

This is a prepublication draft of an article that appears in its final and official form in *The Philosophical Quarterly*.

The Ordinary Language Basis for Contextualism and the New Invariantism

Keith DeRose
Yale University

- 1. Mutually Reinforcing Strands of Evidence**
- 2. Truth/Falsity Asymmetry**
- 3. The Best Cases: Standards Appropriate to Practical Context**
- 4. The Best Cases: Cases Involving No Dispute and No Reversals**
- 5. Problems with First-Person Cases**
- 6. Third-Person Cases**
- 7. Contextualism vs. SSI: A Big, Ugly Tie?**
- 8. The Importance of Arguments from Ordinary Language**
- 9. Methodology, Straightforward Data, Comparative Judgments of Content, Metalinguistic Claims, Belief Reports, and Speech Reports**

The Ordinary Language Basis for Contextualism and the New Invariantism

Keith DeRose
Yale University

The best grounds for accepting contextualism concerning knowledge attributions come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, non-philosophical talk: What ordinary speakers will count as “knowledge” in some non-philosophical contexts they will deny is such in others.

This type of basis in ordinary language provides not only the best grounds we have for accepting contextualism concerning knowledge attributions, but, I believe, is evidence of the very best type one can have for concluding that any piece of ordinary language is context-sensitive. Building on earlier work I’ve done on the topic, but going beyond that earlier work in some key ways, in what follows, we’ll look in more detail at some of the features of the ordinary use of “knows” that make for a compelling case for the contextualist account of that verb (and its cognates), and explain why these features militate so strongly for a contextualist treatment. So we’ll be looking both at what the data from ordinary language are, and at the methodology that takes us from that data to a contextualist conclusion. Thus, this will largely be an exercise in how to do ordinary language philosophy.

1. Mutually Reinforcing Strands of Evidence

Contextualists typically appeal to pairs of cases that forcefully display the variability in the epistemic standards that govern ordinary usage: A “low standards” case in which a speaker seems quite appropriately and truthfully to ascribe knowledge to a subject will be paired with a “high standards” case 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. The contextualist argument based on such cases ultimately rests on the

premises that the positive attribution of knowledge in the low standards case is true, and that the denial of knowledge in the high standards case is true. Why think that both of these claims are true?

Well, first, and most directly, where the contextualist's cases are well-chosen, those are the fairly strong intuitions about the cases, at least where each case is considered individually. Here we appeal to how we, competent speakers, intuitively evaluate the truth values of particular claims that are made (or are imagined to have been made) in particular situations. Our intuitions about such matters can be wrong, of course, but still, they are among our best guides when evaluating semantic theories, especially when we are careful to avoid the types of situations where we are likely to be misled, and there is good reason to think that the intuitions the contextualist appeals to here are among those we should invest much confidence in.

To reinforce these intuitions, the contextualist can appeal to the facts that his cases display how speakers in fact, and with propriety, use the claims in question. As we can sense, perhaps with even more certainty than that with which we judge the truth values of the claims, speakers do in fact use "knows" in the way described, and appropriately so – they will in fact, and with apparent propriety, ascribe "knowledge" in situations like the contextualist's low-standards case, yet will deny "knowledge" when they find themselves in conversational circumstances like the contextualist's high-standards case. This supports the premises that both of the imagined claims are true, since generally (though there are some exceptions), one cannot properly claim to know something that from one's own point of view (given one's beliefs about the underlying matters of fact relevant to the claims in question) is false. So, since the contextualist's cases do not involve speakers who are involved in some mistaken belief about a relevant underlying matter of fact, there is good reason to think that their claims, which are made with perfect propriety, are true, and it's a bad strike against a semantic theory if it rules these claims to be false, as it seems invariantism will have to rule with respect to one or the other of the claims.

To cite these facts concerning common and appropriate uses of "knows" as evidence for the truth of our imagined speakers' claims that is thoroughly independent of our intuitions that their claims are true would be to engage in some illegitimate double-accounting. For it is no doubt largely because we already sense that the speakers in the cases are using "knows" in very common and appropriate ways that we intuitively judge that their claims are true.

But it is possible for claims to appear to be true without being in the relevant sense appropriate, though our intuitions that such claims are true will often be more tentative in such a case. Indeed, if I can be permitted to engage briefly in some personal history, my earliest attempts at formulating a pair of cases to support contextualism suffered from just a defect. My advisor, Rogers Albritton, objected, as near as I can remember, “Nobody would really talk that way!” I replied that it didn’t matter whether people would talk that way. All I needed was that such a claim would be true, and that certainly was my intuition about the truth-value of the claim. He would have none of that, and answered, quite sternly, “Look, if you’re going to do ordinary language philosophy – and that’s what you’re doing here – you’d better do it right.” So I returned to the cases, and adjusted them so that what I had the speakers say seemed quite natural and appropriate, as well as true. And, to my surprise, the changes did not only make the argument read more nicely, which was the only pay-off I was expecting, but actually made it more compelling. Albritton never explained to me why the examples should be constructed so that what’s said is natural and appropriate beyond insisting that that’s how ordinary language philosophy should be done. (He seemed to think it a point too obvious to require explanation, and I was not about to ask!) But I think the reason that helps in supporting the claim that what one’s imagined speaker is saying is true is that it engages the general presumption that, where speakers are not basing their claims on some false beliefs they have about underlying matters of fact, how they naturally and appropriately describe a situation, especially by means of very common words, will be a true description.

2. Truth/Falsity Asymmetry

Partly because they can be so reinforced by the naturalness and appropriateness of the claims involved, intuitions we have that claims are true are more trustworthy than are intuitions we might have to the effect that claims are false. Of course, when a speaker makes a claim that is, and is from her own point of view, false, her claim will be improper as well as false. But this impropriety cannot be used to buttress the intuition that the claim is false with the same security that one can use the appropriateness of a claim to reinforce the intuition that the claim is true, because there is not nearly as strong a presumption that inappropriate claims are false as there is

that appropriate claims are true. As David Lewis points out, “There are ever so many reasons why it might be inappropriate to say something true. It might be irrelevant to the conversation, it might convey a false hint, it might be known already to all concerned...”¹ And though Lewis does not go on to say so, the comparative point he is making depends on, and he therefore implies that, it is not nearly as likely that an appropriately made claim will be false. And that implied point seems right. For it seems that, except where we engage in special practices of misdirection, like irony or hyperbole, we should seek to avoid asserting falsehoods, and we will thus be speaking improperly if we assert what is, from our own point of view, false.

Given that we really do (and with seeming truth and propriety) deny “knowledge” in situations like the contextualist’s high-standards case, the contextualist is wise, then, to construct that case so that the speaker does there deny that the subject “knows,” rather than having the speaker ascribe knowledge in a seemingly false way. Because of the comparative point made in the above paragraph, any intuition we might have that such a speaker’s positive ascription of “knowledge” is or would be false in the high-standards case would be more suspect than is the intuition that contextualists in fact rely on – that the denial of knowledge he actually puts in his speaker’s mouth is true.

For this and some other reasons,² it greatly helps the contextualist’s case that we do (actually, and with seeming propriety and truth) go so far as to deny “knowledge” in situations like his high-standards case. The case for contextualism would be much weakened if the facts were instead only that we are reluctant to positively ascribe “knowledge” in high-standards situations, and/or that if a misguided speaker were to positively ascribe “knowledge” in such cases, the ascriptions would seem false and/or inappropriate.

¹ Lewis, “Causation as Influence,” *Journal of Philosophy* 97 (2000): 182-197; p. 196; ellipsis in original. I advocated the asymmetry urged in this section in correspondence with Lewis in the early 1990s, where I argued that Lewis’s account of the relation between “might” and “would” counterfactual conditionals was in particularly bad trouble because it ruled as false certain claims that intuitively seemed true and that were perfectly appropriate, while its key rival merely had to rule true some claims that seemed wrong. This later came out in my “Lewis on ‘Might’ and ‘Would’ Counterfactual Conditionals,” *Canadian Journal of Philosophy* 24 (1994): pp. 413-418. Lewis might have already accepted that asymmetry at the time. He raised no methodological objection to me in his correspondence.

² For some other reasons, see especially section 3 (pp. 188-194) of DeRose, “Assertion, Knowledge, and Context,” *Philosophical Review* 111 (2002): 167-203.

3. The Best Cases: Standards Appropriate to Practical Context

Contextualists can differ among themselves about what types of features of the context of utterance really do affect the truth conditions of knowledge attributions and to what extent they do so. Thus, there can be pairs of cases, differing only in “disputed” features – features that some contextualists think affect the standards that comprise a truth-conditions of knowledge-attributing sentences, but that other contextualists think have no such effect. Such pairs of cases do not provide the best tests for deciding between contextualism and invariantism. The best case pairs will differ with respect to as many of the features that plausibly affect the epistemic standards, and especially those features which most clearly appear to affect epistemic standards, as possible. It’s about such pairs of cases that the intuitions supporting contextualism will be strongest. And, given the content of their position, invariantists must resist the intuitions supporting contextualism even with respect to such cases. By contrast, it’s of course open to contextualists to hold that there is no difference in truth-conditions between the cases in some of the case pairs.

So, what makes for a good pair of cases? I’ll here go into only a couple of important ingredients. But, first, it makes the relevant pro-contextualist intuitions stronger and more stable if the elevated epistemic standards in the high standards case are tied to a pressing and very practical concern, and thus seem reasonable given the situation, and if the lower epistemic standards in the other case are also reasonable given the different practical situation involved there. And it helps (as we’ll explore more fully in the following section) if all the parties to the conversation accept as reasonable both the elevated standards of the one case and the lower standards of the other. Thus, in the pair of cases I have used – my Bank Cases³ – one character (myself, as it happens), claims to know that a 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

³ DeRose, “Contextualism and Knowledge Attributions,” *Philosophy and Phenomenological Research* 52: 913-929, p. 913.

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 catastrophically 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 that the bank may have changed its hours in the last couple of weeks. This possibility seems well 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.⁴

We can encounter very tricky cases, about which contextualists can disagree among themselves, if the actual practical situation that a speaker faces is importantly different from how he perceives his situation to be. Thus, if there were a variant of the Bank Cases in which it is in fact vitally important that I be right about whether the bank is open on Saturdays, but in which I do not realize that the stakes are so high, perhaps because I have forgotten that I have written the large and very important check, or perhaps because my wife was the one who wrote the check without my being aware of it, so that I claim to “know” the bank is open on relatively meager grounds, some contextualists may think the standards remain low and my claim to “know” is true, while others may think the raised stakes, even when I’m unaware of them, raise the standards for knowledge that govern the truth-conditions of my use of “knows” and its cognates, so that my claim is false. Such cases may be useful for deciding among different versions of contextualism, but are to be avoided in deciding between contextualism and invariantism in the first place. Better cases for that latter purpose, which is the purpose at issue in this paper, will be

⁴ 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.

a low standards case where the stakes are low and are clearly seen to be low, on the one hand, and a high standards case where the stakes are high and clearly seen to be high, on the other.⁵

To avoid confusion, I should here explicitly add that while I believe that the best case pairs for establishing contextualism involve a marked difference in the stakes involved, I myself do not believe that such a difference in stakes is necessary for a difference in the semantic standards for “knows.” In fact, I think that speakers are free to use standards even wildly inappropriate to the practical situations they face – for instance, to use low standards where they face an extremely high-stakes situation in which it would be much wiser for them to employ much higher standards. But when speakers do use inappropriate, and especially wildly inappropriate, standards, our intuitions about the truth-values of their claims become very insecure, partly, I think, because it is easy to confuse the fact that there is something very wrong with their utterance (as will be the case when they are employing standards wildly inappropriate to the situation they face) for their claim’s being false. That’s why it’s best to test contextualism by means of case pairs in which the standards employed in each case are appropriate to the situation of the case in question. And, again, about even such cases, where the intuitions are quite secure, the invariantist will have to rule that at least one of the relevant intuitions is wrong.

4. The Best Cases: Cases Involving No Dispute and No Reversals

Contextualist theories of knowledge attributions are typically developed with an eye toward providing some kind of contextualist response to the problem of philosophical skepticism. But it is important to first develop support for contextualism from ordinary, non-philosophical cases, because, among other reasons, the relevant intuitions are not very strong where the “high standards” case is one involving a philosophical skeptic, who denies “knowledge” in ways that are not tightly tied to appropriate practical concerns, and whose denials seem to be based on

⁵ Thus, it’s worth noting that where Jason Stanley (“On the Linguistic Basis for Contextualism,” forthcoming, Philosophical Studies, section 1) alters my Bank cases in presenting them, collapsing my two separate cases, one of which never involves high stakes, and the other of which does involve high stakes that are clear to the speakers throughout the case, into one temporally extended case in which the speaker first claims to “know”, but then, after being reminded of the high stakes that she has faced all along, concedes that she does not “know”, the change isn’t innocent, but significantly hurts the effectiveness of the cases. For the early, positive claim of knowledge in Stanley’s temporally extended case is of just the problematic type we have been discussing.

standards that are exceedingly high. Many will find their intuitions in revolt against these philosophical denials of “knowledge.”

So, contextualists have wisely appealed to cases which don’t involve philosophical skepticism. It does help if something that functions as a “skeptical hypothesis” is inserted into the conversation in the high-standards case – a hypothesis that is perhaps too far-fetched to be considered in the low-standards case. But these hypotheses are still far more moderate than are the ethereal possibilities on which philosophical skeptics base their hyperbolic doubts (brains in vats, or whatnot), and do seem worth considering under the somewhat heightened scrutiny appropriate to the high-standards case. Thus, in my high standards Bank Case, my wife raises the possibility that the bank has changed its hours in the last two weeks, and in his very similar Airport Case, Stewart Cohen has one of his characters raise the possibility that the itinerary they are relying on to tell whether their flight stops in Chicago might contain a misprint.⁶

But the best cases for contextualism do not just avoid the hyperbolically far-fetched possibilities and the astoundingly high epistemic standards of philosophical skeptics. The best cases avoid there being any dispute among the parties to the conversation over whether the proposition in question is or is not “known,” and avoid any of the speakers reversing what they say on the matter of what is or is not “known.”

When such things are avoided, the intuitions supportive of contextualism are strongest. That’s to be expected, since contextualists themselves can, and some do, have worries about how to handle such cases. In fact, with respect to at least some cases of dispute, my own (contextualist) position is that the speakers involved are not speaking truths. I have worked this out mainly with reference to disputes between philosophical skeptics and their commonsensical opponents,⁷ but the position I advocate there – the “gap view,” on which neither the philosophical skeptic who insists that his opponent does not “know” something nor her commonsensical opponent, who insists that she does indeed “know” the thing in question are making true claims⁸ – can be easily applied to disputes involving more moderate, non-philosophical “skeptics” and their opponents, and I would be inclined to apply that view to many such cases. So it would be a mistake to base contextualism on cases involving such disputes,

⁶ Cohen, “Contextualism, Skepticism, and the Structure of Reasons,” Philosophical Perspectives 13 (1999): 57-89; p. 58.

⁷ See DeRose, “Single Scoreboard Semantics,” forthcoming, Philosophical Studies.

⁸ See the final section of “Single Scoreboard Semantics.”

claiming that both parties are speaking the truth – cases about which even card-carrying contextualists might have their doubts.⁹

Some opponents of contextualism, apparently viewing contextualism as supported mainly by a certain resolution to the problem of philosophical skepticism that it makes possible – a resolution on which the apparently disputing parties are both speaking truthfully – have complained that such a resolution is always available for any dispute, and so the contextualist solution is far too general a problem-solver. We can always claim that a key term involved in a hard-fought dispute is context-sensitive, and thus claim the advantage for our “contextualism” that it dissolves this apparent dispute. As Timothy Williamson sums up his version of this complaint: “Contextualism supplies a perfectly general strategy for resolving any apparent disagreement whatsoever. Since some disagreements are genuine, we should not always follow that strategy.”¹⁰ But this charge seems damaging only if contextualism is being supported exclusively, or at least mainly, by its ability to provide such a resolution to the dispute between philosophical skeptics and their opponents. Wise contextualists, however, as we have seen, seek to support their contextualism by appeals to what transpires in ordinary, non-philosophical uses of “knows,” the most important examples of which are cases that involve no dispute among the parties whatsoever, before applying their contextualism to the problem of philosophical skepticism. Other disputes involving other key terms can be just as well solved by a contextualist approach only if similar evidence of context-sensitivity is available for those terms. Now, contextualists will typically also claim that the treatment of the problem of philosophical skepticism that we are able to supply provides further support for their views. But the solutions there appealed to typically are not dissolutions of dispute between philosophical skeptics and their opponents on which both parties are speaking the truth that are achieved by a one-size-fits-all generic contextualism. Indeed, contextualists typically don’t even accept solutions which so dissolve the dispute: None of Stewart Cohen, David Lewis, or myself are inclined toward what I have called a “multiple scoreboards” view on which both the philosophical skeptic and her opponent are speaking truthfully when they argue, but instead opt for some “single scoreboard”

⁹ Thus, again, Stanley’s alteration of my Bank cases (see note 5, above) hurts their effectiveness, this time by including such a reversal.

¹⁰ Williamson, “Knowledge and Scepticism,” to appear in F. Jackson and M. Smith, ed., The Oxford Handbook of Analytic Philosophy (Oxford University Press); cf. Richard Feldman, “Contextualism and Skepticism,” *Philosophical Perspectives* 13 (1999): 93-115, pp. 103, ff.; and Jason Stanley, “Context, Interest-Relativity, and Knowledge,” draft of 11 January, 2004, p. 46

view on which the dispute is quite real indeed, and it's quite impossible for both disputants to be speaking truthfully.¹¹ (In fact, on the "single scoreboards" view I opt for, and on both of the views that seem to be Cohen's finalists, neither the skeptic nor her opponent are speaking truthfully as they dispute!¹²) We solve (or try to solve, depending on whom you talk to) various skeptical paradoxes that are generated by skeptical arguments not by making both the disputants in a fight turn out to be making truthful claims, but by explaining why we are led into a paradoxical set of intuitions. Our explanations are built upon various standards-changing mechanisms we claim govern knowledge attributions, and so our solutions are not instances of some general contextualist strategy available wherever there is some dispute that is hard to adjudicate.

Unfortunately for the invariantist, the cases where varying epistemic standards are displayed are not limited to cases where there is some dispute over what is "known", or cases in which some speaker is led to reverse what she has said. It is a feature of how we use "knows" that sometimes heightened practical concerns make it seem appropriate for speakers to deny that subjects "know" things so long as those subjects don't meet unusually demanding epistemic standards, and speakers then do deny that the subjects involve "know" propositions that the speakers would ordinarily say those subjects do "know". And sometimes – indeed, ordinarily – when they do this, they meet with no resistance at all from their conversational partners, who have also sensed that unusually stringent standards are appropriate in the context, and have adjusted their use of "knows" to these high standards, themselves saying, with apparent truth and appropriateness, that they and other subjects do not "know" things if those subjects fall short of the extraordinarily high standards in question, even if they meet more ordinary standards. These best high-standards cases involve no element of arm-twisting, arguing or convincing, but only of informing – one speaker informing another that, according to the appropriately heightened standards that obviously govern their conversation, some subject does not "know" something. Paired off against such high-standards cases, it's also a feature of our use of "knows" that often, where unusually heightened practical concerns are not involved, speakers behave in ways that show that they are in a low standards context, informing each other that the same subjects described above, or similarly positioned subjects, do "know" those same things that they alleged

¹¹ See DeRose, "Single Scoreboard Semantics."

¹² See the last two sections of "Single Scoreboard Semantics."

not to “know” in the high standards cases. And where there is no dispute among the parties to the conversation in either of cases; and where none of the speakers in either case is made to reverse herself, saying that something is not “known” that she had earlier claimed was “known,” or vice versa; and where the different standards that seem to govern the high-standards case and the low-standards both seem appropriate, given the practical situations involved in each case; and where none of the claims about who does or doesn’t “know” what are based on false beliefs the speakers have about underlying matters of fact; the intuitions that the relevant claims in each of the cases (the ascription of knowledge in the low-standards case and the denial of knowledge in the high-) are true can become extremely powerful indeed. And these strong intuitions of truth are buttressed by the presumption that what is properly said is true. And even with respect to these very strong strongest cases, invariantists must deny one of these intuitions involved.

5. Problems with First-Person Cases

Thus far, in looking for the best cases to use in support of contextualism, I have been concerned with ways the cases can be constructed to make the contextualist’s premises – that the key claim in each case is true – as secure as possible. But it is also important that the cases be constructed so that those premises support a contextualist conclusion as powerfully as possible.¹³ For a couple of reasons, it helps in this latter respect that the cases involve third-person attributions (and denials) of knowledge, preferably to subjects not involved in the conversation in which the attributions and denials are made, rather than first-person claims to know and first-person admissions that the speaker herself does not know. Thus, contextualists make a mistake when, as I have done, we focus our attention on cases involving first-person claims.

To see one problem with first-person cases, consider my Bank Cases. In the low-standards case, a speaker (myself, as it happens) claims that he himself knows that the bank is open on Saturdays, while in the high-standards case, the speaker admits that he does not know

¹³ In his presentation of the Bank cases, Stanley inserts a “really” into the denial of knowledge; he has the speaker say, “I guess I don’t really know that the bank will be open tomorrow.” This is another (see notes 5 and 9, above) alteration of the Bank cases that hurts their effectiveness, this time not by making it needlessly problematic that the claim in question is true, but rather by making it needlessly problematic that the truth of the claims in question supports contextualism. This becomes problematic because the function of the “really” is quite unclear. The case does not really need to have that “really” in the denial of knowledge, and my original version of the cases includes no such “really.”

that the Bank is open on Saturdays. For reasons we've been considering, the contextualist's premises that these two claims are each true seem quite secure. But do they really support contextualism? I believe so, but that certainly can be questioned. I think these cases support contextualism because the best explanation for why both claims are true, despite their surface opposition to one another, is that the semantic standards for knowledge differ between the cases. And that, in turn, is based on my sense that the subject – who is identical with the speaker in first-person cases like these – is in an equally strong epistemic position with respect to the bank being open on Saturdays in the two cases. What's always seemed to me the most natural explanation for why the speaker can truthfully claim knowledge in the one case but truthfully deny it in the other, then, is that what semantically counts as "knowledge" has changed, so that the subject's same epistemic position in the two cases meets the truth-conditions for "knowledge" in the one case, but not in the other.

But recall the differences between the cases that, according to the contextualist, contribute to the difference in standards between the cases: In the high standards case, the stakes are higher, making it appropriate to worry about possibilities of error that would usually be ignored; one such possibility is actually mentioned by the speaker's conversational partner, and is accepted in the conversation as a legitimate worry; the speaker then is thinking about, and taking seriously, this possibility of error as he makes his claim; etc. The invariantist may claim that it is far from clear that these differences don't affect how well-positioned the subject really is with respect to the bank being open on Saturdays. Perhaps both the positive claim and the denial of knowledge are true, not because the semantic standards for "knowledge" have changed, but because, though the same standards govern both cases, the speaker meets those standards in the so-called "low-standards" case, but fails to meet those same standards in the so-called "high-standards" case.

It's worth noting that contextualism would be validated by mixed explanations, according to which the differences between the two cases induce both a change in what semantically counts as "knowledge" and the strength of the speaker/subject's epistemic position. To make good on the worry currently under consideration, then, the invariantist needs to hold out for the possibility that it is entirely a matter of the speaker/subject's difference in epistemic position between the two cases that explains how both the positive and negative claim can be true. Nevertheless, this

worry considerably weakens the contextualist's argument, so it wise for the contextualist to sidestep this worry if possible.

And there is another problem with first-person cases that is worth considering – and that recent work by John Hawthorne and Jason Stanley on a new form of invariantism makes more pressing. Perhaps different epistemic standards do comprise truth-conditions for knowledge attributions in the contextualist's two first-person cases, but contextualism is still not true, because these differing standards can instead be explained by a view we can call "Subject-Sensitive Invariantism" (SSI). According to SSI, the varying standards that comprise a truth-condition of "I know that P" are sensitive to factors that attach to the speaker as the putative subject of knowledge, rather than as the speaker of the knowledge attribution. That is, according to SSI, these factors of the subject's context determine a single set of standards that govern when the subject himself, or any other speaker, including those not engaged in conversation with the subject, can truthfully say that the subject "knows." Thus, we do not get the result that contextualists insist on: that one speaker can truthfully say the subject "knows," while another speaker, in a different and more demanding context, can say that the subject does "not know", even though the two speakers are speaking of the same subject knowing (or not knowing) the same proposition at the same time. But where the subject faces a different situation, different standards govern whether she, or anybody else, can truthfully say that she "knows." And the backer of SSI can claim that just that is transpiring in the contextualist's two cases: different standards do govern the two cases, but they work in the way SSI claims, rather than as the contextualist would have us believe.

6. Third-Person Cases

Fortunately, then, for the contextualist, we use third-person attributions of and denials of knowledge in describing subjects who are no party to our conversation in a way that demonstrates the same variation in standards that is displayed by our first-person talk of what we ourselves do and do not "know," and, because of this, third-person pairs of cases are available that are as powerful as are the best first-person cases.

In a way that will work as well for other cases, we can convert my first-person bank cases into third-person cases about which the contextualist's premises – that both the positive ascription of knowledge in the low-standards case and the denial of knowledge in the high-standards case – are as secure as they are about the original first-person cases. So, take the low-standards bank case, but make the topic of conversation whether some third party, not present at the conversation, “knows” whether the bank is open, and construct the case so that one of the speakers, but not the other, knows that this absent subject has the same grounds for thinking the bank is open on Saturdays as the speaker has: the absent subject, too, was at the bank just two weeks ago on Saturday, and found it open, etc. Suppose that this far-away subject is herself at the time in another low-standards conversation about whether the bank is open on Saturdays, in which she has counted herself as “knowing” that it is based on her grounds. Now, even though that absent subject, like the speaker, hasn't made any of the special checks that would be needed to rule out such relatively far-fetched possibilities as that the bank has changed its hours in the last two weeks, it's clear that an English speaker, facing such a situation, who is saying that he himself “knows” that the bank is open on Saturdays, will also say that this absent subject “knows” the fact as well. And when the contextualist has the speaker claim that the absent subject “knows,” the intuition that that claim is true will be as powerful as it is in a first-person case, and the intuition will be buttressed by the related considerations that the speaker is speaking as we all in fact do, and with propriety, as much as in the case of the first-person case.

Next, similarly convert the high-standards case so that the topic of conversation is whether some third party, not present at the conversation “knows” whether the bank is open. Suppose again the absent subject is herself at the time in a low-standards setting, but is unaware of the fact that, far away, our speakers are talking about her. But this time, of course, our two speakers are in a high-stakes, high-standards case, and are wondering whether various people, including our absent subject, might “know” that the bank is open. Suppose they are wondering about this because they hope that if some far-away subject does “know” this, our speakers might then be able to use this “knower” as an informant on the matter, and might themselves come to “know” that the bank is open. Give the absent subject in question the same grounds for believing that the bank is open on Saturdays as in the low-standards case, and construct this new case so that, as with the old, one of the speakers, but not the other, knows what kind of epistemic position this absent subject is in at her far-away location. It is clear that an English speaker who

faces this situation, and who is saying that he himself does “not know”, will say that our absent subject, whom the speaker knows has no more evidence than the speaker himself has, and who is being considered as a potential informant, similarly “does not know” that the bank is open, and the contextualist’s premise that this high-standards third-person denial of knowledge is true will be about as secure as is his premise concerning the high-standards first-person case.

Since the contextualist premises we are now considering concern third-person cases, the worries we looked at in the previous section seem not to apply – or at the very least are not as worrisome. To press the first worry, the invariantist will have to claim that the absent party is in a worse epistemic position in our so-called “high standards” case as compared with her position in our “low-standards” case. But we’ve put the absent party in the extremely similar situations in the two cases! As with the first-person cases, she has the same grounds for thinking the bank is open in our two new cases. And in our new third-person cases, she is not confronted with heightened practical concerns and a conversational partner pressing relatively far-fetched possibilities of error in the high-standards case – things which might plausibly be thought to hurt her epistemic position. In fact, our subject’s situation is identical in our two third-person cases except for the difference that, far away from her, different conversations, that she is and will remain oblivious to, are taking place about her. So, to press the first worry, an invariantist must maintain that how strongly our subject is positioned with respect to the bank being open on Saturdays is substantially affected by such factors as what kind of far-away conversation is taking place about the subject. That’s quite implausible on its face. Far better, it seems, to explain how both the positive and negative claims in these two cases can both be true by supposing that the semantic standards, rather than the subject’s ability to meet a single set of standards, differs between these two cases.

With regard to our second worry, these third-person cases provide a powerful objection – to my thinking, a killer objection – to SSI, which predicts that our speaker will apply the lower standards to our subject. For SSI holds that the situation of the subject sets the standards that govern all attributions and denials of knowledge to that subject, and our speaker is aware of the facts that make the subject’s context a low-standards one.¹⁴ Yet it is clear that speakers in

¹⁴ It’s a bit artificial to suppose that our speaker has so much information as I’ve given him about the subject’s conversational situation. I did so to make the argument simpler. In more realistic situations, where the speaker doesn’t know what standards govern the subject’s context, SSI makes somewhat more complicated, but equally false, predictions: It predicts that the speaker will be agnostic about whether the absent subject “knows”, since he

situations like the one our speaker faces in the third-person bank case will apply their own, higher standards, to the subject: If the speaker is counting himself as not “knowing,” then, when he is speaking of the subject as a potential informant in the way we have been imagining, and this subject is no better positioned than is the speaker, the speaker will say that the subject, too, “does not know.” Thus, against very strong appearances, and against the methodology of ordinary language that we have looked at, the defender of SSI will have to deny that the denial of knowledge in our high standards case is true – or that the positive ascription of knowledge in the low-standards case is true.¹⁵

The best defense I can see for SSI here is a projectivist strategy like the one that Hawthorne pursues in defense of an invariantist view much like SSI (pp. 162-166). To explain this strategy, and then why it fails, consider this high-standards case in which a great deal is at stake: Thelma is talking with the police about whether it might have been John who committed some horrible crime. Our speakers have executed conversational maneuvers for driving up the epistemic standards, and Thelma is admitting that she does not “know” various propositions if she is not in an extremely strong epistemic position with respect to them. Thus, even though Thelma has grounds that would usually suffice for claiming to “know” that John was at the office yesterday (she heard from a reliable source that he was in, though she doesn’t know what exactly that report was based on, and she herself saw John’s hat in the hall), she is in her present context

doesn’t know what standards are appropriate to the subject’s context. (In fact, since the standards the speaker applies are extraordinarily high, he should be highly dubious that these are the standards that govern whatever context the subject is in.) But that’s not how speakers in fact behave in our speaker’s situation. If the speaker knows that the subject doesn’t meet the unusually high standards that govern the speaker’s context, the speaker will say that the subject, like the speaker himself, “does not know.”

¹⁵ If the backer of SSI were to accept that the denial of knowledge in our high-standards third-person case is true, while still accepting the truth of the positive ascription of knowledge in the low-standards case, he will have to maintain that, though our speakers are far away from her, and their conversation will never affect her in any other way, the fact that these speakers are thinking of and speaking of the subject as a potential informant about a high-stakes matter in the high-standards case is part of that subject’s “practical context” and thus raises the standards for knowledge that govern whether that subject, or anybody else, whether or not they are talking to our speakers, can truthfully say that the subject “knows” that the bank is open on Saturdays. That seems pretty implausible to me. However, in light of relatively recent developments in the philosophy of mind and language, we should be ready to be surprised by factors quite external to a speaker affecting the content of her thought and speech. Thus, we should seek further evidence against this possible move on behalf of SSI. That evidence is to be found in how the speaker herself, or others talking about her, will describe her. Suppose the subject herself thinks it is quite likely that our far-away speakers are now applying very high standards to her in connection with some very serious matter. If that were enough to greatly inflate the standards that govern whether anybody could truthfully describe our subject as “knowing,” our speaker should then refuse to say that she “knows.” After all, from her own point of view, such a claim is very likely false. But clearly there are many low-standards contexts where the subject’s conversational purposes will call on her to simply apply to herself the low standards appropriate to her own very local context, ignoring what is or might be transpiring in far-away conversations.

admitting that she does “not know” that he was in – though, she adds, she has good reason to think he was in. In this context, she would claim to “know” he was in only if, say, she herself has a clear recollection of having herself seen him in. SSI and contextualism will agree that first-person knowledge claims in such a context are governed by extraordinarily high epistemic standards, and thus that such denials of knowledge are true.

As happens with our bank cases, SSI falls into trouble here when we consider how speakers in such a “high standards context” will use third-person attributions (or denials) of knowledge to describe whether other subjects not present at the conversation “know” various propositions. In many such contexts, speakers will apply to such far-away subjects the same high standards that they are applying to themselves. Thus, for instance, if the police ask Thelma whether Louise, who is not present, but whom they are considering questioning, might know whether John was in, and Thelma knows Louise is in the same epistemic position that she’s in with respect to the matter (she too heard the report and saw the hat, but did not herself see John), Thelma will say that Louise too does “not know” that John was in – though she might again add that Louise too has good grounds to think John was in – even if Thelma is aware that Louise is not herself in a high-standards context. And again, it looks like the SSI theorist’s best option is to deny, against appearances, that this denial of knowledge is true.

According to the projectivist strategy, speakers like Thelma rightly sense that they do not know (they do not meet the standards set by their own high-stakes context), and they then mistakenly project their own ignorance onto absent knowers like Louise (who does meet the standards set by her own context).¹⁶

The problem with this projectivist strategy is that in the relevant cases speakers like Thelma will deny that subjects like Louise “know” even when the speakers have no ignorance to project. To see this, suppose that Thelma does meet even the extraordinarily high standards that govern her context – she does clearly recall herself seeing, and even having a conversation with, John. So, Thelma does know, and has claimed to know, that John was in the office. Still, if the police are wondering whether far-away Louise might also know (so they can have two witnesses), Thelma (who knows and takes herself to know, and thus has and realizes that she has

¹⁶ Hawthorne briefly appeals (p. 164) to some work in psychology that may suggest some mechanisms that might explain our “tendency to overproject our own lack of knowledge to others” (p. 163), but clearly does not want to be dependent here on any particular mechanism. My argument here, though, should work against any projectivist account, whatever the psychological mechanism behind it might be.

no ignorance to project) will still deny that Louise knows if all Louise is going by is having heard the report and having seen John's hat.

7. Contextualism vs. SSI: A Big, Ugly Tie?

For reasons that have become apparent, I think that SSI's fumbling of third-person uses of "knows" provides a quite decisive objection against the view. Defenders of SSI of course see things differently. This is no doubt partly because they don't think the problem is as damning as I think it is, but also partly because they think SSI's rivals, including contextualism, face equally daunting challenges. It's in this spirit that Hawthorne writes: "This is not to say, of course, that there are no counterintuitive consequences to this version of sensitive invariantism. As far as I can see, every candidate story about our puzzle has counterintuitive results. This is no exception."¹⁷

I believe that none of the recent problems alleged by Hawthorne and by Stanley to afflict contextualism are even close to being as powerful as the objection SSI faces in virtue of its mis-treatment of third-person attributions and denials of knowledge, though this is based in several cases on answers to those objections I believe are available. Here is not the place to go through the objections one by one. But there is a potential objection to contextualism that can appear to be extremely strong, is closely related to the objection to SSI that we've just encountered, and that can be answered by an easy, but important, clarification of contextualism. Here is the place to go into that.

In each of the high-standards third-person cases I have discussed here in my attack against SSI, the speaker applies extraordinarily high standards, appropriate to the speaker's own high-stakes context, to an absent subject who is not herself in such a high-stakes context (nor does the speaker take her to be in such a context). And we speakers often will apply standards appropriate to our own contexts to far-away subjects in quite different contexts. Among the situations in which we'll do this are contexts in which we are discussing these absent subjects as potential informants to us – as our speakers are doing in both of the cases I have used.

¹⁷ Hawthorne, p. 162. I'm counting Hawthorne as a defender of SSI, and I think he can be accurately counted as such. But the reader should be aware that Hawthorne does not have a strongly settled view on what is the right theory here; on this, see especially Hawthorne, pp. 187-188.

But there are other conversational situations – situations that Hawthorne rightly pays a lot of attention to – where different conversational purposes are in play, and in which speakers will apply to absent subjects standards that are appropriate to the practical contexts of those subjects. This often happens when, for example, the speakers are discussing practical decisions those distant subjects face: “She should do that only if she knows that...” SSI handles such cases nicely. Indeed, these cases may seem to strongly favor SSI over contextualism, since according to SSI, the subject’s context sets the standards, while it seems that, as Hawthorne writes, according to contextualism, “It is always the ascriber’s standards that call the semantic shots, so to speak” (p. 59). In fact, this can seem to be an advantage for SSI that balances the advantage I’ve claimed for contextualism: Sometimes it seems that the speaker’s context sets the standards; sometimes it seems the subject’s context sets the standards. SSI doesn’t work well on the first kind of case; contextualism on the second. It’s a wash.

If that were the situation we faced, I would consider that to be a tie that renders both sides losers, since I take failures of this magnitude very seriously.

Fortunately, that is not the situation we face, for contextualism can in fact easily handle the second type of case. What is vital to see here is that there is nothing in contextualism to prevent a speaker’s context from selecting epistemic standards appropriate to the subject’s practical situation, even when the subject being discussed is no party to the speaker’s conversation – which is good, because speakers often do select such standards when their conversational purposes call for it. On contextualism, the speaker’s context does always call the shots. Hawthorne is right about that. But sometimes speakers’ own conversational purposes call for employing standards that are appropriate to the practical situation of the far-away subjects they are discussing, and so the shot that the speakers’ context calls can be, and often quite naturally will be, to invoke the standards appropriate to the practical situation faced by the subject being discussed. And one type of conversational situation in which it will be quite natural for speakers to (and perhaps even strange for them not to) employ standards appropriate to subject’s situation are contexts in which the speakers are discussing whether or not the subject “knows” in connection with evaluating a practical decision (including, it’s worth adding, decisions concerning whether to assert something) that the subject faces (or faced, will face, could have faced, etc.): “She should do A only if she knows that P”, “Well, if he knew that P, he could have responsibly done A”, “She should assert that P only if she knows that P.” In such a

context, where whether the subject “knows” is being discussed for such a purpose, it seems strange indeed to employ epistemic standards not appropriate to the practical situation faced by the subject.

But, as is equally evident, and as we have already seen, there are other, very different contexts in which we have quite different reasons for discussing whether a subject “knows,” and in which it is perfectly appropriate to, and in which speakers actually do, apply to subjects standards quite different from those that the subject’s own practical situation calls for (and not because the speakers are mistaken about the nature of the subject’s context). And, disastrously, SSI cannot handle this evident fact, since SSI has the subject’s situation set the standards that govern any speaker’s description of that subject, whatever the speaker’s context.

I should note in passing, without working it out here, that the realization that a speaker’s context may “call the shots” by selecting standards appropriate to the practical situation of the subject being talked about, and that the speaker’s context actually will call such a subject-appropriate shot where what’s being discussed is how the subject should act, is also the key (or at least one of the keys) to answering one of the main criticisms Hawthorne raises against contextualism (at pp. 85-91) – that contextualism implausibly breaks some important ties between knowledge and assertion and practical reasoning. In fact, I think it can be shown that contextualism can accommodate as tight a tie here as the data support, while SSI is hurt by having to make the tie too tight in order to respect it at all.

8. The Importance of Arguments from Ordinary Language

Arguments from ordinary usage – appeals to what speakers of a language do or would say in applying the terms in question to particular situations (both positive and negative claims involving the term), appeals to which simple applications are or would be proper or improper for them to make, and appeals to intuitions as to the truth values of those claims in particular situations – don’t exhaust the considerations that need to be attended to in deciding among semantic theories, but do provide absolutely essential considerations.

This importance is most forcefully impressed upon me when I consider some obviously false theories, and ask myself how (on what basis) we know that they are false, and the closely

related, but perhaps even more important, question of what makes these theories false (in what does their falsehood consist)? So, to take one example of a theory that is obviously too demanding, and one that is obviously not demanding enough, consider, for instance, the crazed theory that I, following others, have discussed, according to which a necessary condition for the truth of “S is a physician” is that S be able to cure any conceivable illness instantaneously, and the theory about the meaning of bachelor like the traditional account, except that it omits the condition that S must be unmarried for “S is a bachelor” to be true, insisting that married men can be truthfully (if perhaps for some reason improperly) called “bachelors.” There are a couple of real clunkers! But in virtue of what is our language in fact such that these strange theories are not true of it? I’m of course not in a position to give a complete answer to this question, but, with respect to the strange theory concerning ‘physician,’ it seems eminently reasonable to suppose that such facts as these, regarding our use, in thought and speech, of the term ‘physician’ are centrally involved: that we take to be physicians many licensed practitioners of medicine who don’t satisfy the demanding requirement alleged; that we seriously describe these people as being ‘physicians’; that we don’t deny that these people are ‘physicians’; that claims to the effect that these people are ‘physicians’ intuitively strike us as true; etc. It’s no doubt largely in virtue of such facts as these that the traditional view, rather than the bizarre conjecture we are considering, is true of our language: The correctness of the traditional view largely consists in such facts. And these facts also seem to provide us with our best reasons or evidence for accepting the traditional, rather than the strange, hypothesis regarding the semantics of ‘physician’. Analogous points hold for why a traditional theory of ‘bachelor’, rather than the bizarre theory we are considering, is true, and for how we know it’s true.

Moving to context-sensitivity, such facts about ordinary usage also provide us with our primary, most important, and best evidence that clearly context-sensitive terms like “tall” are context-sensitive in the way that we take them to be, and are also that in which the context-sensitivity of those terms consist. Ignoring many subtleties, “tall” is context-sensitive in the way we take it to be largely because speakers do (in fact, with propriety, and with apparent truth) seriously say that things are “tall” in “low-standards” contexts when those items meet certain moderate standards for height, even if they don’t meet still higher standards, but will, in “high-standards” contexts, seriously deny that those same things are “tall”, reserving the application of that adjective only for items that meet the more demanding standard for height. In such facts

about ordinary usage we have some of the best possible type of evidence you could ask for that “tall” is context-sensitive.

But likewise then for “knows”. Data of the type we’ve been considering here provides us with the best possible type of evidence you could ask for that “knows” is context-sensitive in at least roughly the way contextualists claim it is. “Knows” is context-sensitive in that way largely because speakers in some contexts do (in fact, with propriety, and with apparent truth) seriously describe subjects as “knowing” propositions when those subjects meet certain moderate epistemic standards with respect to the propositions in question, even if they don’t meet still higher epistemic standards, but will, in other contexts, go so far as to (in fact, with propriety, and with apparent truth) seriously deny that such subjects “know” such things, reserving the ascription of “knowledge” only for subjects that meet some more demanding epistemic standard.

Note that this best kind of case for contextualism is not an indirect argument that takes as its premise that some other term, like “tall,” is context-sensitive, and then argues that because “knows” is so similar to “tall,” “knows” too is context-sensitive.¹⁸ Such an indirect argument would be very insecure, in my opinion, because, while there are very important similarities between the behaviors of “tall” and “knows,” there are also many important differences. Rather, this best type of case for contextualism is direct in that it adduces the same kinds of considerations for the context-sensitivity of “knows” as are available for the context-sensitivity of the likes of “tall.” Now, I will want to ask those who resist the argument on what basis they think we can know that “tall” is context-sensitive, and why they think we don’t have the same types of grounds for thinking “knows” is context-sensitive, but this is not to revert to the indirect argument. Rather, it is an effort to compare the direct argument for contextualism with the direct

¹⁸ In section 2 (“Is ‘Know’ Gradable?”, pp. 123-130) of his “On the Linguistic Basis for Contextualism,” Stanley appears to be attacking just such an indirect argument for contextualism by arguing that in various ways “knows” does not behave like a “gradable” term. I know that some read Stanley as in that section offering an argument for the falsity of contextualism -- an argument of the basic form: If contextualism is true, “knows” would be gradable; but it isn’t; so, contextualism is false. But that Stanley is instead (and more wisely) merely intending to undermine the indirect argument for contextualism comes out fairly clearly when he speculates that “the fact that an expression is gradable is independent of the question of whether it is context-sensitive” (note 7). This independence would undermine Stanley’s case if he were offering the above-considered argument for the falsity of contextualism. It would then be quite odd for him to admit that independence without addressing the glaring problem that it undermines the case he is giving. If, on the other hand, he is merely trying to undermine the indirect argument for contextualism, his procedure here makes perfect sense. He is in that case construing the contextualist as appealing to the various similarities between “knows” and gradable terms. Stanley is then countering that, at least with respect to matters of “gradability,” the similarities aren’t really there. He would then be free to speculate that even if those similarities did exist, that issue is in the final analysis independent of the question of context-sensitivity, anyway.

grounds available for the context-sensitivity of “tall,” and in so doing to give a higher-level argument that the direct argument for contextualism is solid.

As for views like SSI, relative to the types of considerations we have here considered, it is easy to imagine what our language would have been like so that SSI would have been true of it. In short, and crudely skipping lots of details, we would always apply to subjects standards appropriate to those subjects’ contexts (and would be agnostic about whether those subjects “know” where we were problematically unsure about what their practical and conversational, as well as epistemic, situation). That we do not talk about subjects in that way is, in large part, what makes SSI false – it is that in which the falsehood of SSI consists – and also provides us with our best – and quite decisive, in my view – evidence that SSI is false.

9. Methodology, Straightforward Data, Comparative Judgments of Content, Metalinguistic Claims, Belief Reports, and Speech Reports

A certain methodology strongly favors contextualism.¹⁹ This “methodology of the straightforward,” as we may call it, takes very seriously the simple positive and negative claims speakers make utilizing the piece of language being studied, and puts a very high priority on making those natural and appropriate straightforward uses come out true, at least when that use is not based on some false belief the speaker has about some underlying matter of fact. Relatively little emphasis is then put on somewhat more complex matters, like what metalinguistic claims speakers will make and how we tend to judge how the content of one claim compares with another (e.g., whether one claim contradicts another). This methodology favors a contextualist conclusion, because, as I have been urging here, the data concerning what simple positive and negative claims we make involving “knows” do strongly support contextualism, and considerations coming from the other, fancier sources, do not seem as kind to contextualism. One might therefore be tempted to read this paper as an instance of someone following a certain contextualist-friendly methodology to its inevitable contextualist conclusion.

¹⁹ My thinking here was very much helped by Timothy Williamson’s presentation at the Stirling contextualism conference (March, 2004) – though Williamson’s views on the matter are very different from my own. I was also helped here by a post Brian Weatherson had on his weblog in which he complained that contextualists have focused too much on what I’m here calling the straightforward considerations.

Now, I must admit, I do find the methodology of the straightforward somewhat attractive, at least when its favoring of the simple data is not taken too far. This methodology can be based on the thought that what speakers know best about the piece of language at issue is how and when to make simple positive and negative applications of the term in question, and it does seem sensible to me to suppose that that is what speakers would be best at, at least to the extent that it should count as a somewhat worse strike against a theory that it runs afoul of the straightforward, as opposed to the complex, data. And I believe that this sensible thought is often verified by experience – experience of the fancier data misbehaving in various ways. Thus, if push were to come to shove, I suppose I would be inclined to follow the data concerning the straightforward.

But I hold that we can forget about any such general biases here, however sensible they might or might not be, because I think the fancier considerations at least misbehave (and perhaps even work a bit against the invariantist) in the case at hand. Thus, even if we don't relegate such considerations to some secondary, suspect status, they won't speak clearly against contextualism, anyway. So push does not come to shove here.

Where the contextualist has constructed HIGH and LOW (his high-standards and his low-standards cases) properly, I think it's just about as clear as anything in the area gets that one should affirm that the subject "knows" in LOW and deny that the subject "knows" in HIGH. Those are facts that just about any competent speaker can sense, and indeed, most invariantists will admit that that is how we talk, and appropriately so. Invariantists have typically just questioned whether the affirmation in LOW and the denial in HIGH are both true. But the fairly clear intuitions about the cases, where they are considered individually, certainly are that the claims are true, and, as I have already explained, those intuitions about the truth-values of the claims are buttressed here by the very evident facts about proper assertion that most invariantists will themselves admit. That's the straightforward data.

Now, turn to the fancier stuff. First, judgments of comparative content. When we consider the affirmation in LOW together with the denial in HIGH, is it clear that these two claims contradict one another? So far from that being obvious to me, I've always been fairly strongly inclined to think they do not contradict one another. But that's just me, and it is perhaps no accident that I became a contextualist. But when I have presented such cases to large enough groups of students, and then asked about whether the claims contradict one another, I have always found both positive and negative answers strongly represented. Indeed, it has been my

experience that the majority usually answers the question negatively, as I'm inclined to, but that may be due to the way I present the issue – though I try hard to remain neutral until I get the initial answers from the students. At any rate, the answer that the invariantist wants here is very far indeed from being clearly correct. (It is vital to remember here that I am only talking about instances where the contextualist's cases are properly constructed, in line with the guidelines in the first few sections of this paper. Where those guidelines are violated – where, for instance, HIGH and LOW are put in the same conversation, debating each other over whether something is “known” – the intuition that they are contradicting one another can get quite strong.)

What about metalinguistic claims? If we tell each of our two speakers about the conversational situation faced by the other, can/should LOW say, “HIGH's claim is false,” and can/should HIGH assert “LOW's claim is probably false”? (The “probably” is inserted in light of the possibility that HIGH may be facing a situation where the standards for outright assertion are extremely high.) It has always seemed far from clear to me that such metalinguistic statements would be true. In fact, again, I've always been inclined to think they would be false, but I admit this is not intuitively an easy call. But beside the question of whether such metalinguistic claims would be true, even questions about whether they are appropriate are quite vexed. Are those the right things for our subjects to say? My intuitive reaction has always been that, for instance, whether or not LOW can appropriately say, “HIGH does not know,” it really seems a bit too much for her to say “HIGH's claim is false.” Thus, it always intuitively seemed to me that these metalinguistic claims go too far.²⁰ And you will find when you ask enough speakers that the matter seems to a great many of them as it seems to me (at least, again, when you have constructed the cases properly). But at any rate, the answers that favor invariantism here are certainly very far indeed from being clearly correct. Compare this with the straightforward data, where it's about as clear as such things get that the denial in HIGH and the affirmation in LOW are the things to say, and we have very strong intuitive grounds (based on the mutually reinforcing strands of evidence I explained) for thinking such claims are true.

Does LOW's claim contradict HIGH's? Should they each say that the other's claim is (probably) false? Would such a metalinguistic claim be true? As I've noted, I've always been inclined to give the contextualist's answers to these questions. But I have always taken that to be

²⁰ See pp. 925-926 of my “Contextualism and Knowledge Attributions,” Philosophy and Phenomenological Research 52 (1992): 913-929.

a rightly disputed matter. Those are the precisely the questions we need to do some philosophy to answer. I have a hard time relating to an inquiry that takes certain answers to these questions as its starting point, or that makes delivering certain answers to these questions a desideratum of a theory. And to the extent that it is some slight advantage to a theory that it deliver intuitively correct answers to these questions, given my intuitive reactions to the questions, it's the theory that delivers the contextualist answers that gains that slight advantage, at least in my book – though I suppose that would have to the case also for the great number of speakers who have intuitive reactions similar to mine.

Similar points hold for many others bits of fancier data. To take just one more example, consider how “knows” behaves inside of belief reports, and, in particular, whether we use the term in a way that follows a “disquotational schema” like this:

(DSK) If a speaker S1 sincerely utters a sentence of the form ‘A knows that p’, and the sentence in the that-clause means that p, and ‘A’ is a name or an indexical that refers to a, then a speaker S2, also using ‘A’ to refer to a and using the that-clause to mean that p, can truthfully state that S1 “believes that A knows that p.”

If S1’s utterance of “A knows that p” is sincere, then, even if you are in a very different context, you can truthfully assert, “S1 believes what she asserts by the sentence, ‘A knows that p.’” But that is so clunky! DSK tells you that you can (also) come down a level and truthfully state the far simpler, “S1 believes that A knows that p.” Hawthorne thinks our ordinary usage betrays a clear commitment to a principle like DSK;²¹ he writes:

²¹ Hawthorne’s version of DSK is on p. 101. I based my formulation on Hawthorne’s, but modified it by making its consequent meta-linguistic, as seems appropriate to the issues at hand.

Hawthorne goes on to combine his version of DSK with another principle to produce an argument against contextualism (pp. 101-104). But how to understand this argument is unclear to me. At p. 104, he considers this important objection: “But couldn’t an analogue of this argument be raised for, say, all context-dependent comparative adjectives? And wouldn’t this show that the argument proves too much?” He immediately replies, “There are disanalogies here that should not be ignored.” The disanalogy he presses is that “we have very few devices” in common talk by which we can clarify our use of “knows” (p. 105), while we have many such devices by which to clarify our use of comparative adjectives. But Hawthorne does not explain how to modify the principles he uses in the argument at 101-104 so that the a version of his argument generalized so as to be applicable to both the case of “knows” and that of comparative adjectives will apply where there is a shortage of clarifying devices but not where there is an abundance of them. Surely we do not want a disquotation schema that is just like DSK except for its being generalized, and its ending with the proviso, “provided that there are many clarifying devices in ordinary talk by which the key terms of S1’s assertion can be easily modified.” And, at any rate, I think Hawthorne is wrong about the lack of clarifying devices for “knows,” as I believe is shown in Peter Ludlow’s as yet unpublished paper,

Of course, not all words behave disquotationally in this way. If a speaker says ‘It is raining here’, I cannot report the contents of his mind by an ascription of the form ‘He believes that it is raining here’, unless he and I happen to be in the same place. By contrast, it looks very much as if we do adopt something like DSK. If, for example, someone sincerely utters ‘I know that I will never have a heart attack’, we have no hesitation whatsoever in reporting the contents of his mind by claiming that he believes that he knows that he will never have a heart attack. That is how the verb ‘know’ seems to work.

But rather than being clearly right, that just seems wrong, or at the very least, extremely unclear, at least as applied to the relevant test cases, where the contextualist will think there is a marked difference in standards.²²

So, let’s go back to the story of Thelma and Louise from the end of section 6, above. Thelma knows that Louise has just sincerely said, in her “low-standards” context at the tavern, “I know that John was at the office yesterday.” Indeed, we may suppose the Thelma was also at the tavern, and also claimed to “know” that John was in based just on the evidence of the report and having seen his hat. But now Thelma has left the tavern, leaving Louise behind, has encountered the police, and is now saying, as is appropriate to the heightened concerns of her new conversation, and as almost any speaker would, that she herself “does not know” that John was at the office, and that Louise, whom Thelma knows also has only the same meager evidence to go by on the matter, “does not know” that either (though, again, she might do well to add that they

“Contextualism and the New Linguistic Turn in Epistemology.” I agree with Hawthorne that his argument appears to be as strong against comparative adjectives as it is against “knows,” so he is right to worry about proving too much. But, as I have argued in the text here, Hawthorne’s defense of DSK seems to be based on a mistake. Thus, I think he should worry even more about his argument not working in the first place.

²² Admittedly, we are quite liberal in engaging in this type of disquotation where the difference in standards is not severe. This might be thought to be a problem for contextualism: Why would we be so quick to disquote in cases where, if contextualism is true, we should be dubious that the content of the knowledge claim we are disquoting matches exactly the content it bears in our different context? However, one will find the same liberalness in engaging in disquotation with our use of other, clearly context-sensitive terms, like “tall” – or even “here.” The matter of exactly what region of space is designated by a given use of “here” – the matter of exactly how far out “here” reaches – is often a very subtle matter. Yet, except where there is a marked difference in the regions being designated, we disquote (in the sense under discussion) speakers’ uses of “here” very quickly, even in cases where it seems we should be dubious that the area the speaker was designating is exactly the region that “here” picks out in our own context.

both have good grounds for believing that he was in). Now, in this elevated context, will Thelma display “no hesitation whatsoever” in reporting, “Louise believes that she knows”?

Intuitively, it actually seems that it would be wrong for Thelma to say that in her highly-charged context; it is certainly not so clearly the right thing to say as to occasion “no hesitation whatsoever” in saying it. I realize that there are a number of possible explanations for why that is not a good thing, or is at the very least is a problematic thing, for Thelma to say. I’m not trying here to adduce its wrongness as a positive argument for contextualism. I am just pointing out that arguments that attempt to oppose contextualism based on supposed observations of the type Hawthorne uses here are on very shaky ground.

What (in the vicinity) can Thelma say, then? Intuitively, the situation seems to me quite similar to what can be said in an analogous case involving a clearly context-sensitive term like “tall.” Suppose that while discussing the reasons why the 6’5” Reeves is such a great high school basketball player, Lucius has sincerely said, “Reeves is tall,” and suppose that Theo knows that Lucius sincerely said that in such a context. Now, however, Theo is in a context in which various players, including Reeves, are being discussed as potential NBA centers (Reeves is considering skipping college and jumping straight into the NBA), and, as is appropriate when potential NBA centers are being discussed, 6’5” players are definitely not being counted as “tall,” and so Theo has admitted, “Reeves is not tall.” Here, as in the Thelma/Louise/“knows” case, it does not intuitively seem right for Theo to report in his “high-standards” context, “Lucius believes that Reeves is tall.” It seems significantly better (if only because it is less bad), though, for Theo to report, “Lucius said that Reeves is tall” and for Thelma to report, “Louise said that she knows.” But though these “said that” reports seem less bad than the corresponding “believes that” reports we considered earlier, they still seem somewhat problematic; at least a little hesitation seems called for. (Whether such “said that” reports would be true is a very difficult question; they may just be in various ways misleading. I’m here merely reporting intuitions about whether these are the right things for our speakers to say.) What does seem clearly correct is the above “said that” reports with the addition of any of a number of possible follow-ups that at least seem to work precisely by warning the listener against interpreting the content of what Lucius or Louise said as if it had been said in the listener’s new, quite different context (i.e., which seem to work by warning the listener against disquoting in the way in question here): “Lucius said that Reeves is tall, but he was speaking of Reeves as a high school player,” “Louise

said that she knows, but she was speaking casually.” Something like that seems to be the thing to say.

But, again, what is important to current purposes is that the worst thing we have here considered, the simple (i.e., said with no follow-up) “believes that” report, is, at the very least, far from being clearly correct. It certainly is not simply and clearly the right thing to say in the relevant test situations (situations in which contextualists will hold that there has been a marked change in standards). And whether it would be true if it were somewhat improperly said is a very vexed question. So it’s not clear that there even is some shove to be found here in favor of invariantism. If there is, it seems incapable of matching the pro-contextualist push we’ve been investigating in this paper.^{23,24}

²³ If you think that I am wrong, and that there is a strong anti-contextualist shove here to match the stiff contextualist push we have looked at in this paper, one option you might consider is following John MacFarlane in resolving this tension by adopting the radically relativist semantics for knowledge attributions that MacFarlane presents in his “A Relativist Semantics for ‘S knows that p’”; talk at the “Themes in Philosophy of Language” conference at Yale University, November 8, 2003.

²⁴ For very helpful discussion of these matters, I thank Richard Heck, Jason Stanley, Brian Weatherson, Timothy Williamson, the participants in my Spring 2004 Epistemology seminar at Yale, and the participants in and, especially, the organizers of the conference on epistemological contextualism at the University of Stirling in March of 2004.