

Shifting Targets and Disagreements

Robin McKenna

University of Geneva

rbnmckenna@gmail.com

This is a preprint of an article whose final and definitive form will be published in the Australasian Journal of Philosophy; the Australasian Journal of Philosophy is available online at: <http://www.tandf.co.uk/journals/>.

Abstract Many have rejected contextualism about ‘knows’ because the view runs into trouble with intra- and inter-contextual disagreement reports. My aim in this paper is to show that this is a mistake. First, I outline four desiderata for a contextualist solution to the problem. Second, I argue that two extant solutions to the problem fail to satisfy the desiderata. Third, I develop an alternative solution which satisfies the four desiderata. The basic idea, put roughly, is that ‘knowledge’ ascriptions serve the function of recommending good informants, and disagreement over ‘knowledge’ ascriptions is a matter of conflicting recommendations. The upshot is that intra- and inter-contextual disagreement reports provide little - if any - reason to reject contextualism.

1. Introductory Remarks

Epistemic contextualists - henceforth, ‘contextualists’ - think that the extension of the expression ‘knows’ depends on and varies with the context of utterance. Just as sentences containing indexical expressions (‘I’, ‘today’, etc.) and gradable adjectives (‘tall’, ‘flat’, etc.) express different propositions in different contexts, sentences containing the expression ‘knows’ express different propositions in different contexts. I’ll work with a version of the view on which a sentence of the form ‘S knows that p’ (where ‘S’ is a person and ‘p’ a sentence) is true in a context C iff S’s epistemic position with respect to p meets the epistemic standard operative in C [see DeRose 2009: Ch. 1].

While contextualism is usually defended by appeal to certain patterns of linguistic data, it runs into trouble with other patterns. Consider these cases [see DeRose 2009: Ch. 1]:

Setup: Ted, Dougal and Jack are priests. Jack should perform a service this afternoon, but he’s notoriously unreliable. Ted has a long-standing bet with Dougal: If Jack performs the service Dougal wins £10, but if he doesn’t Ted wins £10. When Ted and Dougal go to the church they hear a service in progress, and in the vestry they see Jack’s jacket on the wall. Sister Imelda, who is aware of Ted and Dougal’s bet, winks at Dougal and says ‘Looks like you win today Dougal.’

Satisfied that Jack is performing the service, Ted goes to the pub and Dougal heads home. Assume that Jack performed the service.

PUB: Ted tells Tom that he lost the bet. Tom asks Ted 'How do you know Jack performed the service?' and Ted responds 'Jack never goes anywhere without his jacket, and we saw it in the vestry. Oh, and Sister Imelda seemed to think he was there. So Dougal and I both know.' Satisfied, Tom buys Ted a drink.

STRINGENT POLLY: While Tom is at the bar Polly, who has been eavesdropping, corners Ted. Polly, who notoriously holds demanding epistemic standards, asks Ted why he said that both he and Dougal 'knew' that Jack performed the service when neither of them actually saw him. Ted defends himself, but Polly insists 'No you're wrong. You don't know that Jack performed the service.'

POLICE STATION: A serious crime has been committed and Dougal has been summoned to talk to the police as part of their investigations. They ask Dougal for a statement concerning Jack's whereabouts and Dougal says 'I didn't actually see Jack at the church. We did see his jacket, and he always has that with him. Sister Imelda seemed to think he was there, but she didn't say that she had seen him. So I'm afraid I don't know if he was there or not.' The police ask if Ted 'knows' and Dougal says 'He doesn't know either.' When told that Ted was overheard saying that he did 'know' Dougal says 'No he's wrong. He doesn't know that, for the reason I just gave you.'

On the one hand, these cases provide support for contextualism. Intuitively, Ted, Polly and Dougal all spoke appropriately, and a simple way of explaining this would be by saying that the extension of the expression 'knows' varies with the context: Ted and Dougal 'know' by the epistemic standard operative in PUB, but they don't 'know' by the epistemic standards operative in STRINGENT POLLY and POLICE STATION. On the other hand, they raise two problems for contextualism. First, if Ted satisfies the epistemic standard operative in PUB, it's unclear why Polly and Dougal use the 'disagreement markers' 'No' and 'You're/He's wrong' to negatively assess his 'knowledge' ascription. What, exactly, was Ted wrong about? I'll call this the 'assessment problem.' Second, if 'knows' has different extensions in different contexts then what Ted said doesn't contradict what Polly or Dougal said. But, *prima facie*, Ted, Polly and Dougal

disagree about whether Ted knows that Jack performed the service. What, exactly, do they disagree about? I'll call this the 'disagreement problem.'¹

Many think that these problems give us good reason to reject contextualism; see Hawthorne [2004: 101-104], MacFarlane [2005], Richard [2004] and Stanley [2004: 52-56]. My aim in this paper is to show that this is a mistake. First, I outline four desiderata for a contextualist solution to the assessment and disagreement problems (§2). Second, I argue that two extant solutions fail to satisfy the desiderata (§3). Third, I develop an alternative solution which, at the very least, promises to satisfy all four desiderata (§4). I conclude that the assessment and disagreement problems provide little reason to reject contextualism.

2. Disagreement and Assessment: Four Desiderata

The main aim of this section is to motivate four desiderata for a contextualist solution to the assessment and disagreement problems (§2.2). First, though, I explain why one might think that these problems provide good reason to reject contextualism (§2.1).

2.1. The Dialectic

If we reject contextualism, we need an alternative semantics for 'knowledge' ascriptions. Two alternatives are particularly relevant here. First, one might think that, while 'knowledge' ascriptions express the same proposition in every context, their truth-values vary with the epistemic standard operative in the context in which they are assessed. As applied to our cases, the idea is that Ted's 'knowledge' ascription expresses the proposition *Dougal and Ted know that Jack performed the service*, but that proposition is true as assessed by Ted in PUB and false as assessed by Polly and Dougal in STRINGENT POLLY and POLICE STATION.² I'll call this sort of view 'relativist' [see MacFarlane 2005; Richard 2004]. Second, one might think both that 'knowledge' ascriptions express the same proposition in every context and that their truth-values don't vary with the context of assessment. I'll call this sort of view 'invariantist' [see Brown 2006; Hawthorne 2004: Ch. 4; Stanley 2005: Ch. 5]. The taxonomy is crude, but it will suffice for my purposes.

¹ This assumes a simplistic version of contextualism on which the extension of 'knows' is tied to the epistemic standards of the speaker. While more complex versions are available [see DeRose 2009: Ch. 4] I focus on the simple version because, first, my solution to the assessment and disagreement problems is neutral between simple and complex versions and, second, more complex versions have their own problems [see McKenna forthcoming].

² I'm using italics for propositions here.

Here's a simple argument that, because of the assessment and disagreement problems, we should reject contextualism in favour of invariantism or relativism. Consider this plausible generalisation concerning the behaviour of context-sensitive expressions:

DISAGREES THAT: Let u_1 be a sincere utterance of sentence S by agent A in context C_1 and u_2 a sincere utterance of 'not- S ' by B in C_2 . If from a third context C_3 they cannot be correctly reported by 'A and B disagree about whether S ', then S is semantically context-sensitive.

Meanwhile, if from a third context C_3 they can be correctly reported by 'A and B disagree about whether S ', that is evidence that S is semantically invariant across C_1 , C_2 , and C_3 [Cappelen & Hawthorne 2009: 54-5].

A wide range of context-sensitive expressions behave as DISAGREES THAT predicts. For example, 'left' and 'right' are context-sensitive: things are to the left or the right relative to an orientation, not *simpliciter*. As one would expect, if A says 'The restaurant is on the left' in context C_1 and B says 'The restaurant is on the right' in context C_2 , where C_1 and C_2 differ in orientation, I can't correctly report that A and B disagree about whether the restaurant is on the left in my context C_3 . While DISAGREES THAT perhaps admits of exceptions, it provides good albeit defeasible evidence for or against contextualism about various expressions. If, as invariantists and relativists have argued, one can accommodate the patterns of linguistic data that support contextualism by adopting an invariantist or relativist semantics for 'knowledge' ascriptions [see Brown 2006; Richard 2004; Stanley 2005: Chs. 5 & 6] then the assessment and disagreement problems provide good reason to reject contextualism.

2.2. The Desiderata

What would contextualists need to do in order to deal with the assessment and disagreement problems? In what follows I motivate four desiderata for a contextualist solution.

First, contextualism is motivated by intuitions about the appropriateness of 'knowledge' ascriptions and denials in cases like PUB, STRINGENT POLLY and POLICE STATION. If intuitions about the appropriateness of 'knowledge' ascriptions and denials are part of the motivation for contextualism then, absent some independently plausible explanation why, contextualists can't just deny other intuitions, such as the intuition that Polly and Dougal's uses of the disagreement markers 'No you're/he's wrong' are perfectly appropriate, or the intuition that Ted, Polly and Dougal disagree. Here's MacFarlane:

“One would also need to explain why the data that seems to support [contextualism] ... should be taken so seriously, when the data about third-party assessments, retraction, and disputes are just thrown away. There is no clear reason to favor the “positive” data in this way” [2009: 149].

A contextualist solution to the problems must not discard these intuitions. This is the first desideratum, which I’ll call COMPATIBILITY.

Second, a natural suggestion is to deal with STRINGENT POLLY and POLICE STATION by appeal to conversational implicatures, presuppositions or another pragmatic feature, the idea being that Ted, Polly and Dougal disagree over the pragmatic feature. I discuss this sort of solution below (§3.2; §4). If the contextualist appeals to pragmatics, she must provide independent motivation for thinking that ‘knowledge’ ascriptions generate the relevant pragmatic feature. This is the second desideratum, which I’ll call NOT AD HOC.

Third, in the STRINGENT POLLY and POLICE STATION cases Polly and Dougal use the disagreement markers ‘No you’re/he’s wrong.’ But here are some other equally felicitous disagreement markers:

- (1) Polly/Dougal: What Ted said/thinks/believes is wrong. He doesn’t know that Jack performed the service.

Any contextualist solution to the problems has to deal with the full range of felicitous disagreement markers. This is the third desideratum, which I’ll call RESPECT THE DATA.

Fourth, in the STRINGENT POLLY and POLICE STATION cases Ted, Polly and Dougal seem to disagree. The contextualist has to identify the sense of disagreement in play, and she has to appeal to that sense in explaining the use of the various disagreement markers. This is the fourth desideratum, which I’ll call LOCATE DISAGREEMENT. *Prima facie*, one would expect that contextualists will have a hard time satisfying this desideratum: Ted, Polly and Dougal’s ‘knowledge’ ascriptions and denials may be consistent, so what can they disagree about? But this is too quick. First, it assumes that disagreement must always concern whatever proposition is literally expressed, and therefore Ted, Polly and Dougal can’t disagree because the propositions literally expressed by their ascriptions and denials may be consistent. Consider this case:

- (2) Fred and Wilma are discussing who stole Wilma’s money.
Fred: I was told that Barney stole the money.
Wilma: No that’s wrong. Barney didn’t steal the money.

The most plausible interpretation of (2) is that Wilma disagrees with Fred about whether Barney stole the money, not about whether Fred was told that Barney stole the money. In support of this note that Wilma's objection is that Barney didn't steal the money, not that Fred wasn't told that Barney stole the money. This suggests that disagreement can sometimes concern whatever proposition is most salient, not whatever proposition is literally expressed, where the most salient proposition depends on a range of conversational factors (the purpose of the conversation, etc.). Given the purpose of the conversation, the most salient proposition in (2) is the proposition *Barney stole the money*.

Second, it assumes that disagreement must always concern doxastic attitudes, and therefore that Ted, Polly and Dougal can't disagree because they don't have incompatible beliefs. Consider another case:

(3) Frank: I love haggis.

Stephen: No way! You're wrong. Haggis is disgusting.

I think it's natural to say that Frank and Stephen disagree. In support of this note that Stephen's rejoinder - 'No way! You're wrong' - seems felicitous, but it's hard to see how that could be if Frank and Stephen didn't, in some sense, disagree. This suggests that disagreement can sometimes concern non-doxastic attitudes, and there can be genuine cases of 'attitudinal disagreement'.³ The suggestion makes sense of (3) because Frank and Stephen clearly have incompatible non-doxastic attitudes towards haggis (Frank could not coherently adopt Stephen's attitude without abandoning his own, and vice versa). In §4 I'll say a little more about attitudinal disagreement, and I'll suggest that the contextualist can appeal to a sort of attitudinal disagreement to deal with the disagreement problem.

3. Two Unsatisfactory Solutions

I discuss two extant contextualist solutions to our problems. I argue that neither comes close to satisfying the four desiderata. While this section is largely negative, the discussion prepares the ground for §4.

3.1. Semantic Blindness

Many contextualists have tried to deal with the assessment and disagreement problems by appealing to the observation that competent speakers of a language are sometimes blind to the

³ See Stevenson [1944] for the classic discussion of attitudinal disagreement. For a more recent discussion see Ridge [2013].

context-sensitivity of expressions in that language. The idea is that, because of this blindness, we mistakenly think that the speakers in cases like *STRINGENT POLLY* and *POLICE STATION* disagree, and the speakers in such cases mistakenly take themselves to disagree [see Cohen 1999: 77-9; DeRose 2009: Ch. 5]. Call this the ‘semantic blindness strategy.’ Note that the relevant blindness is of a particular sort, viz. blindness about when speakers disagree. I’ll call this ‘disagreement blindness’, and I’ll briefly return to it in §4.3.

The semantic blindness strategy is a non-starter. The strategy rejects the intuitions of competent speakers about when speakers disagree as a manifestation of some sort of blindness. But this prompts the question why the intuitions about the appropriateness of ‘knowledge’ ascriptions that are supposed to motivate contextualism aren’t also a manifestation of some sort of blindness. The objection should be familiar: The strategy rejects some of the data (violating *COMPATIBILITY* and *RESPECT THE DATA*), and in the process undercuts the original motivation for contextualism.

3.2. Presuppositions of Commonality

It’s perhaps plausible that ‘knowledge’ ascriptions usually carry the presupposition that the conversational participants are similar with respect to their epistemic standards. A presupposition is (roughly) a proposition that is part of the common ground, where a proposition is part of the common ground just in case all those in a conversation accept, it, all believe that all accept, and so on [see Stalnaker 2002]. The idea is that, in contexts where ‘knowledge’ ascriptions are made, it is usually part of the common ground that those involved in the conversation are similar with respect to their epistemic standards. Because this is usually part of the common ground, when we first encounter the *STRINGENT POLLY* case we get the impression that Ted and Polly disagree. If Ted and Polly were similar with respect to epistemic standards, then they would disagree in a fairly straightforward way: Ted would be saying that he knows relative to some epistemic standard, and Polly would be saying that Ted doesn’t know relative to a similar standard. But, *ex hypothesi*, Ted and Polly are not similar with respect to epistemic standards. So our initial impression is mistaken. Like the semantic blindness strategy, this is a sort of error-theory. But, unlike the semantic blindness strategy, the explanation isn’t that we’re blind to the context-sensitivity of the expression ‘knows’. Rather, the explanation is that ‘knowledge’ ascriptions usually carry a presupposition of commonality, and we assume that the *STRINGENT POLLY* case is no different. It’s entirely consistent with this that we’re aware,

or at least aware to some extent, that ‘knows’ is context-sensitive. Call this the ‘presupposition of commonality strategy.’⁴

The presuppositions of commonality strategy faces three objections. First, it violates RESPECT THE DATA. Presuppositions are normally understood in terms of what propositions the participants in a particular conversation all accept, all believe that they all accept, etc. So it’s entirely unclear how the defender of the presuppositions of commonality can deal with cases of disagreement across contexts, viz. cases like POLICE STATION. In my view, this is reason enough to reject the strategy. But, in case the reader is unmoved - perhaps because he or she is unsure about the POLICE STATION case - I’ll offer two further objections.

Second, the strategy has problems satisfying NOT AD HOC. *Prima facie*, this might sound surprising, given that the strategy appeals to a general feature of ‘knowledge’ ascriptions, viz. usually carrying presuppositions of commonality. But why think that ‘knowledge’ ascriptions have this general feature? A natural suggestion is that conversations usually run more smoothly if speakers use their terms to refer to the same things. For example, if Ailsa and Laurie are discussing how tall their friends are, there’s a presupposition to the effect that Ailsa and Laurie associate similar ‘height standards’ with their uses of the word ‘tall’. So it’s hardly surprising that ‘knowledge’ ascriptions usually carry the presupposition that the conversational participants are similar with respect to their epistemic standards.

However, this explanation makes faulty predictions. If Ailsa says that Raibeart is tall, whereas Laurie denies that Raibeart is tall, then Ailsa and Laurie seem to disagree. But imagine that it becomes clear that, while there’s a presupposition that Ailsa and Laurie associate the same standards with their uses of ‘tall’, the presupposition is false. At this point the appearance of disagreement vanishes. In contrast, imagine that it becomes clear that, while there’s perhaps a presupposition that Ted and Polly associate the same epistemic standards with their uses of ‘knows’, the presupposition is false. Even at this point the appearance of disagreement remains. Whatever presupposition explains the apparent disagreement between Ted and Polly, it isn’t the general one that speakers use their terms to refer to the same things.⁵

⁴ For this strategy, albeit applied to ‘funny’ rather than ‘knows’, see López de Sa [2008].

⁵ Objection: If it were pointed out to Polly (or Dougal) that because ‘knowledge’ ascriptions are context-sensitive what Ted asserted was true relative to the epistemic standards in his context, it would be strange if Polly (or Dougal) persisted in thinking that they disagree. Response: I agree that it would be strange for Polly to accept both that what Ted asserted was true relative to his

Third, the evidence suggests that ‘knowledge’ ascriptions don’t usually carry the presupposition of commonality. Consider three standard tests:

TEST ONE: If utterances of a sentence S carry the presupposition that p, it should make sense for someone who was previously unaware of p to respond to an utterance of S by saying ‘Hey, wait a minute. I didn’t know that p!’ [see von Stechow 2004; Yablo 2006]. Compare and contrast:

(4) Raibeart: I’m going to pick up my sister from the airport.

 Catriona: Hey, wait a minute. I didn’t know that you had a sister!

(5) Sorcha: I know that penguins eat fish.

 Morag: Hey wait a minute. I didn’t know that we were similar with respect to epistemic standards!

If ‘knowledge’ ascriptions carry the presupposition of commonality, Morag’s response to Sorcha should make sense on a variety of ways of filling in the context. But, unlike Catriona’s response to Raibeart, it doesn’t.

TEST TWO: If utterances of S carry the presupposition that p, it should be infelicitous to follow an utterance of S with ‘and what’s more, p’ [see von Stechow 2004; Yablo 2006]. Compare and contrast:

(6) #I realise that I forgot to pick my sister up. And what’s more, I forgot to pick my sister up.⁶

(7) I forgot to go to the airport. And what’s more, my sister is waiting for me there.

(8) I know that penguins eat fish. And what’s more, we’re similar with respect to epistemic standards.

Grice’s maxim of quantity requires one to make one’s conversational contributions as informative as required [1989: 26]. Because (6) violates Grice’s maxim (the second conjunct asserts what the first conjunct presupposes), it is infelicitous. In contrast, both (7) and (8) are felicitous (the second conjuncts provide new information).

standards and that they disagree. But the problem for the contextualist is that many have argued that, because Polly and Ted disagree, ‘knowledge’ ascriptions aren’t context-sensitive.

⁶ I’m using ‘#’ to indicate pragmatic infelicity.

TEST THREE: If I utter S and presuppose p, an immediate subsequent denial of p makes my collective utterance very puzzling [see Baker 2012]. Compare and contrast:

(9) #I'm going to pick up my sister from the airport, although I haven't got a sister.

(10) I know that penguins eat fish, although we're not similar with respect to epistemic standards.

The subsequent denial of what is presupposed makes (9) very puzzling. In contrast, even without filling in the context (10) sounds fine.

I conclude that there are good reasons for thinking that 'knowledge' ascriptions don't usually carry the presupposition of commonality. To claim otherwise flies in the face of the best available evidence.⁷

4. Shifting Targets

My aim in this section is to develop solutions to the assessment and disagreement problems that avoid the pitfalls of the two solutions just discussed (§4.1). I argue that these solutions promise to satisfy our four desiderata (§4.2) and I deal with three objections to the solutions (§4.3). The basic idea is that 'knowledge' ascriptions function pragmatically as recommendations.⁸ In rough outline, the solution to the assessment problem has it that the proposition literally expressed by a 'knowledge' ascription is only conversationally relevant insofar as that ascription serves its pragmatic function. If an ascription no longer serves its function another proposition becomes relevant instead, namely the proposition that the ascription would have expressed had it been uttered in the present context. It is this proposition that is 'targeted' by the various disagreement markers. In similarly rough outline, the solution to the disagreement problem has it that disagreement over 'knowledge' ascriptions is a matter of conflicting recommendations.

⁷ Objection: One might think that the presupposition of commonality is like the presupposition that we speak the same language. In most conversational exchanges this presupposition is in place, but it doesn't seem to pass the three tests. Response: While this might help with the third objection, it won't help with the first or second objections. Imagine that, while Ailsa and Laurie speak different languages, they think that they speak the same language. Ailsa and Laurie happen to utter strings of words that are not just meaningful but actually contradictory in both their respective languages. Consequently, Ailsa and Laurie think they disagree. But, upon finding out that they actually speak different languages, the apparent disagreement disappears. In contrast, once Ted and Polly find out that they have different epistemic standards, their apparent disagreement doesn't disappear.

⁸ For a similar idea for contextualism about 'ought' see Björnsson and Finlay [2010].

4.1. The Solution

Recall (2):

(2) Fred: I was told that Barney stole the money.

Wilma: No that's wrong. Barney didn't steal the money.

The thought is that, just as Fred and Wilma disagree about something other than the proposition Fred literally expressed, and just as Wilma assesses something other than that proposition, Ted, Polly and Dougal disagree about something other than the proposition literally expressed by Ted's 'knowledge' ascription and Polly and Dougal assess something other than that proposition. I take what Polly and Dougal assess first (§4.1.1) and what Ted, Polly and Dougal disagree about second (§4.1.2).

4.1.1. The Assessment Problem

Assessment and disagreement can sometimes concern the most salient proposition, where what proposition is most salient depends on the purposes and interests of the conversational participants. But what is the fundamental interest motivating our assessment of 'knowledge' ascriptions? To answer this question I'll sketch an account of the function of 'knowledge' ascriptions.

On one plausible account - based on Edward Craig [1990: Chs. 1-2] - the function of 'knowledge' ascriptions is to identify good informants.⁹ The thought is that humans need information about the world around them, some of which we can obtain ourselves (through perception, introspection and the like), but most of which we need to get from others. We therefore have a further need for a way of keeping track of those who have the required information. A natural suggestion is that we keep track via our ascriptions of 'knowledge': in saying that a subject S 'knows' that p I identify S as a good informant on the matter of p.

Three clarifications: First, while Craig embeds his account of the function of 'knowledge' ascriptions within a larger project concerning the analysis and nature of knowledge, one can go along with everything I'll say about the function of 'knowledge' ascriptions while remaining sceptical about Craig's larger project. The plausibility of his account of the function of 'knowledge' doesn't derive from the plausibility of the larger project (quite the reverse).

⁹ This account has a number of supporters in the contemporary literature; see, for instance, Fricker [2012]; Greco [2009]; Pritchard [2012].

Second, what do I mean by the ‘function’ of ‘knowledge’ ascriptions? An analogy will help here. It’s a commonplace that ‘good’ claims - claims of the form ‘x is good’ - have the function of commending. In saying ‘giving to charity is good’ I *inter alia* commend giving to charity. For familiar reasons, not least the Frege-Geach problem, one shouldn’t conclude that commending is central to the semantics of ‘good’ claims. One might instead conclude that commending is part of the pragmatics of ‘good’ claims. The basic idea is that the function of ‘good’ claims is a matter of what we do with them, viz. commend. To say this is, of course, not to deny that ‘good’ claims also express truth-evaluable propositions. Similarly, I’m suggesting that ‘knowledge’ ascriptions have the function of identifying good informants. For much the same reasons, I conclude that this is part of the pragmatics of ‘knowledge’ ascriptions, not the semantics. Again, the basic idea is that we use ‘knowledge’ ascriptions to do something, viz. identify good informants, and to say this is not to deny that ‘knowledge’ ascriptions also express truth-evaluable propositions.

Third, if the function of ‘knowledge’ ascriptions is to identify good informants then one might think that, because whether a subject is a good informant depends on the context, whether that subject ‘knows’ must also depend on the context. But, as we’ve just seen, that would be too quick. It’s far more plausible that the function of ‘knowledge’ is a matter of pragmatics rather than semantics.

Craig’s account of the function of ‘knowledge’ ascriptions suggests an answer to our original question. The fundamental interest that motivates our assessments of ‘knowledge’ ascriptions is with identifying good informants. This fundamental interest leads to a derived interest in the truth or falsity of particular ‘knowledge’ ascriptions. In any context, a subject who qualifies as a good informant on the matter of p will also qualify as ‘knowing’ that p.¹⁰ But ‘knowledge’ ascriptions are only conversationally relevant insofar as they serve to identify good informants. Combining this with the contextualist thought that the truth-conditions of ‘knowledge’ ascriptions depend on and vary with the contextually determined epistemic standards, we get the result that, if the epistemic standards shift, ‘knowledge’ ascriptions made relative to earlier epistemic standards become conversationally irrelevant. Such ascriptions no longer serve to identify good informants.

If we apply this to our cases we can identify what Polly and Dougal assess in STRINGENT POLLY and POLICE STATION. Ted has offered himself as a good informant on the matter of whether

¹⁰ Note that someone might qualify as ‘knowing’ that p yet not be a good informant on the matter of p, perhaps because they’re unable to share their information.

Jack performed the service. But in STRINGENT POLLY the epistemic standards are unusually high (Polly generally adopts stringent standards). Consequently, the truth or falsity of Ted's 'knowledge' self-ascription is conversationally irrelevant (it doesn't serve to identify Ted as a good informant). However, there is a proposition that would serve to identify Ted as a good informant in this context, and that's the proposition *Ted knows, by Polly's standards, that Jack performed the service*. Consequently, this is the salient proposition for assessment. But, because Ted doesn't know by Polly's standards, that proposition is false, and Polly uses the disagreement markers 'No you're wrong' to negatively assess it. Similar remarks apply to POLICE STATION: the standards are also unusually high (Dougal is in a very high stakes situation), so the truth or falsity of Ted's 'knowledge' self-ascription is conversationally irrelevant (it doesn't serve to identify Ted as a good informant). But there is a proposition that would serve to identify Ted as a good informant, viz. the proposition *Ted knows, by Dougal's standards, that Jack performed the service*. But, because Ted doesn't know by Dougal's standards, that proposition is false, and Dougal uses the disagreement markers 'No you're wrong' to negatively assess it.

Assessments of 'knowledge' ascriptions target the propositions that those ascriptions would have expressed had they been uttered in the context in which they're being assessed. The 'target' of assessment - what I'll call the 'target content' - shifts with the context. I'll call this the 'shifting target strategy.' It's important to distinguish the shifting target strategy from views on which the context in which a speech-act is assessed can affect what proposition is expressed [see Cappelen 2008]. On the shifting target strategy it is the target of assessment - the target content - that depends on the context in which it is assessed, not the proposition actually expressed.

4.1.2. The Disagreement Problem

One might still wonder what Ted, Polly and Dougal disagree about. The proposal is that 'knowledge' ascriptions pragmatically express recommendations. If I ascribe 'knowledge' that p to a subject S I thereby identify S as a good informant on the matter of p. Or, put differently, I recommend S as a good informant on the matter of p. Applying this to our cases, Ted's 'knowledge' self-ascription serves to recommend himself as a good informant on the matter of Jack's whereabouts whereas Polly and Dougal's respective 'knowledge' denials don't serve to recommend Ted as a good informant on that matter. Consequently, Ted, Polly and Dougal have incompatible non-doxastic attitudes. Ted recommends something whereas Polly and Dougal don't recommend that thing. So the relevant sort of disagreement in STRINGENT POLLY and POLICE STATION is attitudinal. This solution to the disagreement problem could fairly be called 'quasi-expressivist.' The solution borrows something from the expressivist - the idea that

disagreement is a matter of incompatible non-doxastic attitudes - but she thinks of this as a matter of pragmatics rather than semantics (hence ‘quasi’-expressivist). So I’ll call this the ‘quasi-expressivist strategy.’

I’d like to pause to deal with two worries about the quasi-expressivist strategy. First, I seem to be assuming that recommendations are just non-doxastic attitudes. If I recommend something I adopt a particular sort of pro-attitude towards that thing, but it’s plausible that there’s a lot more to recommending than this. This worry can be easily dealt with. While I agree that recommendations involve more than non-doxastic attitudes, it can hardly be denied that recommendations involve such attitudes, and I’m primarily interested in the non-doxastic attitude aspect of recommendations.

Second, one might argue that, while it’s plausible that Ted recommends himself as a good informant on Jack’s whereabouts for the purposes of his bet with Dougal, it’s not plausible that Ted recommends himself as a good informant for the purposes of police investigations, or for someone like Polly who adopts stringent epistemic standards. But the quasi-expressivist strategy requires that Ted recommends himself as a good informant for all of these purposes, and more besides. I think this worry is more pressing. To address it I will start with some remarks about recommendations in general.

The worry is that the recommendation expressed by Ted’s ‘knowledge’ ascription has to be qualified or restricted in certain ways, and in particular to some but not all purposes. There are various ways in which one can qualify or otherwise restrict a recommendation (this list is not meant to be exhaustive, and the items on it aren’t meant to be mutually exclusive):

1. Qualification of the thing recommended; for example, I recommend that you try the *blue cheese* on the cheeseboard at a restaurant, as opposed to just recommending that you try the cheeseboard.
2. Qualification of the audience to whom that thing is recommended; for example, I recommend that you try the cheese, *unless you dislike cheese*.
3. Qualification of the circumstances in which that thing is recommended; for example, I recommend that you go for a walk later, *unless it starts to rain heavily*.

While recommendations can be qualified or otherwise restricted in these ways, they are often not qualified. Thus, rather than saying ‘I recommend that you go for a walk later, unless it starts to

rain heavily' I will often just say 'I recommend that you go for a walk later', even though I have no reason for thinking you should go for a walk in the rain. This is no doubt partly for practical reasons: even if one were aware of the ways in which some recommendation had to be qualified or restricted, it would often be tiresome and unnecessary to give the various qualifications and restrictions. More importantly, those issuing recommendations often have no idea whether they need to be qualified, or about how to go about making the relevant qualifications. Again, I might recommend that you go for a walk later, even though I am aware that there are certain circumstances in which you would be best off not going for a walk; for example, if you suffer from a hidden ailment that makes walking a bad idea. Because I have no idea whether those circumstances obtain, or even what those circumstances might be, I can hardly qualify my recommendation accordingly.

Another reason why one might not qualify or otherwise restrict a recommendation is that doing so has the potential to nullify the intended effect of the recommendation. We usually issue recommendations in order to bring about effects; my purpose in recommending that you go for a walk later is to bring it about that you go for a walk later. Imagine that neither you nor I have any reason to think that it will rain heavily later (it's summer, and the weather has generally been good). Saying 'I recommend that you go for a walk later, unless it starts to rain heavily' - as opposed to 'I recommend that you go for a walk later' - makes the possibility that it might rain heavily later salient, and so has the potential to discourage you from preparing to go for a walk ('I'll not go after all. It might rain, and I haven't got a jacket'). Often, if I want to bring about the intended effect, I am well advised to not qualify or restrict my recommendation.

While these aspects of our recommendation-issuing practice might look somewhat problematic - why do we issue unqualified or unrestricted recommendations so readily given our ignorance? - I think it is a fairly accurate description of the practice. For better or worse, when we issue recommendations we often don't restrict them to certain circumstances, or qualify the audience to whom we are making the recommendation.

Turning back to the recommendations expressed by 'knowledge' ascriptions, like recommendations in general they can be restricted or qualified in a range of ways, for instance:

1. Qualification of the subject matter; for example, I recommend Ailsa as a good informant *on the matter of when the train to London leaves* (rather than, say, when the train to Paris leaves).

2. Qualification of the purposes; for example, I recommend Ted as a good informant, *but not for the purposes of police investigations.*

I'm particularly interested in the second way, and I'll focus on it in what follows. Again, even though the recommendations expressed by 'knowledge' ascriptions can be restricted or qualified to some but not all purposes, there are a number of reasons why they will often not be so restricted.¹¹ First, there will often be practical reasons for not restricting them in this way. Consider Ted's situation in PUB. Ted is just having a chat with a friend, and in this sort of situation it would be tiresome and unnecessary to qualify or otherwise restrict his recommendation of himself as a good informant to some set of purposes. More importantly, while Ted may recognise that his recommendation might not be suited to all future purposes, he has little idea which purposes it is suited to, or which it isn't. Because he has little idea what those circumstances are, he can hardly qualify his recommendation accordingly. Faced with a choice between issuing a recommendation not restricted to some set of purposes and not issuing a recommendation at all, and given the point of his conversation with Tom (to explain why he lost the bet), Ted is best off opting for the former.

Another reason why one might not restrict a recommendation of some subject as a good informant to some set of purposes is that doing so has the potential to nullify the intended effect of the recommendation. Again, consider Ted's situation in PUB. Ted's aim in recommending himself as a good informant on the matter of whether Jack performed the service is to identify himself as someone with information that explains why Ted has lost his bet, and to get Tom to accept him as someone with information that explains why he lost. If Ted had restricted his recommendation by indicating that he would not be a good informant for certain purposes, that might well have prevented Tom from accepting Ted as someone with information that explains why Ted lost. Here, as is often the case, if Ted wants to achieve his aim he is well advised to not qualify or restrict his recommendation.

The quasi-expressivist strategy requires that Ted recommends himself as a good informant for a wide range of purposes, including for the purposes of police investigations, and for the purposes of someone like Polly, who generally adopts stringent epistemic standards. For the reasons just

¹¹ I don't deny that one might recommend some subject as a good informant on some matter for certain purposes but not for others, for instance by making the restriction to those purposes explicit. The aim in the main body of the text is to give a number of reasons for thinking that, absent any special reason to think a recommendation of a good informant is being restricted in this way, the recommendation is unrestricted.

given, I think it's plausible that Ted's recommendation isn't restricted to some but not all of these purposes. However, one might worry that, while the recommendation expressed by Ted's 'knowledge' ascription isn't restricted, it doesn't make sense for Polly or Dougal to express disagreement with that recommendation. But consider how recommendations work in general. Suppose Catriona recommends that Ailsa take a walk. If it starts to rain Ailsa might say 'Catriona recommended that I take a walk, but she was wrong - it's raining outside.' Even if we stipulate that Catriona wouldn't recommend Ailsa taking a walk in the rain, Ailsa's reply is felicitous, and it is natural to describe her as disagreeing with Catriona's recommendation. This suggests that, when a recommendation isn't restricted in so as to rule out certain circumstances or purposes, it makes sense to object that the recommendation doesn't hold in those circumstances. Consequently, the quasi-expressivist strategy is supported by general observations about how recommendations work.

This completes the presentation of my solutions to the assessment and disagreement problems. In the remainder of this paper I discuss how the solutions fair with the four desiderata (§4.2) and deal with three further objections (§4.3).

4.2. The Desiderata

In this sub-section I argue that my solutions to the assessment and disagreement problems satisfy the four desiderata outlined in §2.2. I'll take each desideratum in turn.

COMPATIBILITY: The solutions aims to explain both the intuition that Polly and Dougal's uses of the disagreement markers 'No you're wrong' are felicitous and the intuition that Ted, Polly and Dougal disagree. Consequently, the intuitions of competent speakers about these matters are taken into account and not discarded.

NOT AD HOC: My solutions are based on Craig's account of the function of 'knowledge' ascriptions. Insofar as that account is independently plausible, my solution to the assessment and disagreement problems satisfies NOT AD HOC. While I think the account is independently plausible, I can't defend it at length here. I'll just note that the account has many adherents in contemporary epistemology (recall fn. 9).

RESPECT THE DATA: On the shifting target strategy, assessments of 'knowledge' ascriptions target the propositions that those ascriptions would have expressed had they been uttered in the context in which they are assessed (what I called the 'target content'). So Polly assesses the proposition *Ted knows, by Polly's standards, that Jack performed the service* and Dougal assesses the

proposition *Ted knows, by Dougal's standards, that Jack performed the service*. Polly and Dougal could use the disagreement markers 'No you're wrong' to assess these propositions, or they could equally well use the disagreement markers 'What Ted said/believes/thinks is false' to assess them. In each case the disagreement marker is being used to assess the target content.

Consequently, the solutions RESPECT THE DATA.

LOCATE DISAGREEMENT: On the quasi-expressivist strategy, the relevant sort of disagreement is attitudinal. Consequently, the solution to the disagreement problem locates the relevant sense of disagreement.

4.3. The Objections

On the solutions proposed in §4.1 the target content in STRINGENT POLLY is the proposition *Ted knows, by Polly's standards, that Jack performed the service*, but Ted and Polly disagree about whether to recommend Ted as a good informant, not about whether that proposition is true. The solutions 'disconnect' assessment and disagreement: What Polly (or Dougal) assesses is one thing, and what Ted and Polly and (or Ted and Dougal) disagree about another. Contrast (2):

(2) Fred: I was told that Barney stole the money.

Wilma: No that's wrong. Barney didn't steal the money.

The target content here is the proposition *Barney stole the money*, and Fred and Wilma disagree about whether that proposition is true. This invites three objections, which I take in turn.

First objection: I have illegitimately traded on an analogy between (2) and STRINGENT POLLY.

This objection can be easily dealt with. I've only made use of (2) in order to show that disagreement and assessment can concern whatever proposition is most salient, which need not be the proposition that was literally expressed. My argument doesn't rely on an analogy between (2) and STRINGENT POLLY.

Second objection: The disconnect between assessment and disagreement is problematic for two reasons. First, on the shifting target strategy the proposition that is assessed need not be one of the propositions that the speaker expressed (or judged/believed). This, one might think, is odd. Second, recall that the shifting target strategy aims to explain why Polly or Dougal could felicitously say the following:

(11) What Ted said is wrong/what Ted thinks is wrong/what Ted believes is wrong. He doesn't know that Jack performed the service.

The explanation was that Polly or Dougal could use these disagreement markers to assess the target content. But Ted may not assert (or judge/believe) the target content. Again, one might think that this is odd.

In response I'm inclined to appeal to a limited sort of semantic blindness, viz. blindness about how disagreement over 'knowledge' ascriptions works, and in particular about what hearers are assessing when they assess 'knowledge' ascriptions and denials. The idea is that we mistakenly think that assessment of and disagreement about 'knowledge' ascriptions concern one and the same thing. In contrast with the appeal to semantic blindness discussed in §3.1 - what I called 'disagreement blindness' - this sort of blindness doesn't threaten to undercut the original motivations for contextualism. The contextualist who appeals to disagreement blindness is forced to say that apparent disagreement about 'knowledge' ascriptions is merely apparent. In contrast, the contextualist who appeals to the sort of blindness I'm suggesting takes intuitions about when speakers disagree at face value, but adds that we are blind to certain aspects of how disagreement about 'knowledge' ascriptions works. I don't see why a semantic theory should be expected to take intuitions about how disagreement works into account. We don't expect a semantic theory of some context-sensitive expression to take how competent speakers think context-sensitivity works into account. So why expect a semantic theory to take how competent speakers think disagreement works into account?

Third objection: The shifting target strategy holds that Polly and Dougal can felicitously negatively assess Ted's 'knowledge' ascription because, in doing so, they are assessing the proposition that the ascription would have expressed had it been made in their contexts, rather than the proposition that it actually expressed. But it is consistent with this that the following is, strictly speaking, true:

(12) Ted spoke truly/what Ted thinks is true/what Ted believes is true.

The worry is that, far from only being committed to a limited sort of semantic blindness, the shifting target strategy is committed to widespread error. Competent speakers like Polly and Dougal can only felicitously negatively assess Ted's 'knowledge' ascription if they mistakenly think that (12) is false.

I'll make two points in response. First, the error in question - that of speakers disagreeing with those whom they mistakenly take to have spoken falsely - is familiar and widespread. Recall our case of attitudinal disagreement, (3):

(3) Frank: I love haggis.

Stephen: No way! You're wrong. Haggis is disgusting.

A common feature of cases of attitudinal disagreement is that those involved often do mistakenly take their opponents to be wrong about some matter of fact, or to have some false belief. One can fill in the details of (3) such that, because he thinks that whether a food is disgusting or not is an objective matter, Stephen thinks that Frank is wrong about some matter of fact. Assuming that this objectivist view is mistaken, Stephen mistakenly takes Frank to have spoken falsely, and their disagreement to be more than (merely) attitudinal. On the view I'm suggesting, the disagreement between Ted, Polly and Dougal is in this respect like the disagreement between Stephen and Frank.

Second, and more importantly, the objection is most forceful on the assumption that the fundamental interest driving our assessment of and disagreement about 'knowledge' ascriptions is with truth and falsity. If we want to find out whether those around us are making true 'knowledge' ascriptions and denials it would be problematic that we're often wrong about the truth-values of those ascriptions and denials. But, as I argued in §4.1, the fundamental interest driving assessment and disagreement is with identifying good informants. Consequently, we only want to find out whether those around us are making true 'knowledge' ascriptions insofar as doing so helps us identify good informants. In *STRINGENT POLLY* and *POLICE STATION* Polly and Dougal recognise that Ted's 'knowledge' ascription doesn't serve to identify him as a good informant. That he speaks, judges or believes truly isn't important. So, while the shifting target strategy attributes a sort of error to Polly and Dougal, they aren't in error about what's important when it comes to assessing and disagreeing about 'knowledge' ascriptions, viz. the identification of good informants.

5. Concluding Remarks

The usual story is that intra- and inter-contextual disagreement reports provide good reason to reject contextualism. *Contra* the usual story I have argued that contextualists can deal with both sorts of disagreement reports. Those who wish to motivate adopting a noncontextualist semantics for 'knowledge' ascriptions have to look elsewhere.¹²

¹² Thanks to two anonymous reviewers, Sebastian Becker, Matthew Chrisman, Davide Fassio, Patrick Greenough, Allan Hazlett, Torfinn Huvenes, Sebastian Köhler, Joey Pollock, Mark Richard, Tim Sundell and audiences in Aberdeen, Lisbon and Miglieglia. The research for this paper was funded by the Carnegie Trust.

References

- Baker, Carl 2012. Indexical Contextualism and the Challenges from Disagreement, *Philosophical Studies* 157/1: 107-23.
- Björnsson, Gunnar and Finlay, Stephen 2010. Metaethical Contextualism Defended, *Ethics* 121/1: 7-36.
- Brown, Jessica 2006. Contextualism and Warranted Assertibility Manoeuvres, *Philosophical Studies* 130/3: 407-35.
- Cappelen, Herman 2008. The Creative Interpreter: Content Relativism and Assertion, *Noûs* 42/1: 23-46.
- Cappelen, Herman and Hawthorne, John 2009. *Relativism and Monadic Truth*, Oxford: Oxford University Press.
- Cohen, Stewart 1999. Contextualism, Skepticism, and the Structure of Reasons, *Philosophical Perspectives*, 13: 57–89.
- Craig, Edward 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Oxford University Press.
- DeRose, Keith 2009. *The Case for Contextualism: Knowledge, Skepticism and Context, vol. 1*, Oxford: Clarendon
- Fricker, Miranda 2012. Group Testimony? The Making of a Collective Good Informant, *Philosophy and Phenomenological Research* 84/2: 249-76.
- Greco, John 2009. Knowledge and Success from Ability, *Philosophical Studies* 142/1: 17-26.
- Grice, Paul 1989. *Studies in the Way of Words*, Cambridge, MA: Harvard University Press.
- Hawthorne, John 2004. *Knowledge and Lotteries*, Oxford: Oxford University Press.
- López de Sa, Dan 2008. Presuppositions of Commonality, in *Relativising Utterance Truth*, eds. Manuel Garcia-Carpintero and Max Kölbel, Oxford: Oxford University Press: 297-310.
- McKenna, Robin forthcoming. Normative Scorekeeping, *Synthese*.
- MacFarlane, John 2005. The Assessment Sensitivity of Knowledge Attributions, *Oxford Studies in Epistemology* 1: 197–233.
- MacFarlane, John 2009. Epistemic Modals Are Assessment-Sensitive, in *Epistemic Modality*, eds. Andy Egan and Brian Weatherson, Oxford: Oxford University Press: 144-78.
- Pritchard, Duncan 2012. Anti-Luck Virtue Epistemology, *Journal of Philosophy* 109: 247-79.
- Richard, Mark 2004. Contextualism and Relativism. *Philosophical Studies* 119/1-2: 215-42.
- Ridge, Michael 2013. Disagreement, *Philosophy and Phenomenological Research* 86/1: 41-63.
- Stalnaker, Robert 2002. Common Ground, *Linguistics and Philosophy* 25/5-6: 701-21.
- Stanley, Jason 2005. *Knowledge and Practical Interests*, Oxford: Oxford University Press.

- Stevenson, Charles 1944. *Ethics and Language*, New Haven: Yale University Press.
- von Fintel, Kai 2004. Would you Believe it, the King of France is Back, in *Descriptions and Beyond*, eds. Anne Bezuidenhout and Marga Reimer, Oxford: Oxford University Press: 315-41.
- Yablo, Stephen 2006. Non-Catastrophic Presupposition Failure. In *Content and Modality*, eds. Alex Byrne and Judith Jarvis Thomson, Oxford: Oxford University Press: 164-80.