

Contextualism, Contrastivism, and X-Phi Surveys

Keith DeRose, Yale University

draft of Jan. 2011
forthcoming, *Phil. Studies*

The tale that Jonathan Schaffer and Joshua Knobe tell in “Contrastivism Surveyed” (henceforth, “S&K”) is a tragic one for what we may call “standard contextualists” about knowledge attributions. First, they report (word of this has been “out on the street” for a while now) that a recent wave of work in Experimental Philosophy threatens to undermine the intuitive basis that contextualists have claimed for their view. Given the importance of that intuitive basis for the view, this would be very bad news indeed for contextualists. But the saddest stories usually aren’t just bad news from start to finish: True tragedies usually involve some real or apparent upturn in fortune, some glimmer of hope that can then be cruelly stomped out in the tragic ending of the story. So, in Act II, we’re told that these experimental results that seem hostile to contextualism turn out to be driven by what looks like some kind of performance error (though S&K don’t use that term to describe what’s going on) being made by the survey takers: If you ask the survey questions in the right way (as S&K have now done in a new survey), you both get the kind of results contextualists would be looking for, and also the makings of a case that it’s these friendlier results, and not the earlier, hostile ones, that should be heeded. But then we find in Act III that when you look carefully, what the experimental results really turn out to support is a particular form of contextualism, contrastivism, that “standard” contextualists reject (and by “standard” contextualists, I guess I just mean “non-contrastivist” contextualists).

I will here sharply oppose all the phases of the story S&K tell. In Part 1 we will look at the supposed empirical case against standard contextualism, and in Part 2 we will investigate S&K’s supposed empirical case for the superiority of contrastivism over standard contextualism. (I won’t here go into my own reasons for preferring standard contextualism over contrastivism; for that, see DeRose 2009: 34-41.)

1. Experimental Philosophy (X-Phi) and the Empirical Case against Contextualism

1.1 *A wave of bad empirical news for contextualism (and many of its rivals)?*

Yes, word is out that the recent experimental work that S&K report on has undermined the basis for epistemic contextualism. What’s more, as S&K report, many of the contextualists’ opponents—the invariantists—have accepted key elements of that intuitive basis for contextualism, but have tried to account for that basis within an invariantist framework, resulting in forms of invariantism that are designed to accommodate some of the same intuitions that drive contextualism, and those invariantist views are likewise jeopardized. A fairly extensive and robust debate has been raging in

the professional philosophy journals for quite a while now, with almost all participants being at least largely in agreement about what the key intuitions are that should somehow be addressed, but disagreeing about how best to handle them.¹ And I think this statement from S&K well captures the spirit of the word that is spreading through some philosophical circles (largely comprised of younger philosophers, in my experience): “Yet a recent series of empirical studies threatens to undermine this whole debate” (p. 1).

How serious is the threat? I’m a little puzzled about S&K’s thought here. As we’ll discuss later, based largely on their own surveys, S&K thinks something “went wrong” in the previous studies they discuss, and they think you can get results more friendly to contextualism “when the studies are done properly” (p. 16). Nevertheless, in their conclusion, after all that is past, based on the presumably somehow improper earlier studies, they still quite boldly assert (where §2 is the portion of their paper in which they discuss the earlier studies): “The existing debate over contextualism is based in large part on claims about what ordinary speakers will say about bank cases (§1). But these claims do not survive empirical scrutiny. Ordinary speakers do not say what philosophers say they will say (§2)” (p. 29). But, at any rate, S&K seem to think the threat is very serious indeed.

omit s

What is the “wave” of empirical results responsible for this ominous threat?² S&K cite four studies as threatening to “undermine this whole debate”:

- Neta and Phelan (manuscript; henceforth N&P)
- Feltz and Zarpentine (forthcoming, henceforth F&Z)
- May, Sinnott-Armstrong, Hull, and Zimmerman (forthcoming, henceforth MSH&Z)
- Buckwalter (forthcoming, henceforth B).³

But it turns out the intuitive support for contextualism doesn’t really face much of a wave of empirical trouble. Only two of these studies (the last two listed above) even purport to be calling into question any case for any form of contextualism, and, as I will argue, there are some severe problems with taking the results of any of these studies (whatever their aims) as undermining the intuitive basis for contextualism.

1.2 May, Sinnott-Armstrong, Hull, and Zimmerman (MSH&Z), anti-anti-intellectualism, and asking directly whether a character knows

It’s actually a bit generous to say that two of the studies S&K cite target the case for contextualism. One of those two, MSH&Z, takes anti-intellectualism, and mainly, in

¹ There are of course some disagreements about intuitions. For example, in my review of (Stanley 2005), I report a difference of intuition I have with Jason Stanley concerning one of Stanley’s key cases (DeRose 2007: 486). Still, the debate is fairly characterized as one primarily concerned with how to handle certain intuitions, rather than over what the relevant intuitions are.

² At p. 11, S&K describe the earlier studies they discuss as a “recent wave of empirical research calling those claims [made by contextualists and some of their opponents about our intuitions] into doubt.”

³ S&K also discuss a fifth study, Beebe and Buckwalter, but this one concerns another issue and isn’t taken by anyone to undermine the case for contextualism.

particular, Jason Stanley’s anti-intellectualism, as its main target, arguing that their empirical results undermine some of Stanley’s key claims. But they also target Jonathan Schaffer’s contrastivism, claiming to undermine Schaffer’s claims as well, and, as we’ll see, contrastivism can be viewed as a form of contextualism. But however well or poorly MSH&Z engage their intended targets, their results have no tendency to jeopardize the kinds of intuitions generally used to support standard forms of contextualism.

To see why, let’s start with a very basic description of the kind of case I have made for contextualism. My main argument for contextualism is based on pairs of cases, like my bank cases that get discussed in the experimental philosophy literature, that I take to illustrate important features of ordinary speakers’ use of “know(s),” in one of which (the “high standards case,” henceforth HIGH) some speaker says that some subject does *not* “know” some proposition to be the case, and in the other of which (the “low standards case,” henceforth LOW), some speaker says that the subject *does* “know” the thing in question. One difference between HIGH and LOW involves what’s at stake in each case: In HIGH, very much is riding on the speaker being right about the matter in question. Another difference concerns which possibilities for error have been mentioned in the conversations taking place in the stories. The argument is driven by premises claiming that both of the relevant claims – the denial of knowledge in HIGH and the affirmation of knowledge in LOW – are true.⁴ (Contextualism is supported by its providing the best explanation for how both claims are true.) And an important source of support that I claim for these premises is their own intuitive plausibility.⁵ In connection with worries about survey methodology that we’ll briefly discuss in section 1.6, I’ve been warned that presenting my bank cases, as I have formulated them, verbatim to survey takers would be problematic. However, if one were to empirically test the basic type of intuitions my argument for contextualism actually appeals to in any kind of direct way, the way to do so would be to present cases that are like my HIGH and LOW in that they feature a speaker making a denial of knowledge in HIGH and a speaker making an affirmation of knowledge in LOW, and then ask respondents whether they think those claims, made within the stories, are true.

MSH&Z don’t do that. Their surveys concern cases in some ways like the ones I use (including prominently differences in stakes between the different cases they are testing), but they do not ask their survey-takers to evaluate the truth-value of a knowledge

⁴ This argument is described and executed mainly at (DeRose 2005 \cong 2009: Chapter 2), though the argument is completed at (DeRose 2002 \cong 2009: Chapter 3), where I battle against the most popular way of attempting to evade it. For reasons I give at (DeRose 2005: 181-7 \cong 2009: 59-66), it is better for the argument not to be based on cases featuring first-person knowledge claims like my original bank cases, so I actually instead base it on different case pairs that feature third-person knowledge attributions. At (DeRose 2002 \cong 2009: Chapter 3) I also give a second argument for contextualism (the presentation of which is significantly improved in the 2009 version).

⁵ These intuitions, I claim, are one of two mutually supporting strands of support for the key premises of the argument. For further discussion, including an explanation of what the other strand of support is, and of the relation between these two strands of support, see DeRose 2005: 172-4 \cong 2009: 49-52.

claim (or denial of knowledge) made within the story. They simply ask survey-takers whether the subject in the story knows the proposition in question. On the face of it, this is a way to test how intuitively plausible anti-intellectualism is, since anti-intellectualism is the view that the “practical” matter of how much is at stake for a subject is a factor affecting whether or not that subject knows something to be the case. If, as MSH&Z report, you don’t get much of a difference in evaluations of whether the subjects in such case pairs know the propositions in question when these subjects are portrayed situations that differ greatly in terms of how much is at stake for them, that, on the face of it, is bad news for anti-intellectualism. And, indeed, as I’ve mentioned, anti-intellectualism is MSH&Z’s main target. It is also the target of the two studies that don’t even purport to undermine contextualism—N&P and F&Z. But contextualism is not anti-intellectualism. Indeed, though an anti-intellectualist version of contextualism is a coherent option, most contextualists are intellectualists, and, indeed, an explicit or implicit commitment to intellectualism is probably a big part of the reason why they are contextualists.⁶ (That’s certainly true of me.) So if MSH&Z have bad news for anti-intellectualists here, most contextualists, so far from feeling threatened themselves, should take this as good news.

But since such studies are sometimes taken to be a problem for contextualists, it is worth emphasizing just how important is this difference between asking about the truth-values of knowledge claims made within the stories, and asking directly whether a character in the story knows. It may turn out, for all I know, that this difference will not matter much to how survey-takers will tend to answer the questions they are given. But this would mean that they’re probably not very competent in answering at least one of the questions, for the difference here is vital to discerning among the philosophical options on the table.

One way to illustrate this importance is through this question: Assuming that survey-takers will tend to answer truthfully when they respond to prompts like MSH&Z’s that ask directly whether characters in the story know things, what does contextualism predict about how they will answer? General answer: That depends on details of the contextualist position that I haven’t seen worked out by any contextualists. My own contextualist answer: I have no idea! In MSH&Z’s “High Stakes” bank case, for instance, their character Hannah finds herself in a situation in which (of course) the stakes are very high. They then give survey-takers these instructions: “Please check one box to indicate how strongly you agree or disagree with this statement: ‘Hannah knows that the bank will be open on Saturday’.” Though the matter isn’t *entirely* clear, this seems to be asking about whether that attribution of knowledge is true about the story in the context that the survey-takers find themselves in as they take their surveys. If so, how will the contextualist expect them to answer? That depends on what standards we should expect to govern the survey-takers’ context. These survey takers are evaluating the

⁶ For discussion of the relations among contextualism, views like Stanley’s, and intellectualism, see Stanley 2005: 2-3 and DeRose 2009: 185-190.

situation of a character who is in a very high-stakes situation – and a character who would likely use very high standards in evaluating herself. Does that mean the survey-takers’ own context will also be governed by high standards as they evaluate this character? That’s very hard to say. A feature of the contextualist view I have endorsed is that in *some* contexts, speakers apply to subjects the standards that would be appropriate to the practical situations of the subjects being described. This occurs, for instance, when speakers are evaluating whether the subjects “know” things for the purpose of making certain evaluations about the actions or action plans of those subjects. (“She shouldn’t do that, because she doesn’t know that it won’t hurt any bystanders.”) But in other contexts, where other conversational interests are in play, speakers should and do apply epistemic standards that are appropriate to the speakers’ own practical situations, rather than the practical situations of the subjects they are talking about. This happens, for example, when the speakers are interested in whether the subjects “know” things because the speakers are wondering whether these subjects might be sources of information for the speakers.⁷ So, what about this exceedingly odd context in which someone is taking a survey that’s asking them whether a character in a high-stakes story knows? If I’m right that whether speakers will apply to subjects standards appropriate to those subjects’ practical situations, or completely different standards, depends in part on *why* the speakers are discussing the issue of whether those subjects “know” (e.g., are they looking to use a judgment as to whether the subject “knows” in reaching an evaluation of the actions of the subject or are they instead thinking of using the subject as a potential source of information by which they, the speakers, might come to “know” the truth of the matter in question?), then what should we expect when these speakers are taking a survey and they are in a context in which they have no idea what kind of use this judgment is going to be put to? Again, I really have no idea what to expect here. (Though considerable confusion on these survey-taker’s part may be a reasonable guess.) At any rate, I find myself not being in any position to take whatever results emerge as either particularly good news or bad news for my views.

This is not a complaint against MSH&Z: they do not purport to be undermining any case for contextualism like the one I have made. But it is a warning to those who might take results like MSH&Z’s to have some bearing on that case.

1.3 Buckwalter (B) and the problem with versions of HIGH that feature positive attributions of knowledge

The only other study that S&K cite that even purports to undermine any form of contextualism, and the only one that purports to undermine the kind of case I rest my contextualism on, is Buckwalter. Buckwalter explicitly takes my case for contextualism as among his targets (and it seems as his main target), citing DeRose 2005, and in his surveys he does ask whether knowledge claims made by characters in the story are true,

⁷ See DeRose 2005: 188-90 \cong 2009: 238-241.

rather than asking the survey-takers directly whether some character knows something.⁸ However, there are a couple of crucial differences between how I've set up the case for contextualism and how Buckwalter attempted to test it empirically.

First, as I've already explained, my case is based on pairs of cases, in one member of which, LOW, a speaker positively ascribes knowledge to a subject, and in the other of which, HIGH, a speaker says that the subject does not know. I appeal to the intuitions that both the positive attribution of knowledge in LOW and the denial of knowledge in HIGH are true. For some reason, in testing these intuitions, Buckwalter paired a case like my LOW, in which a positive claim to know is made, with cases that in some ways are like my HIGH, but in which the speaker again positively attributes knowledge. When B found that survey takers tend to rule both attributions of knowledge to be true, and the difference how likely they were to so rule to be statistically insignificant, he thought this undermined my case.

But isn't this a fairly innocent change to my cases? I don't think so, for reasons we're about to see. The strange thing is, I didn't just happen to design HIGH so that it included a denial of knowledge; I explained why it was important to the intuitions I was appealing to that I set the case up that way, rather than having the speaker positively attribute knowledge in HIGH and appeal to the intuition that this positive attribution is false. I really had severe doubts that such a judgment of falsehood would be very intuitive. And I explained this in the very paper (2005) that Buckwalter cites. Yet Buckwalter tested my case by means of examples set up in the way I specifically advised against, without explaining what he thought was wrong with my advice, and then goes on to take his results as undermining my case.

As I explained, I think a significant part of what's behind the intuitions I actually appeal to is the realization that my 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 LOW, yet will deny 'knowledge' when they find themselves in conversational circumstances like HIGH" (DeRose 2005: 173 \cong 2009: 50).⁹ This sense that the speakers are using "know(s)" properly yields or buttresses

⁸ An early version of Buckwalter's paper asked directly whether characters in his stories knew things to be the case, and Buckwalter presented his results as undermining my case for contextualism. When Buckwalter posted this earlier version of his paper on "Experimental Philosophy" web site, it inspired a spirited and interesting discussion in the comments to the web blog entry that Buckwalter posted: "Knowledge on Saturday," posted July 26, 2008; currently at: http://experimentalphilosophy.typepad.com/experimental_philosophy/2008/07/knowledge-on-sa.html The current version of Buckwalter's paper now instead asks about the truth-values of claims made within his stories – I presume due to the on-line discussion just mentioned.

⁹ I don't believe x-phi surveys of the type we're discussing here are tests of how speakers tend to use terms. These surveys are set up to test intuitions about the correctness of certain statements. These two issues seem to get conflated in some of the x-phi literature – especially by S&K. Consider, from their "Conclusion":

a judgment that the speakers are speaking truthfully through a certain kind of presumption of truth: “This supports the premises that both of the imagined claims are true, since generally (though there are some exceptions), one cannot properly claim 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” (DeRose 2005: 173 \cong 2009: 50-1). I present this thinking as a second “strand” of support for my premises (that the claims made in the cases are true), where the first strand is the direct intuition that the claims are true. However, I quickly go on to point out that these two “strands” are intertwined, and that “it is no doubt largely because we already sense that the speakers in the cases are using ‘knows’ in common and appropriate ways that we intuitively judge that their claims are true” (2005: 173-4 \cong 2009: 51). And so, it would be no great surprise to me if the intuition I appeal to dissipates when what I think it is partly based on is removed, as it is when Buckwalter changes the cases in the way we are currently considering.

“Well, won’t the contextualist think that the speaker in Buckwalter’s version of the story is using ‘know(s)’ improperly, and so shouldn’t he think that will combine with a general presumption that what’s said improperly is false, so that no real harm is done to the intuition by Buckwalter’s modification?”

No, as I argue (following David Lewis – I think!), there is an important asymmetry here, and the impropriety of a claim doesn’t support a verdict of falsehood for that claim to nearly the same extent that propriety supports a judgment of truth. (See DeRose 2005: 174-5 \cong 2009: 52-3 for discussion.)¹⁰

Indeed, the problems (for the intuitions I appeal to) that are caused by Buckwalter’s switch go beyond the intuitions losing a source of their support to their

The existing debate over contextualism is based in large part on claims about what ordinary speakers will say about bank cases (§1). But these claims do not survive empirical scrutiny. Ordinary speakers do not say what philosophers say they will say (§2).

The empirical data about what ordinary speakers will say suggest that our intuitions are in fact sensitive to contrasts and salient error possibilities, but insensitive to stakes (§§4-5).

Here S&K start by speaking about appeals to how speakers use “know(s),” but think such appeals are refuted by data concerning respondents’ intuitions about how to evaluate certain claims. But how one evaluates a claim one is asked to evaluate is one thing; how one would oneself describe a situation is, at least on the face of it, quite another thing.

¹⁰ There is also the question of whether the speakers in Buckwalter’s versions really would be speaking inappropriately. We cannot assume that because a certain claim is appropriate, the opposite claim would be inappropriate. It could happen that while asserting “P” would be more appropriate than asserting “not-P,” the latter wouldn’t be all that bad. In fact, in some cases, a speaker can appropriately assert either “P” or “not-P”—and in fact, I think it can often happen that either assertion would be *true*, for reasons we’re about to see in the next couple of paragraphs.

having something new lined up against them. In DeRose 2009,¹¹ I have occasion to compare the case for contextualism I actually rely on with what the case would be like if I instead had the speaker in HIGH positively ascribe knowledge and then used a premise that this positive claim is false, and I reached this judgment about the latter kind of case: “I would judge such a case for contextualism to be quite weak, due to the shakiness of the premises that the ascription in [HIGH] is false” (2009: 71). Buckwalter then is testing a premise I judged to be “shaky” enough to render an argument based on it “quite weak.” The problem I have with the switched argument that goes beyond the loss of support from a presumption of truth is that its premise about HIGH will now run afoul of (rather than benefit from, as my actual premise does) “accommodation”: “It may be more difficult than one might think to devise a case that is otherwise suitable to the argument and in which it really does strike us as intuitively clear that the positive ascription in [a switched version of HIGH] is false, because there is pressure on us as interpreters of the ascription to understand it as having a content that makes it true, due to the operation of what David Lewis calls a ‘rule of accommodation’ (Lewis 1979: esp. 346–7)” (2009: 71–2).

I take accommodation to be a *very* important force in the interpretation of context-sensitive language.¹² In fact, I think this force is often responsible for speakers being in a situation in which they’ll be asserting a truth whether they assert either a certain sentence featuring a context-sensitive term or that sentence’s negation! So, for instance, there may be situations in which you and I are both looking at Tom, and we both have a pretty accurate idea of just how tall Tom is, and there’s no particularly strong conversational forces militating either for standards that Tom meets for what counts as “tall” or for standards he doesn’t satisfy, and I will be speaking truthfully if I say to you “Tom is tall” or if I say “Tom is not tall,” because an appropriate rule of accommodation will have it that the standards for “tall” governing our conversation will adjust to make what I say to you count as true.

¹¹ I should point out in fairness to Buckwalter that we are now into material that was new to DeRose 2009, and so was not something Buckwalter had access to as he conducted his study. Any reference to my previous work in this section for which I give only a 2009 reference is material that was new to that work. At the same time, this was material I had written into 2009 before I had ever heard of Buckwalter’s study.

¹² I’ve encountered some philosophers who think I don’t believe in rules of accommodation, apparently because I reject a solution to a skeptical problem that’s based on such a rule at DeRose 1995: 7–13. But I am only arguing there that an appeal to such a rule cannot provide the kind of explanation we should be seeking in solving the skeptical problem I am there wrestling with, not that rules of accommodation don’t exist – or even that they aren’t operating to the skeptic’s benefit in presentations of skeptical arguments. Indeed, though this is (perhaps unfortunately) in a footnote, I write: “None of this is to deny that there is some Rule of Accommodation according to which the standards for knowledge tend to be raised to ‘accommodate’ denials of knowledge. Nor is it even to deny that such Rules of Accommodation help the AI skeptic. In fact, I find it plausible to suppose that many denials of knowledge, including those of AI skeptics, often do exert an upward pressure on the standards for knowledge via some such rule.” (DeRose 2005: p. 10, n. 14).

To change my argument for contextualism so that what was an important and powerful aide to it – accommodation – is instead a force working against it is potentially to damage it quite significantly.

1.4 Buckwalter and the problem of isolating potentially context-affecting factors

A second change Buckwalter makes to my cases is to replace my single HIGH case with two different cases to try to isolate two different ways that my HIGH differs from LOW. My LOW is both a low-stakes case and a case in which certain possibilities of error are not mentioned. I contrast it with HIGH, in which the stakes are high and in which the possibilities of error in question have been brought up. (The character's basis for thinking the bank will be open on Saturday is that he was at the bank just two weeks ago on Saturday, and found that it was open on Saturdays until noon. The possibility of error that gets explicitly mentioned in HIGH is that the bank has changed its hours.)

Buckwalter instead compares a case like LOW with two different cases, one of which features high stakes but no mention of the relevant possibilities of error, and in the other of which the possibilities of error have been mentioned but the stakes are kept low. S&K are bullish about this kind of isolating maneuver, writing that it “removes” a “confound” my cases are guilty of (S&K: 4).

But one person's “confound” is another person's wise testing of a combination of factors. For the purpose of discerning exactly which features of cases are responsible for judgments of knowledge vs. no-knowledge, one may well want to utilize cases that isolate factors in way Buckwalter's do. But if you are trying to test the case for contextualism, you may be doing great harm to the case you are supposedly studying by putting asunder what the case has joined together. I am very explicit of my procedure here and my reasons for it – again, in the paper (2005) that both Buckwalter and S&K cite. I pile on. Why? Well, there can be disagreements among contextualists about which features of context really do affect the standards for knowledge. To discern which form of contextualism is correct, we may want to isolate factors to test them one by one. But in deciding between contextualism and invariantism in the first place, the dialectical situation calls for packing as many potentially important differences as is practicable into our test cases: “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” (DeRose 2005: 176 \cong 2009: 54).

Buckwalter assumes that contextualism predicts that a difference in practical stakes will result in different standards. But that depends on the details of the

contextualist view being discussed. “But don’t the contextualists’ cases typically feature very prominently just such a difference in stakes?” Well, yes, but they also feature other differences, and it may be, for instance, that raised stakes will raise the epistemic standards only when it occurs in combination with other features of context. Or perhaps raising the stakes has no direct impact itself on the standards, but only aides our ability to intuit a raising of the standards that is brought about by some other factor. This is not just an abstract possibility: For my own part, I have long held views which would render me quite unsure whether the standards would go up in cases in which the stakes go up, but in which this happens in isolation from other factors – and, again, I expressed these views in the paper that Buckwalter and S&K cite.¹³ I think that the standards tend to follow (in various ways) the conversational moves made by the speakers (including which possibilities of error are mentioned and taken seriously, but also, via accommodation, whether they’ve claimed or denied “knowledge” and whether these claims have been accepted) in a conversation, rather than, directly, the stakes of the situation they find themselves in. Now, speakers will tend to use higher standards in some of the situations in which they face higher stakes, and there will often be something wrong with using standards inappropriate to the practical situation one faces, but *I* think that if one indicates by verbal maneuvers that one is using certain standards, then even if those standards are inappropriate to one’s practical situation, one’s talk nevertheless really is governed by – its truth conditions are given by – those inappropriate standards. The best cases for testing whether contextualism is true do indeed feature differences in stakes, but, as I’ve stressed, not because the standards directly follow the stakes, but because the relevant intuitions concerning the truth of speakers’ claims are weakened if those speakers are using standards inappropriate to their practical situation, I think, for the reason I gave in 2005:

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) with 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

¹³ To be sure, I haven’t before explicitly expressed doubts about whether differences in stakes alone affect epistemic standards. But as we’re about to see, I have expressed views that would tend to cast some doubt over any confident prediction here.

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" (DeRose 2005: 177-8 \cong 2009: 55-6).

What then should I expect to happen if the stakes are raised, but this happens without being accompanied by any difference in the verbal moves made in the conversation in question (including not only that no new possibilities or error are raised, but also that the speakers positively attribute knowledge)? I don't know, and I wouldn't be at all surprised if intuitions about the truth-values of the claims of such speakers are quite murky. And what should I expect if fairly remote possibilities of error are mentioned and taken seriously by speakers who are in a low stakes situation that does not at all call for such possibilities to be made relevant? I'm on record on that question: I think such talk will in fact be governed by the high standards on which such remote possibilities are relevant, but intuitions that depend on the talk being governed by such standards will tend to be very insecure.

1.5 Feltz and Zarpentine (F&Z)

At least relevant to the concerns I've just raised (in sections 1.2-1.4), the only study that S&K cite that is set up so that it can really test the kind of intuitions utilized in the case for contextualism is one of the studies that does not even purport to be testing contextualism: F&Z. As I've already mentioned, F&Z present themselves as testing anti-intellectualism, especially Stanley's version of anti-intellectualism. But in the process of doing so, they ran a survey on Stanley's version of my bank cases, giving Stanley's descriptions of the cases, "Low Stakes" (LS) and "High Stakes" (HS), verbatim to the survey takers. But Stanley's versions of those two cases are just like mine with respect to the issues I've been discussing here: Stanley has both a high stakes situation and the mentioning of the relevant error possibility built into his High Stakes case, and his Low Stakes case features a lack of any mention of those error possibilities along with low stakes. And Stanley has his speaker, Hannah, positively ascribe knowledge in Low Stakes, while she denies knowledge in High Stakes. And about these cases, F&Z ask the respondents about the truth-values of the claims made by Hannah within the stories, rather than directly asking the survey-takers whether a character in the stories knows.

About LS, F&Z asked respondents to rank, on a scale of 1 to 7 (with 1 indicating "Strongly Agree," 4 indicating "Neutral," and 7 indicating "Strongly Disagree"), the extent to which they agreed with this: "When Hannah says, 'I know the bank will be open tomorrow,' what she says is true." About HS, they asked the same thing about this statement: "When Hannah says, 'I don't know that the bank will be open tomorrow,' what she says is true." Basically, what they found was that speakers were on average extremely neutral about the truth values of both of these assertions. In each case, the mean response was a bit less than 4, very slightly toward the agreement end of their scale: The mean response about HS was 3.74, and about LS was 3.68. But these are so

close to neutral that when you flip the HS score – as F&Z did, since it concerned a denial of knowledge as compared with LS’s positive affirmation of knowledge – the resulting 4.26 is still close enough to LS’s 3.68 that the difference between them is statistically insignificant.

As I’ve indicated, F&Z did not purport to be testing any case for contextualism. Nevertheless, these results (unique among the studies S&K discuss) are *prima facie* trouble for that case.

1.6 X-Phi and survey problems

Much of what flies under the banner of “Experimental Philosophy” (“X-Phi”) consists in checking philosophers’ claims about what is intuitive about this or that against the results of surveys given to ordinary people. The X-Phi work that we, following S&K, have been considering here is of course very much in that vein.

However, there are pitfalls to relying uncritically on survey data, and a general problem with X-Phi, at least as it has so far been practiced, is that it at least mostly seems to have been done without due regard to what is known about survey methodology.¹⁴ In my estimation (though readers should keep in mind that I am a very interested party here), such general problems concerning survey methodology by themselves are pressing enough in the case of the literature we have been looking at to make it unwise for participants in the ongoing philosophical debate over knowledge attributions to significantly alter their (our) debate in reaction to these results – though we should be open to new, better empirical results having a legitimate impact on the debate. But that is a case better made by others, far more expert in survey methodology than I am.¹⁵ What I have been raising are problems of a different sort for this experimental work: problems that seem to stem from misunderstandings of the philosophical view (contextualism)

¹⁴ I have in mind here the kind of general issues about survey methodology that have been prominently discussed by Norbert Schwarz; see Schwarz (1995), (1996), and Schwarz et al. (1991). Simon Cullen presses this criticism against X-Phi (though he does not discuss the works we have been looking at) in his (forthcoming). He argues:

Experimental philosophers can effectively equate survey responses with intuitions only by ignoring the established social and cognitive science literature on survey methodology. In my review of the growing body of experimental philosophy literature I have found just one reference to a serious discussion of survey methodology.* The conclusion to be drawn is that, despite their pretensions, many experimental philosophers have given no serious thought to methodology. This not only undermines their claim to be doing science, as we shall see, it often leaves the philosophical significance of their findings unclear.

In a note that is attached to the point in the above marked *, Cullen reports:

Searching all of the electronic scholarly resources accessible from Melbourne University—which include all major philosophy, psychology, social and cognitive science journals—for the terms [“experimental philosophy” “survey methodology”] turned up only two results, neither of which contained a discussion of survey methodology. Searching for [Schwarz “Experimental philosophy”] yielded two additional results: Goldman and Pust (1998) who cite Schwarz (1995) in their *defence* of intuitions; and Doris et al. (2007) who cite Schwarz (1996).

¹⁵ My methodologically-based pessimism about the bearing of the experimental work we’ve been considering on the philosophical debate is based on intriguing unpublished work, critical of the literature we have been considering, by Jennifer Nagel, who is indeed far more expert in survey methodology than I am.

being tested and that thereby result in problematic assumptions about what the view predicts.

Before moving on in the next section to consider a problem S&K found with the surveys we have been discussing, I will here (1) quickly point out a potential connection between the worries I have been raising and some general methodological worries and (2) briefly discuss the potential bearing of methodological worries on F&Z, the experimental study that raises *prima facie* problems for my case for contextualism.

I have been insisting, among other things, that if you really want to use surveys to test the case for contextualism (at least in anything close to its best form, and in any form I would endorse), then (for reasons I gave in section 1.2) it will not do to directly ask those taking the surveys whether characters in the stories know things to be the case; you must instead have the characters in the stories make claims involving “know(s)” and ask the survey takers about the truth values of those claims. And (for reasons I gave in section 1.3) you really have to make the key claim in one of the stories (HIGH) be a denial, rather than an affirmation, of knowledge. And it may well be that following these insistences of mine can make it difficult to design effective surveys.

One kind of problem that can arise in administering surveys involves shallow processing. S&K explain:

It seems that linguistic input that is not in focus is sometimes only partially processed. For instance, Erickson and Mattson (1981) asked people: “How many animals of each type did Moses put on the ark?” The typical answer they got was two. Few people seemed to notice that the question asks about *Moses*, who was presumably not involved with boarding any animals onto arks.¹⁶

When I ask my readers to consider whether the claims made in my bank cases are true, I count on them to note that I’m asking about the truth values of the key claims being within the stories, and to be sensitive in their responses to the fact that I’m not directly asking them whether characters in the story know things. Based on responses that I’ve received (when I presented the cases in talks to philosophical audiences, and also written discussions, often over e-mail, with those who have read the cases) that go beyond simple one-word responses, but explain my audience’s reasons for their answers, my audiences have had little problem with this. However, my audience has been largely professional philosophers, and otherwise students, with the students usually having had enough experience with philosophical examples to have some idea why they are being asked these strange things, and so what to focus on. But anyone who, like me, has taken a survey when you didn’t have any good feeling for why you were being asked the questions directed at you and so didn’t know what to focus on should be able to appreciate how lost some ordinary person, just being asked about these strange cases on

¹⁶ S&K (forthcoming): 26. S&K note that this possibility was pointed out to them by Jennifer Nagel and they cite Sanford and Sturt (2002) and Sanford et al (2006) as sources on the phenomenon.

some survey, might be.¹⁷ Asked about one of the bank cases, such a survey taker might well suppose that the purpose behind the questions is to find out how much confidence respondents tend to have in the stability of banking hours, and, as apparently can so very easily happen in such surveys, might well be shallow processing right over subtle differences that are of vital importance to the philosophical issues under discussion. I would have thought this problem could be taken care of by putting some preamble like this at the start of the survey: “We are interested in the conditions under which claims about knowledge are either true or false. In connection with that, we’d like to ask you what you think about the following case.” However, I’ve been warned by someone who knows what tends to happen in surveys that I’m being quite overly optimistic in assuming this would solve the problem. In addition to problems brought on by respondents not having a good feel for why they’re being asked these strange questions, and so not knowing what to focus on, there are issues of motivation. When you read about survey results, it may be helpful to actually picture the situation in which those surveys were taken (if you’re told; and if you’re not told, recognize the following as possibilities): for example, students standing at a table at a student center quickly filling out a survey in exchange for a piece of candy, or shoppers at a mall quickly answering a question or two to get a coupon for a Big Mac from a local McDonald’s. Knowing what kinds of questions are likely to get useful responses in such settings is not obvious to the uninitiated, and I have been warned that giving my bank cases as they appear in my philosophical writings to survey takers is likely to run into problems.¹⁸

In connection with that, it is worth recalling that Z&F’s results on Stanley’s version of the bank cases, the results that are *prima facie* problematic to the case for contextualism, feature mean scores that are extremely close to what is labeled on their scale of possible answers as “NEUTRAL” (4 on their scale of 1 to 7). This *could* mean that their respondents are right on top of things (the philosophically relevant details of the case), but just think the matter of the truth values of the claims they are evaluating is a close call. But it could also be that many of their respondents were marking the neutral

¹⁷ In my case, my initially frustrating experience was with an on-line survey I was asked by e-mail to fill out by a former employer. At the beginning, I had no idea why I was being asked the questions on the survey. It seemed to be a fairly long survey (both with many questions, and also in terms of many of the individual questions being a bit complicated), so I didn’t want to mull over each question for very long. I didn’t feel very good about the fairly rushed answers I found myself giving toward the beginning of the survey, and found myself often giving fairly “neutral” answers, (midway between “strongly agree” and “strongly disagree”), largely because I didn’t have a good feel for what the questions were “getting at.” A bit into the survey, I figured out that the institution was interested in reviewing various of their practices, with an eye toward changing them if need be. I suddenly felt much better about the answers I was giving, even when answering fairly quickly, and I was no longer opting for neutral answers out of a sense of confusion. A good hypothesis as to why this change took place is that once I knew the purpose behind the questions, I was better able to focus on and more adequately process the important elements of the questions I was being asked. I remember wanting to – but frustratingly being unable to – go back and redo the early portions of the survey.

¹⁸ This warning came from Jennifer Nagel. Some of the potential problems may come from there being a negation inside the claim being evaluated in HIGH

middle of the scale as a default answer they are giving largely because (perhaps in a way that’s analogous to the confusion I felt while taking the survey described in note 17), they didn’t know what to make of the strange question they were being asked. (This could perhaps be dubbed the “WTF?! neutral response.”)

1.7 S&K’s success at making possibilities salient

S&K believe the results of the studies we’ve been discussing threaten both non-contrastivist forms of contextualism, including my contextualist views, and also the view they defend, contrastivism. (We will consider what these views are and how they differ from each other in the next section.) According to the tables S&K present, at this point in the inquiry, the situation stands as it presented in all but the bottom row of the below:

Table 1

	Sallience for the Ascriber	Stakes for the Subject
Predictions of Contrastivism, according to S&K (p. 14)	Yes	No
Predictions of Non-Contrastivist Contextualism, according to S&K (p. 27)	Yes	No
Results of previous studies, according to S&K (p. 14)	No	No
Corrected results, according to S&K (pp. 21-22)	Yes	No

The top row names various features of cases (and I’ve followed S&K’s wording exactly on these headings of the columns), and the subsequent rows specify what S&K think are the predictions of views or the results of surveys on the matter of whether these factors have or will have an impact on judgments about knowledge. The “previous studies” of row 3 are the studies we have been discussing here. So, with respect to these previous studies, contrastivism and non-contrastivist (or standard) contextualism would seem to be in the same boat, with both predicting that “Sallience for the Ascriber” will have an impact, but “Stakes for the Subject” will not. And the boat they both seem to be in

appears a bit leaky, because, according to S&K, these theories' predictions have not been born out by the studies we've been looking at, since these previous studies suggest that neither factor has a significant impact.

However, as the bottom row shows, S&K believe that they can flip the problematic "No" that the previous studies suggested for "Salience for the Ascriber" to "Yes," bringing the empirical results into line with what they take to be the predictions of the two views under consideration. S&K think that they can "provide an empirically supported explanation of what went wrong in these studies, and show that there is a real salience effect when the studies are done properly" (p. 14). The problem with the earlier studies, according to S&K, is that they "presented the possibility of the bank changing its hours in too abstract and pallid a manner to make it salient" (p. 20). So S&K altered the HIGH version of the bank case they gave to their respondents, having the speaker who raises the possibility that the bank has changed its hours do so in a way that makes that possibility "concrete and vivid, by using a personal anecdote invested with emotional force" (p. 20). They gave "participants in the control condition" this bank case:

Hannah and Sarah are driving home on a Friday afternoon. They plan to stop at the bank to deposit their paychecks. As they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons.

Hannah says, "I was at the bank two weeks before on a Saturday morning, and it was open. So this is a bank that is open on Saturdays. We can just leave now and deposit our paychecks tomorrow morning."

Sarah replies, "Ok, that sounds good. Let's go on Saturday." (pp. 20-21)

"Participants in the salient contrast condition" were given a story exactly like that given to the control group, except for Sarah's reply at the end, which was changed to:

Sarah replies, "Well, banks do change their hours sometimes. My brother Leon once got into trouble when the bank changed hours on him and closed on Saturday. How frustrating! Just imagine driving here tomorrow and finding the door locked." (p. 21)

S&K asked respondents about the degree to which they agreed with the statement that Hannah knew that the bank would be open on Saturday, and they report:

Strikingly, when the error possibility was made more vivid in this way, the results *did* reveal a salience effect, just as all the contextualist models—including the contrastivist model—predict. Participants in the default condition tended to agree with the knowledge attribution (mean rating: 5.54 out of 7), while participants in the salient alternative condition tended to disagree (mean rating: 3.05 out of 7). (p. 21)

S&K think their new results supplant the those of the old surveys so that we should now take the empirical testing to have shown that "Salience for the Ascriber" of error possibilities does have an impact on what respondents will take to be known or not. Though the earlier studies were following the presentations in the philosophical literature in how they had the possibility presented, S&K think that philosophers who have been

persuaded by that literature have been able to pick up on the relevant features of the cases without any special prodding, while non-philosophical respondents may have to be hit over the head more forcefully with error possibilities before they respond properly:

One nice consequence of finding a salience effect is that we need no longer dismiss the intuitions of the epistemologists as theory driven or otherwise defective. For it does seem clear that a wide range of epistemologists, with various theoretical predilections, agreed about the bank cases. We speculate that the epistemologists simply needed less prodding to view the error possibility as salient. (p. 21)

S&K seem to think of respondents' reactions to the earlier, pallid bank cases to represent some kind of performance error – though they don't use that term. But they do indicate that philosophers' reactions to those old cases are not in any way defective, and seem to suggest it's ordinary non-philosophers who have defective reactions to such cases. This reaction is in line with the reservations about x-phi surveys we discussed above in section 1.6.

So at least as they present their new results, S&K seem to think that non-contrastivist contextualism, like their own contrastivism, is vindicated, at least against the old X-Phi studies we've been considering. However, I must note that since S&K don't have their characters claim and deny "knowledge" and then ask their respondents about the truth values of these claims, but instead directly ask their respondents whether Hannah knows, I don't think as they do that any clear predictions of contextualism are vindicated by their survey (for reasons we looked at above in section 1.2). But, of course, this is tied up with some of my reasons for not being too worried about those old studies in the first place. And if S&K are right about the problem they posit for those old surveys, that is further reason, beyond what we've seen so far, to not be so worried about them.

I should also point out that if S&K's interpretation of how their change to the HIGH case works is far from obvious. They claim that what they have done is increase the salience of the relevant error possibility – that the bank's hours have changed. But look again at their version of HIGH. It seems it could just as easily be thought that what they're pounding home to the survey takers is thoughts about the practical consequences or stakes of being wrong. Sarah tells Hannah about poor Leon and the frustration he had to endure when he was wrong about the bank's hours, and then, turning to Hannah's and her own situation, adds, "Just imagine driving here tomorrow and finding the doors locked." Is this not at least hinting that the stakes may be quite high and/or encouraging respondents to focus on the matter of the practical consequences or stakes of being wrong? To be sure, in neither of S&K's cases are participants told what the stakes actually are: The "control" or "default" story does not contain a statement to the effect that "it is not especially important in this case that [the paychecks] are deposited right away," as my original LOW case did (DeRose 1992: 913), and those given the "salient contrast condition" aren't told what, if anything is at stake. (This is in keeping with

S&K's liking the idea of isolating factors, which, for reasons we saw in section 1.4, above, I think can be unwise for our current purposes.) So, that Sarah would carry on as she does in S&K's beefed up HIGH case might well indicate to the survey takers, who up to that point have been told nothing about what stakes are involved, that the stakes are high somehow. At any rate, it's far from clear that what is being forcefully brought to the attention of survey-takers is the particular remote error scenario (the bank's having changed its hours in the last two weeks) as opposed to the practical consequences of being wrong. But in their charts, S&K take their contrastivist view to predict that stakes have no effect on responses, so an account on which their addition to the example affects responses by calling attention to the stakes involved would hurt their position.

2. S&K's Empirical Case for Contrastivism over Standard Contextualism

2.1 Contrastivism vs. standard (non-contrastivist) contextualism

S&K also conduct three new surveys which they argue empirically vindicate contrastivism, both in general and also specifically over non-contrastivist contextualism. S&K explain that on the contrastivist view knowledge is

a three-place relation between a subject, a proposition (the fact), and a contrast proposition (the foil). All knowledge takes the form: *s* knows that *p* rather than *q*.

The contrastivist view can be seen as direct implementation of a classic relevant alternatives view of knowledge (c.f. Austin 1946). The relevant alternatives view holds that whether *s* knows that *p* depends on which alternatives to *p* are relevant. On the contrastivist view, these relevant alternatives show up as an additional third argument of the knowledge relation. (p. 13)

On contrastivism, then, when someone says something of the form "S knows that *p*," what proposition they are expressing by that claim depends on what the (implicit) foil is. It may happen that S knows that *p* rather than *q* (K_{spq}), but S does not know that *p* rather than *r* ($\sim K_{spr}$). It can happen then that one speaker, A1, can truthfully assert, "S knows that *p*," while another speaker, A2 (perhaps engaged in entirely different conversation, as opposed to speaking with A1), can truthfully assert "S does not know that *p*," even though they're both talking about the same S and the same *p* at the same time, if in A1's context, *q* is the implicit foil supplied by context, while *r* is the foil supplied by A2's context. So, this is a form of contextualism, as S&K acknowledge (p. 13).

But what, then, do *non*-contrastivist contextualists hold? As contextualists, we also hold that it can happen that one speaker, A1, can truthfully assert, "S knows that *p*," while another speaker, A2 (who is perhaps not talking to A1) can truthfully assert, of the same S and the same *p* at the same time, "S does not know that *p*," and we agree with the contrastivist that this happens because A2 is not denying the same proposition that A1 is

affirming. However, as non-contrastivists, we do not think the proper analysis of this situation is one on which there is an implicit argument place for a foil proposition in “S knows that p.” On my view, what context instead supplies is an epistemic standard, and what proposition is expressed by a given knowledge attribution depends on what standard is selected by the speaker’s context.

2.2 *The structure of S&K’s argument: their tables and their basic “jewel thief vignette”*

Again pulling various of S&K’s tables into one, but this time adding a new column that wasn’t in our table 1, here’s how S&K see things:

Table 2

	Contrast of the Ascriber	Salience for the Ascriber	Stakes for the Subject
Predictions of Contrastivism, according to S&K (p. 14)	Yes	Yes	No
Predictions of Non-Contrastivist Contextualism, according to S&K (p. 27)	No	Yes	No
Results of previous studies, according to S&K (p. 14)	?	No	No
Corrected results, according to S&K (pp. 21-22)	Yes	Yes	No

Again, the top row names various features of cases (and I’ve followed S&K’s wording exactly on these headings of the columns), and the subsequent rows specify how, in their tables, S&K present what they think are the predictions of views or the results of surveys on the matter of whether the feature in question has an impact on speaker’s judgments about knowledge.

Our new column, not in Table 1, is “Contrast of the Ascriber.” As the question mark in that column indicates in the row for the previous studies, this is something S&K think the earlier studies didn’t test, but S&K’s new studies do test. And as the rows for the predictions of the views make clear, it’s this new column that is crucial, at least on

S&K's views, to the contest between contrastivism and non-contrastivist contextualism, because it's on "Contrast of the Ascriber" that the two views issue different predictions. And as the bottom row indicates, S&K thinks it's contrastivism's prediction that proves correct.

Unfortunately, I don't feel I understand what that crucial new column, "Contrast of the Ascriber," is supposed to be registering or how it's supposed to differ from the old "Salience for the Ascriber" column, much less why my theory (S&K clearly have me pegged as a non-contrastivist contextualist) is supposed to be making different predictions on these two matters.

We will proceed by looking at the three studies S&K did that they believe support a verdict of "Yes" for the impact of "Contrast for the Ascriber," and thereby empirically vindicate contrastivism, both in general and also specifically over non-contrastivist contextualism (which they think predicts "No" for that factor). These studies all involve S&K's below "jewel thief vignette":

Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter forgot to wear gloves. He also forgot about the security camera.

Today, Mary the detective has been called to the scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information. (p. 15)

S&K comment:

Notice that this story is constructed in such a way that Mary's evidence allows her to eliminate certain possibilities but not others. Specifically, her evidence allows her to rule out the possibility that (a) someone other than Peter was the thief, but it does not allow her to rule out the possibility that (b) Peter stole something other than rubies. Given all this, one might well pose the question: "Does Mary know that Peter stole the rubies?" (p. 15)

And indeed, S&K go on to ask in various ways what Mary knows. But before going further, it is important to note that it is very far from clear that S&K's respondents would interpret their vignette in such a way that Mary cannot rule out the possibility that the jewels were something other than rubies, as S&K claim. Much depends on how one reads the last sentence in the story. That Mary has "no further information" is hard to interpret, in large part because it really cannot be taken in full seriousness. Mary was informed there was a theft, we're told. But presumably, she knows much more than we're being explicitly told here. For instance, she seems to know the location of the theft – after all, she seems to be in the right place as she conducts her investigation (or are we

to suppose that happened purely by accident?). So, it's easy to suppose – despite the description saying she has “no further information” – that she actually was given some very basic information, like where the theft took place, that we're not being told about, perhaps by whoever sent her over to investigate, and what she has “no more” of, besides what we've been told, is information that is somehow non-basic—the results of her investigation that go beyond just the basics that she would have been told beforehand. But what does that include? Was Mary also told, in a basic way, roughly when the theft took place (e.g., earlier today)? Hard to say. Is it clear that she wasn't told that it was rubies that were stolen? Is it clear that that nobody in the story knows that, or that nobody who does know that has bothered to mention to Mary what was stolen? Or that being told that it was rubies by someone who knows wouldn't put her in a position to rule out other possibilities? I wouldn't think those matters are clear at all. I don't know about S&K's respondents, but based on the vignette they gave their respondents, *I* was quite confused about just how much information Mary has been given and how, and about which possibilities she is or is not in a position to rule out.

2.3 S&K's “rather than” survey

As we discussed in 2.1, the difference between contrastivism and non-contrastivist contextualism concerns what is supplied by context to determine what proposition is expressed by a given simple knowledge attribution of the form “S knows that p.” Contrastivism says it's a foil proposition that's supplied. Other forms of contextualism will posit that context supplies something else. I think context supplies an epistemic standard. (My reasons for preferring my account over the contrastivist one are explained at DeRose 2009: 34-41.) But in the case of either theory, there are also uses of “know(s)” in which what the theory says is often supplied by context is instead made explicit. A speaker can explicitly say, for instance, “Mary knows that it's Peter, rather than Paul, who stole the rubies.” Here, the speaker is being explicit about a matter – what the foil is – that, according to contrastivism, is left implicit and is supplied by context in the case of simple, unadorned knowledge attributions. (In this case, the foil is specified to be *Paul stole the rubies*.) But one can also say, “Mary knows, by even the most demanding possible standards for knowledge, that Peter stole the rubies,” making explicit a matter that, according to *me*, is left implicit but supplied by context in the case of simple, unadorned knowledge attributions.

The first test S&K cite as telling in favor of contrastivism is one in which they asked one group of respondents whether they agreed with this adorned (“thief contrast”) statement:

Mary now knows that Peter rather than anyone else stole the rubies

but asked another group of respondents whether they agreed with this differently adorned (“jewel contrast”) claim:

Mary now knows that Peter stole the rubies rather than anything else

Quite unsurprisingly, given the evidence the story presents Mary as having, respondents were more inclined to agree with the first of the two statements above. S&K appear to take this as arguing for a “Yes” answer to the question of the impact of the “Contrast of the Ascriber,” and to therefore argue for contrastivism, and for contrastivism over non-contrastivist contextualism.

I find this extremely puzzling. Imagine that some experimental philosophers set out to empirically test *my* brand of contextualism,¹⁹ and they did so by constructing a vignette in which in fact Peter stole the rubies, and in which Mary is described (in concrete terms) as having an amount and quality of evidence that, for most people, would make it a close call whether or not to rule that Mary knows that Peter stole the rubies,²⁰ and they then asked one group of respondents whether they agreed with this (“high standards”) statement:

Mary knows, by even the most demanding possible standards for knowledge, that Peter stole the rubies

and they asked a second group whether they agreed with this (“low standards”) statement:

Mary knows, at least by the most relaxed possible standards for knowledge, that Peter stole the rubies

Imagine that they got truly strikingly different results on these two statements, with those given the low standards statement being much more inclined to agree with the statement shown to them. When I received their excited e-mail, with the subject line “ascriber standards impact observed; standards contextualism vindicated empirically!”, I’m afraid that I’d have to burst the investigators’ bubble in my reply, saying that I find their results quite unsurprising, and, more importantly, of little help to my cause. The dispute my theory is a party to is over how to understand simple, unadorned knowledge attributions (and denials of knowledge), in which whatever (if anything) context contributes to determining what proposition get expressed by such claims is indeed left implicit. What any such theory predicts about how subjects will react to statements in which such matters are made explicit is an extremely complicated matter. Just because other theories

¹⁹ I guess it’s a sign that I’m a true armchair philosopher at heart that I’m here discussing the bearing of the results of an experimental philosophy survey—that I’ve only imagined has been conducted!

²⁰ I presume that when S&K set out to test for “contrast for the ascriber” impact, they tried to construct the case to be such as to make it very likely that such impact could be observed, if the factor in question ever has an impact. Similarly, then, my imagined experimental philosophers would test for “standards for the ascriber” impact by constructing a case designed to most likely register such an impact. I’m thinking this would be done by giving the possible knower in the example a level of evidence that would, for most speakers, put her in the intuitive “grey zone” between knowing and not knowing, and then asking them whether they agree with the statements I have above in the text.

don't single out epistemic standards as context's contribution does not mean that they couldn't explain the results our imagined theorists obtained. In our imagined case, and also in S&K's case, it's easy to come up with explanations rival theories could utilize to explain the results obtained.

I'm not saying the results of such a survey would be completely irrelevant to the evaluation of my brand of contextualism—which is why I say they would be of “little,” rather than of “no” help. If no significant difference or even if only a quite muted difference were observed, I imagine that would be not only quite surprising but also worrisome to me (if the survey were conducted well). That instead there was (we're imagining) a truly striking difference in the direction I'd expect would at least be a sign that I don't have *that* to worry about.

I *will* say, first, that if the impact observed in our imagined study that purported to test my theory were as muted as the impact S&K report in their survey, I actually would be a bit worried about my theory. I mean, when you're able to put important (and at least apparently content-altering) clauses like that explicitly into the sentences respondents are asked to evaluate in order to get different evaluations, shouldn't you be able to get a quite striking difference indeed? S&K write:

Participants rated these sentences on a scale from 1 ('disagree') to 7 ('agree'). The results showed a striking difference in responses between the two conditions. Participants in the thief contrast condition tended to agree with the knowledge ascription they received (mean rating: 4.6), while participants in the jewel contrast condition tended to disagree with the knowledge ascription they received (mean rating: 3.1). This difference was statistically significant. (p. 16)

Statistically significant it is.²¹ But if those were the results of our imagined survey testing for the impact of “ascriber standards,” I'd find these two means too close together, and too close to the middle of the scale, for me to find any vindication, and, in fact, as I've indicated, I'd actually be more worried than cheered by such results. But how strong an impact should we expect, and why? How much is enough; how much is too little? I've said nothing useful about that, I admit, except that it's a very complicated matter: I'm very much just flying by the seat of my pants here. But that's kinda my point: So are S&K. (Or at least they're not sharing any reasons they might have for thinking that the statistically significant but still fairly muted differences they observed are sufficient for vindicating their view.) So I counter their apparent sense that these results are vindicating to their contrastivism with my contrary vague and unsupported sense that they actually count against contrastivism,²² mostly to make the point that neither reported

²¹ The details of S&K's survey are at S&K, p. 31, note 16.

²² I don't want to make too much of these results counting *against* contrastivism. Though I presume (see note 20, above) that S&K designed their story to make an impact to be most likely to show up, there are problems with their story, like the one I mention at the end of section 2.2, above, that may well explain why their results are so muted.

sense here is worth much without some good explanation of why a theory about the meaning of non-adorned knowledge claims should be taken to predict an impact of a certain magnitude on respondents' reactions to explicitly adorned knowledge ascribing claims.

Second, if one is for whatever reason inclined to test theories about simple unadorned knowledge attributions by checking how speakers responded to adorned knowledge claims in the fashion of S&K here, judgments about whether and to what extent the results of such surveys are vindicating or worrying to the theories about simple knowledge attributions purportedly being tested should, at the very least, await results concerning how rival theories fare on analogous tests.

Most importantly, though, I close by reiterating generally that much more needs to be said about why and exactly how theories about unadorned knowledge attributions should lead us to expect various types of results in terms of responses to adorned knowledge attributions.

2.4 S&K's "knowledge-wh" survey

Similar problems affect S&K's second, "knowledge-wh" study. They again gave their "jewel thief vignette" to two groups of participants, and asked those in "the thief contrast condition" about the extent that they agreed or disagreed with this sentence:

Mary knows who stole the rubies,

while asking the other group – those in the "jewel contrast condition" – about the extent to which they agreed or disagreed with this sentence:

Mary knows what Peter stole.

Unsurprisingly, given the evidence the vignette describes Mary as having, the first group of participants were more agreeable to the sentence they were given. S&K report:

Again, the results showed a marked difference between conditions. Participants in the thief contrast condition tended to agree with the statement that Mary knew who stole the rubies (mean rating: 4.91), while participants in the jewel contrast condition tended to disagree with the statement that Mary knew what Peter stole (mean rating: 2.62). This difference was statistically significant. (pp. 17-18)

It's quite unclear to me why and how these results about how speakers react to these "knowledge-wh" constructions bear on the issue of whether to accept the contrastivist's account of the meaning of the quite different "knowledge that" claims. Here's S&K's explanation (given before they tell us the results of their survey) for why these results would (and so do, once we see what they are) support contrastivism:

If knowledge were simply a relation between a person and a proposition, all these differences in contrasts should have no effect on the truth values of the knowledge ascriptions themselves. Ultimately, the whole issue would just come down to whether or not Mary stood in the right sort of relation to the proposition that Peter stole the rubies. If

she did, both knowledge ascriptions would be true; if not, both would be false. On the other hand, if contrastivism is correct, these differences in contrast can actually make a difference to the truth values of the knowledge ascriptions. Since the two knowledge ascriptions discussed here pick out two different contrasts, it can turn out that one of them is true and the other false. (p. 17)

But again, the dispute between contrastivism and its rivals that we are addressing here is over the meaning of “knowledge that” sentences. And it is only on *some* of theories that are rivals to contrastivism – the non-contextualist (a.k.a. invariantist) ones – that “knowledge that” claims, like “Mary knows that Peter stole the rubies,” express the same relation between Mary and the proposition that Peter stole the rubies in whatever context they are uttered (so long as they are about the same Mary, the same Peter, and the same rubies, of course). And none of the various theories about the meaning of “knowledge that” claims (including contrastivism, insofar as we construe it as just a theory about such sentences) by themselves tell us anything about the meaning of the very different “knowledge-wh” sentences like the ones S&K had the participants in this study evaluate, and none of them in any obvious way issue any predictions about the acceptability of “knowledge-wh” claims. To take these results as having a bearing on the issue of the meaning of “knowledge that” claims would seem to require at least much more explanation than S&K provide here.

2.5 “Epistemic Aspectism”

Before moving to S&K’s final survey – and the one that has the best chance of having a real bearing on the debate between contrastivism and its rivals – it may be helpful for me to say a few words about what we might call “epistemic aspectism.” As I’m using that term, it refers to a series of theses about constructions involving “know(s)” that seem closely related to the sentence form “S knows that P.” Assume that Peter stole the rubies: That it was indeed Peter who did the deed, and that it was indeed rubies that he did the deed to, and that stealing (rather than begging or borrowing or etc.) is indeed what he did. An epistemic aspectist of one variety or other will hold that the claims in one or the other of the following groups need not all have the same truth value (even though they all contain references to the same Mary, the same Peter, and the same rubies, and are all about the same time period):

another

- 1a. Mary knows who stole the rubies
- 1b. Mary knows what Peter stole
- 1c. Mary knows what Peter did to the rubies

- 2a. Mary knows that it is Peter who stole the rubies
- 2b. Mary knows that it is rubies that Peter stole
- 2c. Mary knows that it is stealing that Peter did to the rubies

- 3a. Mary knows that Peter, rather than someone or something else, stole the rubies
- 3b. Mary knows that Peter stole the rubies, rather than something else
- 3c. Mary knows that Peter stole the rubies, rather than doing something else to them

- 4a. Mary knows that Peter stole the rubies
- 4a. Mary knows that Peter stole the rubies
- 4a. Mary knows that Peter stole the rubies

- 5a. Mary knows that Peter stole the rubies [said in a context in which the focus is on the question of who stole the rubies]
- 5b. Mary knows that Peter stole the rubies [said in a context in which the focus is on the question of what Peter stole]
- 5c. Mary knows that Peter stole the rubies [said in a context in which the focus is on the question of what Peter did to the rubies]

So, for instance, a type-1 epistemic aspectist (or a “knowledge-wh aspectist”) will hold that it can happen that 1a is true, while 1b is false. And such an epistemic aspectist will likely think that a (perhaps suitably filled in) scenario like S&K’s jewel thief vignette is just the kind of situation in which such a divergence in truth values will be realized.

Each of the above groups of sentences employ some device or other to focus on some aspect of “Peter stole the rubies”: the a sentence in each group focuses attention on who did the deed, the b sentence on what the deed was done to, and the c sentence on what deed was done. Aspectists think that the epistemic requirements for the truth of these knowledge-ascribing sentences will also somehow be focused on the aspect of Peter’s stealing of the rubies that the sentence is focused on. I suppose a very strong form of what we should still call “epistemic aspectism” would hold that, for instance, it needn’t even actually be rubies that were stolen for “Mary knows who stole the rubies” to be true. (Though in that case, we can’t stipulate that in all cases it was indeed rubies that were stolen when we set up the sentences.) But more interesting (because more plausible) forms of epistemic aspectism hold that, while it must indeed be Peter who did the deed, and rubies that the deed was done to, and stealing that was done, for any of the sentences in the relevant group to be true, the further epistemic requirements for the various sentences to be true attach only to (in the case of what we might call “moderate epistemic aspectism”), or more strongly to (in the case of what we might call “weak epistemic aspectism”) the aspect of “Peter stole the rubies” that is somehow or other being focused on.

So, for instance, an evidentialist, who thinks that, beyond the truth of what is believed, questions of knowledge turn entirely on whether the believer has sufficient evidence for her beliefs, will hold that, beyond Mary’s believing that Peter stole the rubies and its being true that Peter stole the rubies, 1a requires for its truth only that Mary

have sufficient evidence for the who-done-it aspect of “Peter stole the rubies,” if this theorist is a *moderate* type-1 aspectist. If she is instead a *weak* type-1 aspectist, she will hold that Mary must have some level of evidence for all the aspects of “Peter stole the rubies” (the who-done-it, the what-was-done, and the what-it-was-done-to aspects), but will hold that a higher level of evidence is required for the focus aspect. Or consider a “safety” theorist, who thinks that, beyond Mary’s believing that and it being true that Peter stole the rubies, questions of what Mary knows here turn on whether there are possible worlds that are too close to the actual world in which Mary holds this belief but it is false. (Very similar remarks would apply to my own, closely related “double safety” account of knowledge, on which knowledge is also killed if there is a too-close possible world in which Mary disbelieves that Peter stole the rubies even though he did steal them.) A safety theorist who is a *moderate* type-1 aspectist will hold that the only too-close non-actual worlds that make 1a go false (in the actual world) are ones in which Mary’s belief is wrong because it is not really Peter who did the deed; non-actual possible worlds in which she goes wrong about a non-focus aspect of “Peter stole the rubies” don’t kill her knowledge of “who stole the rubies” in the actual world. A *weak* type-1 aspectist could instead hold that while 1a can go false because Mary believes that Peter stole the rubies in a too-close non-actual world in which it’s something else that Peter stole, worlds in which Mary’s belief goes false due to a failure of a non-focus aspect of “Peter stole the rubies” have to be closer to be knowledge-killers than worlds in which Mary’s belief is false because she’s going wrong about the focus aspect have to be.

When we get to type-5 aspectism, we are at a view with at least considerable overlap with contrastivism, as we’ve been construing it here: as a view about the meaning of sentences of the form “S knows that P,” the very sentence form involved in group 5. Usually, contrastivists go beyond mere type-5 aspectism. (Whether to take it as definitional of “contrastivism” to go beyond this, or whether to take mere type-5 aspectism to be a limiting case of contrastivism is a terminological matter that isn’t clear to me.) “Contrastivists,” I take it, will tend to hold what we might call “full-blown contrastivism,” according to which “Mary knows that Peter stole the rubies” can have different truth values in different contexts, even when describing the same situation *and even in different contexts which are focused on the same aspect* of “Peter stole the rubies”. For instance, it can happen that “Mary knows that Peter, rather than Paul, or Mary, or any of the other well-known thieves in the region, stole the rubies” is true, while “Mary knows that Peter, rather than a super-advanced robot, made to look and act just like Peter, stole the rubies” is false (while both sentences are describing the same situation).²³ So, the simple sentence, “Mary knows that Peter stole the rubies” expresses

²³ At least this divergence in truth values can happen according to typical contrastivist verdicts about the truth-values of the relevant “rather than” knowledge-ascribing sentences, which seem to be guided by something like what we can call this ‘contrastive-sensitivity’ condition: If S would have believed that p even if q had been the case, then it is not the case that S knows that p rather than q. Whether these verdicts are actually correct is a very tricky issue. For discussion, see DeRose 2009: 39-40, n. 37. I am here

a truth about this situation when the implicit foil is the one about the other ordinary thieves, while that same simple sentence expresses a false proposition about the same situation when the implicit foil proposition concerns the super-advanced robot, even though the two contexts in question are both ones focused on the “who-done-it” [or I guess this should be broadened to “who-*or-what*-done-it”] aspect of “Peter stole the rubies.”

Still, type-5 aspectism is at least quite close to contrastivism as we’ve been thinking of it, and all these different forms of aspectivism are likely quite close to the various facets of “contrastivism,” insofar as it is construed as a more general theory about the meaning of various constructions involving “know(s).” And this may explain S&K’s interest in surveys which seem to be directly testing various forms of aspectism while they are seeking to defend contrastivism.

So a couple of points are in order here. First, while the two surveys of S&K’s that we have looked at so far are relevant to evaluations of the relevant forms of aspectism (their “rather than” being relevant to type-3 aspectism and their “knowledge-wh” survey being relevant to type-1 aspectism), it’s hard to say exactly what the impact of these results should be on the debates over the relevant versions of aspectism, and it would be premature at best to say that these surveys vindicate even the forms of aspectism they seem set up to test. For my part, ever since I encountered a well known paper by Fred Dretske²⁴ that prompted me to consider the views, I’ve been intrigued by various forms of what we’re here calling epistemic aspectism. And I’ve always assumed (without the backing of careful empirical testing) that all five forms of aspectism we’re discussing here (as well as a couple of other forms I haven’t brought up here) have some intuitive plausibility behind them, of the rough type that S&K’s surveys seem to be set up to pick up on. A main stumbling block to accepting these versions of aspectism has been the thought that the intuitions that drive it invariably seem to be in conflict with other intuitions that seem about as strong, or often stronger, than they are. So, for instance, about some situations like S&K’s jewel thief vignette, 1a, “Mary knows who stole the rubies,” can to some extent seem true, while 1b, “Mary knows what Peter stole,” can to some extent seem false, I suppose. But, about those same situations, these conditionals also each seem correct, at least to me (for the first conditional, remember that we’re taking the situation to be one in which Peter in fact stole the rubies):

accepting such verdicts for the sake of argument. To reject them would be to alter contrastivism’s applications quite drastically from what contrastivists take them to be.

²⁴ Dretske 1972, which I didn’t encounter until quite a few years after it came out; see especially pp. 435-437 for Dretske’s discussion of knowledge. Dretske didn’t consider all the forms of epistemic aspectism we’re discussing here; his focus was on what we’re here calling type-4 aspectism. But once one sees Dretske’s discussion, it doesn’t take much imagination to start wondering about other forms of epistemic aspectism.

If Mary knows who stole the rubies, then Mary knows that Peter stole the rubies

If Mary knows that Peter stole the rubies, then Mary knows what Peter stole

Indeed, to me at least, the judgments that the conditionals directly above are true seem if anything to be more intuitively appealing than is the judgments that 1b is false (though this of course depends on how exactly the story is told) – and they seem to *at least* seriously compete with the judgment that 1a is true in terms of intuitive plausibility. So what we have on our hands are four judgments, all of which have some individual plausibility behind them, but which at least seem to be such that they can't all be true: that 1a is true, that 1b is false, that the first conditional above is correct, and that the second conditional is correct. The issue has always seemed to me to be one of finding the best resolution of such conflicts. (Similar conflicts arise in evaluating other types of aspectism.) Key to the construction of resolutions to these conflicts is not just measuring the degree to which the various claims are intuitively plausible, but also, vitally, accounts of the meanings of and connections between the various relevant forms of knowledge claims, by which the intuitive plausibility of the various claims can be explained – and sometimes explained away. How do the results of S&K's "knowledge-wh" survey bear on all this? That's hard for me to say. Even supposing they've got things set up to provide a good test of how intuitive sentences like 1a and 1b are in the relevant situations, this doesn't tell us much about the comparative matter of how these supposed intuitions stack up against the other intuitions that they seem to be in competition with. Though I'm again admittedly flying by the seat of my pants here, I am inclined to think that the results that S&K obtained are weak enough that their effect on me is to make me somewhat *less* inclined to accept type-1 aspectism than I previously was – though this is based in large part on what my prior expectations were. (I thought the positive intuitive support for type-1 aspectism was stronger than what seems to be indicated by S&K's knowledge-wh survey.)

Second, there's no obvious path from type-1 and/or type-3 aspectism to type-5 aspectism, much less to full-blown contrastivism.

2.6 S&K's "context" survey

Fortunately, S&K's third survey does seem designed to test type-5 aspectism, which is the type most directly related to contrastivism as we're here construing it. Here's their description of the study and its results:

All participants received the very same jewel thief vignette used in the two prior studies. Participants in the thief-contrast condition then read:

Everyone is now asking the big question: Who stole the rubies? The news reporter is about to write a story about Mary. He is wondering if Mary now knows who stole the rubies. He writes: "Mary now knows that Peter stole the rubies."

Please tell us whether you agree or disagree with the news reporter's claim, "Mary now knows that Peter stole the rubies."

Participants in the thief-contrast condition instead read:

Everyone is now asking the big question: What did Peter steal? The news reporter is about to write a story about Mary. He is wondering if Mary now knows what Peter stole. He writes: "Mary now knows that Peter stole the rubies."

Please tell us whether you agree or disagree with the news reporter's claim, "Mary now knows that Peter stole the rubies."

The results showed the usual difference between conditions. Participants tended to agree with the knowledge attribution in the thief-contrast condition (mean rating: 5.24 out of 7) but to disagree with the knowledge attribution in the jewel-contrast condition (mean rating: 2.97 out of 7). This difference was statistically significant. (pp. 18-19)

Before we get to what I take to be the main problem with this survey, we should pause to notice two important things. Note first the big difference between this study and S&K's other jewel thief studies. Those other studies concerned only their original jewel thief vignette (which I reproduce below), which did not contain any character inside the story making any attribution or denial of knowledge. When S&K then asked the survey participants to what extent they agreed with some statement to the effect that Mary knew something, they were not evaluating a claim made within the story, but were presumably evaluating whether the claim in the (rather weird) survey-taking context was correct *about* the story. The attributers, then, were the survey participants themselves, and the "attributor's context" was the rather odd context of taking a survey. Now in this last study S&K do have a knowledge claim being made within the story and seem to be asking their survey participants to evaluate that inside-the-story claim. This changes things drastically. Now the attributor is the reporter inside the story, and so when we speak of the "salience for the attributor" or the "contrast of the attributor," we would seem to be speaking of the reporter's context of writing his article.

Second, note that even if these results supported what the surveys seem to be set up to be most directly testing the positive intuitions behind–type-5 aspectism–this would not establish what I’m here calling full-blown contrastivism. For what it’s worth, I’ve always thought I could buy into the various types of aspectism or not, so far as my standards contextualism went. I was somewhat inclined *not* to buy accept these forms of aspectism, for reasons of the type I indicate in the second-to-last paragraph of the above section, but the decision of whether to go for the various forms of aspectism always seemed to me fairly independent from whether to go in for my style of contextualism. This even when we got to type-5 aspectism. It’s only when one crosses into what I’m here calling “full-blown contrastivism” that I thought of the contrastivist approach as being in serious competition with standards-contextualism. (And at that point, I’ve always opted for standards contextualism, again for the reasons now articulated at DeRose 2009: 34-41.)

But how strongly do these results support even type-5 aspectism? The main problem here is the danger that – indeed, I’d say the great likelihood that – the participants in the two groups are construing the situation in relevantly very different ways.

Both groups first read the original jewel thief vignette, so let’s look at that again:

Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter forgot to wear gloves. He also forgot about the security camera.

Today, Mary the detective has been called to the scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and identified Peter’s fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information. (p. 15)

Now, recall our discussion of this from section 2.2. Though the last sentence above stipulates that Mary “has no further information,” that statement will be hard for readers to interpret, in large part because they cannot take it in full seriousness. Can or will readers safely assume that beyond having been told that a theft has taken place and having the evidence described above Mary literally has *no* further information? Can they safely assume that she wasn’t told where the theft took place? If so, how did she end up, as she seems to have done, at the right place to conduct her investigation? Or will readers assume that that was pure luck? And if she was told where the theft took place, is it safe to assume that she was not told roughly when it took place – that whoever told her a theft took place didn’t tell her, for instance, that it took place “earlier today”? Or, to get to the issue crucial to our discussion, is it safe to assume that she wasn’t told what was stolen? So, as I would interpret the above vignette, I wouldn’t take it that Mary has literally *no* information beyond what I’m being told, but rather that she was told *some* basic

information (including where the theft took place), and I'd interpret that last sentence as saying she has no further information beyond such quick basics that she would have been told by whoever sent her over to investigate. It would be unclear whether this basic information would include that rubies had been stolen, as opposed to the vaguer information that some jewels or other had been stolen, or, I suppose (though this would appear far less likely) that nothing at all was said about what was stolen.

Now look at the new vignette endings that our two groups were given. What's told to those in the "thief contrast condition" would encourage them to construe the evidential situation such that it's completely unproblematic that it was rubies that were stolen. Remember that these folks don't read the alternative, "jewel contrast" ending; they just get the common, old vignette beginning plus their own ending. So they've read nothing that gives them any indication that the evidential situation is such as to render it in any way problematic that it was rubies that were stolen. Nobody seems to be questioning that, and indeed, everyone around Mary, including the reporter, seems to be happily presupposing that it was rubies that were stolen. By marked contrast, what's told to those in the "jewel contrast condition" will encourage them to think that the evidential situation is such as to render it problematic indeed that it was rubies that were stolen. Why would everyone be asking about what was stolen if the evidence left no question about that?

If the two groups of survey participants are likely construing what the evidential situation within the story is like quite differently in a very relevant way, that blocks the inference from their different reactions to a conclusion that they are interpreting the reporter's knowledge claim as having different contents. Their different evaluations can very well indeed be due to different construals of the situation the reporter is describing, rather than different meanings for the description itself.²⁵

²⁵ Thanks to the organizers of the 2010 Oberlin Philosophy Colloquium, where I delivered this paper, and to the participants at the conference, especially to Patrick Rysiew, my commentator. Thanks to Jennifer Nagel and to Jonathan Schaffer for very helpful comments on an earlier draft.

REFERENCES

- Austin, J. L. 1946. "Other Minds," *Proceedings of the Aristotelian Society* 20: 149-87.
- Beebe, James and Wesley Buckwalter *forthcoming*. "The Epistemic Side-Effect Effect," *Mind & Language*.
- Buckwalter, Wesley *forthcoming*. "Knowledge Isn't Closed on Saturdays," *Review of Philosophy and Psychology*.
- Cohen, Stewart 1988. "How to be a Fallibilist," *Philosophical Perspectives* 2: 91-123.
- 1999. "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13: 57-89.
- Cullen, Simon *forthcoming*. "Survey-Driven Romanticism," *European Review of Philosophy*.
- DeRose, Keith 1992. "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52: 913-29.
- 1995. "Solving the Skeptical Problem," *Philosophical Review* 104: 1-52.
- 2005. 'The Ordinary Language Basis for Contextualism and the New Invariantism', *Philosophical Quarterly* 55: 172-98.
- 2007. Review of Stanley (2005), *Mind* 116: 486-9.
- 2009. *The Case for Contextualism*, Oxford: Oxford University Press.
- Doris, John M., Joshua Knobe, and Robert L. Woolfolk 2007. "Variantism about responsibility," *Philosophical Perspectives* 21:183-214.
- Dretske, Fred I. 1972. "Contrastive Statements," *Philosophical Review* 81: 411-437.
- Feltz, Adam and Chris Zarpentine *forthcoming*. "Do You Know More When it Matters Less?" *Philosophical Psychology*.
- Goldman, Alvin and Joel Pust 1998. "Philosophical Theory and Intuitional Evidence," in W. Ramsey and M. DePaul, eds., *Rethinking Intuition*, Totowa: Rowman & Littlefield.
- Lewis, David 1979. 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* 8: 339-59.
- May, Joshua, Walter Sinnott-Armstrong, Jay G. Hull, and Aaron Zimmerman *forthcoming*. "Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study," *Review of Philosophy and Psychology*.
- Neta, Ram and Mark Phelan *manuscript*. "Evidence that Stakes Don't Matter for Evidence."

Schaffer, Jonathan and Joshua Knobe forthcoming. "Contrastive Knowledge Surveyed," *Noûs*. Page references are to a prepublication version of the paper available on Wiley's web site for *Noûs*, accessed 6 January 2011.

Schwarz, Norbert 1995. "What respondents learn from questionnaires: The survey interview and the logic of conversation," *International Statistical Review/Revue Internationale de Statistique* 63: 153–168.

Schwarz, Norbert 1996. *Cognition and communication: Judgmental biases, research methods, and the logic of conversation*. Mahwah: Erlbaum.

Schwarz, Norbert, Fritz Strack, and Hans-Peter Mai 1991. Assimilation and contrast effects in part-whole question sequences: A conversational logic analysis. *The Public Opinion Quarterly* 55: 3–23.

Stanley, Jason 2005. *Knowledge and Practical Interests*, Oxford: Oxford University Press.