

Reply to Nagel

~~5/23, 18bot+end~~

The key test cases for deciding between my brand of contextualism and Jennifer Nagel's brand of invariantism are the third-person examples. As matters currently stand, first-person cases, like my original Bank cases (pp. 1-2), are pretty useless here. Nagel can agree that the speaker's claim to "know" in Case A and his admission that he doesn't "know" in Case B are both true; she just accepts a different account of why it is that both assertions can be, and are, true, according to which it is because in B the speaker doesn't meet the attitude requirement for knowledge, while he does meet that requirement in A. Perhaps not at all unexpectedly, I find my proposed explanation more plausible; presumably, Nagel sees things differently. We can try to hash this out, but extremely tricky questions, and no stable enough answers, about how to understand the attitude needed for knowledge derail this attempt to decide between accounts. But no worries: Without having to first decide these issues, third-person cases provide the natural test cases to decide between the views. We can use them to decide between theories, and then go back to the first-person cases and *apply* what we've learned from the third-person cases to guide our handling of the trickier first-person cases.

Well, that may be a bit too quick, so a couple of points of clarification. First, I should emphasize the part about first-person cases being useless *as matters currently stand*. Mainly by developing a secure approach to the tricky questions about attitude, someone could come to be in a position to make a good, direct (not going through considerations about third-person examples) case for preferring one approach to the first-person cases over the other. I take it Nagel is not presenting that kind of case here. The bulk of her paper (the part before her last three paragraphs, where we turn to the third-person cases) defends the viability of her account of the first-person cases, but she doesn't seem to be claiming to be showing she has some important advantage here. Nagel is presenting and defending the viability of an alternative approach. Second, I want to emphasize that I'm not saying the first-person cases, or

the tricky issues that arise in trying to handle them, are unimportant. In fact, I am extremely interested in these issues (most of which seem to be such that results of empirical investigations should and will loom large in addressing them) – and very much hope to be learning important things from Nagel, as well as others, about them. I am just saying that as things stand, the way to decide between the theories is to look to the natural test cases, which are third-person cases.

The problem with would-be arguments for contextualism that are like my argument (in Chapter 2), but are based on cases involving first-person knowledge claims rather than, as I have it, third-person knowledge attributions, is that they are open to the worry that the features of the situation of the subject/speaker that the contextualist claims raise the epistemic standards in his so-called high-standards case might instead be working by weakening the subject/speaker's confidence and/or the strength of her epistemic position. So, for example, in Bank Case B, the subject's wife reminds him about how much is at stake and also mentions that banks do change their hours from time to time. The contextualist may hold that the differences in Case B work to make the denial of knowledge in that case true, at least in part, by raising the epistemic standards that govern that scenario. But how can he be so sure that they don't instead operate by weakening the subject's epistemic position? (Might the fact that his wife is bothering to bring up this possibility be some reason for our subject to think that the bank in question might well have recently changed its hours?) Or, to move to a worry directed related to Nagel's approach, how can the contextualist be so sure that the dire situation, plus the subject's wife's dwelling a bit on the direness of the situation, hasn't lowered the subject's confidence to the point that, for purely "attitudinal" reasons, he can no longer be said to know that the bank will be open?

So the wiser contextualist does not offer such an argument, but sidesteps these worries by moving to third-person cases, like my third-person bank cases, or my Office cases that Nagel discusses toward the end of her comments. Here, too, fairly remote possibilities of error are brought up and worried about by our speakers, but

there is no serious worry that these conversational developments are hurting either the subject's strength of epistemic position or, more relevant to Nagel's account, her level of confidence, because she is not on the scene as these developments unfold: as our speakers realize, our subject is far away and quite oblivious to what is transpiring in their conversation.

So such third-person examples are the natural test cases for deciding between the views in question. And it seems that it is contextualism that wins this test. Nagel seems to agree that in my third-person Office case, speakers at least seem to correctly and truthfully say what the cases have them say (Thelma saying in the tavern that Lena "knows" and Louise denying that Lena "knows" while talking with the police). Her strategy is to explain this away. Nagel suggests that her view might actually have an advantage here, since, she suggests, it might provide a "simpler explanation of why Louise's claim seems true." Here Nagel is suggesting (given what immediately precedes it) that her proposed explanation may be simpler than the contextualist explanation, on which Louise's claim seems true because it is true (and, as competent speakers, we're able to recognize that Louise is choosing the correct thing to say here). I'm not really understanding what inclines Nagel toward that comparative judgment. But however complicated the contextualist account of the appearance of truth turns out to be (and my presentation above seems far from complete), Louise is speaking as we would expect her to on contextualism, and the intuition that she is speaking truthfully is one a good form of contextualism validates.

By contrast, Nagel's proposed form of invariantism would at least at first look lead us to expect Louise *not* to say that Lena doesn't "know."¹ After all, Thelma says

¹ Wouldn't Nagel's invariantism lead us further, to expect that Louise will, like Thelma, go so far as to positively assert that Lena does "know"? No, it can account for why Louise would not go that far, at least given the knowledge rule for assertion. For on this invariantist account, Louise herself does not know that John was in, and so does not know that Lena knows that John was in. Thus, by KAA, she is in no position to assert that Lena knows. But on this invariantist account, Louise should think that Lena *probably* does "know," and so she certainly shouldn't say that she doesn't know.

that Lena “knows,” and Lena appears the same from Louise’s perspective as from Thelma’s with respect to all the matters that on Nagel’s view are relevant to whether Lena knows.² That Louise is now being interviewed by the police concerning some very high-stakes matter, in particular, clearly is doing nothing to trouble Lena’s *confidence* that John was in, since Lena is completely oblivious to these developments. In these types of cases – the very cases that provide the natural tests for deciding between contextualism and Nagel’s brand of invariantism – speakers don’t behave as Nagel’s theory would at first look predict.

Nagel proposes to handle that problem by means of the suggestion that “when the police ask her about Lena’s epistemic position, Louise gives an evaluation of the epistemic position that she naturally expects the police would find in Lena were they to interview Lena.” Nagel proposes two forms such an explanation can take, on one of which Louise’s own thinking about whether Lena “knows” is distorted by these thoughts of whether Lena would know if she were interrogated, and on the other of which Louise is just telling the police what they are really interested in. (What is directly important to the police is whether Lena would know if asked by them so that they can use her as a witness or informant; it’s pretty useless to them that she knows while home alone, except insofar as that has implications for what she will know when they’re questioning her.)

I believe the considerations Nagel raises can account for why on her theory Louise should not simply (without further explanation) tell the police that Lena does

² Other invariantist views on which the epistemic standards providing the truth-condition for speaker’s claim are set by the practical situation of the subject can claim here that the fact that the police wish to interview Lena about such an important matter is in fact an important feature of Lena’s practical situation, whether she, or Thelma, realizes it or not. They can then with some plausibility argue that Thelma is speaking falsely in saying that Lena knows, and explain why Thelma would make a false claim by means of a claim that she is ignorant of a fact that is crucially important to the issue of whether Lena knows. To assess such a suggestion, we need to modify our cases a bit; see 2009: 231-2.

“know.”³ Consider this analogous case. George is going to try to find Frank, whom he only knows by description, in a big, crowded ballroom, and is talking to Lou on his cell phone as he approaches the room. In an attempt to get one last bit of information to help him hopefully identify Frank, right before he terminates the call and enters the ballroom, George asks Lou whether Frank is sitting down. Suppose that, as Lou realizes, Frank is now standing up, but will be sitting down in a few moments, before George will be in any position to spot him. Here I would agree that it would seem misleading and somehow wrong for Lou *simply* to tell George, “No, he’s standing up.” (Though it’s worth noting that it would still seem that he would then be speaking a truth, even if a misleading one.) Because the truth that Frank is standing up is not in line with what is directly important to George (whether Frank will be sitting when he sees him), Lou shouldn’t simply tell George that truth without further explanation. For somewhat analogous reasons, in our original case involving what Lena knows, it’s very plausible to suppose that if Nagel’s theory were correct and it were simply true that Lena “knows,” it would still be wrong for Louise to tell this truth to the police, at least without some further explanation.

But note that in our ballroom case, it is wrong for Lou to tell George that “Frank is standing up” or that “Frank is not sitting down,” and Lou’s judgment as to whether it’s now true that Frank is sitting down is not likely to be distorted by thoughts of how he might most helpfully communicate with George. In these situations, it seems to me, we as speakers do not have license to assert falsehoods, despite their practical usefulness. And if we do, our assertions of these falsehoods do not give the appearance of truth. So, the truth (“He’s standing”) is misleading; the falsehood (“He’s sitting”) is a more useful thing for the listener to believe, given his present purposes, but, of course, has the annoying drawback of being false. What’s a poor speaker to do? We run into these situations all the time, and handle them with

³ Remember that we already have another account for why on Nagel’s theory Louise should not tell the police that Lena “knows,” and this other explanation implies that she shouldn’t tell them this even with some further explanation, but should instead say something like that Lena “probably knows”: See note 1, above.

ease, it seems to me. One can opt for the truth together with further explanation: “He’s standing, but he’ll be sitting in a moment.” Alternatively, we often answer with an assertion that does not directly respond to the question we were asked. Lou is asked whether Frank *is* sitting, but he may respond with an assertion that speaks to Frank’s future, rather than present state: “He’ll be sitting in just a moment”, “He’ll be sitting by the time you get to him.” There’s probably much that can be said about why we would choose slightly more complicated useful truths over simpler useful falsehoods. It probably has something to do with the fact that our listeners’ purposes may easily go beyond what is apparent to us, or even if their current purposes *are* limited to what we take them to be, they can always develop other purposes later on, while still relying on the “information” we have given them. That we should stick to the truth in such situations, even when that means getting a little wordier than false-speaking would allow us to be, seems a very sensible policy. But for whatever reason, that does seem to be the policy in cases seriously analogous to my third-person cases. For, as I said, it does seem wrong in these circumstances to assert the falsehoods. And, in a connected matter, in these analogous situations, we are not typically misled into thinking the helpful falsehood is true.

Thus, while the considerations Nagel raises can account for why on Nagel’s theory Louise cannot simply tell the police the truth (given Nagel’s theory) that Lena “knows,” they don’t seem very promising for motivating what Nagel needs here: that Louise will go so far as to tell the police the falsehood (given Nagel’s theory) that Lena doesn’t “know.” It seems what we should expect on Nagel’s theory is that Louise will either tell the police the truth together with further explanation: “She (probably) knows, but she won’t know when you ask her” (see note 1 for why the “probably” is needed); or perhaps just the truth that doesn’t directly respond to their question: “She won’t know when you ask her.” These actually don’t sound to me very good. But the important point here is that we don’t have good reason to suppose that on Nagel’s theory speakers will assert the relevant falsehoods here. Or at least we don’t *yet* have such a reason. Perhaps by finding analogous situations where we do happily assert falsehoods, or by finding other further considerations to make this kind

of explanation plausible, Nagel or someone else can make good on this attempt to explain away these results of our crucial test. What I'm suggesting is that the proper attitude for us to take to this attempt is the one I propose at (DeRose 2009: 235-6) to a different invariantist attempt to explain away seemingly pro-contextualist data: As of now, it doesn't succeed in defeating the force of the pro-contextualist considerations in play.

Reply to Ludlow

Peter Ludlow discusses a couple of cases that concern just how uses of “know(s)” are connected to practical reasoning. While these are very important types of cases for theories of knowledge attributions to address, I should begin by clarifying the nature of my arguments against SSI. Ludlow focuses his discussion on my arguments against SSI in Chapter 7, and seems to present the two cases—an “ignorant high stakes” case and my “multi-tasking” example—as playing a major role in my attack, when in fact they are not of a type that form the basis of my anti-SSI argument. I do write of the latter case that it “may be something of an embarrassment” for some versions of SSI, but that's because it's tough case and seems a bit of an embarrassment for *anybody*, so I acknowledge that it “does not in the final analysis provide a very strong objection” to the versions of SSI in question (2009: 272). Its purpose is to set up one of my real arguments against SSI – the one that immediately follows in the next section, and is based on a quite different example.⁴ Ludlow is right that third-person cases play a key

⁴ The argument in question is in section 12 of Chapter 7, and is based on the example of Judith, an agent who is sleeping and not really engaged in any practical matter in particular right now. The case it utilizes can in an extended sense be called a “multi-tasking” argument, as it is based on the observation that Judith, like other agents, can be evaluated with respect to several different practical situations, which for all the speakers know she might face in the coming days, even as nothing changes about the agent except the interest speakers happen to be taking in her. But Ludlow writes “DeRose thinks this makes a good case against SSI” right after discussing my “walking talker” example, which is the one I say *doesn't* make a strong case against SSI. See especially (2009: 275-6) for a comparison of the two cases, explaining how the Judith provides a strong argument against SSI, while the walking talker example does not.

role in my case against SSI, but he seems to get off the track in discussing “ignorant high stakes.” My methodology concerning which kinds of intuitions about cases should be given most credence would direct me away from leaning on intuitions about such cases,⁵ and when I do discuss cases similar to the ones Ludlow presents, it is to set them aside as not being the best kinds of test cases (p. 55). The argument by means of cases that I do rely on against SSI (at pp. 230-4) utilizes an example in which the speakers are under no illusions about what’s happening with the far-away subject they’re talking about.

Well, even if “ignorant high stakes” cases don’t play a role in my anti-SSI crusade, what do I say about them? In Ludlow’s case, Frank asks Angie whether she “knows” that some O-rings won’t become brittle in low temperatures. Though Frank is in fact asking this on his boss’s behalf for purposes having to do with a high-stakes, life-or-death matter to which very high epistemic standards are appropriate, Angie believes he’s asking on his own behalf and in connection with a low-stakes endeavor, and so responds that she does indeed “know.” Frank then reports to his boss that “Angie knows” that the O-rings won’t become brittle. I take Frank’s ascription of knowledge to Angie, while talking with his boss, to be false: Both he and his boss would be thinking in terms of very high standards, which Angie does not meet. The first-person knowledge claim Angie makes to Frank is more interesting. I’m inclined to judge it true, since Angie would be intending low standards, and then to hold that Frank misinterprets Angie’s claim. But there seems room for other contextualists to let the very high standards assumed by Frank, the hearer of the claim, to get into the act (by which I mean: to affect the truth-conditions of the speaker’s claim) in various

⁵ This methodology is discussed in sections 2-5 of Chapter 2. Most relevant here is that I trust intuitions to the effect that a claim a speaker has made is true, where the claim is also appropriate, and where the speaker is not mistaken or ignorant about some relevant matter where this mistake or ignorance might be able to explain why their claim may have been appropriately made even if it were false; see esp. pp. 50-1.

ways that could render Angie’s claim truth-value-less, or perhaps even false.⁶ My own call here—that what standards Frank is assuming don’t have an affect on the truth-conditions of Angie’s claim as she talks to him—is part of my general inclination for how to treat context-sensitive terms (including uncontroversially context-sensitive terms) in analogous situations where there are misunderstandings about the purposes that are behind various questions and to which the answers will be put. (And I should perhaps add that it’s a call I’m very open to revising.) I don’t take my verdicts here to be intuitively clearly correct (though I also don’t take myself to be in intuitive trouble here), and wouldn’t wield them as premises in an anti-SSI argument.

Ludlow makes some extremely interesting suggestions for how such cases can be handled. He may well be pointing to a fine (and the best) way for SSI to handle such cases. But I didn’t doubt that SSI could work out a plausible enough way of handling such cases. My worries about SSI stem from other cases, and my grounds for preferring a contextualist over an SSI approach to these “ignorant high stakes” cases is not that the contextualist has an important intuitive advantage here, but that the contextualist approach works better on some other cases, and we can then apply what we learn from those other cases to these more problematic cases.

Ludlow is right to notice how much critical attention I give to SSI, and reasonable in suspecting that it is the rival I am “*most* concerned” about. And it *may* indeed be the one I’m most concerned about—in the sense that I think it’s the alternative that it’s most important to engage with (though I also spend a lot of energy combating CI, and think important points in philosophy of language emerge from that battle). But I’d like to quickly clarify my attitude here, because SSI is *not* the rival

⁶ In Chapter 4, I consider several different ways to handle the truth-conditions of knowledge attributions in cases where two speakers are disagreeing with one another about whether or not someone “knows,” and seem to pushing the “conversational score” in different directions. These options (and others, too) will also be available to those who opt to let the epistemic standards the hearer is assuming affect the truth-conditions of the speaker’s claims in the quite different situations (involving no disagreement) like the example we’re considering here. They are ways of adjudicating different “personal” contents into common contents, in cases where a common content are appropriate.

I'm most concerned with in the sense of finding it the most threatening. As I note in the book, "I have long thought that how we use third-person knowledge attributions show that SSI cannot be right, and have therefore long thought of classical invariantism as the real threat to contextualism."⁷ True, I side with SSI against CI in the verdict I issue on many important cases. However, I think the verdicts about the third-person cases I use against SSI are about as solid as are the verdicts we both use against CI, and the awkwardness I note at (DeRose 2009: 234) involved in the situation of SSI – that of needing to reject the maneuvers of CI in defending itself against the intuitions about various first-person cases, while SSI itself needs to make some such maneuvers in defending itself from what seem to me to be intuitions just about as secure, but concerning third-person cases – seems to me to render SSI the less threatening rival. CI may be from my point of view more often wrong about cases than is SSI, but it seems to have a more coherent defensive strategy available to it.

Why then all the attention to SSI? For one thing, of course, it came to be a "hot view," with some very able advocates whose defense of it caused it to be widely discussed. This made it professionally prudent for me to deal extensively with the view: It is considered more of a mark against one's work that one has failed to adequately deal with an opposing view that is quite popular than it is that one has neglected a rival that nobody sees as very promising anyways. But beyond (and more important than) that kind of prudential calculation, is this: If philosophers as able as, and as knowledgeable about the subject-matter as, Fantl and McGrath, Stanley, and Hawthorne, find SSI very attractive (though Fantl and McGrath have moved to a form of subject-sensitive contextualism), despite being well aware of the basis of my own pessimism about the view, some wise intellectual humility required me to face the possibility that the view might just be more promising than I was (and remain) personally inclined to judge it to be.

⁷ (DeRose 2009: 108-09). Also relevant: (DeRose 2009: 230, n. 5).

But perhaps most important of all is that I share Ludlow's sense of where the exciting and important progress is likely to be made in this long-standing discussion of the semantics of knowledge attributions. Advocates of SSI have made explicit some very important issues concerning the connections between knowledge attributions (and denials of knowledge) and action that any theory of knowledge attributions, including mine, should address. (And the later sections of my Chapter 7 hopefully represents a good start toward a contextualist approach to these issues.) If in his closing paragraph Ludlow is right (and I'm bold enough to hope that he is) that the debates I'm engaged in with rivals will likely "lead to some very interesting and largely unexplored territory in epistemology," this will be in large part because these arguments will lead to advances on just those issues that have been most clearly spotted as important applications of epistemological theories by advocates of SSI.

Fantl and McGrath: Loose Use

One of Jeremy Fantl and Matthew McGrath's (F&M's) main motivations (perhaps *the* main motivation they have) for moving from what they describe as my "cases-as-data strategy" to their proposed "cases-as-conclusions" methodology is that my principal argument for contextualism is open to the objection that the positive claim to know in my LOW case is an example of loose speaking. They suggest I could complete my argument if I would "develop tests for loose speech" and then argue that they're not met in my LOW cases – though they think I would be better off changing strategies altogether. The work of developing this "loose speech" approach to the data contextualists try to use has been done for F&M, principally by Wayne Davis (2007), whom they can and do cite. By contrast, the critical points I believe need to be made about the "loose speech" approach are not out there (that I know of) in a form I can just cite. Since the viability of this loose speech approach is important to the methodological differences between F&M and me, I will take this section to briefly make some observations about that approach.

As Davis notes (Davis 2010: 1155), Peter Unger’s theory of the mid-to-early switch 1970s (“early Unger”) that I do argue against is an example of a loose usage invariantist theory—an invariantist account on which the intuitions about the truth-values of knowledge claims that go against the theory are explained away as examples of loose speech. If one takes this type of approach, one will naturally opt for a version of invariantism that sets the standards for knowledge very high, and early Unger certainly did that: His semantics for various terms (including “knows,” “flat,” and many others) was *maximally* demanding: for a surface to be such that it could truthfully be called “flat,” for example, it must be such that it is impossible for any surface to be flatter than it is; for someone to be such that one can truthfully describe them as “knowing” that p, they must be so well-positioned with respect to p that it is impossible for any subject to be better positioned with respect to any proposition; etc. This makes all or almost all of our positive uses of the terms in question come out false, which, interestingly, Davis does recognize as a problem. But he proposes a solution, which seems to be to ease up just a bit on the severity of the semantics. At the end of my third chapter (DeRose 2009: 117-27), I argue against what I took to be the two best attempts to defend an invariantist approach: both that of early Unger, and the much later and very different invariantist theory of Patrick Rysiew, whose semantics were so forgiving that it’s actually the denial of knowledge in my HIGH cases that he argues is false. Davis thinks I make it too easy on myself by picking two easy invariantist targets, one of which employs truth-conditions that are too demanding and the other of which is too lax. Davis feels his own approach to “knows” and other terms, which employs “strict but satisfiable truth conditions” (Davis 2010: 1156), provides the “Goldilocks” (not his term) form of invariantism that’s just right (Davis 2010: 1155-6).

A problem with such an approach is that it’s hard to see just how very demanding to make one’s semantics. If you’re not going to go all the way with early Unger, but after siding with the skeptics (even *extremely* annoying everyday, as opposed to philosophical, skeptics) and writing off what seem to be even *extremely* serious positive uses of “know(s)” (including, I presume, even some uses in high-stakes court

testimony) as examples of false, loose speech, one then seeks to suddenly draw a line in the sand and say, “Wait a minute, skeptics, now *that’s* getting *too* demanding,” it’s easy to worry that one is trying to close the barn door after far too many of the horses have already escaped, and with little in our linguistic practices to look to for guidance in deciding just how few of the last horses should be kept.

But to seriously evaluate such an approach, we need to know just how demanding its semantics is, so we can evaluate just how much systematic falsehood it implicates us in on the one hand (how much truth-telling it’s able to save in our positive uses), and on the other hand whether it can account for all the apparently high-standards denials of “knowledge” we should want to account for. I’ve never seen a way to plausibly thread that needle in general, but I must admit that I really don’t feel I have a good grasp of just how demanding Davis’s semantics is – just what it takes to know on his view. (Much depends here on what it takes on his view to be “completely justified” in believing something, but that’s precisely what I don’t feel I understand, in Davis’s hands.) Readers are encouraged to read at and around esp. (Davis 2007:426-7), where we learn, for instance, that “When she arrives at the bank Saturday morning after nine o’clock and finds it bustling with activity, I think it is literally and strictly speaking true that Hannah now knows the bank is open.” So Davis seems no early Unger. But he certainly does seem to get *extremely* demanding.

F&M agree with Davis that his approach has the advantage over early Unger that it is not “wedded to an ‘impossibly demanding semantics.’” But though F&M are quoting the phrase “impossibly demanding semantics” from my depiction of Unger’s theory, it’s not as if Unger semantics for them implied that *all* positive uses of the relevant terms (the “absolute terms”) are false. Unger’s favorite example was ‘flat’, and though he held to a maximally demanding semantics for that term, according to which a surface must be one flatter than which no surface can be conceived (to use some appropriate Anselmian terminology to describe the view) before it can be truthfully described as ‘flat’, he always threw in some important qualifications when describing the results of such a demanding account, constantly writing such things as

that, for instance, “We do not speak truly, at least as a rule, when we say *of a real object*, “That has a top which is flat”” (Unger 1975: 49, emphasis added) – the implication being that we can truthfully say that an ideal plane we are discussing while, say, in working on a geometry problem, is “flat.” (So his semantics only proves “impossibly demanding” for real physical objects of the type we’re familiar with.) Perhaps Unger’s theory (whether or not Unger thinks so) will allow true ascriptions of “knowledge” to God? And on the other side, there seems to be no effective limit on how demanding Davis will get on claims like “It’s three o’clock.”⁸

At any rate, a main thrust of my third chapter is to argue that ~~that~~ pragmatic explanations of semantic intuitions can be carried too far, and to start on the task of discerning when such maneuvers are legitimate and when not. I conclude that things aren’t looking up for invariantist attempts to explain away the data supportive of contextualism, since they seem to get on the wrong side of the reasonable tests that seem to emerge from a look at some fairly clearly successful and also clearly unsuccessful uses of such maneuvers. While F&M seem to write as if arguing against pragmatic theories like Davis’s would be something new added to my case, Davis himself sees that a pragmatic theory like his would be among my intended targets in the argument I have given,⁹ and he seeks to explain why my argument does not work against *his* loose speech proposal. Key to Davis’s attempted escape is the claim that

⁸ Davis reports: “[Stewart] Cohen has objected that if all sentences ascribing knowledge of precise measurements are strictly speaking false, then no one ever knows what time it is, or when anything occurs, or how long anything is.” Davis’s response, that follows immediately, very much surprised me: “That conclusion would be absurd indeed.” Perhaps not coincidentally, Davis thinks he can meet this challenge: “That conclusion would be absurd indeed, but it does not follow. I can know that the stick is about 3 ft. long, and that suffices to know how long it is” (Davis 2007: 417). But if claims that someone know what time it is always have to a hedge like “about” or “approximately” in them to be true, isn’t that pretty bad already? If it would not just be a cost, not just damaging, or even very damaging, but positively “absurd” to have to hold that no one ever knows what time it is, isn’t it at least a fairly serious cost to have to hold that no one ever knows such unhedged things as that it’s three o’clock?

⁹ (Davis 2007: 412). Davis is here responding to the argument as it appeared in (DeRose 1999) and (DeRose 2002), which later formed the basis of Chapter 3 of the book currently under discussion (DeRose 2009).

there are clear examples of loose speech that he can use as the model for what he claims is going on when speakers claim to “know” things in my LOW cases.¹⁰

But isn't Davis right? *Isn't* it clear that, as F&M confidently assert in their second paragraph, “Loose speech is pervasive”? Well, certainly in *some* sense, yes, but I suggest caution here, because it might not be the sense that the “loose usage” theorist needs as examples of appropriately asserted *falsehoods*. Let's begin by looking at one of Davis's presumably clear examples, his “Time Measurement Case”:

A. Wondering how hard the final exam was, I ask Mike how long he took to finish. He answers ‘Two hours.’ B. When Nora says that she took two hours and four minutes to finish the exam, Mike responds ‘You took even longer than me. It took me two hours and two minutes.’

Immediately after presenting the case, Davis asserts, “What Mike said in A is false if what he said in B is true” (Davis 2007: 407). But why think that? I don't see any problem with ruling both claims as true (I'm assuming A and B represent distinct conversations), and indeed, those are my verdicts about the cases. If there's supposed to be some connection between these judgments that rules out both claims being true, that reasoning (and readers are encouraged to examine the brief paragraph in which Davis explains this verdict) escapes me, for reasons that may become apparent as we proceed.

But Davis may well mean not to be so much making some connection here, but just to be registering his sense or intuition that Mike's answer in A is false if he

¹⁰ Davis writes, “DeRose (1999, p. 200) knew of no case in which what a speaker says appears true just because of a true implicature,” and he replies, “Loose use provides countless examples” (Davis 2007: 413). Now, what I had written was that none of the “*clearly successful*” pragmatic maneuvers of the relevant type explain away apparent truth by appeal to the fact that a true implicature is generated. Davis seems to get that qualification, because he has just made a claim that there are “clear cases” of loose use, writing, “On the contrary, I am accounting for the observed behavior of ‘know’ in terms of the general phenomenon of loose use. I use the close similarity between cases of contextual variation in the use of ‘know’ and *clear cases* in which diverse terms are used with variable strictness as evidence that the variation is due in part to its being used loosely or strictly” (Davis 2007: 412, emphasis added).

took two hours and two minutes to finish. And if Davis were to just take it to be obvious beyond the need for argument that what Mike says in A is false if it took him 122 minutes to finish, he certainly wouldn't be alone. In fact, he would be following the lead of perhaps the most important recent work on loose speech (not cited by Davis, but I assume he's aware of it), (Lasersohn 1999), which begins with a very similar case about time. Lasersohn's abstract begins with the claim,

It is a truism that people speak 'loosely'—that is, that they often say things that we can recognize not to be true, but which come close enough to the truth for practical purposes.

And the body of his paper begins as follows:

People speak with varying degrees of precision, and often speak quite loosely. Suppose, for example, I tell John that Mary arrived at three o'clock. In certain relatively unusual circumstances, the exact second of her arrival might be important, but most of the time this level of precision is not required. So if John finds out later that Mary didn't arrive at three but at fifteen seconds after three, it would be unreasonable of him to complain 'You said she came at three!'

But whether or not John is acting unreasonably in this situation, I think we have to concede that he is, strictly speaking RIGHT: when I told him that Mary arrived, I said something that was literally false, not true. (Lasersohn 1999: 522)

What I wish to suggest here is that while it is certainly true that “people speak with varying degrees of precision,” and that it is clearly true that in *some* good sense that we often speak loosely (and I'll explain in what sense I think that happens a few paragraphs down), it is very far indeed from clear, much less a “truism,” that these (or other) examples are cases of “loose speech” as *Lasersohn analyzes that term*, and as Davis needs it to be understood, because it is very far indeed from clear that our speakers are asserting falsehoods.

Indeed, left to my own devices, I would have thought that it was at least somewhat clear in the relevant cases that the speakers were speaking truthfully. And I wouldn't be alone in that, either. As it happens, when I first discussed Lasersohn's

example with a linguist I know who was clearly going along with Lasersohn's verdict, I happened to be carrying in my backpack John Hawthorne's then fairly recently published (2004), so I was able to pull that book out and show where Hawthorne claims that in many relatively undemanding contexts (no doubt including those like the context Lasersohn imagines), one speaks truthfully when one asserts "It is three o'clock," even when it's a full minute (and not just 15 seconds, as Lasersohn has it) after three o'clock! And Hawthorne gives no argument for this verdict. He issues it and then uses it in the argument he's giving (Hawthorne 2004: 99). So it must have seemed to him clear enough to be so used.

I suppose that, given this sharp disagreement in how things seem, I should no longer take it to be so clear that these "loose uses" are true. I'm just a little hesitant to do that, however, due a suspicion that those on the other side are being misled by faulty reasoning. Space is drawing short, so I'll quickly state my suspicion in one of its main forms. I suspect that those who favor an invariably very demanding semantics in these cases are being (mis)led by such considerations as that

a) "It's three o'clock, but it's actually a few seconds after three o'clock"

would seem to be just wrong, in just about any context they're said. To the extent that such considerations are responsible for their judgments, I believe those judgments are based on a bad underestimation of the ability of contextualist accounts to explain such data. Suppose that "It's three o'clock" gets used with varying degrees of precision in this (contextualist/indexicalist) way: how close to three o'clock it must be for such a claim to be *true* varies with context, and in less demanding contexts, it can certainly be true when it's, say, 15 seconds, or even a full minute, after exactly three o'clock. A "loose use" of such a claim, then, would be one governed by a relatively low standard of precision in that sense. The prospects for such a view to account for why a) sounds so bad seem to me bright, since it can be based on some claim to the rough effect that when you bring up the matter of a few seconds and treat it as relevant, you shouldn't/can't combine that with a low-standards-of-

precision claim that “It’s three o’clock,” where a few seconds give-or-take need to be deemed irrelevant. The rough idea could be that even in a context that starts out as governed by low standards of precision (that would render the first half of (a) true) would tend to *become* a high-standards context when (a) is asserted. And that could I think explain why (a) sounds so bad, in just about any context you imagine it being asserted.

I have similar reactions to Davis’s other examples. Those wishing to pursue these issues might do well to start with section of Davis’s paper entitled “Semantic treatments of loose use” (Davis 2007: 415-17) where Davis reports suggestions made by Stewart Cohen (not in any work published yet, but in correspondence with Davis) for how to treat Davis’s examples in the contextualist/indexical way alluded to in the above paragraph. (It’s of course no surprise that Cohen and I have similar views here.) Davis argues against Cohen’s ideas, but readers may well get good ideas of how to respond to Davis’s arguments. I’ll just report that I’m *still* finding it very far indeed from clear that the “low-standards” claims that serve as the examples of “loose speech” are false. A main problem with the account of early Unger was that, although he provided a lot of company for the misery he visited on our use of “know(s),” there was an alternative contextualist approach to all of the relevant terms that avoided the misery quite generally, and so Unger’s general approach had no clearly correct applications and it was difficult to see where the pressure was coming from to make us put up with all the false-speaking the theory implicated us in. If the examples of Davis and Laserssohn were clear cases of false, loose speaking, that would have solved one of the main problems of Unger’s approach.

(as later Unger realized)

not right: need something like: "would have solved, *for the loose usage approach,* one of...

Reply to Fantl and McGrath

So I’m a lot less worried about the threat of LOW turning out to be an example of false “loose speaking” than F&M seem to be. On the other side, I seem to worry a lot *more* than F&M do about the truth of general principles like their

(Actionability) If you know that p, then p is actionable for you, i.e., you can count on p's being the case,

which plays a crucial role, at Step 3, in their suggested “cases-as-conclusions” strategy for supporting contextualism. I won't here pursue the question whether my type of argumentative strategy is better than F&M's alternative proposal, especially since, as F&M note, these strategies needn't crowd each other out, but can be pursued together.

But I do have to express my skeptical outlook on *Actionability* in another connection, for though F&M propose that principle as part of their suggested route to contextualism, a destination I like, that principle takes us to subject-sensitivity, which I resist. (F&M then are proposing a trip to subject-sensitive contextualism, which to my thinking is a very mixed bag.)

A reason for being initially skeptical about principles like *Actionability* is precisely that they are powerful principles of very general application, whose plausibility should be judged in large part by how correct they seem in their most problematic applications. By comparison, if someone claims that all the starters for a certain college basketball team are good enough to play in the WNBA, you'll want to focus on the worst of those starters in evaluating the plausibility of that general claim. Likewise, if someone claims that in all possible cases, if an agent knows some proposition to be true, then that proposition is actionable for that subject, you should focus on the most problematic cases for that principle in evaluating it. Now, if it turns out that it seems quite likely that even the worst of those starters in the basketball case really is good enough for the WNBA, then the general statement can turn out to be pretty plausible. (And of course, I don't reject all general principles.) And, as I admit, and even argue a bit for, even the most problematic cases for *Actionability*—which I take to be cases like Jessica Brown's that I discuss at (DeRose

2009: 265-9)¹¹—don't seem to be *clear* counter-examples to it. On the other hand, it's also quite plausible that those are cases where knowledge is not actionable (and so it's plausible that they are cloudy counter-examples to the principle). Those applications of the principle, and therefore the principle itself, at least insofar as we are judging the principle by its applications to cases, are quite problematic.

But our evaluation of this principle needn't be confined to just the plausibility of its applications to particular cases, and F&M cite reasons for accepting it that don't just consist of pointing to its plausible simple applications to cases.¹² I can briefly explain my more skeptical attitude toward *Actionability* by discussing the dialectical situation at a fairly high level of abstraction and saying that (in Chapter 7) I try to account for the kinds of grounds F&M cite for their principle by appeal to other things, often relativized versions of one-way knowledge-action principles like *Actionability*, where these relativized principles (and/or other materials I use in my accountings) don't lead to anti-intellectualism the way that *Actionability* itself does. F&M's strategy is not to dispute that these alternative attempts to account for the relevant data succeed or are promising. (We should not, and I don't, take their silence

¹¹ At one point, immediately after explaining why they think a certain combination of considerations “tips the balance in favor of an SS-view,” F&M write, “DeRose begs off because he ‘couldn’t reach a conclusive enough verdict’ about such principles. (267).” I should make clear that what I was not able to reach a stable verdict about was whether even relativized versions of these one-way knowledge-action principles were correct. This is why I took the fairly nuanced stance that if there are sufficient grounds for accepting a one-way knowledge-action principle, those grounds can likely be accounted for by a relativized principle. About unrelativized principles like *Actionability* itself (though the specific one-way principles I treated in the book were slightly different), I was able to reach at least a negative verdict, due to considerations I’m presenting here: mainly, that they lead to anti-intellectualism. I hope this presentation will make what I take to be the basic structure of case against such principles clearer, even as it can’t pursue any details to the extent they are discussed in the book.

¹² In this respect, F&M seem better off than Hawthorne, whose case for a very powerful principle he starts wielding consists of pointing to a case where its application is plausible. Given that (as I point out) there are also some very problematic applications of his principle, this really does seem much like the person who would point to one of the better starters on the college team, plausibly claim that that player is a shoe-in for the WNBA, and then infer the general claim, while all the while you know that there are those other starters who seem very implausible as WNBA candidates. See (DeRose 2009: 161-6) for discussion.

here to be an implicit endorsement: When working within space limitations, they no doubt had to be very selective in what they included.) Their strategy seems rather to claim that, *even if* the accounts I attempt work, there are still important relative disadvantages for a position like mine. But the advantages F&M cite (their five bullet points toward the end) are advantages that the accounts in question, if successful, are supposed to *neutralize*. The goal, for instance, to take their third point, is to explain why certain forms of reasoning would appear to be valid even though they aren't. Or, to take their fifth point, to explain why a certain unspoken sentence saying a subject knows something may be true in a certain context, even though it can seem false. It may be a bit better to accept that things are as they seem than to explain away the appearances, but if the explaining away does really work, that takes the sting out of denying that things are as they seem.

It's important for me that this explaining away of various appearances is done against the backdrop of Chapter 6, where I try to establish the attractions of intellectualism and some of the relative costs of anti-intellectualism. F&M may be able to use the contextualist aspect of their view to mitigate or remove some of those costs incurred by their anti-intellectualism, but intellectualism still presents a very attractive account of the types of factors that can matter to knowledge.¹³ It's this that motivates an intellectualist account of the kind of phenomena F&M point to, even if this account will at key points involve explaining away rather than validating certain appearances.

neutralize

As F&M set things up, a crucial difference between my theory and theirs is how we rule on their "LOW-on-HIGH" case: That's the sole difference between the theories that makes their "scoreboard" chart, and it proves decisive on that scoreboard because it's there that I get my only frownie-face, while they get smiley-faces straight across the board. So I should register that I reverse those rulings. My

¹³ And very interestingly, when I motivate the intuitive attractiveness of this account toward the beginning of Chapter 6 at pp. 189-90, I do so with F&M's effective help, quoting (Fantl and McGrath 2007).

crucial inclination is to say that what Thelma says there is true, so it's DeRose's contextualism that gets ~~the~~ smiley-face. But I'll also say that I don't think that intuitions about cases like these, where a speaker's ignorance of facts that potentially play a big explanatory role in why it's appropriate for her to say what she says, are worth much. We can try to adjudicate this difference by considering related cases. So consider what happens when we suppose Thelma realizes that Louise is or may well be talking to the police in a high-stakes discussion. (In F&M's "LOW-on-HIGH," it's Louise's "knowledge" that's at issue.) I think that if Thelma's conversation at the tavern remains focused on the low-stakes matter of whether Louise will "know" to pick up her winnings on the \$2 bet, she can quite appropriately, and will in fact if she's a normal speaker, continue to describe Louise as "knowing" that John was at the office, and I think that this suggests that the fact that Louise is in a high-stakes talk with the police doesn't count against the truth of Thelma's attributions of knowledge to her in contexts where low-stakes purposes hold sway (since one cannot appropriately say what from one's own point of view is likely false), and that therefore Thelma's attribution in the unmodified LOW-on-HIGH is true. But I realize there is ~~certainly~~ room to dispute my ruling on the modified case and/or its bearing on the unmodified case.

References

Davis, Wayne A. 2007. "Knowledge Claims and Context: Loose use," *Mind* 132: 395-438.

Philosophical Studies

_____. 2010. Review of DeRose (2009), *Mind* 119: 1152-7.

DeRose, Keith. 1999. "Contextualism: An Explanation and Defense," in J. Greco, & E. Sosa (eds.), *Blackwell Guide to Epistemology*, Malden, Mass.: Blackwell Publishers, 187-205.

_____. 2002. "Assertion, Knowledge, and Context," *Philosophical Review* 111: 167-203.

_____. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*, Oxford: Oxford University Press.

Fantl, Jeremy and McGrath, Matthew. 2007. "On Pragmatic Encroachment in Epistemology," *Philosophy and Phenomenological Research* 75: 558-589.

Hawthorne, John 2004. *Knowledge and Lotteries*, Oxford: Oxford University Press.

Lasersohn, Peter 1999. "Pragmatic Halos," *Language* 75: 522-551.

Unger, Peter 1975. *Ignorance: A Case for Scepticism*, Oxford: Oxford University Press.