

Chapter 8

Contextualism:
An Explanation and Defense

Keith DeRose

from Greco & Sosa, eds.
The Blackwell Guide to
Epistemology (1999)

In epistemology, “contextualism” denotes a wide variety of more-or-less closely related positions according to which the issues of knowledge or justification are somehow relative to context. I will proceed by first explicating the position I call contextualism, and distinguishing that position from some closely related positions in epistemology, some of which sometimes also go by the name of “contextualism.” I’ll then present and answer what seems to many the most pressing of the objections to contextualism as I construe it, and also indicate some of the main positive motivations for accepting the view. Among the epistemologists I’ve spoken with who have an opinion on the matter, I think it’s fair to say a majority reject contextualism. However, the resistance has to this point been largely underground, with little by way of sustained arguments against contextualism appearing in the journals,¹ though I have begun to see various papers in manuscript form which are critical of contextualism. Here, I’ll respond to the criticism of contextualism that, in my travels, I have found to be the most pervasive in producing suspicion about the view.

1. What Is Contextualism?

As I use it, and as I think the term is most usefully employed, “contextualism” refers to the position 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) vary 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) in order for such a statement to be true. In some contexts, “S knows that P” requires for its truth 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 very 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. Thus, the contextualist will allow that one speaker can truthfully say "S knows that P," while another speaker, in a different context where higher standards are in place, can truthfully say "S doesn't know that P," though both speakers are talking about the same S and the same P at the same time. The "invariantist" – Peter Unger's good name for one who denies contextualism² – will have none of this. According to her, there's a single, invariant set of standards which, at least as far as truth-conditions go, govern the use of knowledge attributions regardless of the context in which they're uttered; thus, our two speakers can't both be speaking a truth.

I am not alone in endorsing contextualism so construed; other contextualists prominently include Stewart Cohen, David Lewis, Gail Stine, and Peter Unger.³ Now it should be, and I think it largely is, fairly uncontroversial that in different conversational contexts, quite different standards govern whether ordinary speakers *will say* that someone knows something: What we're happy to call knowledge in some ("low-standards") contexts we'll deny is knowledge in other ("high-standards") contexts. The invariantist needn't deny this, and if she's wise she won't deny it. Nor need she or should she deny that this is a very useful feature of our use of the relevant sentences. What she must deny is that these varying standards for when ordinary speakers *will* attribute knowledge, and/or for when they're in some sense *warranted* in attributing knowledge, reflect varying standards for when it is or would be *true* for them to attribute knowledge, for, again, according to the invariantist, the *truth-conditions* of the relevant sentences do not vary in the relevant way.

Contextualism, so understood, then, is a position about knowledge attributions (sentences attributing knowledge to a subject) and denials of knowledge – precisely, a thesis about their truth-conditions. This has been known to give rise to the following type of outburst: "Your contextualism isn't a theory about *knowledge* at all; it's just a theory about *knowledge attributions*. As such, it's not a piece of epistemology at all, but of the philosophy of language."

Of the many things that can be said in response to this type of charge, let me limit myself here to just this. To the extent that contextualism/invariantism is an issue in the philosophy of language, it's a piece of philosophy of language that is of profound importance to epistemology. How we should proceed in studying knowledge will be greatly affected by how we come down on the contextualism/invariantism issue. For contextualism opens up possibilities for dealing with issues and puzzles in epistemology which, of course, must be rejected if invariantism is instead the correct position. And how could it be otherwise? Those who work on the problem of free will and determinism should *of course* be very interested in the issue of what it means to call an action "free." If that could mean different things in different contexts, then all sorts of problems could arise from a failure to recognize this shift in meaning. If there is no such shift, then that too will be vital information. In either case, one will want to know what such claims mean. Likewise, it's important in studying knowledge to discern what it means to say someone knows

something. If that can mean different things in different contexts, all sorts of problems in epistemology, and not just in philosophy of language, will arise from a failure to recognize such shifts in meaning. If, on the other hand, there is no such shift, then we're bound to fall into all sorts of error about knowledge, as well as about "knows," if we think such shifts occur.

2. Contextualism regarding Other Epistemic Terms

Contextualism, as described above, is a thesis about knowledge attributing and denying sentences. But, since there are other terms with analytic ties to the concept of knowledge, we should expect that if contextualism about knowledge is true, there should be corresponding shifts in the content of sentences containing those other terms. To use David Lewis's words (though he wasn't writing about "knows" when he used them), we should expect the content of knowledge attributing sentences and the sentences containing the other terms to "sway together."⁴ For instance, to a first approximation, at least, "It's possible that P_{ind}"⁵ is true if and only if the speaker of the sentence doesn't know that P is false.⁶ Given contextualism, then, we should expect that, as the standards for knowledge go up, making it harder for belief to count as knowledge, it should become easier for statements of possibility to be true.⁷ And, since "It's certain that P" is the dual of "It's possible that P_{ind} is false", we should expect that as the standards for knowledge go up, making it harder for knowledge attributions to be true, it should also become harder for such expressions of impersonal certainty to be true.⁸

We can construe contextualism regarding *justification* as an analogue of what we're calling contextualism about knowledge: According to the contextualist about justification, the standards for justified belief that a subject must meet in order to render true a sentence describing a belief of hers as justified vary with context. The relation between knowledge and justification is controversial, and neither of these forms of contextualism clearly implies the other. If one holds that a belief's being justified is a necessary condition for its being a piece of knowledge, then one may believe that the two contextualisms are closely related: Perhaps it's *because* the standards for justification vary with context that the standards for knowledge so vary. However, it's widely accepted today that more is needed for knowledge than simply justified true belief, and it may be varying requirements for that something more that's reflected in the varying standards for knowledge – in addition to, or instead of, varying standards for justification.

In what follows, when I write simply of "contextualism," I will mean contextualism regarding knowledge; when I mean to refer to contextualism about justification or some other epistemic term, I will explicitly state so. To the best of my knowledge, neither I nor any of the contextualists mentioned above in section 1 have either endorsed or rejected contextualism about justification.⁹

3. Contextualism is Not a Thesis about the Structure of Knowledge or Justification

In his influential paper, "A Contextualist Theory of Epistemic Justification,"¹⁰ David Annis presents what he calls "contextualism" as an alternative to both foundationalism and coherentism in the issue of the structure of justification. Now, I think that even Annis's "contextualism" is not a structural alternative to those two theories, but is rather best construed as a form of foundationalism. But the vital point to be made now is that "contextualism" as I am construing it here, is certainly not a thesis about the structure of knowledge or of justification. It is, in fact, consistent with either foundationalism or coherentism.

If you're a foundationalist, then if you're also a contextualist, you may well come to think of the issue of which beliefs are properly basic (i.e., the issue of which beliefs are justified to a degree sufficient for knowledge independent of any support they receive from other beliefs), and/or the issue of how strongly supported a belief in the superstructure must be in order to count as knowledge or as a justified belief, to be matters that vary according to features of conversational context. And if you're a coherentist, then if you're also a contextualist, you'll probably want to hold that how strongly beliefs must cohere with one another in order to count as knowledge (if they're true), or to count as justified, is a contextually variable matter. So contextualism will certainly color the theories of either structural camp. But contextualism is not itself a structural alternative. Nor does it in any obvious way favor one structural alternative over the other.

4. "Subject" vs. "Attributor" Contextualism

Some distinguish between "subject contextualism" and "attributor contextualism," or use slightly different titles to mark the same distinction.¹¹ The basic issue here is whether the varying standards a subject must live up to to count as knowing are relative to the context of that subject, or rather to the context of the attributor – the person describing the subject as a knower or a nonknower. I should be clear that what I am calling "contextualism" is a form of what these folks call "attributor contextualism." But it is worth briefly describing what these views call "subject contextualism" to distinguish it from what I am here calling "contextualism."

Some "subject contextualists" point to examples in which features of the surroundings of the putative subject of knowledge which don't constitute any part of his evidence for the belief in question, and which the subject may even be completely oblivious to, impact on whether the subject knows. For instance, in Carl Ginet's much-discussed fake barn example,¹² if a subject is driving through a region teeming with fake barns, deceptive enough that they would have fooled him if he had come across them, but is luckily encountering the only real barn anywhere about, and he confidently believes that he is seeing a barn, most respond that the subject doesn't know that what he's seeing is a barn. Presumably,

though, he would have known it was a barn in a normal situation in which there are no fakes about. Here the presence of fakes in the region seems to rob the subject of knowledge, even if the subject doesn't know about the fakes and hasn't even encountered one. Now, I endorse the intuitions appealed to here: You don't know with the many fakes about, but do know in the normal situation. So, in this way, questions of knowledge are relative to the subject's "context," where "context" is being used to describe such extraevidential features of the subject's situation. So if accepting this is tantamount to being a subject-contextualist, then sign me up. But I think it's better to reserve the label "contextualism" for the more controversial thesis described in section 1, above.

Other "subject contextualists" point to features of the subject's *conversational* context: How important is it that the parties to the discussion in which the subject may be engaged be right on the matter in question (how high are the stakes for them)? What has transpired in the conversation in which the subject may be engaged? In what "community" is the subject operating, and what epistemic standards are appropriate to that community's intents and purposes? Now, I think that the conversational features that these subject contextualists point to are the kinds of features which affect whether attributions of knowledge are true, but I think it's these features of the *attributor's* context that are important. Suppose that you and I are in a discussion in which we're employing very high epistemic standards. Perhaps we're discussing a matter of great importance which calls for great caution, and we've made it clear by what we've said that we're employing such very high standards. So, though we have enough evidence for the belief in question that we'd have claimed to know it if we were in a more ordinary context, we've each denied that we know the item in question. Now, suppose I raise the question of whether Mary, a friend of ours who is not present at our conversation, knows the item in question. And suppose that you know that Mary has precisely the same evidence for the belief that we have, but that the issue isn't important to her or to those with whom she's presently speaking, and that she is in a context in which quite lax epistemic standards are being employed. Should you describe Mary as a knower? It seems that if you're not willing to call *us* knowers, then you shouldn't call Mary, who's in possession of the same evidence we have, a knower either. Here, the standards are set by the features of the attributor's setting. Of course, for certain purposes, we may wish to evaluate a subject's belief relative to standards set by features of her context. But there's nothing in attributor contextualism to rule this out: among the many standards a speaker's context may select are those relevant to the subject's context.

5. A Brief History of Contextualism

Theories according to which there are two senses of "know" – a "low," "weak," or "ordinary" sense on the one hand, and a "high," "strong," or "philosophical" sense, which is much more demanding, on the other – can be viewed as limiting

cases of contextualist views. Such a view was prominently defended by Norman Malcolm in his 1952 "Knowledge and Belief."¹³ For reasons I'll touch on in section 7, below, current contextualist theories don't hold that there are just two different sets of epistemic standards governing the truth-conditions of knowledge attributions, but rather posit a wide variety of different standards.

In important work on knowledge and skepticism in the early and mid-1970s, which culminated in his 1975 book *Ignorance*,¹⁴ Peter Unger argued that, in order to really know something, one must be in a *very* strong epistemic position with respect to that proposition – so strong, in fact, that it would be impossible for anyone ever to be better positioned with respect to any matter than you are now with respect to the matter in question. Though the terminology wasn't in place yet, largely because the contextualist alternative to it wasn't in place yet, what Unger was there defending was *skeptical invariantism*. It was a form of invariantism because, so far as their truth-conditions go, Unger claimed that a single set of epistemic standards governed attributions of knowledge, in whatever context they were uttered. And it was *skeptical* invariantism because those standards were held to be very demanding. (Nonskeptical invariantism, then, is invariantism that keeps the standards governing the truth-conditions of knowledge attributions constant, but meetably low.) And Unger drew the skeptical conclusions that were naturally implied by such a stance. Importantly, Unger *did* admit that varying standards for knowledge govern our use of sentences of the form, "S knows that P," but did not endorse contextualism, because Unger claimed that these varying standards were only standards for whether it was *appropriate* to say that S knows; the truth-conditions for the sentence, as I've already noted, were, according to Unger, constant, and *very* demanding. Thus, the skeptic is right when she says we don't know, and we are saying something false (though perhaps appropriately) when, even in ordinary, nonphilosophical discussions, we claim to know this or that. This position, invariantism about truth conditions (whether this invariantism is skeptical like Unger's or nonskeptical), combined with variable standards for warranted assertability, is the great rival to contextualism. The "rival" came first, however: It was largely in response to this "invariantist" theory of Unger's that the early contextualist views of the late 1970s and early 1980s – like that expressed by David Lewis in a short section of his 1979 "Scorekeeping in a Language Game" (see note 3, above) and in contextualist versions of the Relevant Alternatives approach – were developed.

Later, Barry Stroud, in chapter 2 of his prominent 1984 *The Significance of Philosophical Scepticism*,¹⁵ while not advocating skeptical invariantism, did seek to defend the view by appealing to the idea that the varying standards which can govern the truth conditions of knowledge attributions might instead just govern their conditions of warranted assertability.

According to the very prominent "Relevant Alternatives" (RA) account of knowledge, which came to the fore during the middle and late 1970s, the main ingredient which must be added to true belief in order to yield knowledge is that the believer be able to rule out all the relevant alternatives to what she believes.

The range of alternatives to, or contraries of, what one believes that are relevant is held to be sensitive to a variety of factors. Many relevant alternative theorists held that the matter of which alternatives are relevant can be sensitive to the conversational context of the attributor of knowledge. This yields a contextualist version of the Relevant Alternatives theory. Among the most prominent examples of RA in the mid-1970s, Stine clearly held a contextualist version of RA, and in his important "Discrimination and Perceptual Knowledge," Alvin Goldman, while he sketched both a contextualist and an invariantist version of RA, expressed his preference for the contextualist version. Later, Stewart Cohen developed a version of RA that was explicitly a contextualist one, though Cohen used a different term: He called his an "indexical" version of RA.¹⁶ But while the most frequently held and probably most defensible manifestation of the RA approach is the contextualist version of the theory, RA can be held in an invariantist form. If you hold that the range of relevant alternatives is not sensitive to the conversational context of the attributor of knowledge, but only to factors about the putative subject of knowledge and her surroundings, the result is an invariantist form of RA. (Some will call such versions of RA instances of "subject contextualism," but, as noted in the previous section, what I am calling "contextualism" is only what makers of such distinctions call "attributor contextualism.")¹⁷

Unger's 1984 book *Philosophical Relativity* contained what was at that time – and for some time to come, for that matter – easily the most complete exposition of the contextualist view. But while this book represented a change of mind for Unger from his skeptical writings of his *Ignorance* period, he was not advocating contextualism in *Philosophical Relativity*. Instead, he defended the "relativist" conclusion that contextualism and his earlier invariantist views which led to skepticism were equally good theories, and that there simply is no fact of the matter as to which view is correct. Unger's relativism, defended, as it is, by parity considerations, according to which the advantages and disadvantages of contextualism and invariantism balance each other out in such a way that there is no winner, is a precarious view to defend: Any contextualist who succeeds in defeating invariantism will conquer Unger's relativism as an automatic corollary, and the same will happen for any invariantist who produces a successful argument against contextualism. But here Unger laid out very carefully the invariantist rival to contextualism, together with an argument that it was, while not superior to contextualism, at least an equal of it. With his 1986 "The Conc Model of Knowledge" (see note 3), Unger finally joined the ranks of the contextualists, but did not counter his earlier arguments that invariantism is superior to or at least the equal of contextualism. Struggling against the invariantist rival that Unger set up remains a main task of contextualism.

6. Contextualism and Skepticism

Contextualist theories of knowledge attributions have almost invariably been developed with an eye toward providing some kind of answer to philosophical

skepticism. For many of the most powerful skeptical arguments threaten to show that we just plain don't know what we ordinarily think we know. They thus threaten to establish the startling result that we never, or almost never, truthfully ascribe knowledge to ourselves or to other mere mortals.

But, according to contextualists, the skeptic, in presenting her argument, manipulates the semantic standards for knowledge, thereby creating a context in which she can *truthfully* say that we know nothing or very little. Once the standards have been so raised, we *correctly* sense that we could only falsely claim to know such things as that we have hands. Why then are we puzzled? Why don't we simply accept the skeptic's conclusion and henceforth refrain from ascribing such knowledge to ourselves or others? Because, the contextualist continues, we also realize this: As soon as we find ourselves in more ordinary conversational contexts, it will not only be true for us to claim to know the very things that the skeptic now denies we know, but it will also be *wrong* for us to *deny* that we know these things. But then, isn't the skeptic's present denial equally false? And wouldn't it be equally true for us now, in the skeptic's presence, to claim to know?

What we fail to realize, according to the contextualist solution, is that the skeptic's present denials that we know various things are perfectly *compatible* with our ordinary claims to know those very propositions. Once we realize this, we can see how *both* the skeptic's denials of knowledge *and* our ordinary attributions of knowledge can be correct.

Thus, it is hoped, our ordinary claims to know can be safeguarded from the apparently powerful attack of the skeptic, while, at the same time, the persuasiveness of the skeptical argument is explained. For the fact that the skeptic can install very high standards which we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in more ordinary conversations and debates.

The success of such an approach to skepticism hinges largely on the contextualist's ability to explain how the skeptic raises the standards for knowledge in the presentation of her argument. If such an explanation can be successfully constructed, such a solution to skepticism can be very attractive, and can provide a powerful motivation for accepting contextualism.¹⁸

7. The Contextualist Approach to Skepticism and to What Goes On in Conversation

But while philosophical skepticism has drawn much of the attention of contextualists, support for contextualism should also – and perhaps primarily – be looked for in how “knows” is utilized in nonphilosophical conversation. For as I've already noted, we do seem to apply “knows” differently in different contexts, a phenomenon that, at least on the surface, seems to promise significant support for contextualism. And the contextualist's appeal to varying standards for knowledge in his solution for skepticism would rightly seem unmotivated and *ad hoc* if we

didn't have independent reason to think such shifts in the content of knowledge attributions occur.

Additionally, the contextualist solution to the skeptical puzzles might fail to do justice to the impact the skeptical arguments can have on us if it were only in the presence of such arguments that the standards for knowledge were shifted upward. Why? Well, it is a fairly natural reaction to the skeptical arguments to suppose that they induce us to raise the standards for knowledge. In most classes of any size in which I've first presented skeptical arguments to introductory students, some of them will pursue such an analysis of the argument's force. Usually, they propose a version of the “Two Senses of ‘knows’” theory I mentioned in section 5. Students sometimes label the two senses they posit “weak” and “strong,” or sometimes “low” and “high,” and once “regular” and “high octane.” But though many will suspect that the skeptic is somehow “changing the subject” on us, she certainly isn't doing so in any very obvious way – as is shown by the fact that some students in most introductory classes will reject the suggestion that any such thing is going on. Yet, if there were just two senses of “knows” – one normal and quite common, and the other very strong and brought on only in contexts in which philosophical skepticism is being discussed – it would probably be quite clear to us that the skeptic was doing something fairly new and different when she started using “knows” in the “high octane” sense, and it would probably be pretty obvious she was “changing the subject” on us, and the arguments wouldn't seem to be a threat to our ordinary knowledge. On the other hand, if, as current contextualists hold, the standards for knowledge vary with context even in ordinary, nonphilosophical conversations, and the skeptic is utilizing mechanisms for the raising of epistemic standards that we're familiar with from such ordinary conversations, then it would seem much more likely that the skeptic's argument would strike us as threatening our knowledge ordinarily so-called, since the skeptic's use of “knows” would much more likely pass for what ordinarily goes on with the use of the term. As Stine writes: “It is an essential characteristic of our concept of knowledge that tighter criteria are appropriate in different contexts. It is one thing in a street encounter, another in a classroom, another in a court of law – and who is to say it cannot be another in a philosophical discussion? . . . We can point out that some philosophers are very perverse in their standards (by *some* extreme standard, there is some reason to think there is an evil genius, after all) – but we cannot legitimately go so far as to say that their perversity has stretched the concept of knowledge out of all recognition – in fact they have played on an essential feature of the concept” (Stine, p. 254).

8. The Warranted Assertability Objection

As already mentioned in section 1, the issue dividing invariantists and contextualists is not whether in different conversational contexts, quite different standards govern whether ordinary speakers *will say* that someone knows something. Of

course, what we're happy to call knowledge in some ("low-standards") contexts we'll deny is knowledge in other ("high-standards") contexts. The issue is whether these varying standards for when ordinary speakers *will* attribute knowledge, and/or for when they're in some sense *warranted* in attributing knowledge, reflect varying standards for when it is or would be *true* for them to attribute knowledge. The contextualist will appeal to pairs of cases where the standards for knowledge seem to vary: Low-standards cases in which a speaker seems 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 truthfully to describe that same subject as a nonknower. If the contextualist has chosen her pair of cases well, there will be a quite strong intuition about each assertion, at least when it is considered individually, that it is true. The invariantist, of course, cannot accept that both of the speakers' assertions are true, and so must deny a quite strong intuition. But it is often argued, the idea of varying standards for the warranted assertability of knowledge attributions can help the invariantist explain away the intuition that is hostile to her.

How so? Well, it has proven generally fruitful in philosophy to explain away certain intuitions by means of what we may usefully call *warranted assertability maneuvers* (WAMs). Such maneuvers involve explaining why an assertion can seem false in certain circumstances in which it's in fact true by appeal to the fact that the utterance would be improper or unwarranted in the circumstances in question. The idea behind the maneuver is that we mistake this unwarranted assertability for falsehood. Alternatively, but less commonly, 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-conditions and conditions of warranted assertability are quite different things, but that we can easily mistake one for the other. According to perhaps the most common objection to contextualism – the warranted assertability objection – the contextualist has confused a variance in the warranted assertability conditions of knowledge for a variance in their truth conditions.

To assess the power of this important objection, we should try to develop guidelines for the proper use of WAMs. Toward that end, in the next section, we'll examine a successful and then a patently unsuccessful example of such maneuvers. In section 10, then, we'll seek to draw some general lessons about the conditions under which such maneuvers should be taken seriously. We'll then be in a position, in section 11, to assess the power of the invariant's use of the warranted assertability objection against the contextualist.

9. Warranted Assertability Maneuvers, Good and Bad

When a speaker knows that P, it can seem somehow wrong, and to some it can seem downright false, for her to say "It's possible that P_{ind}." Suppose, for instance, that Kelly knows full well that a certain book is in her office. Suppose

Tom wants to borrow the book, and he asks Kelly whether the book is in her office. Here, it seems somehow wrong for Kelly to assert, "It's possible that the book is in my office." Indeed, pre-theoretically, there's some tendency to think she'd be saying something false. Such tendencies, if unchecked, could tempt one toward a "Don't Know Either Way" (DKEW) account of "It's possible that P_{ind}," according to which:

DKEW: S's assertion, "It's possible that P_{ind}" 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 think, for the correct account lies down the "Don't Know Otherwise" (DKO) path:

DKO: S's assertion, "It's possible that P_{ind}" is true iff (1) S doesn't know that P is false.¹⁹

According to DKO, Kelly is asserting the truth in our example. The tendency to think she's saying something false can be explained away as follows. Both "P" and "I know that P" are stronger than – they imply but are not implied by – "It's possible that P_{ind}," 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 Kelly knows that P, then they're usually in position to assert that they know that P, and they're always in a position to assert P itself. Thus, by the "Assert the Stronger" rule, they should assert one of those two things rather than the needlessly weak "It's possible that P_{ind}." To say the weaker thing is to make an inappropriate assertion, and it's this unwarrantedness of the assertion that we're mistaking for falsehood. In the following section, we'll see reasons for thinking that this example of a WAM is quite credible. But even before such reasons are given, I hope you will be able to sense that this maneuver is more successful than the WAM we're about to consider.

Suppose a crazed philosopher of language were to defend the view that the truth-conditions of "S is a bachelor" do not contain any condition to the effect that S be unmarried; rather, the sentence is true iff S is a male. You might think this view is refuted by such facts as that we speakers of English have strong intuitions that assertions of the form "S is a bachelor" are false when they're said of married men, and that, in our linguistic behavior, we'll refrain from saying of any male that we believe is married that he's a bachelor, and in fact will go so far as to deny that he's a bachelor. But our imagined philosopher, though crazed, is not without resources. Suppose he attempts to explain away such facts as follows: "When S is married, it is inappropriate or unwarranted to assert 'S is a bachelor.' We mistake this unwarranted assertability for falsehood. That explains why we find such assertions false when made of married men and why we won't make such assertions about married men."

10. Conditions for Successful Uses of Warranted Assertability Maneuvers

The second WAM above – the one made in defense of the weird theory about the meaning of “bachelor” – is, I hope you can sense, wildly unsuccessful. That’s no doubt partly because it’s being offered in defense of such a loser of a theory – a theory for which it’s difficult to imagine what possible positive support it might have. But it’s important for our present purposes to notice the deeper reason why the defensive maneuver has no force here. It’s an instance of a general scheme that, if allowed, could be used to far too easily explain away the counterexamples marshalled against *any* theory about the truth-conditions of sentence forms in natural language. Whenever you face an apparent counterexample – where your theory says that what seems false is true, or when it says that what seems true is false – you can very easily just ascribe the apparent truth (falsehood) to the warranted (unwarranted) assertability of the sentence in the circumstances problematic to your theory. If we allow such maneuvers, we’ll completely lose our ability to profitably test theories against examples. It would be disastrous to let theories off the hook with respect to putative counterexamples so easily.

Does that mean that we should generally stop putting any faith in WAMs? No. A number of contrasts between the first and the second example of WAMs from the above section will provide some guidelines for when such maneuvers can be accorded some legitimacy. Consider, then, the following features of the first example – the maneuver made on behalf of the DKO approach to “It’s possible that P_{ind} .”

First, where a speaker knows that P , while it does indeed seem wrong, and may even seem false to some, for her to assert “It’s possible that P_{ind} ,” it seems just as bad – in fact, worse – and certainly seems false, for her to instead say “It’s impossible that P_{ind} ” or “It’s *not* possible that P_{ind} .” 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. So 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), 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. So there appears to be no problem. Given just this, we can see already that the intuition of falsehood that some have about “It’s possible that P_{ind} ” where the speaker knows that P is a better candidate for explaining away via a WAM than are the intuitions of falsehood regarding the application of “ S is a bachelor” to married men.

Second, and closely related, the maneuver used in defense of DKO can appeal to the generation of a *false implicature* to explain the appearance of falsehood. Given that there is an “Assert the Stronger” rule, asserting “It’s possible that P_{ind} ”

will generate what, following H. P. Grice, has come to be called an *implicature* to the effect that the speaker doesn’t know that P .²⁰ An implicature is not part of “what is said,” to use Grice’s favorite phrase, in making an assertion, but it is something conveyed by the making of the assertion. In the case of “It’s possible that P_{ind} ,” the listener can, on the assumption that the speaker is following the “Assert the Stronger” rule, calculate that the speaker does not know that P , since if he did, he would have said that he knew that P , or at least would have said that P , rather than the needlessly weak “It’s possible that P_{ind} .” Thus, if a speaker breaks that conversational rule, and asserts “It’s possible that P_{ind} ” where he knows that P , while he won’t have said anything false, he will have conveyed a false implicature to the effect that he doesn’t know that P .

Most of the clearly effective WAMs involve explaining away apparent falsehood of an assertion by appeal to the generation of a false implicature. This is not surprising. 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.

What of the defense of the crazed theory about “bachelor”? As it stands, it doesn’t appeal to the generation of a false implicature to explain away apparent falsehood. It’s rather an instance of what we may call a “bare warranted assertability maneuver” – a WAM that simply explains away the problematic intuitions of falsehood by claiming the assertions in question are unwarranted or explains away intuitions of truth by appeal to the warranted assertability of the relevant assertions, without further explaining *why* the true assertions are unwarranted or the false ones warranted.

But that’s just the explanation as it now stands. It could be beefed up to appeal to the generation of implicatures. Suppose our crazed philosopher were to get even more resourceful, and argue as follows: “There’s a conversational rule to the effect that you shouldn’t assert ‘ S is a bachelor’ where S is married. Thus, when you make such an assertion, your listener, assuming that you’re following the rule, will gather that S is unmarried. Thus, making such an assertion will generate an implicature to the effect that S is unmarried. Where S is married, this implicature is false. That explains why we find such assertions to be false when they’re made of married men: We mistake the falsehood of the implicature that the assertion generates for the falsehood of the assertion itself.” So the crazed theory regarding “bachelor” can be defended by appeal to implicatures.

But there are two problems with this strategy. The first problem is closely associated with our first contrast between our two WAMs. Our crazed philosopher doesn’t only have intuitions of falsehood to explain away, but also intuitions of truth which are problematic to his theory. When we say of a married man that he is not a bachelor, that seems true, but is false according to the crazed theory. Even if you can come up with a good explanation for why the assertion would generate some true implicature, this wouldn’t seem to help much. For don’t we want to avoid falsehood both in what we implicate and (especially!) in what we actually say? So, it would seem that it would be unwarranted to assert a falsehood,

even if doing so generates a true implicature. Thus it's no wonder that most clearly successful WAMs involve explaining away apparent falsehood by appeal to the generation of a false implicature; none I know of involve explaining away apparent truth by appeal to the generation of true implicatures.

Finally, a vitally important contrast between our two WAMs involves *how* they explain the generation of the implicatures to which they appeal. They defense of DKO utilized a very general rule of conversation – “Assert the Stronger” – which applies to assertions of any content.²¹ This general rule, together with DKO's account of the content of “It's possible that P_{ind}” generates the implicature that S doesn't know that P. By contrast, the defense of the crazed theory of “bachelor” resorted to positing a *special* rule attaching only to assertions involving the term “bachelor.” Again, the danger of making it too easy to sweep away intuitions emerges. *Any* theory which omits what is in fact a truth-condition for a type of assertion could just admit the missing condition, not as a truth-condition, but rather as a condition for warranted assertability that's generated by some special rule that attaches only to assertions containing the relevant term, and then go on to explain away the intuitions of falsehood that the theory will inevitably fall prey to. If such moves are allowed, it's difficult to see how we could ever discern truth-conditions from conditions of warranted assertability. But it's not so easy to generate the implicatures you need to deflect the apparent counterexamples to your theory by means of *general* conversational rules. If your theory is subject to apparent counterexamples where the relevant sentences seem false though your theory says they're true, but if your proposed truth-conditions, together with general rules of conversation, which can be tested on very different sentences, predict that false implicatures will be generated in the circumstances that generate the apparent counterexamples, that seems to have significant potential to mitigate the damage of the apparent counterexamples.

In summary, then, the rather successful defense of DKO starts out with a better candidate for WAMming: An intuition that an assertion is false, where the opposite assertion also seems false, indicating that some intuition here has to be explained away. It then explains away the apparent falsehood of “It's possible that P_{ind}” where the speaker knows that P by means of an appeal to the generation of a false implicature. And it explains how this implicature is generated by means of a very general rule of conversation together with the DKO account of the content of the assertions in question. By contrast, the lame defense of the crazed theory of “bachelor” starts with a bad candidate for WAMming: An intuition that an assertion is false, where the opposite assertion appears to be true, so that our intuitions about both the assertion in question and its opposite would have to be explained away. To the extent that it appeals to the generation of implicatures in its explanations, it has to generate these implicatures by means of special rules that apply only to assertions involving the terms in question. And even then, it runs into trouble with intuitions of truth, where it seems to involve itself in claims that false assertions can be warranted if they generate true implicatures.

11. Evaluation of the Invariantist's Warranted Assertability Maneuver

We can now evaluate the warranted assertability objection to contextualism. If our above investigation of the conditions under which to give credence to WAMs is at all on the right track, then this objection to contextualism is a prime example of a WAM we should *not* give any credence to, since it fails every test we discerned above.

First, in the “high-standards” contexts, we don't just refrain from ascribing knowledge to the same subjects we're happy to call knowers in more lenient contexts; rather, we go so far as to appropriately *deny* that those subjects know. 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 they know, while in the “high standards” context, it seems appropriate and true to deny that similarly situated subjects know and it seems inappropriate and false to say they do know. 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. Like our imagined philosopher who defends the crazy theory about the truth conditions of “bachelor” sentences, our invariantist about “knows” is trying to employ a WAM on a set of data that doesn't seem a good candidate for WAMming.

Next, it's difficult to see how the invariantist could appeal to general rules of assertability to explain why the misleading appearances she alleges are generated. Truth be told, the warranted assertability objection against contextualism usually 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 invariantists go beyond such bare maneuvers, with an exception I'll present in the next paragraph, they seem to appeal to *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.” 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.

The exception I alluded to is Unger's defense of skeptical invariantism in *Ignorance*, which does a bit better by our criteria for successful WAMs. While Unger didn't appeal there to a fully general conversational rule like “Assert the Stronger,” he did attempt to treat “knows” as an instance of a *fairly* wide group of terms which he called “absolute” terms. A precise exposition of where and how Unger drew the line between absolute and other terms is beyond the scope of the present short paper. It will suffice for our purposes to note that it was

indeed a fairly wide group of terms, that included among their number such terms as "flat," "dry," "straight," "empty," and "square." Thus, while Unger held the skeptical invariantist position that for any assertion of "S knows that P" to be true, S must be in such a strong epistemic position with respect to P that it is impossible that anyone should be better positioned with respect to any proposition than S is positioned with respect to P, with the result that none or almost none of our attributions of knowledge are true, he did not thereby make "knows" an "isolated freak" of our language, as he put it. For he held similar views about other absolute terms like "flat," according to which surfaces had to meet very stringent, absolute standards of flatness for a sentence describing them as "flat" to be true. In general, then, according to this view, positive assertions containing absolute terms (like "S knows that P" or "X is flat") have incredibly demanding truth-conditions, making most of our uses of such assertions false, but, in general, when the person, object, or situation being described comes *close enough* for present conversational intents and purposes to satisfying the demanding criteria for the application of the relevant absolute term, it is warranted for one to falsely assert that the surface is flat, that the person knows, etc.

The Unger of *Ignorance*, then, is not subject to the charge of appealing to a special rule for the warranted assertability of "knows." Consequently, Unger's defense of invariantism is on firmer ground than the WAMs that are typically leveled against contextualism. However, several features of Unger's old invariantism work together to make it unattractive.

First, Unger's invariantism is a *skeptical* invariantism. Most who reject the contextualist's varying standards, I think, don't imagine that the constant standards they endorse will be so demanding as to be unmettable by mere mortals. However, it's unclear, to say the least, that a general invariantism about absolute terms of the kind Unger mounts could be used in defense of a moderate or nonskeptical invariantism, according to which the constant standards that govern the truth-conditions of sentences containing absolute terms will be measurably low, for such a moderate account will not be able to utilize Unger's account of assertability going by whether the belief in question is *close enough* to being knowledge to be appropriately called such. Second, the cost of the generality of Unger's account is that the impossibly demanding standards and the resulting systematic falsehood in what we say spreads to a large stretch of our language. Thus, we end up speaking falsely whenever we describe a physical surface as "flat," or when we apply any number of "absolute" terms in the way we're accustomed to applying them. I'm fairly confident that most would find a general contextualist approach to absolute terms far more plausible than such a relentlessly demanding invariantist approach. Finally, by not utilizing a thoroughly general rule which has clearly correct applications like "Assert the Stronger," the Unger of *Ignorance* loses a lot of leverage in advocating his view. His rule of assertability would be something like: "When x is close enough, for present conversational intents and purposes, to satisfying the semantic requirements for 'F,' where 'F' is an absolute term, it is appropriate to describe x as being F." But a general contextualist account of the

use of what Unger calls "absolute" terms which avoids the systematic falsehood is available, and so it's difficult to see where the pressure to accept a demanding invariantist account will come from. By contrast, a WAM that utilizes a rule like "Assert the Stronger," will have the advantage of being based on a rule that we are independently motivated to accept. Thus, I don't think many will be tempted by Unger's old approach.

Most who let considerations like those embodied in the Warranted Assertability Objection dissuade them from accepting contextualism don't see themselves as accepting anything so radical as Unger's general approach of relentless stringency and systemic falsehood for large stretches of our language. But then, as far as I can see, they are being moved by nothing more than a bare WAM that miserably fails to meet all the reasonable criteria we can discern for what it would take for a WAM to be successful. Of course, some invariantist may be able to more successfully explain away the facts supportive of contextualism in a way that I can't see. But to date, there's been no such explanation that I am aware of. And even if a general rule can be found to do *some* of the damage control that the invariantist needs done, the generation of implicatures seems to hold promise only for explaining away the intuitions of falsehood that the invariantist must reject, and seems powerless against the intuitions of truth that any invariantist will find lined up against his position, for reasons we have already seen. The Warranted Assertability Objection to contextualism, in short, seems to be a paradigm case of a powerless WAM.

Notes

- 1 A notable exception to this is Stephen Schiffer's "Contextualist Solutions to Scepticism," *Proceedings of the Aristotelian Society* 96 (1996): 317-33. I do not here respond to Schiffer's criticisms. Schiffer's criticisms raise important issues about the viability of what he calls "hidden indexical" accounts of various types of sentences. I am eager to defend not only contextualism about knowledge attributions, but also hidden indexical accounts of attitude ascriptions, from Schiffer's attacks (see esp. pp. 510-19 of Schiffer, "Belief Ascription," *Journal of Philosophy* 89 (1992): 499-521), but such a defense, if it were to be at all adequate to the task, would be far too extensive for the present paper.
- 2 See Peter Unger, *Philosophical Relativity* (Minneapolis, MN: University of Minnesota Press, 1984).
- 3 See Cohen, "Knowledge, Context, and Social Standards," *Synthese* 73 (1987): 3-26; Cohen, "How to be a Fallibilist," *Philosophical Perspectives* 2 (1988): 91-123; DeRose, "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52 (1992): 913-29; DeRose, "Solving the Skeptical Puzzle," *Philosophical Review* 104 (1995): 1-52; Lewis, "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8 (1979): 339-59, esp. Example 6, "Relative Modality," pp. 384-5; Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (1996): 549-67; and Unger, "The Cone Model of Knowledge," *Philosophical Topics* 14 (1986):

- 125-78. Gail Stine advocates a contextualist version of the "Relevant Alternatives" theory of knowledge in Stine, "Skepticism, Relevant Alternatives, and Deductive Closure," *Philosophical Studies* 29 (1976): 249-61. Fred Dretske is another prominent relevant alternative, but whether he's a contextualist is extremely difficult issue. I won't here go into a long, inconclusive story about this matter, but will content myself with pointing the reader to the most enlightening material on it: See pp. 191 (starting with the paragraph that begins "This brings me to the very important distinction . . .") -196 of Dretske, "Dretske's Replies," in Brian P. McLaughlin, ed., *Dretske and His Critics* (Cambridge, MA: Blackwell, 1991). *Editor's note:* versions of contextualism are also endorsed by Williams and Longino in their essays in this volume.
- 4 Lewis, *Counterfactuals* (Cambridge, MA: Harvard University Press, 1973), p. 92; the full sentence, and the sentence which precedes it, which are about the link between similarity and counterfactual conditionals, read: "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."
 - 5 The subscript "ind" indicates that the embedded P is to be kept in the indicative mood: very different possibilities are expressed where the P is subjective. The subjunctive, "It's possible that I should not have existed," is just plain good sense, while the indicative "It's possible that I don't exist" is bizarre.
 - 6 A more exact analysis of the sentence would still involve the concept of knowledge in a way that, though more complicated, doesn't ruin the point I'm making. See DeRose, "Epistemic Possibilities," *Philosophical Review* 100 (1991): 581-605 for such a more exact analysis.
 - 7 For more on contextualism and statements of possibility, see ch. 2, section V ("Contextualism and Epistemic Modal Statements") of my *Knowledge, Epistemic Possibility, and Skepticism* (Ph.D. diss., University of California, Los Angeles, 1990; University Microfilms International order number 9035253), and my "Simple Might's, Indicative Possibilities, and the Open Future," *Philosophical Quarterly*, 1998, especially section 3 ("The Second Problem, the Context-Sensitivity of Knowledge, and the Methodology of Flar-Footed, 'All-in-One-Breath' Conjunctions").
 - 8 Though "It's certain that P" is the dual of "It's possible that P_{ind}," it doesn't mean just that the speaker knows that P, because the provisional analysis of "It's possible that P_{ind}" as simply being a matter of the speaker not knowing otherwise is just a first approximation, and so isn't exactly correct. For a more exact account of the meaning of "It's certain that P," see my "Simple Might's," esp. section 8 ("It's Certain That"). Expressions of the form "S is certain that . . .", where S is a subject - expressions of *personal* certainty - of course, are a completely different matter.
 - 9 In "Elusive Knowledge," Lewis explicitly denies that justification is even necessary for knowledge. It's still open to him, of course, to be a contextualist about justification, but if he were to go that route, this contextualism would be independent from his contextualism regarding knowledge.
 - 10 David Annis, "A Contextualist Theory of Epistemic Justification," *American Philosophical Quarterly* 15 (1978): 213-19.
 - 11 As of now, I haven't seen this in print, though I've read several papers in manuscript that mark this distinction, some of them with exactly those labels. (When those exact labels are used, their use *may* derive from my distinction between "subject

factors" and "attributor factors" in part II of DeRose, "Contextualism and Knowledge Attributions.")

- 12 This case appeared prominently in Alvin Goldman, "Discrimination and Perceptual Knowledge," *The Journal of Philosophy* 73 (1976): 771-91. On p. 252 of Stine, "Skepticism, Relevant Alternatives, and Deductive Closure," it is reported that Goldman attributes the example to Ginet.
- 13 Norman Malcolm, "Knowledge and Belief," *Mind* 51 (1952): 178-89.
- 14 Peter Unger, *Ignorance: A Case for Skepticism* (Oxford: Oxford University Press, 1975). This book incorporates, with some improvements, Unger's important journal articles from the early 1970s, while adding new material as well.
- 15 Barry Stroud, *The Significance of Philosophical Skepticism* (Oxford: Oxford University Press, 1984).
- 16 See especially Cohen, "How to be a Fallibilist."
- 17 For more on the relationship between RA and contextualism, see part II of DeRose, "Contextualism and Knowledge Attributions," and also DeRose, "Relevant Alternatives and the Content of Knowledge Attributions," *Philosophy and Phenomenological Research* 56 (1996): 193-7.
- 18 For much more on the contextualist approach to skepticism, see DeRose, "Solving the Skeptical Problem."
- 19 This account is on the right path (the DKO path), but is not correct as it stands. See DeRose, "Epistemic Possibilities," for my best attempt at the truth conditions of "It's possible that P_{ind}."
- 20 See esp. "Logic and Conversation," in H. P. Grice, *Studies in the Way of Words* (Cambridge, MA: Harvard University Press, 1989).
- 21 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 in section 2 of Frank Jackson, "On Assertion and Indicative Conditionals," *Philosophical Review* 88 (1979): 565-89.