

HOW TO LINK ASSERTION AND KNOWLEDGE WITHOUT GOING CONTEXTUALIST:
A REPLY TO DEROSE'S "ASSERTION, KNOWLEDGE, AND CONTEXT"

The knowledge account of assertion takes two forms. According to one, in asserting that p one (standardly) represents oneself as knowing p .¹ According to the other, "assert only what you know" is a constitutive norm of the practice of assertion.² Keith DeRose has recently maintained that in either form, the knowledge account of assertion – when combined with some empirical considerations about our assertive practice – yields a straightforward argument for the contextualist view that the truth conditions for knowledge ascriptions vary with features of the context of utterance; as he puts it, "The knowledge account of assertion demands a contextualist account of knowledge and is simply incredible without it" (2002, p. 182). However, this claim is incorrect. One could accept both the knowledge account of assertion and DeRose's considerations about our assertive practice, and still maintain the invariantist view that the standards one must meet in order for a knowledge ascription to be true are not determined by *anyone's* practical or conversational circumstances.³ As I will argue, DeRose's argument fails because of a subtle mistake concerning the notion of *warranted assertability*. One consequence of this mistake is that while DeRose aims to respond to the threat to contextualism presented by so-called "warranted assertability maneuvers", the threat and his response pass each other by. Another is that it can come to look as if invariantism is much more easily dismissed than is in fact the case.

I.

In a slightly enthymematic presentation, DeRose's argument runs as follows.

The knowledge account of assertion provides a powerful argument for contextualism: If the standards for when one is in a position to warrantably assert that P are the same as those that constitute a truth condition for “I know that P,” then if the former vary with context, so do the latter. In short: The knowledge account of assertion together with the context sensitivity of assertability yields contextualism about knowledge (187).

This argument can be fleshed out as follows.⁴

1. Whether one is warranted in asserting p depends in part on one’s epistemic relation to p: in order to warrantably assert that p, one must be epistemically positioned well enough with regard to p.
2. According to the knowledge account of assertion, one is warranted in asserting that p *only if* one knows that p. This rule is the only rule specific to assertion which concerns when one is epistemically well-enough positioned to assert a specific proposition.⁵ So, assuming the correctness of the knowledge account of assertion, “one is positioned well enough to assert that P iff one knows that P” (180).
3. Consequently, the epistemic standards which one must meet in order to warrantably assert that p in a given conversational context may be *identified* with the standards for knowledge in that context.⁶
4. So if the epistemic standards for warranted assertion vary with the conversational context, then the standards for knowledge vary with the conversational context as well.

5. The epistemic standards for proper assertion *do* vary with the conversational context. (“how well-positioned one must be with respect to P to be able to properly assert that P is a variable and highly context-sensitive matter” (p. 181).)
6. So, by 4 and 5, the standards which must be met in order for a knowledge attribution (made in a particular context) to be true will vary with the conversational context.

In short: One is epistemically warranted in asserting *p* iff one knows that *p*. But whether one is warranted in asserting *p* varies with the context of utterance. So, too, then do the conditions for knowledge.

I propose not to question step (5). Numerous examples suggest that it is at least roughly correct. Suppose, for instance, that we are sitting in my back yard idly chatting over drinks after a hard day. Someone asks, “So, who is going to win the election?” “Kerry,” I assert. “The voters are tired of Bush’s immaturity.” Other conversational participants might disagree with my comment, but nobody would judge it inappropriate or improper. It was perfectly relevant and polite. Most importantly, nobody in this context would think it appropriate to respond, “But you don’t know that,” or anything along those lines. Compare this case, then, with one in which (my epistemic position having remained unchanged) my six-year-old son asks me, really wanting to know, “Daddy, who will win the election?” In this context, the appropriate response would be something like, “I don’t know, but I think that Kerry will win.” It would be inappropriate for me to reply, “Kerry will;” here someone could fairly object, “But you don’t know

that!” Thus, whether one’s epistemic position is adequate for proper or warranted assertion depends upon features of the context.

Still, DeRose’s argument won’t go through unless the same thing is meant by “epistemically warranted assertion” when we say that the standards one must meet in order to make an epistemically proper or warranted assertion vary with the context and when we say, in step (3) of the argument, that the standards for proper or warranted assertion may be identified with the standards for knowledge. Otherwise, the argument will fail because of an equivocation; (5) will not enable us to discharge the antecedent of (4).

Let’s consider, then, in what sense the knowledge account of assertion yields an account of warranted assertability. According to the Moorean view, asserting that *p* involves representing oneself as knowing that *p*. On the face of it, the most this view generates is that one will not be (epistemically) warranted in asserting that *p* unless one is (epistemically) warranted in representing oneself as knowing that *p*. But the view is silent on when that is.⁷ It is therefore doubtful that the view yields an account of warranted assertability at all. For this reason, it is best to focus on Williamson’s version of the view, where matters are clearer.

According to Williamson, the rule “Don’t assert that *p* unless you know that *p*” is a “constitutive norm” of the practice of assertion. A constitutive norm is a norm to which any act (or, more generally, norm-evaluable thing) of a certain type is *answerable*, as a condition of being an act (etc.) of that type. So the claim is that in virtue of being an instance of *assertion*, a speech-act is defective – for reasons internal to the practice of asserting -- if the speaker does not possess knowledge of the truth of the proposition

asserted. This is not to say that the speech-act in question is not an assertion: it is an assertion, but a defective one. Likewise, this is not to say that the speaker is appropriately blamed or chastised for having issued a faulty assertion: for instance, he might quite reasonably but incorrectly have taken himself to have knowledge; or it may be that issuing a fully non-deficient assertion is not very important in his circumstances. We'll return to this last point shortly. The important point for now is that on Williamson's view, any such assertion, however reasonable or excusable, is defective in an important sense.

An account of assertion along Williamson's lines directly yields a norm regarding the epistemic position required for fully non-defective assertion. It is tempting to put this by saying that the account tells us when an assertion would be (epistemically) warranted, and it would then be natural to take this result to provide a (partial) account of the theoretical notion of warranted assertability as it figures in familiar discussions in the contextualism literature. However, the notion of warranted assertability which is yielded by Williamson's account cannot play the standard theoretical role. Here's why.

A familiar objection to contextualism involves what DeRose calls a "Warranted Assertability Maneuver (WAM)". The proponent of this objection *accepts* that the standards which govern our ordinary knowledge attributions exhibit contextual variability, but suggests that this contextual variability may not reveal anything about the *truth conditions* for knowledge attributions, because the variability may only be "in the conditions under which it is appropriate to claim to know" (p. 170). The objection thus depends upon a distinction between the conditions which must be met in order for a knowledge attribution to be true ("truth conditions") and the conditions which must be

met in order for it to be appropriate or warranted (“assertability conditions”); the suggestion is that the latter do not directly reveal the former, since a knowledge attribution might be true but not warranted or warranted but not true. (The latter tack is that of the skeptical invariantist who thinks that because the standards for knowledge are *always* unmeetably high, no knowledge attribution is ever true.) More generally, the distinction between truth conditions and assertability conditions, and the suggestion that the two might come apart, is essential to the strategy of all WAMs.

A WAM involves explaining why an assertion can seem false (or at least not true) in certain circumstances in which it is in fact true by appeal to the fact that the utterance would be improper or unwarranted in the circumstances in question. Going the other way, an intuition that an assertion is true can be explained away by means of the claim that the assertion, while false, is warranted, and we mistake this warranted assertability for truth. Either way, the maneuver is based on the correct insight that truth and warranted assertability are quite different things, but that we can easily mistake one for the other (171-2).

As DeRose emphasizes here, the notion of assertability that is relevant is one on which an assertion may be *appropriately asserted* but *false*.

Consider, then, the notion of warranted assertability which we derived from Williamson’s version of the knowledge account of assertion (hereafter, “warranted assertability_{KAA}”). According to this account, the assertability condition for a proposition *P* is that *the speaker knows that P* (186). A speaker can’t know that *P* unless *P* is true. So on this account of assertability conditions, assertability conditions and truth conditions

cannot come apart in the way that they must if the WAM strategy is to make any sense at all.⁸ So warranted assertability_{KAA} must not be the notion of warranted assertability that figures in the standard objections to contextualism and other familiar WAMs in the philosophy of language literature. In *this standard sense of “warranted assertion”*, DeRose is wrong to identify the standards for warranted assertion that P with the standards “that constitute a truth condition for ‘I know that P’” (187). It is this *standard sense* of “warranted assertion” which is at stake in step (5) of DeRose’s argument. So far, then, it looks as though the argument is indeed guilty of equivocation.⁹

That DeRose identifies these distinct notions of warranted assertability is clear from passages such as the following.

It is difficult to deny that the matter of how well positioned one must be with respect to a matter to be able to assert it varies with the context: What one can flat-out assert in some “easy” contexts can be put forward in only a hedged manner ... when more stringent standards hold sway. Even invariantists, who deny that the *truth* conditions of knowledge attributions are sensitive to varying standards, tend to agree that the *warranted assertability* conditions of knowledge attributions vary with the context. (Indeed, as we’ve seen, they attempt to use that fact in explaining away the allure of contextualism.) And it’s clear that this is true not only of knowledge attributions, but of assertion generally... . And no “Warranted Assertability Maneuver” can be wielded against us here, for the strong intuitions we are utilizing are themselves judgments concerning warranted assertability (pp. 187-8, italics in original).

As should be clear, DeRose's claim here is that the conditions for warranted assertion vary with the context – in the very sense of “warranted assertion” that is at stake in the dispute over warranted assertability maneuvers. But that's not the notion of warranted assertability which derives from the knowledge account of assertion.

It might be thought that DeRose does not *need* to equivocate here, because (it might be suggested) even if we work only with warranted assertability_{KAA}, there is contextual variability in the conditions one must meet in order to make a warranted_{KAA} assertion. However, once we explicitly identify the relevant conditions for warranted_{KAA} assertion — namely, that the speaker knows the truth of p — it is hardly obvious that these conditions *are* contextually variable. Indeed, it is no more obvious than that the conditions for knowledge are contextually variable themselves. Perhaps it is rather the case that in *this* sense of “epistemically warranted assertion,” many of our ordinary assertions are not warranted.

II.

In fact, that many of our ordinary assertions are not epistemically warranted_{KAA} is just what the invariantist proponent of the knowledge account of assertion should suggest. To bring this out, I want to consider a line of thought which DeRose might try. We can distinguish two components of the conditions for warranted assertability_{KAA}: first, the truth requirement; second, everything else that is required in order for one to know that p. For convenience, call the latter the “justification component”. (This is just to give it a label; I don't mean to import any substantive epistemological theory here.)

And use “epistemic position” as a label for the features of a person which are relevant to whether s/he has satisfied the justification component. One might then argue as follows:

1. Whether a given epistemic position regarding P is sufficient for proper assertion of P varies with contextual factors. (Empirical fact)
2. Necessarily, one can properly assert that P if and only if one’s epistemic position regarding P is good enough to satisfy the justification component of knowledge. (Knowledge account of assertion)
3. So the standards one must meet to satisfy the justification component of knowledge are contextually variable. (By 1, 2)
4. So contextualism is correct.

This argument appears to evade the objection offered so far.¹⁰ But it is far from airtight.

To help make the argument fully explicit, I want to characterize an intuitive notion of epistemically “appropriate” or “proper” assertion. Say that an assertion is *epistemically acceptable* in a given context just if it would not be appropriate (in a broad, undifferentiated sense independent of considerations of politeness) to object to it on epistemic grounds in that context, e.g., by saying something like, “But you don’t know that!” or “But you don’t have adequate evidence for that!” This is the sense of epistemically “appropriate” or “proper” assertion in which, in the earlier example, my assertion that Kerry will win the election is clearly appropriate. I will reserve the term *warranted assertion* for the notion of epistemically non-defective assertion which is yielded by the knowledge account of assertion (that is, warranted assertion_{KAA}). As before, I will use the term *epistemic position* to refer to what is evaluated in order to

determine whether a person satisfies the conditions for knowledge other than the truth requirement (the “justification component”, as I previously termed it).

Using this terminology, the argument can be formulated as follows.¹¹

1. There are epistemic positions with the following feature: in some conversational contexts an assertion would be epistemically acceptable when the speaker is in that epistemic position regarding the content of that assertion, but in other situations it wouldn't.
2. So there can be contextual variability in whether a given epistemic position is adequate for epistemically acceptable assertion.
3. Whether a given epistemic position is adequate for epistemically acceptable assertion is a matter of whether one has satisfied the justification component of the conditions for warranted assertion_{KAA}.
4. So there is contextual variability in the standards which one must meet in order to satisfy the justification component of the conditions for warranted assertion_{KAA}.
5. But the standards which one must meet in order to satisfy the justification component of the conditions for warranted assertion_{KAA} are identical to the standards which one must meet in order to satisfy the justification component of the conditions for knowledge.
6. So the latter standards are contextually variable as well.
7. So contextualist approaches to knowledge are correct.

Notice that this argument requires that *three* sets of standards be identified: (i) the standards for epistemically acceptable assertion, (ii) the standards for satisfaction of the justification component of the conditions for warranted assertion_{KAA}, and (iii) the standards for satisfaction of the justification component of the conditions for knowledge. On the assumption that the knowledge account of assertion is right to hold that knowledge is *the only* constitutive norm for the practice of assertion, the identification of standards (ii) and (iii) (step 5 of the argument) is acceptable. The crucial question, then concerns step (3) of the argument – the identification of standards (i) and (ii). If it is successfully challenged, then the invariantist could accept steps (1) and (2) as well as the knowledge account of assertion *without being forced to accept claim (4)*. The invariantist could thus avoid DeRose’s contextualist conclusion.

There is a gap here which can be exploited by the invariantist. As Williamson insists, the notion of warranted assertion that can be derived from his account is a *technical notion*: the notion of an assertion which satisfies the constitutive rule of the practice of assertion. An assertion satisfies that rule when (and only when) certain epistemic conditions are met. However, when the rule is not satisfied because those conditions are not met, it may not be appropriate or acceptable, in the particular practical and conversational context, to criticize or object to the assertion on account of this deficiency. This allows slack between the notion of “epistemically acceptable” assertion – an assertion which is not appropriately objected to or criticized on epistemic grounds in a particular conversational context -- and the notion of warranted assertion generated by Williamson’s account.

Consider, then, the uncontroversial data: whether one's assertion is epistemically acceptable (in the technical sense defined above) depends on features of the context. Once we recognize the conceptual gap between the two sets of standards, the question becomes one of the best theoretical explanation of the variability in the standards for epistemically acceptable assertion. Within the framework of the knowledge account of assertion, there are two places where contextual factors could have an affect that would yield the uncontroversial data. First, they could influence the standards one must meet in order to satisfy the justification component. The uncontroversial data would then result if in our conversational practice we tried to follow the dictates of contextually set standards for warranted assertion_{KAA}. (This is, in effect, DeRose's suggestion.) Second, contextual factors (such as the conversational participants' practical goals and circumstances) might affect the conversational participants' attitude towards the standards for warranted assertion_{KAA} – the extent to which they take them seriously or think that meeting them matters. On this explanation of the variability, we would regard an assertion as epistemically acceptable – even if we don't judge that the speaker satisfies the justification component -- when, and because, the speaker's being in a strong epistemic position *isn't important to us*. If this were the correct explanation, then the standards for warranted assertion_{KAA} might well be context invariant; an invariantist could thus appeal to this potential source of contextual variability in our assertive practice in order to explain the uncontroversial data.

The latter explanation is consonant with a broader phenomenon relating to constitutive rules. As Williamson observes, considerations about the purposes for which we are engaging in an activity sometimes lead us to be lax in applying the constitutive

rules of that activity.¹² Suppose, for example, that my son and I are playing a game of chess. The way *we* play chess, he (unlike me) is allowed not to move a piece after touching it, to “take back” bad moves, and even to go back several moves in order to avoid a checkmate which has in fact occurred. An observer might object, “But those are all clear violations of the rules of chess. If you allow all that, then you aren’t playing chess at all!” But this is wrong-headed. We are playing chess, but incorrectly. And that’s okay, given our purposes. There is no reason, in these circumstances, to insist that the rules be followed to the letter. So when my son fails to move one of his pieces after touching it, what he does is – in the context of our game – reasonable, appropriate, and warranted, in the sense that it is not appropriately subjected to criticism or correction in that context. This is not because he reasonably thought that what he was doing was in accordance with the rules. He might know that it isn’t. Rather, it is because for our purposes here and now, in this chess game, some of the rules of chess don’t matter much at all. Likewise, such considerations can sometimes lead us to require *more* than the constitutive rules require. When I’m playing chess with my son, we sometimes decide that I will start the game without my queen. Again, we’re playing chess, but incorrectly. We have good reason to do so. Any objection would be out of place.

A proponent of the knowledge account should opt for an analogous explanation of our assertive practice if there are cases in which the conversational participants would legitimately regard an assertion as epistemically acceptable – as not appropriately criticized on epistemic grounds – even though they would *not* regard the speaker as being in an epistemic position that is adequate for knowledge. For such cases would indicate two things: first, that the conditions for epistemically acceptable assertion are not

extensionally equivalent to the conditions for warranted assertion, and, second, that there are cases in which – despite the contextual variability in the standards for epistemically acceptable assertion -- the conversational participants do not regard the standards for the justification component as shifting with the standards for epistemically acceptable assertion. It would be plausible, in relation to such cases, to suggest that the conversational participants are rather *ignoring* the constitutive norms of assertion. And in the absence of some further evidence and explanatory machinery, it would therefore be perfectly plausible for the invariantist to maintain that the standards for knowledge remain fixed, while the contextual variability in the standards for epistemically acceptable assertion is to be explained in the suggested terms.

This line is in fact the more plausible position, if one accepts the knowledge account of assertion. For there are clear examples of precisely the relevant sort. The previous example of my assertion that Kerry will win the election is a case in point. No one in the context thinks that I know that Kerry will win, and no one thinks that I am in an adequate epistemic position to know it. Nonetheless, they would all agree that my assertion is epistemically appropriate in that context; it would be completely inappropriate, given the purposes of our conversation, for someone to object, “But you don’t know that!” This appropriateness is not the “secondary warrant” that comes from reasonably taking oneself to meet the conditions for warranted assertion, since in this sort of case one *doesn’t* take oneself to meet those conditions.¹³ The contextualist approach thus completely fails for this sort of case. What a proponent of the knowledge account of assertion should say about such a case is that the standards for epistemically acceptable

assertion have dropped because satisfying the constitutive norm of assertion doesn't matter very much in that context.¹⁴

Of course, it remains possible for the contextualist to utilize this explanation for this sort of case but to claim that the contextualist explanation remains the best available for *other* cases of shifting standards for epistemically acceptable assertion. However, in order for this move to be more than an abstract possibility, the contextualist will need to specify the relevant cases. In particular, the contextualist will need to provide some non-controversial examples for which contextualism can account while invariantism cannot, examples in which the standards for epistemically appropriate assertion differ while: 1) the difference cannot plausibly be explained in terms of differences in the degree to which adherence to the knowledge rule for assertion matters in the different conversational settings, and 2) the difference cannot plausibly be explained in terms of more general conversational principles in interaction with differences in the conversational settings. However, it is doubtful that such examples are available. For instance, it might be proposed that the invariantist cannot explain the way in which the contemplation of uneliminated doubts can sometimes shift the standards for epistemically acceptable assertion. But this proposal seems to be incorrect. In some cases, the invariantist can plausibly reject the supposed datum by arguing, for instance, that the mere contemplation of the possibility that I am a brain in a vat does not in fact induce any tendency to withhold the assertion that I have hands. In other cases, the invariantist can appeal to specifically epistemological considerations to explain the data, as when — in a situation in which meeting the knowledge rule is important — a speaker is reminded or learns of considerations which provide a good reason for doubting the claim in question.

In still other cases, the invariantist can appeal to shifts in the extent to which adherence to the knowledge rule matters, claiming that the appeal to reasons for doubt can function to suggest, impose, or enforce strict adherence to the rule.¹⁵

There is room for debate regarding these proposals. The crucial lesson here is just that these empirical issues are what needs to be discussed. There isn't a direct argument to the truth of contextualism from the knowledge account of assertion and the fact that the standards for epistemically acceptable assertion shift with the conversational context.

III.

I have suggested that DeRose's argument is vulnerable to something very much like a WAM. DeRose maintains that the contextual variability in whether an assertion is epistemically acceptable shows that there is contextual variability in the epistemic position required for knowledge. The argument is facilitated by a use of the term "warranted assertion" which is supposed to span the gap between the conditions for epistemically acceptable assertion and the conditions for knowledge. But I have argued that an invariantist can plausibly charge that the notion of warranted assertion that derives from the knowledge account of assertion doesn't do the needed work. The undeniable contextual variability is merely a matter of what I have termed "epistemically acceptable assertion"; it does not entail contextual variability in the standards one must meet in order to make an epistemically non-defective assertion or in the standards for knowledge.

A WAM functions by suggesting that an account of the conditions under which we appropriately use a given term cannot be directly transformed into an account of the truth conditions (semantics) of ascriptions of that term. My objection is not a WAM,

because there is no particular term whose appropriate use is under scrutiny. Rather, my objection suggests that one cannot directly move from an account of when we would appropriately regard an assertion as not open to objection or criticism along epistemic lines to an account of the conditions under which an assertion is epistemically non-defective. So the fact that an account of the former will involve contextual-variability does not indicate that an account of the latter should do so as well. The objection is not merely that there is a gap “in principle”. First, the objection is supported by examples: cases in which we would appropriately regard an assertion as not open to appropriate epistemic criticism even while granting that the assertion is not epistemically non-defective. Second, the proposed anti-contextualist account of the data is an instance of a much more general phenomenon relating to constitutive rules: how seriously we take the constitutive rules of a practice on any given occasion depends on our purposes in engaging in the practice and our circumstances, and in some cases we may demand even more than is demanded by the rules of the practice. Thus, our actual behavior – acceptable and appropriate behavior – in engaging in the practice will deviate from the constitutive norms in ways that can be characterized in terms of more general considerations not specific to the practice in question. This applies not just to our linguistic behavior, but to our participation in any rule-governed practice. Thus the objection meets DeRose’s desiderata for a successful WAM, to the extent that they are applicable to this case: the claimed deviation between the rules governing the practice and what we see in our (appropriate and acceptable) behavior can be accounted for in terms of general principles; the objection is not simply an ad hoc maneuver involving special principles tailored to the specific case.

I haven't addressed all of DeRose's challenges to the invariantist.¹⁶ My concern here has not been to defend invariantism on all fronts, nor to establish its correctness, but rather to determine whether there is a direct argument for contextualism from the knowledge account of assertion. I conclude that there is a plausible way to combine invariantism with the knowledge account; whatever difficulties the combined view encounters derive from challenges which face the invariantist anyway.¹⁷

References

Author. [year] Publication 1.

Blackson, Thomas A. 2004. "An Invalid Argument for Contextualism," *Philosophy and Phenomenological Research*, vol. 68 no. 2, pp. 344-5.

DeRose, Keith. 2004. "The Problem with Subject-Sensitive Invariantism," *Philosophy and Phenomenological Research*, vol. 68 no. 2, pp. 346 – 50.

2002. "Assertion, Knowledge, and Context," *The Philosophical Review*, vol. 111, no. 2, pp. 167-204.

1992. "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research*, vol. 52, pp. 913 – 929.

Rysiew, Patrick. 2001. "'The Context-Sensitivity of Knowledge Attributions", *Noûs*, Vol. 35, No. 4 (December 2001): 477-514.

Williamson, Timothy. 2001. *Knowledge and Its Limits*, Oxford University Press.

¹ DeRose traces this view to G.E. Moore (1962), p. 277. For other references, see DeRose (2002), fn. 20, p. 199.

² Williamson (2001), chapt. 11.

³ One lacuna in DeRose's argument has recently been noted (Blackson, 2004): as explicitly stated, it does not yield a reason to prefer *contextualism* over *subject-centered invariantism*, which holds that the truth conditions of knowledge ascriptions vary with the practical and conversational context of the *subject* of the knowledge ascription. DeRose has now filled the gap, explaining why he finds subject-centered invariantism unappealing (DeRose, 2004). However, a more basic issue still remains, as I try to bring out here.

I defend a form of invariantism in my (author, year of publication).

⁴ In this presentation of the argument, I ignore the distinction between contextualism and subject-centered invariantism.

⁵ In particular, there is no rule specific to assertion which would ever make it the case that in order to make a warranted assertion, one must be in a *better* epistemic position than that required for knowledge. DeRose does not explicitly defend this claim. As he writes, "I will join Williamson in holding that this is the only rule governing assertion that has to do with asserting only what one is positioned well enough with respect to – or, when using the other form of the knowledge account of assertion, that the strength of the position that one represents oneself as being in when one asserts that P is just that of knowing that P, nothing more or less" (180). It should be clear that without some claim along these lines, DeRose's argument is a non-starter, since such a claim is necessary in order to yield the biconditional which underwrites step (3) of the argument.

⁶ Strictly speaking, it is incorrect by the contextualist's lights to talk simply in terms of "the standards for knowledge" or to say that the standards for knowledge shift with the

context. Rather, the issue should be put metalinguistically. I ignore this complication here and in the remainder of this paper, since it is irrelevant to the issues under discussion and requires unnecessarily cumbersome formulations.

⁷ DeRose writes: “If one represents oneself as knowing that P by asserting P, then, to avoid falsely representing oneself, one should follow the rule of asserting only what one knows” (p. 180). This is true enough. But the consequent yields a norm characterizing *warranted assertability* only if warranted assertion requires not falsely representing oneself. Given the claim that one represents oneself as knowing P by asserting P, the weakest premise needed to produce the desired result (i.e., that one is never (epistemically) warranted in asserting P unless one knows that P) is this: One is never (epistemically) warranted in representing falsehoods as being the case. But it is quite unclear whether this premise is true, since it is unclear just what is meant by “represent” in this context. Since what is represented is being treated as distinct from what is asserted, the necessary premise must be broader than “One is never warranted in asserting what is false” (which is itself controversial). If “representing” is a species of implicature, however, then the premise would seem to be false, since one can be epistemically warranted in generating false implicatures. Moreover, there is no obvious reason why someone who is committed to the knowledge account of assertion must also be committed to this premise. So it is hardly obvious that this version of the knowledge account straightforwardly yields a characterization of epistemically warranted assertion. More work needs to be done.

⁸ Which DeRose evidently thinks it sometimes does, since he thinks that some WAMs are successful.

⁹ One further drawback of DeRose’s proposal is worth noting as well. If we identify the notion of warranted assertion that derives from Williamson’s account with the familiar notion of appropriate or warranted assertion – the notion which is supposed to roughly correspond to our ordinary judgments of appropriate assertion – then we will be in possession of a quick argument against the skeptical invariantist, as follows:

1. It is sometimes appropriate or warranted to assert that p.
2. If one is warranted in asserting that p, then one knows that p.
3. So one sometimes knows that p.
4. So skepticism is false.

But this argument seems *too* quick. The problem, it seems, is an equivocation between two quite different senses of “warranted assertion”. Here, however, the equivocation does not concern whether *truth* is required for warranted assertion, but rather whether the norms of warranted assertion could be known to be satisfied *even if skepticism were true*. On the familiar notion of warranted assertion invoked by Stroud and Unger in defense of the skeptic, it is assumed that even if skepticism were true, one could know that one’s assertion was warranted or appropriate. But that can’t be the notion of warranted assertion that derives from the knowledge account of assertion.

¹⁰ It is worth noting that if one accepts Williamson’s view of knowledge, then one cannot offer this argument. According to Williamson, knowledge is a non-analyzable, factive mental state. Consequently, on his view the truth requirement and the “justification component” cannot be distinguished in the relevant way, so one cannot hive off the truth component in the way this line of argument requires. (Williamson himself, however,

rejects contextualism.) Likewise, one can't argue in this fashion if one accepts the view that knowledge requires an epistemic position that is truth-entailing.

¹¹ Again, my presentation of the argument will ignore the distinction between contextualism and subject-centered invariantism.

¹² Williamson (2001), pp. 256, 258-9.

¹³ DeRose introduces this notion of secondary warrant in the course of arguing (quite correctly) that the knowledge account of assertion is not threatened by the fact that when a speaker reasonably, but incorrectly, takes herself to know something, there is a sense in which her assertion is appropriate.

In many other cases, a speaker reasonably thinks she knows, though in fact she doesn't, for what she believes is not even true. The knowledge account of assertion would lead us to expect that though such speakers are breaking the rule for assertion, their assertions are warranted in a secondary way, since they reasonably take themselves to know what they assert. Thus, our sense that such speakers are at least in some way asserting appropriately does not falsify the knowledge account of assertion, which leads us to expect that what speakers properly assert in one way (at least relative to the matter of asserting only what they're positioned well enough with respect to) is what they in fact know, and that what they properly assert in another way is what they reasonably take themselves to know (p. 199 fn. 23).

In addition to the example discussed in the main text above, there are a variety of cases of epistemically acceptable assertion to which this notion of "secondary warrant" does not apply. Many predictions and retrodictions are of this sort. So, too, is Williamson's

example of a person who yells out, “There’s your train,” knowing he doesn’t know it, because it is probably your train and you have only moments to catch it (Williamson, p. 256). Urgency and convenience can lead us to ignore the constitutive rule of assertion just as much as can the allure of idle chatter.

¹⁴ Cf., Williamson’s comment: “When assertions come cheap, it is not because the knowledge rule is no longer in force, but because violations of the rule have ceased to matter so much” (2001, p. 259).

¹⁵ The idea here would be simply this. In a great many ordinary conversational contexts, even ones in which imparting information is at stake, strict adherence to the knowledge rule is not demanded; an assertion often can be epistemically appropriate if the speaker comes close enough to possessing knowledge for current conversational purposes.

Sometimes, however, strict adherence to the knowledge rule is demanded by the practical and conversational setting, and the speaker can be reminded of this fact by the appeal to possibilities of error which the speaker recognizes he would have to be able to rule out in order to possess knowledge. For instance, imagine a situation in which a great deal rides on a check’s being deposited before Monday morning. The speaker recently found the bank open on a Saturday and has no indication that the bank would be changing its hours. In such a situation, the speaker might be reminded that the practical and conversational setting demands strict adherence to the knowledge rule by a conversational participant’s noting that after all, banks do sometimes change their weekend hours suddenly and without notice. Once this has been noted, and strict adherence to the knowledge rule has been enforced, it would indeed seem epistemically unacceptable for the speaker to assert, on the grounds available to him, that the bank will be open on Saturday. (This is, of

course, a modified version of DeRose's well-known "Bank Case", which originally appeared in his (1992).)

¹⁶ In particular, I haven't addressed his extremely important question, Why (if nonskeptical invariantism is correct) it is not only *not* acceptable to assert "I know that P" in a high standards context when one *does* know that P (and even knows one knows it), but also *acceptable* to assert "I don't know that P" (which is false and known to be such). Any invariantist view must answer this question, or it must deny that there is contextual variability in the standards for appropriate knowledge ascription. For an attempt in the former direction, see Rysiew (2001). I myself favor the latter path. (See [author, year of publication].)

¹⁷ [acknowledgements]