

# Assertion, Knowledge, and Context

forthcoming, *Philosophical Review*

Keith DeRose

*Yale University*

This paper brings together two positions that for the most part have been developed and defended independently of one another: contextualism about knowledge attributions and the knowledge account of assertion.

The positions under discussion are both located in the area of overlap between epistemology and the philosophy of language and they have both received a good deal of attention in recent years. But there is further reason for surprise that more has not been done to bring them together. The chief bugaboo of contextualism has been the concern that the contextualist is mistaking variability in the conditions of warranted assertability of knowledge attributions for a variability in their truth-conditions. This strongly suggests that a serious assessment of contextualism will demand a discerning look at the question of what it takes for a speaker to make a warranted assertion. And it turns out that the knowledge account of assertion — according to which what one is in a position to assert is what one knows — promises to provide a large and important part of the answer to this question.

It also turns out that the knowledge account of assertion dissolves the most pressing problem confronting contextualism, and, on top of that, provides a powerful positive argument in favor of contextualism. Or so I will argue.

It will take a bit of work to uncover contextualism's most pressing problem, since the critics of contextualism have not themselves been very proficient in identifying it. In Part I, then, after briefly explaining contextualism, I will present its main problem — the “Generality Objection”, as I will call it — and will show why this is such a serious concern. In Part II, after the knowledge account of assertion and the grounds that support it have been briefly explained, I will show its need for contextualism and present the positive argument for contextualism that it provides. In Part III, I will use the knowledge account of assertion to squash the Generality Objection and will also cast other shadows over the prospects for anti-contextualism.

## **1. Contextualism and its Main Problem: From the Warranted Assertability Objection to the Generality Objection**

1.1. Contextualists and Their Cases. Contextualists hold that the 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) fluctuate in certain ways according to the context 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) for such a statement to be true. In some contexts, “S knows that P” requires that S have a true belief that P and also be in a very strong epistemic position with respect to P, while in other contexts, the same sentence may require for its truth, in addition to S’s having a true belief that P, only that S meet some lower epistemic standards.<sup>1</sup>

Contextualist accounts of knowledge attributions have been almost invariably developed with an eye toward providing a response to philosophical skepticism. For some skeptical arguments threaten to show that we know nothing (or little) of what we think we know, and thus we are wrong whenever (or almost whenever) we say or think that we know this or that. But according to typical contextualist analysis, the skeptic, in presenting her argument, executes conversational maneuvers by which she raises, or at least threatens to raise,<sup>2</sup> the standards for knowledge to a level at which we count as knowing nothing or little, and according to which she can truthfully say that we don’t know. Thus, the contextualist hopes to explain the persuasiveness of the skeptic’s attack, but in a way that makes it unthreatening to the truth of our ordinary claims to know. For the fact that the skeptic can install (or threaten to install) very high standards for knowledge that we do not meet has no tendency to show that we do not meet the more relaxed standards that govern more ordinary conversations. That contextualists have so consistently and so quickly turned their attention toward providing such a response to skepticism would lead one to conclude that the predominant recommendation of contextualism is its ability to provide such a response.

However, support for contextualism should also — and perhaps primarily — be looked for in how “knows” and its cognates are utilized in non-philosophical conversation. If our non-philosophical usage of the relevant terms did support contextualism, that would seem to be important evidence in favor of the theory, and contextualist treatments of skepticism could

proceed with a lot more leverage behind them. On the other hand, the contextualist's appeal to varying standards for knowledge in his solution to skepticism would rightly seem unmotivated and ad hoc if we didn't have independent reason from non-philosophical talk to think such shifts in the content of knowledge attributions occur.

But it can seem that there is no real problem for the contextualist here. It's an obvious enough observation that what we will count as knowledge in some non-philosophical contexts won't pass as such in others. As J.L. Austin observed, in many ordinary settings we are easy, and say things like (Austin's example): "I know he is in, because his hat is in the hall." But, even with no philosophers in sight, at other times speakers get tough and will not claim to know that the owner was present based on the same evidence; as Austin notes: "The presence of the hat, which would serve as proof of the owner's presence in many circumstances, could only through laxity be adduced as a proof in a court of law."<sup>3</sup> We needn't invoke anything as unusual as a high-stakes court case to find such variation — as I'm sure Austin realized. A wide variety of different standards for knowledge are actually used in different ordinary contexts. Following Austin's lead, the contextualist will appeal to pairs of cases that forcefully display this variability: "Low standards" cases in which a speaker seems quite appropriately and truthfully to ascribe knowledge to a subject will be paired with "high standards" cases in which another speaker in a quite different and more demanding context seems with equal propriety and truth to say that the same subject (or a similarly positioned subject) does not know.

To make the relevant intuitions as strong as possible, the contextualist will choose a "high standards" case that is not as ethereal as a typical philosophical discussion of radical skepticism: a "skeptical hypothesis" may be employed, but it will be much more moderate than the playthings of philosophers (brains in vats, evil geniuses, or whatnot).<sup>4</sup> And it makes the relevant intuitions more stable if the introduction of the more moderate skeptical hypothesis and the resulting raise in epistemic standards are tied to a very practical concern, and thus seem reasonable given the situation. It also helps if this is all accepted as reasonable by all the parties to the conversation. Thus, in the pair of cases I have used — my Bank Cases<sup>5</sup> — one character (myself, as it happens), claims to know that the bank is open on Saturday mornings in the "low standards" case. This belief is true, and is based on quite solid grounds: I was at the bank just two weeks ago on a Saturday, and found that it was open until noon on Saturday. Given the practical concerns involved — my wife and I are deciding whether to deposit our paychecks on

Friday, or wait until Saturday morning, where no disaster will ensue if we waste a trip to the bank on Saturday only to find it closed – almost any speaker in my situation would claim to know the bank is open on Saturdays. And, supposing “nothing funny” is going on (there has not been a recent rash of banks cutting their Saturday hours in the area, etc.), almost all of us would judge such a claim to know to be true. But in the “high standards” case, 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. (And, of course, the bank is not open on Sunday.) Given all this, my wife seems reasonable in not being satisfied with my grounds, and, after reminding me of how much is at stake, in raising, as she does, the possibility (the “skeptical hypothesis”) that the bank may have changed its hours in the last couple of weeks. This possibility is fairly far-fetched by ordinary standards, but is downright moderate when compared with the possibilities employed by philosophical skeptics, and seems worth worrying about, given the high stakes we are dealing with. Here I seem quite reasonable in admitting to her that I “don’t know” that the bank is open on Saturdays, and in endeavoring to “make sure.” Almost everyone will accept this as a reasonable admission, and it will seem true to almost everyone.<sup>6</sup>

1.2. The Warranted Assertability Objection. When in 1986 I first formulated this pair of cases and presented them to my sympathetic but critical teacher, Rogers Albritton, he admitted that they reflect how “knows” is used, but immediately raised the concern that the variability in epistemic standards that my cases display may be only a variability in the conditions under which it is appropriate to claim to know.

Such has been my experience ever since. Wisely, defenders of invariantism — Peter Unger’s good name for the denial of contextualism<sup>7</sup> — have accepted the evident facts about typical and, in the relevant sense, appropriate linguistic behavior that such case pairs display: Indeed, that is how we talk, and appropriately so. What is controversial is whether these varying standards for when ordinary speakers will attribute knowledge, and for when they’re warranted in attributing knowledge, reflect varying standards for when it is or would be true for them to attribute knowledge, and the contextualist’s claim is the more controversial one concerning a variation in truth conditions.<sup>8</sup>

Accordingly, the most influential source of resistance to contextualism is the warranted assertability objection — the charge that what the contextualist takes to be a variation in the truth-conditions of knowledge attributions is in reality only a variation in the conditions for the warranted assertability of those claims.

If the contextualist has chosen his pair of cases well, there will be a quite strong intuition about each of the assertions (both the positive ascription of knowledge in the low standards case and the denial of knowledge in the high standards setting), that it is true, in addition to being warranted. The invariantist cannot accept that both of the speakers' assertions actually are true, and so must deny a quite strong intuition. The warranted assertability objection is an attempt to explain away one or the other of these intuitions, which together are lethal to the invariantist. Perhaps the positive ascription of knowledge in the low standards case is really false, but seems true because, due to the low standards for warranted assertability that are in place there, the positive ascription of knowledge is quite appropriate, and we mistake the warranted assertability of the claim for truth. Or maybe it's the denial of knowledge in the high standards case that's false but appropriate: Due to the high standards for the warranted assertability of knowledge in place there, a positive claim that the subject knows would be unwarranted (though true), and it's the denial of knowledge that is appropriate (though false). The invariantist wielding the warranted assertability objection may often sensibly decline to say which of our intuitions is wrong. About some cases, she might admit that it's hard to say whether the subject knows or not. She will merely claim that the fact that a fairly clear appearance of truth attaches to both the positive ascription of knowledge and the denial of knowledge is due to the warranted assertability that both of the claims enjoy, each in its own context.

The warranted assertability objection to contextualism, then, is an example of what we can call warranted assertability maneuvers (WAMs). Such a maneuver 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.

Though that general insight is sound, and though some instances of this type of maneuver have been correct, there is a distinct danger that WAMs can be carried too far and misused, as can be illustrated by the below fiction.

1.3. The Myth of Jank Fraction: A Cautionary Tale. Suppose that, for some reason we needn't concern ourselves with, some philosophers toward the middle of the 20<sup>th</sup> Century took a liking to the notion that all it takes to know a proposition is that one believe it.<sup>9</sup> We may suppose that their view came to be called the "Equivalence Thesis" because, according to it, "S knows that P" is equivalent to — has the same truth-conditions as — "S believes that P."

You don't have to be a veteran of the Gettier wars to be immediately struck by counter-examples to this sorry theory. Some accounts of knowledge face the problem of Lucky Louie: They counter-intuitively count as knowers subjects who believe a proposition in some baseless way, say, by means of a wild guess, but by sheer luck, turn out to be right. The Equivalence theorists faced not only that, but also the problem of Unlucky Ursula, whose belief is not only based on a pure guess, but is also false! Indeed, we may suppose that among the many counter-examples hurled at the hapless Equivalence theorists, ones involving subjects with false beliefs were thought to be the most pressing.

Enter Jank Fraction, who defended the Equivalence Thesis against its most pressing problem by adding to the Thesis the auxiliary hypothesis that an assertion of "S knows that P" generates an implicature to the effect that P is true, and thus a warranted assertability condition of "S knows that P" is that P be true. This, however, is only an implicature and a warranted assertability condition of "S knows that P," according to Fraction, not a truth condition. Thus, according to Fraction, the Equivalence Thesis is correct about the truth conditions of "S knows that P." Fraction argued that, since his "Supplemented Equivalence Thesis" — the Equivalence Thesis regarding the truth-conditions of knowledge attributions together with his additional claim about those sentences' warranted assertability conditions — predicts that "S knows that P" will be unassertable where P is false, it can explain why we think such a knowledge attribution is false where the relevant P is false, and thus our intuitions that false beliefs don't amount to knowledge really don't hurt his theory. We may finally suppose that Fraction closed his 1979 defense of the Equivalence Thesis, "On Assertion and Knowledge Attributions," with these words:

In my view this puts a very different complexion on certain putative counter-examples to the Equivalence thesis. We saw, for instance, how the Equivalence theorist must hold that “Unlucky Ursula knows that Carter was re-elected” is true. But what is it that is immediately evident about this putative counter-example? Surely that it has very low assertability. But the Supplemented Equivalence theory explains this, and what a theory well explains cannot be an objection to that theory.<sup>10</sup>

1.4. Lame WAMs and the Warranted Assertability Objection to Contextualism. There’s something fishy about Fraction’s maneuver, as I hope you can sense. Fraction’s defense should not be allowed to mitigate the negative verdict we are inclined to reach against the Equivalence Thesis. The sense that this WAM is unsuccessful is partly due to the fact that it is being offered in defense of such a loser of a theory — a theory for which it is difficult to imagine what positive support it might have.<sup>11</sup> But it is important to notice the deeper reasons why this defensive maneuver should be allotted no force here. For now, we’ll focus on this important reason: Fraction’s WAM is an instance of a general scheme that, if allowed, could be used to far too easily explain away the counterexamples marshaled against any account of the truth-conditions of sentences in natural language. Whenever your theory seems to be wrong because it is omitting a certain truth-condition — as Fraction’s theory of “S knows that P” seems to be wrongly neglecting to include a condition to the effect that P must be true for a sentence ascribing knowledge of P to a subject to be true — you can simply claim that assertions of the sentences in question generate implicatures to the effect that the condition in question holds. Thus, you claim, the alleged condition is a warranted assertability condition of the relevant sentences: to assert such a sentence where the condition does not hold will be unwarranted because the speaker will generate a false implicature. Your critic, you claim, is mistaking the falsity of an implicature that is generated by the sentences in the relevant circumstances for the falsity of what the speaker says in uttering the sentence, and is thus mistaking what is only a warranted assertability condition of the sentence for a truth-condition of it.

So armed, you are ready to defend, say, an account of “S is a bachelor” according to which S’s being unmarried is not a truth condition of the sentence. Of course, real and imagined

married men will provide a slew of apparent counter-examples to your besieged theory. But rather than modifying your account to remedy your apparent mistake, you can follow Fraction's example and supplement your theory with the auxiliary claim that S's being unmarried is an implicature generated by an assertion of "S is a bachelor," and is therefore a warranted assertability condition for the sentence, though it is not a truth-condition. Thus, your supplemented theory explains why "S is bachelor" seems false when said of a married man. And what a theory well explains...

"But wait! Aren't there WAMs that we rightly do give credence to?"

Yes. Here's one that I have used<sup>12</sup> – I believe effectively. When a speaker knows that P, it can seem somehow wrong, and to some it will seem downright false, for him to say "It's possible that P<sub>ind</sub>."<sup>13</sup> Suppose, for instance, that Ringo wants to borrow a certain book, and he asks Paul whether the book is in Paul's apartment. If he knows full well that the book is there, it would be somehow wrong for Paul to answer, "It's possible that it's there." Indeed, pre-theoretically, many feel some tendency to judge that Paul would then be saying something false. (Many will find this tendency in competition with an opposing intuitive pull toward the verdict that the statement is true — though somehow misleading and wrong.) Such tendencies could tempt one toward a "Don't Know Either Way" (DKEW) account of "It's possible that P<sub>ind</sub>", according to which:

DKEW: S's assertion, "It's possible that P<sub>ind</sub>" is true iff (1) S doesn't know that P is false and (2) S doesn't know that P is true.

But this temptation should be resisted, I believe, for the correct account lies down the simpler "Don't Know Otherwise" (DKO) path:

DKO: S's assertion, "It's possible that P<sub>ind</sub>" is true iff S doesn't know that P is false.<sup>14</sup>

According to DKO, Paul is asserting the truth in our example. To the extent that one does have the intuition that Paul is saying something false, the backer of DKO can attempt to explain away that intuition, and other intuitions to the effect that speakers who know that P speak



falsely in saying that P is possible, as follows. Both “P” and “I know that P” are stronger than — they imply but are not implied by — “It’s possible that P<sub>ind</sub>,” according to DKO. And there’s a very general conversational rule to the effect that when you’re in a position to assert either of two things, then, other things being equal, if you assert either of them, you should assert the stronger. Now when someone like Paul knows that P, then they’re in a position to assert that P — and they’re often in a position to assert even that they know that P. Thus, by the “assert the stronger” rule, they should assert one of those stronger things rather than the needlessly weak “It’s possible that P<sub>ind</sub>.” To assert that weak possibility statement is unwarranted and generates the false implicature that the speaker doesn’t know that P by the following Gricean reasoning, which is based on the assumption that the speaker is following the “assert the stronger” rule: “If he knew that P, he would have been in a position to assert something stronger than ‘It’s possible that P<sub>ind</sub>’, and thus would have asserted some stronger thing instead. But he did assert ‘It’s possible that P<sub>ind</sub>’, not anything stronger. So he must not know that P.” We can easily mistake the falsehood of the implicature generated by such an assertion of “It’s possible that P<sub>ind</sub>” and the resulting unwarrantedness of the assertion for the falsehood of the assertion itself.

Or so the backer of DKO can claim. And plausibly so. My purpose here is not to argue for the ultimate effectiveness of this WAM, or for the correctness of the DKO approach (though the maneuver is effective and the approach is correct), but to use this WAM as an illustration of one that, unlike the ones we considered above, has at least some force, as I hope the reader can sense. We will later see further reasons for thinking that this WAM is quite credible. But for now, note this central reason for taking it seriously — at least more seriously than Fraction’s lame maneuver and our imagined WAM used in defense of the strange theory of “bachelor”. A vitally important contrast here involves how the wielders of these WAMs explain the generation of the implicatures to which they appeal. The defense of DKO utilizes a very general rule of conversation — “Assert the Stronger” — which applies to assertions of any content.<sup>15</sup> This general rule, together with DKO’s account of the content of “It’s possible that P<sub>ind</sub>”, generates the implicature that S doesn’t know that P, by reasoning of a familiar, Gricean style. By contrast, Fraction’s WAM and the defense of the crazed theory of ‘bachelor’ each resort to positing a special rule attaching only to assertions involving the relevant term — ‘knows’ or ‘bachelor’. This is what rightly arouses our suspicions that any theory which omits what is in fact a truth-condition for a type of assertion could just as well execute this maneuver. But it’s

not so easy to generate the implicatures you need to deflect the apparent counter-examples to your theory by means of general conversational rules that can be tested on very different sentences. If your theory of the truth conditions of the relevant sentences, together with such general rules, really does predict the occurrence of apparent counter-examples to your theory, that would seem to legitimately mitigate the force of those apparent counter-examples.<sup>16</sup>

The problem anti-contextualists face — or, as we will see later, just one of the daunting problems they face — as we turn back to the warranted assertability objection against contextualism, is that, as this objection has so far been formulated, it does no such thing. No defender of invariantism has proposed an account on which general rules of conversation together with their proposed invariantist account of the truth-conditions of knowledge attributions predict the pattern of varying conditions of warranted assertability of knowledge attributions that we encounter in natural language.<sup>17</sup> Truth be told, the warranted assertability objection against contextualism often takes the form of a bare warranted assertability objection: It's simply claimed that it's the conditions of warranted assertability, rather than of truth, that are varying with context, and the contextualist is then accused of mistaking warranted assertability with truth. To the extent that defenders of invariantism go beyond such bare maneuvers, their hints point in the direction of special rules for the assertability of “knows”, like “If someone is close enough, for present intents and purposes, to being a knower, don't say that she doesn't know, but rather say that she knows.” And of course, if he's allowed to appeal to the bare possibility that warranted assertability is being confused with truth or to special rules about the term in question, even our theorist about “bachelor” can rebut the evidence against his theory.

Indeed, if we consider contrasts between forceful and forceless WAMs different from the one we have just contemplated and attempt to discern various features WAMs should possess in order to be accorded any mitigating force, and we then apply these results to the case at hand, as I've done elsewhere,<sup>18</sup> and as we'll briefly see in section 3.2, we find that the warranted assertability objection against contextualism miserably fails to meet all manner of other reasonable criteria we can discern for what it would take for a WAM to be successful.

1.5. The Generality Objection. The contextualist may simply take this all as good news and declare victory, as I did in the earlier work already alluded to (see note 18). But one can't help but think that we haven't gotten to the bottom of things yet. For very smart philosophers, who

would not be at all tempted by our imagined lame defense of the crazed theory of ‘bachelor’, nevertheless do find the warranted assertability objection to contextualism quite attractive, and I must confess that I can feel the attraction, too, though I ultimately want to reject the move. Can it really be such a loser of a maneuver? What is it about that particular appeal to warranted assertability considerations that makes it so attractive?

An explanation for that appeal is not that far to seek, especially when we keep in mind the concerns about generality just raised. Those concerns should lead us to ask: Is the phenomenon we’re dealing with particular to knowledge attributions, or do the varying epistemic standards we’ve discovered affect other, quite different, assertions as well?

Here the invariantist should be pleased to note that this variation is ubiquitous, affecting the standards, at least for warranted assertability, of assertions whose truth-conditions obviously don’t vary as epistemic standards change. Indeed, varying standards do in some way govern our use of “S knows that P”, but what about that embedded “P”, which can be just about any proposition? Are not the “high standards contexts” in which it is very difficult to warrantably assert “S knows that P” also contexts in which it becomes wrong to assert the simple “P”? In the “high standards” Bank Case, for instance, where it seems I can’t claim to know the bank is open on Saturdays, it also seems it would be wrong for me to assert, flat-out, “The bank is open on Saturdays.”

In the case of first-person knowledge attributions — “I know that P” — the assertability of the knowledge claim and of the simple “P” seem to fade away together as the epistemic standards governing a conversation become more demanding. The knowledge claim may lead the way in the move toward unassertability: It will seem more problematic than the assertion of “P” in various intermediate contexts. But the simple “P” follows closely enough behind that it’s difficult to come up with cases in which “I know that P” has become clearly unwarranted while “P” is still clearly warranted. Things are quite different with second- and third-person knowledge attributions. Here, of course, the speaker will often be in a position to assert “P” but will clearly be in no position to assert that a certain S knows that P — sometimes simply because the speaker doesn’t know anything about S. Still, if we start with a second- or third-person case in which “P” and “S knows that P” actually are both warranted, and then raise the standards in our imagination, we find that the two claims move toward unassertability together — though, again, the knowledge attribution may lead the way.

These observations give rise to a very serious challenge to contextualism that I will call the Generality Objection. There is some very general rule of conversation to the effect that one should assert something only if one is well-enough positioned with respect to that proposition to properly assert it and thus a condition on warranted assertability that one be well enough positioned with respect to what one asserts.<sup>19</sup> We may of course hope for a more specific rule along these lines – one that specifies how well positioned one must be with respect to a proposition to be able to assert it. (And we are about to discuss such a more specific rule below in Part 2.) But at this point, the Generality Objector is seeking to defend an invariantist account of the context-variability in the assertability of knowledge attributions by appealing just to the fact that there is some such requirement on warranted assertion and to the observation that in general, and not just with respect to knowledge ascribing sentences, how well positioned one must be with respect to something to be able to properly assert it is a context-variable matter. For as the Generality Objector will correctly point out, when the simple “P” becomes unassertable because we are moving into more demanding contexts, this is generally not due to any change in the truth conditions of “P”. “P” can be just about anything, and most of our assertions are obviously insensitive in their content to what epistemic standards happen to be governing their use. For example, though the content of an assertion of “The bank is open on Saturdays” may be context-sensitive in various other ways, the conditions under which it is true (as opposed to assertable) clearly don’t depend at all on what epistemic standards are in play when the sentence is asserted. But, as is generally agreed, the truth of “P” is a necessary condition of the truth of “S knows that P.” Thus, it is the reverse of surprising that “S knows that P” should become unassertable due to high epistemic standards when (or roughly when) “P” thus becomes unassertable: If one is not well enough positioned to assert that “P”, then of course one won’t be well enough positioned to assert the stronger “S knows that P”. Since “P” becomes unassertable in high-standards contexts even though there is no change in its content as we move into high-standards context, and since the drift toward the unassertability of “S knows that P” as we move into more demanding contexts is just what we would expect given that “P” displays a similar drift, why suppose the unassertability of the knowledge claim in high contexts is due to a change in content it undergoes as we move into such contexts? According to the Generality Objection, there is no good reason to suppose there is such a variation in truth-conditions of knowledge attributions.

That's a tough question, and a tough objection. To be able to answer it, we will employ the knowledge account of assertion, to which we now turn.

## **2. The Knowledge Account of Assertion – and Contextualism**

2.1. The Knowledge Account of Assertion. What is the content of the rule that prohibits us from asserting what we are not well enough positioned to assert? How strong a position must one be in with respect to a proposition to be able to properly assert it? A little reflection on these questions can quickly lead to pessimism about there being any sufficiently general answer that is any more informative than “Assert only what you are well-enough positioned with respect to,” since it's fairly clear that, as the Generality Objector observes, one is here shooting at a moving target: In some conversational contexts, one must be extremely well-positioned with respect to a proposition to be able to properly assert it, while in other contexts one can properly assert something with respect to which one is only moderately well-positioned.

Nevertheless, an impressive answer to our question has emerged, and there are good reasons to think it is correct. Appropriately enough, according to the knowledge account of assertion, one must know that P in order to be well-enough positioned with respect to P to assert it.

The knowledge account of assertion has been packaged in two basic forms. Following the basic direction of comments by G.E. Moore, most advocates tend to formulate the account in terms of a principle to the effect that when one asserts that P, one represents it as being the case that one knows that P.<sup>20</sup> But more recently, Timothy Williamson, in defending the knowledge account, has expressed it as a claim, not about what one represents as being the case in making an assertion, but about what rule governs the practice of assertion: According to Williamson, the constitutive rule of the practice of assertion is that one should assert only what one knows.<sup>21</sup> For our purposes, these are just two sides of the same coin: 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; and if assertion is governed by a rule that one should assert only what one knows, then one will represent oneself as knowing that P when one asserts that P. We can leave it open whether to follow Williamson in holding that this rule is the single constitutive

rule specific to the practice of assertion, all other rules governing assertion being consequences of this single rule, together with more general norms not specific to assertion.<sup>22</sup> I will join Williamson in holding that this is the only rule governing assertion that has to do with asserting only what one is well-enough positioned 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. What we leave open here is whether all of the other, quite different rules of assertion that don't have to do with asserting only what one is well-enough positioned with respect to — like those enjoining us to assert only what's conversationally relevant — can be derived from the knowledge rule together with other rules not specific to assertion. In either of its forms then, the knowledge account of assertion says that one is well-enough positioned to assert that P iff one knows that P.

As happens with other rules, a kind of secondary propriety/impropriety will arise with respect to this one. While those who assert appropriately (with respect to this rule) in a primary sense will be those who actually obey it, a speaker who broke this rule in a blameless fashion (one who asserted something she didn't know, but reasonably thought she did know) would in some secondary sense be asserting properly, and a speaker who asserted something she thought she did not know, but in fact did know (if this is possible) would be asserting improperly in a secondary sense.<sup>23</sup>

The case for the knowledge account of assertion has been (powerfully) made elsewhere,<sup>24</sup> and I will not here rehearse that whole case, or even any part of it in any great detail. I will, however, quickly point out what to my thinking is one of the most important recommendations of the account: that it provides a nice handling of the knowledge version of Moore's paradox and other troubling conjunctions. Famously, Moore noted the oddity of assertions of the form "P, but I don't believe that P" — this is the standard version of Moore's paradox. Less famously, but equally insightfully, Moore also noted the oddity of

1. P, but I don't know that P.

Moore's example: "Dogs bark, but I don't know that they do."<sup>25</sup> Such conjunctions are clearly consistent: Of course, P could be true without my knowing it. Yet they "clash"; they sound like contradictions. Moore's handling of this clash is built on the claim that (as it was put by later

writers in the Moorean tradition) in asserting that P, one represents oneself as knowing that P, or, as Moore himself put it: “By asserting p positively you imply, though you don’t assert, that you know that p.” (p. 277). Thus, when one asserts the first half of Moore’s conjunction, one is representing it as being the case that one knows that dogs bark. Consequently, when one goes on to say in the second half of the sentence that one does not know that dogs bark, one is saying something inconsistent with what one represented as being the case in asserting the first half. This explanation is attractive because it supports our sense that some inconsistency is responsible for the clash involved in asserting the conjunction, while, at the same time, happily removing that inconsistency from the realm of what’s asserted: The conjunction asserted is itself perfectly consistent, but in trying to assert it, one gets involved in a contradiction between one thing that one asserts, and another thing that one represents as being the case. And the knowledge account can also be used to explain other troubling conjunctions that I can’t see how to handle without the account.<sup>26</sup> These and other grounds (see note 24, above) give us strong reason to think that the knowledge account of assertion is correct.

2.2. The Knowledge Account of Assertion Contextualized. If so, how can we reconcile the correctness of the knowledge account of assertion with our earlier reason for being pessimistic about any such an account being generally right — namely, that 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? That’s no problem — for the contextualist. The context-variability in what we are positioned to assert is just what the knowledge account of assertion would lead us to expect if what counts as knowledge is a context-variable matter. The contextualist about knowledge who also accepts the knowledge account of assertion welcomes the context-variability in what we can assert, and indeed should have been quite worried if that variability was not found.

Given contextualism, the knowledge account of assertion naturally takes a relativized form: To be positioned to assert that P, one must know that P according to the standards for knowledge that are in place as one makes one’s assertion. Following David Lewis (though Lewis was writing of a connection other than that between assertability and knowledge), the contextualist can say:

I am not one of those philosophers who seek to rest fixed distinctions upon a foundation quite incapable of supporting them. I rather seek to rest an unfixed distinction upon a swaying foundation, claiming that the two sway together rather than independently.<sup>27</sup>

What of the advocate of the knowledge account of assertion who does not accept contextualism? Such a character is in serious trouble. Given invariantism about knowledge, the knowledge account of assertion is an untenable attempt to rest a madly swaying distinction upon a stubbornly fixed foundation. Less metaphorically, it is an attempt to identify what is obviously a context-variable standard (the standard for the warranted assertion of “P”) with what one claims is a context-invariable standard (the relevant truth-condition of “S knows that P,” according to the invariantist). The knowledge account of assertion demands a contextualist account of knowledge and is simply incredible without it.<sup>28</sup>

“But wait! Even the invariantist accepts that varying epistemic standards govern knowledge attributions, including, most relevantly here, first-person claims to know. Her only difference with the contextualist is that she holds that it is only the warranted assertability conditions, rather than the truth conditions, of such attributions that vary with context. So can’t the invariantist accept something at least very much like the knowledge account of assertion, but escape trouble by tying the assertability of P to the assertability, rather than the truth, of ‘I know that P?’”

No. The relevant warranted assertability condition for “P” cannot be plausibly equated with the relevant warranted assertability condition for “I know that P”. (Of course, there are many warranted assertability conditions for assertions. We are here interested only in those that pertain to one’s being well-enough positioned with respect to the proposition asserted to be able to assert it, and are ignoring other conditions – e.g., that the assertion be conversationally relevant.) The connection between knowledge and assertion that works is one that ties the relevant warranted assertability-condition of “P” with the relevant truth-condition of “I know that P”. That identification provides a strong argument for contextualism, as we will discuss below in section 2.4. But alternative potential connections, including that suggested above on behalf of the invariantist, hover about us, and are only subtly mistaken. So before putting the knowledge



account of assertion to use, we should look carefully at the connection between knowledge and assertability to make sure we have it right.

2.3. Assertability and Knowledge: Getting the Connection Right. In section 2.4, we will follow a fascinating (but largely overlooked) earlier paper by Robert Hambourger<sup>29</sup> in arguing for contextualism about knowledge from the context-variability of assertability. Unfortunately, Hambourger does not locate a secure path to get him from his solid premise to his desired conclusion, instead getting entangled in subtly false claims about the relation between knowledge and assertability. Since these misfiring attempts to connect knowledge with assertability are tempting, it's worth discussing them.

According to Hambourger, the standard, in terms of how well-positioned one is with respect to P, that one must meet to be able to properly assert that P is the same as the standard, in terms of position with respect to P, that one must meet to properly claim to know that P,<sup>30</sup> which, in turn, is equated with the standard one must meet for one's claim to know to be true.<sup>31</sup> Both equations of standards — 1) those for properly asserting that P with those for properly asserting that one knows that P, and 2) those for properly asserting that one knows that P with those for actually knowing that P — are mistaken, as I trust the below considerations will show to anyone who has deliberated over close calls over whether one is well enough positioned to claim to know that P, or should cool one's heels and only assert that P.

We are often unsure about whether we know. Sometimes, beyond merely being unsure, we positively go wrong about what we know because we are basing our judgment that we do know on some false belief about some underlying factual matter relevant to the issue of whether we know. Example: Henry quite reasonably thinks that he knows that he's seeing a barn, because he thinks he is in a very normal situation; unbeknownst to him, he's in a region teeming with fake barns and in fact doesn't know.<sup>32</sup> But we can be unsure whether we know even without being mistaken about our factual situation: The factual suppositions on which you're operating are all correct, but your grasp of your factual situation is a bit incomplete, and it's a very close call whether or not you know, as you yourself will say if asked whether you know. You just can't tell whether you know. Well, then, you're in no position to assert that you do know. But on some of those occasions, of course, you really do know: Are we to suppose that all such close calls between knowing and not knowing occur in situations where in fact one does not

know? (If so, then realizing that all close calls are actually cases of non-knowing would make it all too easy to correctly handle a close call!) But if you know, you meet the epistemic standards for actually knowing. And if, as we've already determined, you are not positioned to assert that you know, then you don't meet the standards for asserting that you know. Hence, equation 2) is invalidated.

There is a certain kind of counter-example to equation 2) that I can't give here: I can't describe in more concrete terms a situation where a subject clearly knows but clearly is in no position to assert that she knows. I had to describe the situation as involving close calls. But if the calls are close, they won't be clear. Nevertheless, we can see by the above argument that there must be some unclear cases where one knows but is not yet well enough positioned to assert that one knows.

Such unclear scenarios also give the lie to equation 1), but let us proceed in more picturesque terms this time. Imagine a situation in which you're clearly well-enough positioned with respect to P to assert both that P and that you know that P: You see a cat just a few feet in front of you, the situation is favorable, and you are clearly able to properly assert that you are seeing a cat and that you know that you're seeing a cat. Then consider a series of cases just like the one just described, except that in each case you are a bit further from, and therefore get a little worse look at, the cat. At the end of the series, you are clearly in no position to assert even that the thing you're seeing is a cat, much less that you know it to be a cat. Toward the middle of this series of cases, you get many close calls, both for whether you're able to assert that P and for whether you can properly claim to know that P. I don't think we will find any particular distance at which we get a clear counter-example to equation 1) — a distance at which it's clear that you can properly assert that P but clearly cannot properly claim to know that P. Nevertheless, as we can all sense, throughout the intermediate cases, the knowledge claim is more problematic than is the simple assertion that P, and it seems clear that tougher standards govern the propriety of the claim to know that P than govern the simple claim that P, even if the standards are too close to one another to yield any cases where a speaker clearly meets one but clearly falls short of the other. (Compare: I'm grading philosophy essays on a 100 point scale. My standards are clearly tougher for an essay meriting a score of at least 77 than for earning a score of at least 75, despite the fact that there won't be any essays which to me both clearly meet those lower standards and clearly fail to meet the higher.)

It's often tough to properly assert. Even where you have a reasonable belief that P, it's far from automatic that you're in any position to flat-out assert it. And if you feel that the issue of whether you know that P looms large in deliberating over whether you are in a position to assert that P, then, of course, being an advocate of the knowledge account of assertion, I sympathize with that feeling. Nevertheless, as I hope we can all sense, going beyond the simple assertion of P to the claim that one knows that P is, as Austin put it, "taking a new plunge" ("Other Minds," p. 99), and, contrary to equation 1), is evidently subject to higher epistemic standards.

What is true is the equation that 1) and 2) together imply: that of the standards for knowing that P with those for being in a position to properly assert that P. That equation is the knowledge account of assertion: One must know that P to be well-enough positioned to properly assert that P.

While, like equations 1) and 2), it maintains an important tie between knowledge and proper assertion, it also can explain why claiming to know is to take a greater "plunge" than simply asserting P, why it's hard to come up with clear counter-examples to equations 1) and 2) despite those equations' evident falsehood, and can perhaps even help to explain the appeal of the KK Principle.

Simple as it is, I've found the below chart nevertheless helps in making the needed points.

<u>In asserting:</u>	<u>S (of course) asserts that:</u>	<u>And S represents it as being the case that:</u>
"P"	P	S knows that P
"I know that P"	S knows that P	S knows that S knows that P

Alternatively, we can re-label our columns in terms of the truth-conditions and the relevant warranted assertability condition:

<u>S's assertion:</u>	<u>Truth-condition of assertion:</u>	<u>Warranted-assertability condition of assertion:</u>
“P”	P	S knows that P
“I know that P”	S knows that P	S knows that S knows that P

As the charts show, the knowledge account of assertion does equate what S asserts with “I know that P” with what S represents as being the case by asserting “P”, and it equates the relevant warranted assertability condition of simple assertion with the truth-condition of the knowledge claim. However, despite these equations, when you compare apples with apples and oranges with oranges, the knowledge claim (on the bottom of each chart) is stronger than the simple assertion (directly above it), both with respect to what you assert (the middle column), and with respect to what you represent as being the case by making the assertion and the relevant warranted-assertability condition of your assertion (the right column). Thus, the account respects the fact that making the knowledge claim really is taking more of a plunge than is simply asserting that P.

Why are counter-examples to the mistaken equations tough to come by? 1) mis-equates the warranted assertability conditions for “P” with the warranted assertability conditions for “I know that P”. 2) mis-equates the truth-conditions of the knowledge claim with the knowledge claim’s warranted assertability conditions. In each case, according to the knowledge account of assertion (see the second version of our chart), this is a misfiring attempt to equate something stronger (S’s knowing that S knows that P) with something slightly weaker (S’s knowing that P). But, given that it’s difficult to come up with a case where a subject clearly knows some P, while clearly failing to know that she knows that P, the knowledge account of assertion can then explain why it’s difficult to construct clear counter-examples to equations 1 and 2.

Also, according to the Account, one represents it as being the case that one knows that one knows that P when one claims to know that P. Given a general tendency to confuse what a speaker merely represents as being the case in making an assertion with what the speaker actually asserts, this feature of the Account, together with the already noted difficulty in coming up with clear cases where a subject knows without knowing that she knows may go a long way toward explaining the attraction that the “KK Principle” — the claim that if S knows that P, then

S knows that S knows that P — holds for some.<sup>33</sup> (Here I must tread lightly, since the Principle has never held much allure for me.)

2.4. The Argument for Contextualism. 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 comprise 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<sup>34</sup> yields contextualism about knowledge.

One thing that one looks for in an argument is that its premise be more secure than is the conclusion it is being used to establish. This points to a question: If the knowledge account of assertion is right to equate the standards for the warranted assertability of “P” with those for the truth of “I know that P”, then why is the claim that the former are context-variable any more secure than the claim that the latter are? What makes the first claim a suitable premise for the second?

It is clear and uncontroversial that knowledge attributions are in some way governed by varying epistemic standards. What makes contextualism controversial are warranted assertability concerns — the worry that these varying standards are only conditions for warranted assertability, rather than the truth, of the attributions. This worry about our conclusion does not apply to our premise, which itself already concerns warranted assertability.

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 context: What one can flat-out assert in some “easy” contexts can only be put forward in a hedged manner (“I think...”, “I believe...”, “Probably...”, etc.) when more stringent standards hold sway. Even invariantists, who deny the truth-conditions of knowledge attributions are sensitive to varying standards, tend to agree that the warranted assertability conditions of knowledge attributions vary with 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 assertions generally — clear enough that I trust that illustrative examples are not needed here. And no “Warranted Assertability Maneuver” can be wielded against us here, for the strong intuitions we are here utilizing are themselves judgments concerning warranted assertability.

Given how secure is the premise that assertability is context-variable, the knowledge account of assertion, which provides the only other premise needed to establish contextualism, is lethal to invariantism. It turns against the invariantists the very context-sensitivity in warranted assertability that they themselves habitually appeal to. Unfortunately for invariantists, the knowledge account of assertion is not only fatal, but also very well supported.

### **3. “I Don’t Know”, the Demise of the Generality Objection, and Other Problems for Invariantism**

3.1. The Generality Objection Defeated. According to the Generality Objection, the fact that knowledge attributions tend to become unassertable as we move to contexts governed by higher epistemic standards is just one part of a very general phenomenon. In such contexts, it becomes more difficult to assert anything. There’s nothing special about knowledge attributions here. And since this rise in the conditions for warranted assertability are not generally reflective of any change in truth-conditions, there’s no reason to be found here for thinking the truth-conditions of knowledge attributions vary with context in the relevant way.

As I indicated at the start of this paper, the knowledge account of assertion yields an answer to the Generality Objection. It does so by providing the contextualist with an alternative, and, as it turns out, superior, explanation for the harmony the Generality Objector points out: Since, as the knowledge account of assertion tells us, the relevant warranted assertability condition for the assertion of the simple “P” is that the speaker know that P according to the epistemic standards that govern her assertion, as those standards for knowledge go up, making “I know that P” go false and unassertable, the standards for the warranted assertability of the simple “P” will rise, making it unassertable too. Thus it’s no surprise that the assertability of the knowledge claim and that of the simple claim fade away together as we move to contexts governed by higher standards.<sup>35</sup> But on our new account (based on contextualism and the knowledge account of assertion), there is something special about knowledge here. The variability in the epistemic standards governing knowledge is not just one more aspect of a general variability in the epistemic standards for assertion, and certainly does not derive from the more general variability. Rather, as the standards go up, it is because less and less counts as

knowledge that fewer and fewer things become assertable. The general variability in the standards for warranted assertability is explained by the variability in the standards for knowledge. And, according to our new account, when the epistemic standards go up past the speaker's ability to meet them, "I know that P" not only becomes unassertable, as the simple "P" does: It becomes false. This suggests a test for deciding between the two accounts on hand: If, as our new account would have it, "I know that P" becomes false, and not just unassertable, then we should not just be reluctant to claim to know, but should go so far as to deny that we know.

Consideration of the relevant conversational phenomena reveals that our new account is superior to the Generality Objector's account. First, consider what happens to a simple "P" when we move from a context in which one counts as being in a position to assert that P to a context in which one does not so count. Not well-enough positioned to assert that P, what can one say instead? Often, some form of hedged claim is a good choice in such a situation: Instead of flat-out asserting P, one can, for instance, say, "Probably P", "I'm pretty sure that P", "I think that P", etc. Or perhaps one can assert some weaker proposition less specific and less ambitious than P that one is still well enough positioned with respect to, despite the elevated standards. And in some situations silence is a good choice. But one thing one cannot do is to assert that not-P! In the "high standards" Bank Case, for example, the epistemic standards seem too elevated for me to properly declare flat-out, "The Bank is open on Saturdays." But, of course, neither am I in any position to say "The Bank is not open on Saturdays"! I am not in a position to say that even according to the more relaxed standards governing the "low standards" Bank Case, and, unsurprisingly, the rise in standards does not make it easier to make the assertion.

If the Generality Objector were correct that the context-sensitivity displayed by knowledge attributions is just more of the same, we would expect similar behavior from "I know that P." But in fact knowledge claims behave very differently. Often, when we move into more demanding contexts where "I know that P" becomes unassertable, the speaker will do something the Generality Objector should find shocking: She will go so far as to deny that she knows, admitting, "I don't know that P."

Indeed, the pairs of cases contextualists have typically used to motivate their view have displayed just such verbal behavior: "Low-standards" cases, where an attribution of knowledge seems perfectly in order (is just what a normal speaker would/could say) are paired with "high-standard" cases in which speakers, describing similarly positioned subjects, not only seem

prohibited from properly describing similarly positioned subjects as knowers, but, with complete propriety go so far as to describe those subjects (often themselves) as non-knowers.<sup>36</sup> As I've already noted, defenders of invariantism have wisely accepted that these cases reflect how we use "knows". They have merely attempted to claim that this behavior can be explained just as well without the contextualist's appeal to varying truth-conditions for knowledge attributions.

But the Generality Objection seems able to handle only our reluctance to claim knowledge, and seems ill-suited to explain why we go so far as to deny that we know. The Generality Objection would have us expect the former, since it predicts that, like all other assertions, claims of knowledge will become harder to make as the epistemic standards go up. But it is at a complete loss to explain the evident fact that "I do not know" often becomes assertable as the standards go up, when it was unassertable at lower standards.

The invariantist may claim that when "I know that P" becomes unassertable, we mistake its unassertability for falsehood, and, overreacting to the claim's unassertability, we go overboard and go so far as to assert its negation.<sup>37</sup> But, as I've already pointed out, we do not do this with the simple "P". And it's no close call, either: It seems crazy to assert "Not-P" in cases where we are not in a position to assert that P due to elevated standards, though we would meet more ordinary standards for being positioned to assert that P. The Generality Objection leaves it quite a surprise that when it comes to the knowledge claim, we do go so far as to deny it.

But our alternative explanation, based on contextualism and the knowledge account of assertion, handles all this nicely. It too explains why both "P" and "I know that P" become unassertable as the standards come to clearly exceed those the speaker can measure up to. But it is also no surprise on our alternative account, based on the knowledge account of assertion, that as the standards go up to a level at which the subject's belief that P clearly does not meet them, she will be described as not knowing that P.

3.2. Check the Negations! Where it is questionable whether a claim in a situation would be false, or merely inappropriate, a good piece of general methodological advice (like the general advice to "Follow the money" in a certain type of criminal investigation) is: Check the negations (of the relevant claims).

As we've seen, following that advice shows us where the Generality Objection stumbles: While "I know that P" and "P" behave similarly enough to make the Objection look promising,



with both becoming unassertable, and roughly together, as the epistemic standards go up, the glaring difference between “I don’t know that P” (which often becomes assertable when the standards go up) and “not-P” (which doesn’t) spells the doom of the Generality Objection (especially given the presence of an alternative account of the relevant phenomena that outperforms it here).

Our willingness to admit in “high standards” contexts that we don’t know things we do claim to know in less demanding contexts also casts other imposing shadows over the prospects of any warranted assertability maneuver working for the invariantist here. Beyond wiping out the Generality Objection, which is the only form of such a maneuver (at least that I know of) that looked at all promising, it also presents a formidable general obstacle to the success of WAMs in this case, because this aspect of the behavior of “I know” / “I don’t know” puts any such a WAM on the failing side of further criteria (beyond the generality criterion we detected above in section 1.4) that we can discern for successful WAMs.

Consider first a reasonable WAM — the one we considered earlier, in defense of the DKO account of “It’s possible that  $P_{ind}$ ,” will do. Recall that the problem cases for DKO are cases where the speaker knows that P is true. About such cases, many have some inclination to think that it would be false for the speaker to say, “It’s possible that  $P_{ind}$ ” (though, again, many will instead or in addition find themselves with some opposing inclination to think the claim would be true, though inappropriate and misleading). We have already seen how the defender of DKO can use the very general “Assert the Stronger” rule to disarm or at least mitigate the force of this intuition that is hostile to his theory. Auxiliary considerations provide further leverage for this maneuver. First, of course, is the fact that the hostile intuition (that “It’s possible that  $P_{ind}$ ” is false in the relevant situation) is for many not very strong and/or is balanced by a conflicting intuition.

And if we “Check the Negation,” we find two further good-making features motivating the WAMing of this intuition. First, where a speaker knows that P, while it does indeed seem wrong, and may even seem to some false, for her to assert “It’s possible that  $P_{ind}$ ,” it seems just as bad — in fact, considerably worse — and certainly seems false, for her to instead say “It’s impossible that  $P_{ind}$ ” or “It’s not possible that  $P_{ind}$ ”.<sup>38</sup> But it seems quite unlikely that both “It’s possible that  $P_{ind}$ ” and “It’s not possible that  $P_{ind}$ ” are false, so we have good reason to believe that something is not as it seems here; we’re going to have to explain away the misleading

appearance of falsehood of something. By contrast, in the problem cases for the crazed theory of bachelor (cases involving married males), or for Fraction's theory of knowledge (cases involving false beliefs), there's no such pressure to have to explain away any appearances: It seems false to say of married men that they are bachelors, and it seems true to say of them that they are not bachelors; it seems false to describe false beliefs as cases of knowledge, and it seems true to say that believers of falsehoods do not know what they believe. So there appears to be no problem in these cases, making it all the more problematic to claim that in fact there is a big problem here, with both the intuitions of falsehood and of truth being mistaken.

Second, and closely related, having to explain away intuitions of truth as well as of falsehood not only increases the number of intuitions that our lame maneuvers must explain away, but also places on them an explanatory burden of a different and more problematic kind.

As we have seen, the plausible WAM used in defense of DKO, like all of the clearly effective WAMs I know of,<sup>39</sup> appeals to the generation of a false implicature to explain away the appearance of falsehood in the problem cases for the theory. It is not surprising that a true assertion will be inappropriate, and may seem false, if it generates a false implicature: We should avoid conveying falsehoods, whether the falsehood conveyed is part of truth-conditional content of the assertion or is an implicature generated by the assertion, and where something false is conveyed by the making of an assertion, it's not surprising that we might mistake that for the assertion's itself being false.

By contrast, it seems much more problematic to claim that an assertion that seems true is in fact false by means of a claim that, though the assertion itself is false, it generates a true implicature and is therefore a warranted assertion. For, except where we engage in special practices of misdirection, like irony or hyperbole, don't we want to avoid falsehood both in what we implicate and (especially!) in what we actually say? So, it would seem that a false assertion will remain unwarranted, despite whatever true implicatures it might generate.

Putting all our criteria together, we can see how much our proposed WAM in defense of the DKO account of "It's possible that  $P_{ind}$ " has going for it. It starts out with the advantage of being aimed at an intuition that is a good candidate for being WAMed: an intuition that an assertion ("It's possible that  $P_{ind}$ " in cases where the speaker knows that P) would be false that is opposed by an intuition that the opposite assertion would also be false in the relevant circumstances. It then explains away the relevant apparent falsehood by appeal to the fact that a

false implicature would be generated by the assertion in the relevant situations – which, as we have seen, is less problematic than is trying to explain away the appearance of truth by appeal to the generation of a true implicature. And finally, it explains how the false implicature in question gets generated by appeal to very general rules, rather than to ad hoc rules that attach only to sentences of the type in question.<sup>40</sup> By contrast, Fraction’s defense of the “Equivalence thesis” and our imagined defense of the crazed theory regarding ‘bachelor’ falls on the bad side of the distinctions we have been using.

Unfortunately for invariantists, it looks like their attempts to handle the conversational patterns that seem to support contextualism do no better by our reasonable criteria than the desperate and lame maneuvers we have been imagining. As has emerged as we checked the relevant negations, invariantists do not begin with a good candidate for WAMing, and they have to explain away as misleading intuitions of truth as well as intuitions of falsehood. For in the “low standards” contexts, it seems appropriate and it seems true to say that certain subjects know and it would seem wrong and false to deny that they know, while in the “high standards” context, it seems appropriate and true to say that similarly situated subjects don’t know and it seems inappropriate and false to say they do know.<sup>41</sup> Thus, whichever set of appearances the invariantist seeks to discredit — whether she says we’re mistaken about the “high” or the “low” contexts — she’ll have to explain away both an appearance of falsity and (much more problematically) an appearance of truth. And, as we have seen, the warranted assertability maneuver, as it has been put forward by its advocates, fails our test of generality, and, the Generality Objection, which is our attempt to help out the invariantists on this score, while initially promising, ultimately fizzles when we “check the negations” and discover that the explanation on which the Generality Objection is based is inferior to a rival contextualist explanation.

It is very easy for defenders of invariantism to “chalk it all up to pragmatics” – to claim that the varying epistemic standards that govern our use of “knows” are only varying standards for when it would be appropriate to say we know. But it turns out to be no easy matter to make good on such a claim, and we now have little reason to think such claims credible. Not only does such a maneuver fail the tests we have discerned, but, as we saw above in Part 2, what a close look at the conditions for warranted assertability actually leads to is a powerful argument in favor of contextualism.

---

This paper is dedicated to the memory of Rogers Albritton, my teacher and dissertation advisor at UCLA, in discussion with whom my thinking on the topics of this paper began and began to take shape. I am very fortunate to be among the many philosophers who have benefited so much from long discussions with the Socrates of our age. For helpful comments, I am grateful to the editors of the Philosophical Review and to audiences at Syracuse University, Yale University, Rutgers University, the University of Vermont, Franklin & Marshall College, Notre Dame University, Brown University, and the University of Miami, where I floated some of the ideas presented here.

<sup>1</sup> For more on contextualism, including a brief history and references to much of the literature, see my “Contextualism: An Explanation and Defense,” in J. Greco and E. Sosa, ed., The Blackwell Guide to Epistemology (Malden, Mass.: Blackwell Publishers, 1999), 187-205.

<sup>2</sup> Under what conditions does a speaker succeed in raising the epistemic standards, especially when confronted by an opponent determined not to let her get away with raising them? I leave that open here, though I address that in another paper I am now working on. In the meantime, as in earlier work (see especially the second “important point” I raise at p. 6 of my “Solving the Skeptical Problem,” Philosophical Review 104: 1-52), I appeal to the claim that the skeptic executes conversation maneuvers that at least threaten to raise the standards – whether or not she succeeds in raising them – to explain her persuasiveness.

<sup>3</sup> Austin, “Other Minds,” J.L. Austin, Philosophical Papers, Third ed. (Oxford: Oxford University Press, 1979), 76-116; p. 108, fn. 1.

<sup>4</sup> For a variety of reasons, some will suspect that there is something problematic about the thought that philosophical skeptics raise the epistemic standards to such an extreme level that nothing (or almost nothing) counts as knowledge. Some, for instance, will think there is something illegitimate with the thought of installing standards according to which one can truthfully say that we “don’t know” even “hinge propositions” (as those working in the wake of

---

Wittgenstein sometimes call them). Some of these thinkers may be receptive to contextualist treatments of the more moderate cases of elevated standards, like the Bank Case I'm about to present, while balking at the thought that philosophical skeptics succeed in installing extreme standards, some of them instead thinking that the most the skeptic can do with respect to our most secure beliefs is create a context in which claims to know (or not to know) them are neither true or false. My own view on this matter is that while attempts to exploit the relevant rules to raise epistemic standards are not always successful (as I've indicated in note 2, above), and while participants to a conversation for good reasons may be especially motivated to resist attempts to install the most extreme standards, there is no bar in principal to such standards coming into force. (Reasons for thinking there is such a bar in principle may include various accounts of knowledge which tie it to such things as the satisfaction of norms of epistemic cooperation, as in Richard W. Miller, "The Norms of Reason," Philosophical Review 104 (1995): 205-245, if one thinks the tie cannot hold where the most extreme standards are in place; other accounts of the concept of knowledge which tie it to a certain function the concept plays, where extreme standards conflict with that function; and also (very) broadly verificationist reasons having to do with the inconceivability of anyone's ever meeting the extreme standards.) In this paper, I am just defending contextualism per se, and so will not address this issue which is important to the application of contextualism to philosophical skepticism. However, while I am noting (admitting) the fact that contextualist intuitions (that denials of knowledge are true in "high standards" cases) are not as strong in the philosophical cases, it is worth also noting (admitting) that this fact may support and encourage the thought that philosophical skeptics cannot succeed in creating contexts in which they can truly say of (just about) anything that we "don't know" it.

<sup>5</sup> DeRose, "Contextualism and Knowledge Attributions," Philosophy and Phenomenological Research 52: 913-929, p. 913. While I had thought about Austin's cases, and they were in the background, the immediate inspiration for my Bank Cases were Peter Unger's cases involving whether his character John knows that there is milk on the rug (Unger, Philosophical Relativity. (Minneapolis: University of Minnesota Press, 1984), pp. 50-51). Unger contrasted an ordinary ascription of knowledge to John with a case in which philosophical skepticism is being discussed, and the skeptical possibility that "There is an evil demon deceiving me into false

---

belief in the matter of whether or not there is milk on the rug” is haunting poor John. To make the cases less philosophical and more convincing, my “high standards” Bank Case involved a more moderate hypothesis, the taking seriously of which seemed reasonable given practical, non-philosophical concerns not present in the “low standards” case. These Bank Cases were featured later in colloquium talks I gave in the Winter of 1989-1990, beginning at UCLA in late 1989, and in my dissertation (“Knowledge, Epistemic Possibility, and Scepticism,” Ph.D. diss., University of California, Los Angeles, 1990), before appearing in the paper referenced at the top of this note. Among others who use such cases in discussions of contextualism, most notable is Stewart Cohen, who employs a very similar case – his Airport Case – in his “Contextualism, Skepticism, and the Structure of Reasons” (Philosophical Perspectives 13 (1999): 57-89, p. 58). The material immediately following the case contains a very nice, pro-contextualist discussion of it.

<sup>6</sup> Of course, we may begin to doubt the intuitions above when we consider them together, wondering whether the claim to know in the first case and the admission that I don’t know in the second can really both be true. But when the cases are considered individually, without worry about the other case, the intuitions are quite strong, and, in any case, the linguistic behavior displayed in the cases quite clearly does accurately reflect how “I know” / “I don’t know” is in fact used.

<sup>7</sup> See Unger, Philosophical Relativity.

<sup>8</sup> The controversy, then, is over a crucial issue about which Austin was (unwisely, I believe) quite unconcerned. In “A Plea for Excuses” (Philosophical Papers (third ed.): 175-204), attached to the (to my thinking) wonderful passage, “Certainly, then, ordinary language is not the last word: in principle it can everywhere be supplemented and improved upon and superseded. Only remember, it is the first word,” Austin drops this (to my thinking) troubling footnote: “And forget, for once and for a while, that other curious question ‘Is it true?’ May we?” (p. 185). As H.P. Grice points out, in connection with this troubling footnote, Austin “consistently avoids the words ‘true’ and ‘false,’ using instead such expressions as ‘it will not do to say’...” and seeks “to

---

put on one side the question whether such applications would be inadmissible because they were false, because they would lack a truth-value, or for some other reason” (Grice, Studies in the Way of Words (Cambridge: Harvard University Press, 1989), p. 13). Austin displays a similar attitude in his epistemological writings, avoiding issues of whether our epistemological claims (especially claims about what is and is not known) are true or false. Thus, though Austin is something of a granddaddy of contextualism, in that he put on display a lot of the linguistic phenomena contextualists appeal to, he himself should not be classified as a contextualist. All grant that in some way our use of “knows” is governed by different standards in different contexts. The question is: Are the varying standards we can discern those that govern the truth-conditions of the sentences, or merely their warranted assertability conditions? Insofar as I can discern an answer to this in Austin, it is (a peculiarly aggressive): “I don’t care.”

<sup>9</sup> I don’t know of any philosophers who have endorsed such a view. However, the idea that knowledge requires only true belief has been accepted; see Crispin Sartwell, “Knowledge is Merely True Belief,” American Philosophical Quarterly 28 (1991): 157-165; and his “Why Knowledge Is Merely True Belief,” Journal of Philosophy 89 (1992): 167-180. I think this cannot possibly be right as a general account of knowledge, but cases in which it intuitively seems that a subject is truthfully credited with knowing despite the fact that her belief has little more than being true to recommend it as knowledge have long tempted me toward the contextualist position I thought was suggested by Saul Kripke in his “Nozick-bashing” lectures (circa 1985), and has now been taken by John Hawthorne (“Implicit Belief and A Priori Knowledge,” Southern Journal of Philosophy 38 (2000), Spindel Conference Supplement: 191-210), that in certain contexts the standards for knowledge are so low that little or perhaps nothing more than true belief is required for knowledge in those contexts. Kripke presented such a case in the above-mentioned lectures: a case in which we’re asking “What did Nixon know, and when did he know it?,” and where we are unlikely to let Nixon off the hook if his defense is, “Well, at that point, I merely had true belief. I don’t know if I satisfied any of these other requirements for knowledge.” Sartwell (see esp. p. 160 of his 1991 paper) and Hawthorne (pp. 201-206) also present some tempting cases. This contextualist view is easily accommodated (though not necessitated) by the account of knowledge used in my “Solving the Skeptical Problem” (esp. pp. 36-37), where I picture knowledge of P as requiring that P be true, and that the subject’s belief as

---

to whether P is true match the fact of the matter, not only in the actual world, but in the “sphere of epistemically relevant worlds” centered on the actual world (the worlds similar enough, in the relevant respects, to the actual world), where the size of the sphere varies with context, being larger in contexts governed by higher standards. The lowest standards, then, may involve a limiting case where the sphere collapses to a point – the actual world – and in such a context, all that’s required for knowledge is true belief. But in cases governed by “ultra-low” standards of some type, does knowledge really collapse all the way down to true belief? I wonder. Even in such cases, “S knows that P” may require for its truth some further conditions, like, for instance, that S not have strong reasons for thinking P is false.

<sup>10</sup> The character Jank Fraction, of course, is based on the actual Frank Jackson, and this passage is a simple transformation of the closing paragraph of Jackson’s “On Assertion and Indicative Conditionals,” Philosophical Review 88 (1979): 565-589.

<sup>11</sup> By contrast, the actual philosopher Frank Jackson, on whom my fiction of Jank Fraction is based, utilized a maneuver in some ways similar to Fraction’s, but in defense of a theory for which there are genuine positive reasons. The “Equivalence” theory that the actual Jackson defends claims that indicative conditionals are equivalent to material conditionals (i.e., they are true iff their antecedent is false and/or their consequent is true). A reason for accepting that Equivalence theory is to be found in the apparent validity of arguments of the form “P or Q; therefore, If not-P, then Q”. Example (due to Robert Stalnaker): Either the butler did it or the gardener did it. Therefore, if the butler didn’t do it, then the gardener did” (Stalnaker, “Indicative Conditionals,” Philosophia 5 (1975): 269-286, p. 269).

<sup>12</sup> DeRose, “Contextualism: An Explanation and Defense,” pp. 196-7. The next couple of paragraphs, though slightly modified, derive from that discussion.

<sup>13</sup> The subscript “ind” indicates that the embedded P is to be kept in the indicative mood: very different possibilities are expressed where the embedded P is subjunctive.

<sup>14</sup> Though this is the right path, DKO is not correct as it stands. For my best attempt to get the details right, see DeRose, “Epistemic Possibilities,” Philosophical Review 100 (1991): 581-605.



---

<sup>15</sup> Don't mistake generality for exceptionlessness. Though the rule is very general, there are occasions on which one should not assert the stronger, as is pointed out by Jackson ("On Assertion and Indicative Conditionals," pp. 569-73).

<sup>16</sup> This preference for explanations generated by general rules over those that utilize special rules amounts to, in Grice's terminology (see Studies in the Way of Words), a preference for uses of conversational implicatures over conventional implicatures. I join Kent Bach ("The Myth of Conventional Implicature," Linguistics and Philosophy 22 (1999): 327-66) and Jason Stanley ("Modality and What is Said," Philosophical Perspectives, 16 (2002): 321-44) in being wary of conventional implicatures, but the character of my wariness is more like the wariness that Grice himself displayed. I am largely sympathetic with Stanley's statement (made appealing to the arguments in Bach) that "the category of conventional implicature is simply suspect. The arguments for the existence of conventional implicature rest entirely on rather unconvincing appeals to intuitions. Conventional implicature is also theoretically artificial, and sits uneasily with other Gricean distinctions" (p. 333). However, some intuitions that support the existence of conventional implicatures strike me as worthy of some credence, even if they are not compelling. I have in mind here the intuitions like those that support the favorite example of conventional implicature: that "P but Q", while equivalent to "P and Q" in its truth conditions, does generate a conventional implicature, not carried by "and": that there is some contrast between P and Q (see, for example, Jackson, "On Assertion and Indicative Conditionals, pp. 573-74). I am at least open to this being a genuine example of conventional implicature (but see Bach's discussion), so my skepticism is not really over the existence of conventional implicatures. Rather, my worry is that even if some conventional implicatures do exist, it is problematic to appeal to them in certain explanations. This is like the wariness of Grice, who accepted some cases as being examples of conventional implicatures, but who, after displaying a supposedly model case of conventional implicature – the case of "but" – rather warily writes, "In any case, the nature of conventional implicature needs to be examined before any free use of it, for explanatory purposes, can be indulged in" (Studies in the Way of Words, p. 46).

Note that even if "but" provides us with an example of conventional implicature, our access to this conventional implicature is through a "strange but true" intuition: We find that, where P and Q are both true, but there is no contrast between them, "P but Q" is somehow wrong

---

to assert, but would still be true. (Example: ?“The Cubs scored no runs at all, but they lost the game.”) These “strange but true” intuitions, that are needed to establish the existence of conventional implicatures in the first place, would at the same time undermine the methodological presupposition behind Fraction’s – and Jackson’s – procedure: namely, that our immediate access is only to warranted/unwarranted assertability, and our intuitive access to truth/falsehood is only through warranted/asserted assertability (see the imagined Fraction quotation at the end of section 1.3, above, and the actual Jackson passage, on which Fraction’s is modeled, at p. 589 of “On Assertion and Indicative Conditionals”). For if our only intuitive access to truth/falsehood came indirectly through warranted/unwarranted assertability, how could we get the “true” part of “strange but true” intuitions here, where the sentences are unwarranted?

<sup>17</sup> Patrick Rysiew now attempts to take up the challenge and work out, on an invariantist semantics, how one can explain the relevant behavior of knowledge attributions – that they are somehow subject to varying epistemic standards – in his “The Context-Sensitivity of Knowledge Attributions,” *Noûs* 35 (2001): 477-514. While, as Rysiew points out, his account is “underwritten” by the very general conversational rule, “be relevant,” his is not the type of account I have urged is to be desired. On Rysiew’s relevant-alternativist theory of knowledge attributions, an assertion of “S knows that P” carries two separate but related meanings – though these will sometimes have at least roughly the same value. A knowledge attribution semantically imparts that S has a true belief that P and can rule out all the relevant alternatives to P, where the standard for relevance is invariant (see p. 488), but also conveys the second, pragmatic meaning that S has a true belief that P and can rule out all of the salient alternatives to P, where the salience of alternatives is a highly context-sensitive matter. Where the contextually salient alternatives diverge significantly from the relevant alternatives, the injunction to “be relevant” gets the listener to fasten on the second, pragmatic meaning, since the semantic content of “S knows / doesn’t know that P” is clearly irrelevant to the purposes at hand. Thus, on Rysiew’s treatment (p. 490) of my “high-standards” Bank Case, I utter the falsehood that I “don’t know” that the bank is open on Saturday, but all is well, for my wife can be counted on to ignore the falsehood I strictly said, since mere “knowing” / “not knowing” is not what is relevant to our there unusually heightened concerns, and focus instead on the second, pragmatic meaning that

---

my assertion carries – that it’s not the case that I have a true belief and can rule out all of the salient alternatives to P. I am highly suspicious of accounts that help themselves to two such meanings in the way Rysiew’s does, disliking not only the loss of economy in explanation, but also worrying that we will not be able to combat all manner of absurd theories about the truth-conditions of various sentence if defenders of those theories are able to posit separate pragmatic meanings that do the work of accounting for usage in the troublesome cases, allowing their account of the truth-conditions to sit safely off in the corner, above the fray. Of course, in some sense, we all must hold that our assertions carry meanings beyond what goes into their truth-conditions. As we’ve seen, I hold that “It’s possible that P” conveys that the speaker doesn’t know that P, though that thing conveyed is not part of the truth-conditions of the utterance. But I find such a claim unobjectionable when the second, pragmatic meaning is generated by the semantic meaning together with general principles. In such a case, it seems that the semantic meaning is responsible for handling the phenomena, working through the general principles. I am open to hearing evidence that certain assertions carry two meanings in the way Rysiew posits, but I think it is rational to prefer explanations that appeal to only one “basic”, we may say, meaning – the other, “non-basic” propositions conveyed being generated by that one meaning and general principles.

<sup>18</sup> DeRose, “Contextualism: An Explanation and Defense,” pp. 196-203.

<sup>19</sup> One might question whether assertion is governed by a rule concerning how well-positioned one must be with respect to a proposition in order to be able to properly assert it. What any such rule seems ultimately aimed at securing is that, at least often enough, only the truth is asserted. Thus, in the simplest scenario — and perhaps therefore the scenario we should presume, absent any reason to think matters are more complicated — assertion would not be governed by any rule about how well positioned one must be in order to assert, but rather by the simple rule most directly tied to that ultimate aim: Assert only what’s true. Thus, Timothy Williamson, in defending his version of the knowledge account of assertion, expends much effort in battling the rival “truth rule,” according to which one must assert only what is true (Williamson, Knowledge

---

and its Limits (Oxford: Oxford University Press, 2000), pp. 244-249). See also the second paragraph of note 23, below.

The rule here – that one should assert only that with respect to which one is well enough positioned to assert – is basically Grice’s second maxim of Quality – “Do not say that for which you lack adequate evidence” (Studies in the Way of Words, p. 27) – correcting for Grice’s overly narrow focus on evidence. I think a speaker should be counted, in Gricean terminology, as having generated a conversational implicature to the effect that she is in a position to assert that P when she does assert that P, though Grice himself might have declined to call this an “implicature” at all, since it simply specifies that the speaker is obeying the second maxim of Quality, and Grice stipulates that “implicature” excludes commitments that are trivially generated by the supposition that the speaker is observing the Cooperative Principle and its associated conversational maxims (pp. 41-42). For relevant discussion of the issue of whether such commitments should be counted as implicatures, see “Conditional Assertions and ‘Biscuit’ Conditionals,” by Richard E. Grandy and myself, Noûs 33 (1999): 405-420, n. 13, pp. 417-418 and n. 19, p. 419.

<sup>20</sup> See Peter Unger, Ignorance: A Case for Scepticism (Oxford: Oxford University Press, 1975), Chapter VI; Michael Slote, “Assertion and Belief,” in J. Dancy, ed., Papers on Language and Logic (Keele: Keele University Library, 1979); and DeRose, “Epistemic Possibilities,” sections G-H. Williamson (Knowledge and Its Limits, p. 252) reports that Lloyd Humberstone suggests that this talk of “representing oneself” in certain ways by making assertions may have begun with Max Black: “In order to use the English language correctly, one has to learn that to pronounce the sentence ‘Oysters are edible’ in a certain tone of voice is to represent oneself as knowing, or believing, or at least as not disbelieving what is being said” (Black, “Saying and Disbelieving,” Analysis 13 (1952): 25-33, p. 31).

<sup>21</sup> See Williamson, Knowledge and Its Limits, Chapter 11, which updates Williamson, “Knowing and Asserting,” Philosophical Review 105 (1996): 489-523.

---

<sup>22</sup> Knowledge and Its Limits, p. 241. That there is such a unique rule is little more than an item of faith for Williamson, with no justification offered other than that a simple account consisting of such a single rule would be “theoretically satisfying” (p. 242).

<sup>23</sup> An example of someone reasonably but falsely thinking they know something is provided by Ginet’s famous “barn case,” which can be (and in fact most naturally is) imagined so that the character is reasonable in thinking that he knows: Henry quite reasonably thinks that he knows that he’s seeing a barn, because he reasonably thinks he is in a very normal situation; unbeknownst to him, he’s in a region teeming with fake barns and in fact doesn’t know, though in fact he happens to be seeing a real barn right now. (This case appeared prominently in Alvin Goldman’s “Discrimination and Perceptual Knowledge,” Journal of Philosophy 73 (1976): 771-91, pp. 772-73. Gail Stine reports (“Skepticism, Relevant Alternatives, and Deductive Closure,” Philosophical Studies 29: 249-61, p. 252) that Goldman attributes the example to Carl Ginet.) 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 well enough positioned 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. I believe the knowledge account fits the data well, and, in fact, provides the best account of the full range of data, though I am not here pursuing a full-fledged defense of the account. What would show that the knowledge account of assertion is too strong are cases in which it’s apparent that a speaker in some way properly asserts what she doesn’t even reasonably take herself to know.

Similarly, the truth account of assertion (according to which assertion is governed by the “truth rule,” mentioned above in note 19) is not refuted by the existence of just any example in which a speaker asserts a falsehood, but we can sense that she is in some sense warranted in making her assertion, for it may be that the speaker reasonably believes that she is speaking the truth. What would, and does, spell trouble for the truth account (by showing it to be too weak)

---

are cases in which a speaker is not warranted in asserting some P that both is true and is such that she very reasonably takes it to be true. In many such cases, the speaker is in a position to say something like “P is very likely,” “It’s almost certain that P,” “I believe that P” or even, “I reasonably believe that P,” but, as we can all sense, is in no position to flat-out assert the simple “P.” (See also Williamson’s discussion of the truth rule, cited above in note 19.)

<sup>24</sup> While I still like and recommend Unger’s argument in Chapter VI of Ignorance, Chapter 11 of Williamson’s Knowledge and Its Limits is now the state of the art. Williamson discusses and expands on some (but not all) of the grounds given by Unger and Slote (pp. 252-255), adds further positive arguments of his own, answers objections to the account (pp. 255-260), and argues against rival accounts. An important rival account is the probability account, on which one should assert only what is (in the relevant sense) sufficiently probable. V.H. Dudman uses lottery cases to forcefully attack the probability account in his “Probability and Assertion,” Analysis 52 (1992): 204-211, though he goes too far, I think, in his further claim that “Assertability goes out of the window as soon as the underlying thought is reduced to relying on ‘mere’ probability” (p. 205); for discussion, see section X (pp. 576-78) of my “Knowledge Assertion, and Lotteries,” Australasian Journal of Philosophy 74 (1996): 568-80.

<sup>25</sup> Moore, Commonplace Book: 1919-1953 (London: Allen & Unwin, 1962), p. 277.

<sup>26</sup> In other work, I’ve used the claim that in asserting something one represents oneself as knowing it, together with independently supported analyses of the relevant modal statements, to also explain these other “clashes” (all given in general schematic form, except for 4, which is an example of the type of conjunction, involving a “simple ‘might’”, that I mean to indicate):

2. It’s possible that not-P, but P.
3. It is possible that X V’s, but it is not possible for X to V.
4. My car is in the Main Street parking lot, but it might not be there.
5. If P had been the case Q would have been the case, but if P had been the case Q might not have been the case.

We have strong reason to believe that each of 1-5 is consistent. The consistency of 1 is addressed in the text, above, and is obvious in any case. The consistency of 2 and 4 are just

---

about as obvious: Surely the mere possibility of not-P is consistent with P. If these were inconsistent, you could infer that not-P from the mere possibility of not-P, which seems absurd. For reasons for thinking 3 is consistent, see “Epistemic Possibilities,” pp. 602-603. Following David Lewis (see Counterfactuals (Cambridge: Harvard University Press, 1973), pp. 2, 21-24), many subscribe to the “duality thesis” of the relation between “might” and “would” counterfactual conditionals, and would therefore hold that instances of 5 are genuinely inconsistent. However, the duality account is wrong, and instances of 5 are consistent, as I show in my “Lewis on ‘Might’ and ‘Would’ Counterfactual Conditionals,” Canadian Journal of Philosophy 24 (1994): 413-418 and my “Can It Be That It Would Have Been Even Though It Might Not Have Been?” Philosophical Perspectives 13 (1999): 385-413. As I have argued, the hypothesis that one represents oneself as knowing something when one asserts it provides the best treatment of these “clashing conjunctions” — a treatment that happily does not make any of 1-5 out to be a genuine inconsistency: I address 1-3 at “Epistemic Possibilities,” pp. 596-605; the treatment of 2 is expanded to cover statements like 4 at in my “Simple *Might*’s, Indicative Possibilities, and the Open Future,” The Philosophical Quarterly 48 (1998): 67-82, pp. 72-73; 5 is addressed in my “Can It Be...?”, pp. 388-390).

<sup>27</sup> Lewis, Counterfactuals, p. 92. Lewis is there writing about the link between similarity and counterfactual conditionals. Though Lewis was a contextualist (see especially his “Elusive Knowledge,” Australasian Journal of Philosophy 74 (1996): 549-567), he would not have tied assertability to knowledge, as I advocate, because he tended to gravitate toward a probability, rather than a knowledge, account of assertion. However, the probability account Lewis used was well-suited to accommodate the context-variability of assertability. Lewis held that one could properly assert something if its probability was “sufficiently close to 1” (Lewis, “Probabilities of Conditional and Conditional Probabilities,” Philosophical Review 85 (1976): 297-315, p. 297). It would be easy, and indeed quite natural, to then claim that the matter of just how close is “sufficiently close” varies with context, which would allow assertability and sufficient probability to “sway together.”

---

<sup>28</sup> Williamson mentions the possibility of contextualizing the knowledge account of assertion (Knowledge and Its Limits, pp. 254-55), but is non-committal about it, writing only that his account “permits such contextual variation.”

<sup>29</sup> Hambourger, “Justified Assertion and the Relativity of Knowledge,” Philosophical Studies 51 (1987): 241-269.

<sup>30</sup> This is the thesis of Part II (pp. 248-255) of Hambourger’s paper. Focusing too much on the notion of evidence, as he does throughout his paper, Hambourger puts that thesis as follows: “The amount of evidence needed to be epistemically justified in claiming to know a proposition (relative to a given standard of caution) simply is the amount needed to be epistemically justified in asserting the proposition itself straight-out (relative to that standard)” (p. 249). I have converted Hambourger’s talk about how much evidence one has into the more general terms of how well-positioned one is.

<sup>31</sup> This is one of the main conclusions of Part III (pp. 255-262) of Hambourger’s paper. He writes, “I conclude, therefore, that the amount of evidence needed to know a proposition is not fixed once and for all, but that it varies with standards of caution, and, in particular, that the amount needed to know a proposition under any given standard is just that needed to be epistemically justified in claiming to know it. . .under that standard” (p. 260). Again (see the immediately preceding note), Hambourger expresses these as standards for how much evidence one must have, and I have again converted that to the more general matter of standards for how well positioned one must be.

<sup>32</sup> This again is Carl Ginet’s famous “fake barns” case; see note 23, above.

<sup>33</sup> What might also help in explaining the attraction of the KK principle is locating a principle in the vicinity of KK, but which is weaker and more defensible than KK itself. The principle that knowing is a “weakly luminous” condition, in the sense I describe in my review of Williamson’s Knowledge and Its Limits (British Journal for the Philosophy of Science 53 (2002): 573-577) may serve this purpose. Briefly, to say that knowing is weakly luminous, in this sense, is to say



---

that whenever one safely knows (whenever one knows and is not close to not knowing), one is in a position to know that one knows.

<sup>34</sup> Again, I am here speaking of a particular kind of warranted assertability: that having to do with whether one is in a position to make the assertion one is making. I take it as obvious that there is context-variability in this aspect of the conditions for warranted assertability.

<sup>35</sup> And, as we have seen, to the extent that the assertability of the knowledge claim and the simple claim don't fade away exactly together, our account, by making the warranted assertability condition for the knowledge claim a bit more stringent than that of the simple claim, does a good job of handling this.

<sup>36</sup> This is true, for instance, of both Cohen's Airport Case (referred to in note 5, above) and of my Bank Cases (described in section 1.1, above).

<sup>37</sup> There is a type of case where speakers will say "Not-P", where it seems that they are merely denying the assertability of "P", but such cases are special in that they involve unusual emphasis and a certain follow-up: they are followed by what should instead be said. A special case of this that is quite common is where speakers (at least appear to) deny, with emphasis, something weaker and then go on to say something stronger: "He's not big; he's huge"; "She's not smart; she's brilliant". Here, it seems, what's being conveyed is that, for instance, he isn't only big, and since something stronger – "He's huge" – should instead be said, the weaker thing – "He's big" – should not be asserted. I take it then that these should not be understood as really being denials of the weaker claims ("he's big", "she's smart") or as assertions of the negations of the weaker claims, but rather as denials that the weaker claim should be made. And in general, where a flat-footed (unemphasized) and unfollowed denial seems wrong, but the same "denial" can appropriately be said with the right emphasis and follow-up, we should conclude that it is really only the assertability, rather than the truth, of what appears to be denied, that is really being denied in the special cases (of emphasis and follow-up). The classic treatment of this general phenomenon, known as "metalinguistic negation," is in Laurence R. Horn, A Natural History of Negation (University of Chicago Press, 1989).

---

<sup>38</sup> But recalling the previous note, it seems that “It’s not possible that P” can be said where the speaker knows that P, under the special circumstances that the “denial” of the possibility is said with the right emphasis, and is immediately followed by a stronger statement: e.g., “It’s not possible that the book is there; it’s certain that it is” or “It’s not possible that the book is there; it is there.” Since these fit the general pattern mentioned in the preceding note, I conclude that they are not after all real denials of the possibility (note too that “It’s impossible” still seems wrong), but are just ways of saying that the mere possibility is not what should be asserted.

<sup>39</sup> I am limiting my attention here to WAMs that are aimed at what might otherwise really be tempting treatments of sentence types. I take it the defense of DKO we are considering falls into this camp: One might really be tempted by the DKEW account of “It’s possible that P<sub>ind.</sub>” Likewise, to give just one (well-known) example of a clearly successful WAM, initially one might really be tempted to think that it is not true to say of a bright red (British) mailbox that is only two feet away from one in bright sunlight that it “looks red to me,” and one might thereby seek an account of “looks red” that respects this appearance (see Grice, Studies in the Way of Words, esp. pp. 227-230). However, I take it that Grice has successfully WAMed the intuition that “looks red” cannot be truthfully applied in such a case. Note that this WAM gets on the good side of all my tests: It seeks to explain away an appearance of non-truth; it does so by means of a general conversational rule; and the appearance of non-truth that it seeks to explain away is balanced by a similar appearance of non-truth that attaches to the opposite statement: In the circumstances under discussion, it also seems wrong to say that the fire engine “does not look red” – except, to recall the previous two notes, in the special case where one says this with the proper emphasis and a proper follow-up: “It doesn’t look red; it is red!”

<sup>40</sup> Rysiew argues that even the “doesn’t know otherwise” condition (the condition that the speaker not know that not-P) is merely something conveyed by an assertion of “It’s possible that P<sub>ind.</sub>”, and is not part of the truth-conditional content of the modal statement. After noting that I accept that “I don’t know that P” is merely conveyed by, and is not part of the truth-conditional content of, “It’s possible that P<sub>ind.</sub>,” he writes: “Curiously, however, [DeRose] seems not to see that an exactly parallel argument establishes that the speaker’s not knowing that not-p isn’t part

---

of the truth-conditions of such statements either” (Rysiew, “The Context-Sensitivity of Knowledge Attributions,” p. 510, n. 30). However, so far from being “exactly parallel,” the arguments are very far apart indeed. My case for WAMing the condition I reject, both here and in the paper Rysiew addresses (“Contextualism: An Explanation and Defense”), is based on three considerations I put forward that I argue make this WAM credible. At least two of the those three considerations clearly fail to apply to the “parallel” case Rysiew proposes: (1) Where the condition I reject as a truth-condition (that the speaker not know that P) fails to hold (where the speaker does know that P), while it can seem false for her to say “It’s possible that P<sub>ind</sub>”, it also seems false for her to make the opposite assertion (that P is “not possible” or is “impossible”). This makes this condition suspect, and a good candidate for WAMing. The same cannot be said for the condition I accept – that the speaker not know that not-P. Where it fails (where the speaker does know that not-P), it both seems false for the speaker to assert “It’s possible that P<sub>ind</sub>” and it seems true for her to deny that P is possible. (2) Thus, in WAMing the “suspect” condition, I only have to explain away intuitions of falsehood, while Rysiew will have to do not only that, but also explain away intuitions of truth, which, for the reason I’ve given, seems more problematic. I have a hard time evaluating whether any maneuver of the type Rysiew proposes will be able to successfully generate the result he needs based on his account of truth-conditions of “It’s possible that P<sub>ind</sub>” together with general conversational rules, and thus get on the “good side” of my third consideration, for I can’t see what he thinks the truth-conditions of “It’s possible that P<sub>ind</sub>” are. But no attempt to explain away the condition I endorse can be even close to “exactly parallel” to the account by which I battle the condition I reject, for, as we have seen, my account relies on the claim that “P” and “I know that P” are stronger than (imply but are not implied by) “It’s possible that P.” A parallel account targeting the condition I endorse would then have to rely on the very implausible claim that “Not-P” or “I know that not-P” imply “It’s possible that P<sub>ind</sub>.” Also, especially in “Epistemic Possibilities,” I use the account of “It’s possible that P<sub>ind</sub>” that I accept – which includes as a truth-condition what Rysiew rejects as such – to explain many aspects of the behavior of the sentence-type, and I cannot see how all this can be explained if this account is rejected.

---

<sup>41</sup> Here it is crucial to keep in mind that in this paper, when we speak of “high contexts,” we are not thinking of contexts in which the hyperbolic doubts of philosophical skepticism are in play. There the relevant intuitions are not that strong. We are rather thinking of cases like my “high standards” Bank Case – cases where doubts a bit more remote than are usually taken seriously, but nearly so remote as the playthings of philosophers, are in play, and reasonably so, given practical concerns. About such cases, I submit, the relevant intuitions really are quite strong.