnish's Maxim of Quantity-Quality: *Make the strongest relevant claim justifiable by your evidence* (see Harnish, 'Logical Form and Implicature', p. 362). Harnish's maxim mentions both strength and relevance, which suggests that he favors a distinction such as ours.

relevance and submaxims 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.

Two things become clearer once we accept such a distinction: (1) Brown leads us toward the Gricean maxim that is appropriate for the provision of an invariantist WAM, and (2) her WAM stumbles because it leans on the wrong submaxim of that maxim. Perhaps by relying on a submaxim of strength, the invariantist can provide a WAM that does everything it is supposed to do without stumbling.

Here, then, is a different invariantist explanation—an explanation that appeals to a submaxim of strength—of the unwarrantedness of DeRose's true assertion. In tough, the conversation between DeRose and his wife concerns the practical question of whether they should wait in the queue to deposit their paychecks or wait until Saturday. Further, in tough, DeRose's wife reminds him how important it is that the checks are paid in before Monday, and raises the possibility that the bank has changed its hours since the last time he was there. Given the stakes and the mentioned error possibility, it seems that what is salient to the conversation in tough is a strong epistemic position, one strong enough to rule out even unlikely errors, such as the possibility that the bank has changed its hours. If DeRose were to assert 'I know that the bank will be open', he would implicate that he occupies the salient epistemic position, a position, moreover, that he does not occupy. Thus, his assertion, even though it would be true, is too strong and hence unwarranted.

DeRose's assertion in tough that he knows that the bank will be open on Saturday would be *stronger* than other assertions in the sense that it would be *more informative* than other

assertions, 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."[31] Were DeRose to assert in tough that he knows that the bank will be open, he would convey the information that he can eliminate even the salient alternatives to the bank's being open. He does not convey this information, however, if he makes a different assertion, for example, the assertion that he doesn't know that (or whether) the bank will be open. The former assertion—that he knows that the bank will be open—conveys the information that he can eliminate possibilities left open by the latter assertion. Thus, the assertion that he knows that the bank will be open is more informative—and hence stronger—than other assertions.

Moreover, his assertion in tough would be *too* strong because it would implicate that he can eliminate all of the salient error possibilities.[32] But since he's not in the salient epistemic position in tough, he can*not* eliminate all of those error possibilities.

But what about DeRose's assertion that he *doesn't* know that the bank will be open? This assertion, in accordance with the submaxim of strength, is warranted in tough. Once again, in that context, what is salient to the conversation is a position that's strong enough to eliminate even unlikely error possibilities. Moreover, DeRose's assertion 'I don't know that the bank will be open' conveys the information that he cannot eliminate all of those salient possibilities. Thus, since he is in fact in no position to eliminate all of those possibilities, his assertion is made in accordance with the appropriate submaxim of strength—it does not exceed the threshold of informativeness.

[31] Levinson, *Presumptive Meanings*, p. 115. Compare Levinson, *Pragmatics*.

[32] DeRose's assertion would have this (false) implicature primarily because those participating in the conversation (tend to) mistake the salient epistemic position for the relevant one.

How, then, does the warrantedness of DeRose's assertion that he doesn't know that the bank will be open create the appearance that it's *false* that he *does* know? In accordance with the Gricean submaxims of strength, he is warranted in asserting in tough that he doesn't know that the bank will be open. This warranted assertion can generate the false implicature that his belief that the bank will be open does *not* meet the standards for knowledge. This false implicature can in turn create the appearance that it's false that he knows that the bank will be open. In this way, then, the warrantedness of his assertion that he doesn't know that the bank will be open can lead us mistakenly to suppose that it's false that he does know.

3. The Adequacy of the Strength-Based WAM

Is our WAM satisfactory? So far as I know, only DeRose has proposed criteria of adequacy for WAMs.[33] He suggests that

(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;

(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); and

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

[33] See K. DeRose, 'Contextualism: An Explanation and Defense', in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology* (Blackwell, 1999), especially pp. 201-203. Brown, too, is concerned to provide a WAM that can either meet or plausibly resist each of these three conditions; see Brown, 'Adapt or Die', pp. 283-285.

Our strength-based WAM meets all three of these conditions. First, our WAM begins with intuitions of falsity that conflict within one and the same context. Recall tough, which is characterized by the fact that it's governed by unusually high epistemic standards. In that context, it seems false that DeRose knows that the bank will be open. After all, he can't rule out the possibility that the bank will be closed on Saturday. But doesn't it also seem false that he *doesn't* know that the bank will be open? As DeRose himself says, while speaking of radical skeptical contexts, "[m]ost of us feel some ambivalence" here.[34] For we realize that DeRose "was at the bank just two weeks ago on a Saturday, and found that it was open until noon."[35] We also realize that he has no reason in tough to think that the bank has changed its hours—all that's happened in tough is this: DeRose's wife *mentions the possibility* that the bank has changed its hours. I submit, then, that we should maintain that both of the following hold, simultaneously and intracontextually, in tough: (A) It seems false to say that DeRose *knows* that the bank will be open, and (B) it seems false to say that he *doesn't* know that the bank will be open. The invariantist, then, must explain away the appearance in (A). But since this is an appearance of falsity, she satisfies DeRose's first condition.

Second, our invariantist WAM explains away this appearance of falsity by appealing only to the generation of false implicatures. The warrantedness of DeRose's assertion that he doesn't know that the bank will be open generates the false implicature that his belief fails to meet the standards for knowledge, which in turn causes us mistakenly to suppose that it's false that I know that I have hands.

Finally, in explaining why DeRose's assertion that he doesn't know that the bank will be open is warranted, our WAM appeals to very general conversational rules, namely, to Grice's

[34] K. DeRose, 'Solving the Skeptical Problem', in K. DeRose and T. A. Warfield (eds.), *Skepticism: A Contemporary Reader* (Oxford UP, 1999), pp. 208-9. (Compare DeRose, 'Solving the Skeptical Problem', p. 185.)

[35] DeRose, 'Assertion, Knowledge, and Context', pp. 169-70.

maxims of Quantity and to submaxims of strength. So, since our WAM meets all three of DeRose's conditions, both he and Brown should count it as satisfactory.[36]

4. Conclusion

I have endeavored to maintain, as has Brown, that we can provide a WAM against contextualism and in favor of invariantism. Still, as I have shown, Brown's WAM is less effective than it might be, for it relies on a rule of relevance. I have argued that the invariantist's WAM should rely instead on a submaxim of strength, which will allow us to provide a WAM that is both effective and satisfactory. Invariantists should therefore welcome a WAM of this sort.[37]

[36] Once again, Brown, too, is concerned with DeRose's three conditions; see note 33.

[37] For helpful criticisms and comments, I thank Jessica Brown, Albert Casullo, Duncan Pritchard, the audience at the University of Stirling's conference on Epistemological Contextualism, and an anonymous referee for *The Philosophical Quarterly*.