

Epistemic Contextualism: A Normative Approach

Robin McKenna

University of Edinburgh

rbnmckenna@gmail.com

Abstract In his *Knowledge and Practical Interests* Jason Stanley argues that the view he defends, which he calls interest-relative invariantism, is better supported by certain cases than epistemic contextualism. In this paper I argue that a version of epistemic contextualism that emphasizes the role played by the ascriber's practical interests in determining the truth-conditions of her 'knowledge' ascriptions - a view that I call interests contextualism - is better supported by Stanley's cases than interest-relative invariantism or other versions of epistemic contextualism. My main aim is to show that interests contextualism is a viable, if often over-looked, alternative to the usual positions in the contemporary debate.

1. Introduction

Epistemic contextualism (henceforth, *contextualism*) is the view that the truth-conditions of 'knowledge' ascriptions and denials depend upon the context in which they are uttered.¹ Contextualists have often appealed to intuitions about the truth or falsity of 'knowledge' ascriptions and denials in arguing for their view. The argument usually goes something like this. Intuitively, in some contexts 'knowledge' ascriptions of the form 'S knows that p' (where 'S' is the name of a person and 'p' is a sentence) are true. However, equally intuitively, in other contexts 'knowledge' denials of the form 'S does not know

¹¹ When presenting epistemic contextualism one has to be careful about issues of use and mention. Contextualists think that the expression 'knows' means different things in different contexts, so when, in presenting the view, one talks about whether a subject 'knows' or has 'knowledge' it's important to make it clear whether one intends to be using or mentioning the expression. If one is using it, the meaning is tied to one's own context; if one is mentioning it, some other context. In this paper I've adopted the following conventions. First, when presenting contextualist views, or describing what contextualists say about particular cases, I'll always put quotation marks around 'knows' and its cognates. Second, when presenting non-contextualist views - views which deny that the expression 'knows' means different things in different contexts - or describing what these views would say about particular cases, I won't use quotation marks. In cases where I'm doing both, or neither, I've erred on the side of caution and used quotation marks.

that p ' (where 'S' names the same person, in the same epistemic position, and ' p ' is the same sentence) are also true. By positing contextual variation in the truth-conditions of 'knowledge' ascriptions and denials the contextualist offers a straightforward explanation of contextual variation in our intuitive judgements about their truth-values. Call these sorts of arguments *context-shifting arguments* (CSAs, for short).²

In this paper I argue for a view of the features of context that determine the truth-conditions of 'knowledge' ascriptions that I'm going to call *interests contextualism*.³ The basic idea behind interests contextualism is, in my view, highly intuitive. A prominent contextualist idea is that a 'knowledge' ascription of the form 'S knows that p ' is true in an ascriber A's context C iff S's evidence eliminates every not- p alternative relevant in C, where what alternatives are relevant varies with the context. The interests contextualist says that the relevant alternatives are just the alternatives that those in the context ought to consider, whether those in the context happen to actually consider them or not, and where what alternatives those in the context ought to consider depends on their practical interests (explaining this dependence will be the task of §3). On my view, the features of context that determine the truth-conditions of 'knowledge' ascriptions are facts about what alternatives those in the context ought to consider. That's why the view I defend can be thought of as normative.

The main thesis of this paper is that interests contextualism is better supported by the range of cases discussed in Jason Stanley's *Knowledge and Practical Interests* - which I'm going to call *Stanley's cases* - than some of its competitors. In particular, interests contextualism is better supported by these cases than Stanley's view, which he calls *interest-relative invariantism*. The interests contextualist, like the interest-relative invariantist, thinks that practical interests play a central role in determining the truth-values of 'knowledge' ascriptions. But, unlike the interest-relative invariantist, the interests contextualist can deal with certain cases in which the ascriber of 'knowledge' and the

² Cf. Stewart Cohen's airport cases (Cohen, 1999) and Keith DeRose's bank cases (DeRose, 2009).

³ I've previously defended this view in McKenna, 2011. In McKenna, 2011 I was primarily concerned with the differences between interests contextualism and other versions of contextualism in the contemporary literature. In this paper my main aim is to show that, because it emphasizes the role of the ascriber's practical interests, interests contextualism does better than Jason Stanley's interest-relative invariantism in certain cases.

subject of the ‘knowledge’ ascription differ in their practical situation. I’ll argue for the main thesis in §4. Before doing so, I’ll say a little more about interests contextualism (§3) and about its competitors (§2). I’ll finish by saying a little about the further advantages of interests contextualism over IRI and standard contextualism (§5).

It’s worth noting that, while a number of people have indicated that there is conceptual space for a contextualist view that focuses on the practical interests of ‘knowledge’ ascribers, there haven’t been many attempts to develop such a view in any detail.⁴ For instance, in his recent book *The Case for Contextualism* Keith DeRose briefly considers what he calls a ‘pure reasonableness view’, on which the truth-conditions of ‘knowledge’ ascriptions depend upon whatever epistemic standard is ‘most reasonable’ in the relevant context.⁵ Interests contextualism is a sort of pure reasonableness view, where what it is for a standard to be reasonable is understood in terms of practical interests (again, more on this in §3). DeRose’s main complaint with this sort of view is that what makes a standard reasonable isn’t at all obvious (DeRose, 2009, p. 142). My aim in §3 is to address this complaint. In arguing for my main thesis - that interests contextualism is better supported by Stanley’s cases than some of its competitors - I aim to show not just that interests contextualism is a possible position but that it’s a viable, if often overlooked, alternative to the usual positions in the contemporary debate.

2. Standard Contextualism and Interest-Relative Invariantism

CSAs

I’m going to quickly comment on why I think CSAs are relevant to views about the semantics of the expression ‘knows’ and its cognates.⁶ Here are two ways one might go about defending their relevance. In DeRose’s view ‘the best grounds for accepting contextualism come from how knowledge attributing (and knowledge denying) sentences are used in ordinary, non- philosophical talk’ (DeRose, 2009, p.47). DeRose should, I think, be understood as holding that CSAs provide good but not, of course, decisive grounds for accepting contextualism. Our intuitions about the truth-values of ‘knowledge’ ascriptions and denials are data that a view about the semantics of ‘knows’

⁴ Henderson, 2009 is a possible exception.

⁵ Stanley, 2005a also briefly discusses this sort of view (see pp. 25-8).

⁶ Henceforth I’ll omit references to cognates of the expression ‘knows’, but what I have to say about ‘knows’ will apply to its cognates as well.

should take into account. But a particular view about the semantics of ‘knows’ may give a good explanation of our intuitions yet face several objections. This, I take it, is the position that the contextualist is in.⁷

Another way to defend the relevance of intuitions about the truth-values of ‘knowledge’ ascriptions and denials to semantic theorizing would be to argue that these intuitions follow directly from our understanding of the propositions that these ascriptions and denials express. I have the intuition that an assertion of the sentence ‘some animals have fur’ is true. A straightforward explanation of why I have that intuition is that, in general, intuitions about the truth-values of utterances of ordinary sentences of a natural language follow directly from our understanding of the propositions they express, along with our knowledge of certain empirical facts. Our intuitions about the truth-values of ‘knowledge’ ascriptions and denials should be no different.⁸ I’m not particularly concerned with which sort of defense you prefer. But I do need intuitions about the truth or falsity of knowledge ascription and denials to be relevant to semantic theorizing. My argument in what follows just assumes that they are.

The Competitors

What I’m going to call *classical invariantism* is the view that the expression ‘knows’ is context-insensitive and only truth-conducive factors, such as the extent to which the subject is justified, play a role in determining whether that subject knows a given proposition that she believes. What I’m going to call *interest-relative invariantism* (I’ll often

⁷ A number of philosophers have claimed that the linguistic data shows that ‘knows’ behaves very differently from other context-sensitive expressions. See Stanley, 2005a, Chapter 2 for the argument that ‘knows’, contra what some contextualists claim, isn’t analogous to ‘tall’ because it isn’t gradable and Hawthorne, 2004, pp. 104-8 for the argument that ‘knows’, unlike, say, ‘tall’, doesn’t accept modifier phrases. I’d like to note two things here. First, these aren’t knockdown objections; if successful, they show that the contextualist cannot appeal to an analogy between ‘knows’ and, for instance, ‘tall’. As Stanley acknowledges, they leave open the possibility that “the alleged context-sensitivity of instances of ‘know that p’ could be modeled on the context-sensitivity of some other kind of expression” (Stanley, 2005a, p. 72). See Blome-Tillmann, 2008 for a proposed analogy between ‘knows’ and the expression ‘snores’. Second, Stanley openly acknowledges that the arguments summarised above don’t dilute the force of CSAs. He says they show that ‘except for the intuitions we have about [e.g. bank cases] there is no further evidence that knowledge ascriptions are context-sensitive’ (Stanley, 2005a, p. 33). He then goes on to argue that certain cases - Stanley’s cases - support interest-relative invariantism rather than contextualism. I’m going to argue that this claim is false. These cases support interests contextualism rather than interest-relative invariantism.

⁸ For a defense of this view see Stanley, 2005b.

abbreviate this as IRI) is the view that 'knows' is context-insensitive but non-truth-conducive factors, such as the subject's practical interests, can play a role in determining whether that subject knows a given proposition that she believes.⁹ In this paper I'm going to work with a rough-and-ready characterization of the view on which whether a subject knows depends upon her practical situation (while the view defended in Stanley, 2005a is rather more sophisticated, the details don't matter here). When the subject is in a high stakes situation we can say that the epistemic standards are high, and she doesn't know very much, if anything. When the subject is in a low stakes situation we can say that the epistemic standards are low, and she knows quite a lot.

In this paper I'm going to work with a relevant alternatives contextualist framework on which 'S knows that p' is true in A's context C iff S's evidence eliminates every not-p alternative relevant in C. In terms of this framework, what I'm going to call the *standard contextualist* thinks that the relevant alternatives in a context are often just the alternatives that are *actually* considered in that context.¹⁰ Here's DeRose committing himself to what I just called standard contextualism:

'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 "know(s)". 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 (and even realize that they face) an extremely high-stakes situation in which it would be much wiser for them to employ much higher standards' (DeRose, 2009, p. 55).

'It's good for speakers to use reasonable standards ... But if they opt for unreasonable standards, I'm inclined to think the truth-conditions of their claims then reflect those

⁹ The name is from Stanley, 2005a. For a similar position see Hawthorne, 2004.

¹⁰ I should note that nothing of substance relies on my use of the relevant alternatives framework. An alternative framework would be one on which 'S knows that p' is true in A's context C iff S's epistemic position with respect to p is strong enough to meet the epistemic standard operative in C. In terms of this alternative framework, the standard contextualist thinks that the epistemic standard operative in a context is often just the standard that is actually used in that context, and the interests contextualist thinks that the operative standard is always the standard that those in the context ought to use, given their practical interests.

unreasonable standards that they are indicating' (DeRose, 2009, p. 142).

It's important to be clear about the relationship between interests and standard contextualism. What I take the above quotes to show is that someone like DeRose - a prominent standard contextualist - thinks that the epistemic standard that plays a role in determining the truth-conditions of 'knowledge' ascriptions is often going to be the epistemic standard that those in the context *actually* use. In contrast, the interests contextualist thinks that the epistemic standard operative in a context is *always* going to be the epistemic standard that those in the context *ought* to use, given their practical interests. Presumably the standard contextualist doesn't deny that practical interests can play a role in determining the truth-conditions of 'knowledge' ascriptions (most obviously, if the stakes are high then that's going to incline you to use high epistemic standards). But the standard contextualist certainly denies that the relevant epistemic standard is always the one that those in the context ought to use, whether they actually use it or not. In §4 I'll argue that interests contextualism is better supported by Stanley's cases than standard contextualism precisely because, for the interests contextualist, the relevant epistemic standard is always the one that those in the context ought to use.¹¹

The Cases

It will be helpful in what follows to have the details of what I called Stanley's cases to hand:¹²

Low: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit a check. It's not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long. Realizing that it isn't very important that the check is deposited right away, Hannah says, 'I know the bank will be open tomorrow, since I was there just

¹¹ I want the characterisation of standard contextualism to be broad enough that people such as Cohen, DeRose and Lewis all count as standard contextualists. I've already discussed DeRose, and I'll mention Cohen below (§3). As for Lewis, he certainly thinks that alternatives that are actually considered will be relevant (cf. Lewis, 1996, pp. 559-60) and none of the other rules he lists governing whether an alternative is relevant make reference to what alternatives those in the context ought to consider, so it seems fair to classify him as a standard contextualist.

¹² Because of the dialectical strategy of this paper I've taken these cases more-or-less word for word from Stanley, 2005a. The cases are originally from DeRose (cf. DeRose, 2009, Chapter 1).

two weeks ago on Saturday morning. So we can deposit our check tomorrow morning.’

High: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit a check. Since they have an impending bill, and very little in their account, it’s very important that they deposit their check by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, ‘I guess you’re right. I don’t know that the bank will be open tomorrow.’

Ignorant High: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit a check. Since they have an impending bill, and very little in their account, it’s very important that they deposit their check by Saturday. But neither Hannah nor Sarah is aware of the impending bill, nor of the paucity of available funds. Looking at the lines, Hannah says to Sarah, ‘I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our check tomorrow morning.’

High Ascriber–Low Subject: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit a check. Since they have an impending bill, and very little in their account, it’s very important that they deposit their check by Saturday. Hannah calls up Bill on her cell phone, and asks Bill whether the bank will be open on Saturday. Bill replies by telling Hannah, ‘Well, I was there two weeks ago on a Saturday, and it was open.’ After reporting the discussion to Sarah, Hannah concludes that, since banks do occasionally change their hours, ‘Bill doesn’t really know that the bank will be open on Saturday’

Low Ascriber–High Subject: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit a check. Since they have an impending bill, and very little in their account, it’s very important that they deposit their check by Saturday. Two weeks earlier, on a Saturday, Hannah went to the bank, where Jill saw her. Sarah points out to Hannah that banks do change their hours, and Hannah responds, ‘That’s a good point. I guess I don’t really know that the bank will be open on Saturday.’ Coincidentally, Jill is thinking of going to the bank on Saturday, just for fun, to see if she meets Hannah there. Nothing is at stake for Jill, and she knows

nothing of Hannah's situation. Wondering whether Hannah will be there, Jill utters to a friend, 'Well, Hannah was at the bank two weeks ago on a Saturday. So she knows the bank will be open on Saturday.'

Here are my intuitions about these cases, and I take it that they will be widely shared. In *Low* Hannah's self-ascription of 'knowledge' is true, in *High* Hannah's self-denial is true, in *Ignorant High* Hannah's self-ascription is false and in *High Ascriber-Low Subject* Hannah's denial of 'knowledge' to Bill is true. I'm not alone here. Stanley takes it as read that we have this set of intuitions (cf. Stanley, 2005a, p. 5).¹³ (What about *Low Ascriber-High Subject*? Stanley claims that Jill's ascription of 'knowledge' to Hannah is false but, as we'll see in §4, I'm not so sure about that).

In the next section I'm going to say a little more about interests contextualism. I'll then argue that Stanley's cases provide better support for interests contextualism than either IRI or standard contextualism. But I want to make it clear that nothing I say in this paper is going to show that Stanley's cases provide better support for interests contextualism than classical invariantism. Unlike both contextualists and interests-relative invariantists, classical invariantists don't take the set of intuitions above at face value. Rather, classical invariantists argue that these intuitions can be explained away as a pragmatic phenomenon.¹⁴

3. Interests contextualism

On the interests contextualist view, an alternative is relevant in a context iff those in the context ought to consider that alternative, where what alternatives those in the context ought to consider depends on their practical interests. But what determines their practical interests? What they take their interests to be, or some more 'objective' factor? In this section I'll address these questions.

¹³ In the literature these intuitions have often been taken as read. However, see May et al., 2010 for some contrary data and DeRose, 2011 for a contextualist response.

¹⁴ Cf. Brown, 2006, Hazlett, 2009 and Rysiew, 2001. The basic idea is that, instead of explaining contextual variation in intuitions about the truth-values of 'knowledge' ascriptions by positing contextual variation in their truth-conditions (contextualism) or truth-values (IRI) we can just posit contextual variation in their assertability conditions. I'd like to note that, while this is a popular move, there are a number of problems with the various pragmatic accounts that have been proposed. See, for example, Baumann, 2011 and DeRose, 2009, Chapter 3.

I'll start by considering an illustrative example.¹⁵ Saskia and Marie are trying to decide when to leave their flat to go to the cinema. They want to see the film *Troll Hunter*, which is showing at 6pm. It's currently 5.30pm and, in normal traffic conditions, it would take them 15 minutes to get there. However, this is a weekday and rush-hour traffic can get pretty bad where they live. Because they're not aware that the rush-hour traffic gets so bad, neither Saskia nor Marie are considering this, so they're planning on leaving in 10 minutes. In an objective sense of 'having a reason', Saskia and Marie have a reason to consider the possibility that they will be held up in traffic. However, in a subjective sense, Saskia and Marie have no such reason.

I'd like to pause to say a little about what I mean by 'subjective' and 'objective' senses of having a reason. The idea I'm drawing on here should be familiar from Bernard Williams' classic paper "Internal and External Reasons". Williams' gives an example involving an agent - call him Bernard - who wants to drink from a glass if it contains gin, but has no such desire if it contains petrol. Imagine that the glass contains petrol. Williams says that, in one sense - a sense connected with explanations of action - Bernard has got a reason to drink from the glass but, in another sense, Bernard hasn't got such a reason, and actually has a reason to not drink from the glass (Williams, 2001, pp. 78-9).

What I'm claiming is that the reason Saskia and Marie have to consider the possibility that they will be held up in traffic is of the same sort as the reason Bernard has to not drink from the glass. So when I say that Saskia and Marie have a reason, in the objective sense, to consider the possibility that they will be held up in traffic, I certainly don't mean that Saskia and Marie have what Williams called an 'external reason'. Part of Williams' aim in his paper is to argue for a conception of internal reasons on which Bernard has an internal reason to not drink from the glass. On Williams' view, the fact that the glass contains petrol gives Bernard an internal reason to not drink from it just so long as he would be motivated to not drink from the glass were he aware that it contained petrol. The basic idea is that an internal reason must be the sort of thing that would motivate an agent to act if that agent were aware of it. Saskia and Marie have an internal reason to consider the possibility that they will be held up in traffic because, if they were aware that the traffic is going to be bad, they would consider it. (If, in contrast to the original case, Saskia and Marie didn't care about getting to the film on time then they wouldn't have

¹⁵ Just to be clear: the purpose of this example is entirely illustrative. I'm not using the case I'm about to describe to motivate interests contextualism.

such an internal reason. In that case, even if they were aware that the traffic is going to be bad they still wouldn't consider it).

The distinction I'm making between 'subjective' and 'objective' senses of having a reason doesn't line-up with Williams' distinction between internal and external reasons, but it does line-up with a familiar distinction in epistemology. There's a difference between the possibilities one ought to consider given what one knows or believes about one's present situation and the possibilities one ought to consider given the facts. If Bernard were presented with misleading evidence that the glass contains gin - an open, half-empty gin bottle right next to the glass, say - then, given what he knows or believes about his present situation, he has got a reason to drink from the glass. Further, given what he knows about his situation, he hasn't got a reason to consider the possibility that the glass contains petrol. However, given the facts - the fact that the glass contains petrol - he has a reason to not drink from the glass. Further, given the facts, he has a reason to consider the possibility that the glass contains petrol. The first sort of reason - the reason Bernard has to drink from the glass - is what I mean by 'subjective'. The second sort is what I mean by 'objective'.¹⁶

Here's the basic idea I want to take from this. We attach different levels of practical importance to the various projects we engage in. When we engage in these projects there are certain things that we should consider. Saskia and Marie are engaging in the project of getting to the cinema in time to see a film, and this is a project to which they attach a decent amount of practical importance (they'd far rather get there on time than not). For the purposes of this project, there are certain possibilities that will almost always be relevant, and the possibility of getting caught in traffic is plausibly one of them. These possibilities are relevant in virtue of the practical importance of the project Saskia and Marie are engaged in, not in virtue of the fact that they consider them. There are also certain possibilities that will almost never be relevant, for example the possibility that the world might end in the next five minutes. These possibilities are irrelevant in virtue of

¹⁶ I acknowledge that, to some, this objective sense might be mysterious, or even controversial. I'd like to offer two points in response. First, as I've said, what I'm talking about aren't external reasons in the sense of Williams, 2001. Second, the target of this paper - Stanley, 2005a - appeals to a notion of what possibilities are epistemically probable for a subject that is objective in the same sense that the reason Saskia and Marie have to consider the possibility that they'll get stuck in traffic is objective (see Stanley, 2005a, Chapter 5).

the practical importance of the project Saskia and Marie are engaged in, not in virtue of the fact that they don't consider them. Projects differ in their practical importance, and the more important the project, the greater the number of relevant possibilities.

Interests contextualism is the combination of the basic contextualist semantic thesis - that the truth-conditions of 'knowledge' ascriptions depend upon the context in which they are uttered - with an account of which features of context determine the truth-conditions of 'knowledge' ascriptions. On the interests contextualist view the relevant features are facts about what alternatives those in the context ought to consider, where the 'ought' is understood in the objective sense discussed. So, for example, in *Low* Hannah and Sarah are trying to decide whether to stop at the bank. It isn't vitally important that they cash their check before Saturday, and that's why it's not the case that they ought to consider the alternative that the bank has changed its opening hours. That means it's irrelevant, so, assuming Hannah can rule out all relevant alternatives, she speaks truly in self-ascribing 'knowledge'. However, in *High* it is vitally important to Hannah and Sarah that they cash their check before Saturday, and that's why they ought to consider the same alternative. That means it's relevant, so Hannah speaks truly in denying that she 'knows'. For the interests contextualist, what alternatives one ought to consider in a context depends upon the practical importance of the project one is engaging in. The relevant alternatives are relevant in virtue of the importance of the project, not in virtue of the fact that they are actually considered.

In the next section I'm going to argue that Stanley's cases provide better support for interests contextualism than IRI or standard contextualism. Before doing so, I'm going to consider two objections to the view that I've just sketched.

Upon reading that, on my view, an alternative is relevant in a context iff those in the context ought to consider it, where what alternatives those in the context ought to consider depends upon their practical interests, one might wonder what I say about cases where those in the context attend to practically irrelevant possibilities. Do I say that, in all such cases, these alternatives are irrelevant? I could, but I don't have to. It may well be that, in certain cases, attending to an alternative can affect what alternatives those in the context ought to consider. A good example of a case where this might be going on is

Jonathan Vogel's parked-car case.¹⁷ Imagine that I've parked my car in what I think is a fairly safe area and, while presumably I'm aware that it's possible that my car might have been stolen, I don't take that possibility seriously and I happily ascribe 'knowledge' of where my car is parked to myself. Now imagine that a friend starts to insist that it's possible that my car might have been stolen, and she makes this possibility particularly vivid to me by explaining how massively important it must be to me that my car not be stolen (imagine I've only just bought it). It may well be that my friend doing so affects my practical situation. Before I wasn't worried about my car, whereas now I am. If my friend doing so affects my practical situation, then that will expand the range of possibilities it's in my practical interest to consider. So, sometimes, attending to possibilities might make them relevant. But it will only do so by affecting my practical situation. (Of course, it's also perfectly possible that I might regard my friend's raising these possibilities as a nuisance, in which case they wouldn't be relevant).

Another worry you might have about interests contextualism is that in some cases it isn't going to be entirely clear *what* the practical interests of those in the context are. In contrast to the cases I've discussed in this paper, real life is complex, and real people have a far wider range of interests and concerns than Saskia and Marie, or the characters in *Low*, *High* and the rest of Stanley's cases. This worry might take an epistemic form, viz. in a particular case how am I to determine the practical interests of an ascriber in order to evaluate the truth or falsity of her 'knowledge' ascription? The worry here is that, on the interests contextualist view, the facts I need in order to evaluate her 'knowledge' ascription - facts about her practical interests - might not be available to me. Further, and perhaps more pressingly, one might worry that, in some cases, an ascriber's practical interests won't just be hard to determine but actually indeterminate, perhaps because she has conflicting interests. The worry here is that, on the interests contextualist view, the facts that are relevant to the assignment of a truth-value to her 'knowledge' ascription - facts about her practical interests - might not even be determinate.

I think that these are hard questions for the interests contextualist, and answering them properly would be a task for another paper. In what follows I pursue the following strategy. I will give some reasons for thinking that the interests contextualist can avoid

¹⁷ Cf. Vogel, 1999.

the second worry, and then I will argue that something much like the first worry is also a problem for some of interests contextualism's rivals, in particular IRI and standard contextualism.

First, it's not clear what would have to be the case for an ascriber's practical interests to actually be indeterminate. Say that I have a bunch of conflicting desires of roughly similar strength: a desire for a drink, a desire to continue my work, a desire to go for a walk, and so on. If I'm trying to decide what to do, it may well be that I have a hard time deciding which of these desires to act on. After all, each of the options has something to be said for it. Sometimes, indeed often, what it would be in my best interests to do isn't entirely clear to me. But that doesn't mean that there isn't one thing that it would be in my best interests to do. It just means that figuring out what that thing is can be difficult. So why think that, because 'knowledge' ascribers might have conflicting interests, their practical interests may be indeterminate? After all, we don't think that my conflicting desires might mean that there isn't a single thing that it would be in my best interests to do. For this reason, I would suggest that the interests contextualist can avoid the second worry.

Second, the epistemic worry isn't just a problem for the interests contextualist. Interest-relative invariantists are going to face the same question - the question of how to determine practical interests - given that, on their view, the truth-values of 'knowledge' ascriptions depend upon the *subject's* practical interests. In a particular case, in which I'm neither the ascriber nor the subject, I'm not going to find it any easier to determine the practical interests of the subject of a 'knowledge' ascription than the practical interests of the ascriber. Further, standard contextualists are going to face a similar, although somewhat different, question (actually, two questions). On the standard contextualist view, the truth-values of 'knowledge' ascriptions depend upon a number of features of the conversational context. However, standard contextualists are usually rather vague as to *which* features of the conversational context are relevant. A common move here is to give a list of some of the relevant features. Here's Stewart Cohen:

'How precisely do the standards for [epistemic predicates] get determined in a particular context of ascription? This is a very difficult question to answer. But we can say this much. The standards are determined by some complicated function of speaker

intentions, listener expectations, presuppositions of the conversation, salience relations, etc., by what David Lewis calls the conversational score' (Cohen, 1999, p. 61).

I take it that the standard contextualist faces two questions here. First, can she say anything more about this complicated function? Second, if the function is so complicated, how can I, in a particular case, determine the conversational score in order to evaluate the truth or falsity of a 'knowledge' ascription? So the standard contextualist seems to face something like the epistemic worry.¹⁸

4. Interests Contextualism vs. IRI and Standard Contextualism

In this section I'll argue for three claims. First, unlike standard contextualism, interests contextualism and IRI can accommodate our intuitions about *Ignorant High*. Second, unlike both standard and interests contextualism, IRI has problems accommodating our intuitions about *High Ascriber-Low Subject*. Third, both standard and interests contextualism can accommodate our intuitions about *Low Ascriber-High Subject*. Taken together, these three claims establish that interests contextualism does a better job of accommodating our intuitions about Stanley's cases than IRI or standard contextualism.

Ignorant High

(A reminder: Hannah and Sarah have a bill due, and very little in their account. Neither Hannah nor Sarah is aware of the bill, or their perilous financial situation. Based on her memory of it being open on a recent Saturday, Hannah says that she 'knows' the bank will be open on Saturday. Intuitively, Hannah speaks falsely). Stanley has argued that IRI, unlike contextualism, can accommodate our intuitions about this case (Stanley, 2005a, p. 25). In *Ignorant High*, Hannah and Sarah are in a high stakes situation, even though they're not aware of it. The defender of IRI says that the epistemic standards are high, which makes Hannah's knowledge self-ascription false, so IRI accommodates our intuitions. In contrast, standard contextualists will have to say that Hannah speaks truly. The alternative on which the bank has changed its opening hours isn't considered by Hannah or Sarah, which means it's irrelevant. The standard contextualist has to deny our intuitions because, on her view, an alternative that isn't actually considered isn't relevant.

¹⁸ Just to be clear: my point isn't that standard contextualists can't answer these questions. My point is just that standard contextualists, interests contextualists and interest-relative invariantists all face tough questions, and it isn't immediately clear how they should go about answering them.

However, the interests contextualist can accommodate our intuitions about *Ignorant High*. For the interests contextualist, an alternative is relevant in a context iff those in the context ought to consider it. In *Ignorant High* Hannah and Sarah are trying to decide whether to stop at the bank and, even though they're not aware of it, it's vitally important that they do so before Saturday. That means that they ought to consider the alternative that the bank has changed its opening hours, so it's relevant. So Hannah speaks falsely, assuming that her evidence doesn't rule out the alternative that the bank has changed its opening hours. The interests contextualist can deal with *Ignorant High* because, on her view, an alternative is relevant iff those in the context ought to consider it, whether they actually consider it or not.

High Ascriber-Low Subject

(A reminder: Hannah and Sarah have a bill due, and very little in their account. Bill tells Hannah that he remembers being in the bank on a recent Saturday, so it will be open this Saturday. Hannah says that, since banks do change their opening hours, Bill doesn't 'know' that it will be open this Saturday. Intuitively, Hannah speaks truly). The usual story is that IRI, unlike contextualism, has problems with these sorts of cases. The contextualist can say that the alternative that the bank has changed its opening hours is relevant, whether because Hannah and Sarah consider it or because, given their practical situation, they ought to consider it. But, *prima facie*, IRI cannot accommodate our intuitions. Bill isn't in a high stakes situation, which means that the epistemic standards aren't high, so Hannah's denial that Bill knows is false.

In response, the defender of IRI needs to explain why our intuitions about *High Ascriber-Low Subject* are mistaken. In what follows I'll focus on Stanley's explanation.¹⁹ We need a general account of what goes on in cases where we have an ascriber in a high stakes context - call her High - and a subject in a low stakes context - call her Low - and High denies that Low knows that p. Stanley argues that the reason why we have mistaken intuitions about these cases is that ascribers tend to *project* their epistemic standards onto the subject. We have this tendency, says Stanley, because of what we're doing when we inquire into whether a subject knows, for example, if a certain bank is open on Saturdays. Generally, when inquiring into whether a subject knows that p we are using that subject as a source of information. So what we want to establish is whether, if the subject had

¹⁹ Cf. Stanley, 2005a, Chapter 5.

our interests, that subject would know that p. As a consequence, we project our epistemic standards onto the subject and evaluate whether her evidence suffices for knowledge by our, rather than the subject's, epistemic standards. This explains why High denies that Low knows that p. High projects her epistemic standards onto Low, and Low doesn't meet those standards, even though she meets her own standards. In what follows I'll call this the *projectivist strategy*.

The projectivist strategy has it that we *mistakenly* project our epistemic standards whereas the contextualist has it that we correctly project our standards. It might look like there's a stalemate here: on what grounds do we decide whether the projection of epistemic standards is mistaken or not? In what follows I'm going to argue that there are good reasons for thinking that the contextualist is better off here.

First, it's plausible that there should be a defeasible default assumption to the effect that, when speakers choose to apply a particular epistemic standard, they aren't mistaken in doing so. This default assumption can, of course, be defeated by a whole host of considerations. For example, if the way speakers apply some standard is incoherent then they must be making some sort of mistake. If this is correct, Stanley needs to give us a reason why the projection of epistemic standards is mistaken. In contrast, in the absence of a reason to think the projection is mistaken, the contextualist doesn't have to give a reason to think we project correctly.

I think this is right, but it doesn't go very far. The default assumption in question is, presumably, highly defeasible. In particular, if it were the case that contextualism were worse off than interest-relative invariantism in certain other respects, then that would give us good reason to reject the default assumption. This, I take it, is something that the defender of the projectivist strategy - in particular, Stanley - takes himself to have argued for.

Second, the projectivist strategy is a sort of error-theory, so the defender of the strategy has to say something about why the error in question arises and, in particular, whether the error is a result of some more general tendency. Ideally, an error-theoretic explanation will appeal to a general tendency speakers have to go wrong that not only

explains the error in question but much more besides. I think the following passage from DeRose makes this point nicely.

‘While the projectivist account Stanley suggests is *possible*, he provides, and I can see, little reason to actually accept it. Stanley does not explain how misleading tendencies of overprojection that we might display elsewhere can be applied to the case at hand to promote the thought that our tendency to deny that subjects ‘know’ when we are in such situations is a mistake’ (DeRose, 2009, p. 235) (DeRose’s emphasis).

I also think this is right, but it has more of the look of a challenge than an objection, and I’m not sure it’s a challenge that will be all that hard to meet. What I said in response to the first point seems to apply here as well. If, as Stanley argues, contextualism is worse off than interest-relative invariantism in other respects, that would give us a reason to accept the projectivist strategy.

The real problem with the projectivist strategy is that there are good reasons for thinking that the feature of our epistemological practice that it appeals to - that, in evaluating whether a subject ‘knows’ that *p*, we’re evaluating that subject as a source of information - reveals something important about the social role or function of ‘knowledge’ ascriptions.²⁰ It’s plausible that one of the social roles or functions of ‘knowledge’ ascriptions is to, as it’s sometimes put, ‘tag’ good sources of information.²¹ When you or I ascribe ‘knowledge’ that *p* to a subject *S* we, in effect, recommend *S* as an authority on the matter of *p*.

Here are two reasons why one might think that ‘knowledge’ ascriptions play this social role. First, we clearly have a need for a way of identifying good sources of information on matters that are of interest to us. It’s not possible to obtain all the information that we need first-hand, so there’s a need for a way of keeping track of those who have the relevant information. Second, if I’m told that someone ‘knows’ something, that the train

²⁰ Cf. Henderson, 2009, pp. 123-5 for a similar objection.

²¹ The picture I’m drawing on here is from Craig, 1990. Note, though, that for Craig *the* social role or function of knowledge ascriptions is to tag good sources of information. All I need is that *a* social role of knowledge ascriptions is to tag good sources of information. I submit that, while there may be good reasons to doubt the stronger claim, there are good reasons to endorse the rather weaker claim (some of which I discuss in what follows). For discussion see Kelp, 2011.

leaves in 10 minutes, say, then it seems right to say that I'm entitled to treat that person as a source of this information. They have been recommended to me as someone with the information that I need. So 'knowledge' ascriptions are what meet the need for a way of keeping track of good sources of information.

If, as I've just argued, the feature of our epistemological practice that Stanley appeals to reveals something important about the social role of 'knowledge' ascriptions, then the defender of the projectivist strategy has to say that the social role of 'knowledge' ascriptions involves some sort of error or mistake. This, I think, is a real cost for anyone who wants to adopt it.²²

Low Ascriber-High Subject

(A reminder: Jill is thinking of going to the bank on Saturday on the off chance that she'll meet her friends Hannah and Sarah there. Nothing is at stake for Jill, but Hannah and Sarah have an impending bill due, and very little in their account. Based on her meeting Hannah at the bank on a recent Saturday, Jill says that Hannah 'knows' the bank will be open on Saturday). The usual story is that contextualism, unlike IRI, has problems with these sorts of cases. Hannah is in a high stakes situation so, for the interest-relative invariantist, Jill's knowledge ascription is false. However, Jill doesn't consider the alternative that the bank has changed its opening hours so, for the standard contextualist, that alternative is irrelevant. Assuming that Hannah's evidence suffices to rule out all relevant alternatives, Jill's 'knowledge' ascription is true. This, so the usual story goes, is a problem for the standard contextualist. In this section I'm going to argue that the interests contextualist has no problem here (I should note that DeRose, 2009, Chapter 7 argues that the standard contextualist can also deal with this sort of case, so I'm not claiming that interests contextualism is at an advantage over standard contextualism here).

²² Why can't the defender of the projectivist strategy say that while some, or maybe most, 'knowledge' ascriptions have the function of tagging sources of information, there isn't a necessary connection between tagging information and 'knowledge' ascriptions? It could be that the cases where 'knowledge' ascriptions don't have the function of tagging sources of information are just the cases to which the projectivist strategy applies. I don't claim to have a decisive reason for rejecting this response, but I do think it looks ad hoc. In particular: is there an explanation *why* the cases to which the projectivist strategy applies are the cases where 'knowledge' ascriptions don't have their usual function? If not, it would be a remarkable coincidence.

Rather than discussing the original *Low Ascriber-High Subject* case, I want to distinguish two sorts of cases in which the ascriber is in a low stakes practical situation and the subject in a high stakes practical situation. Once I've explained how the interests contextualist deals with these two sorts of cases, I'll explain why this helps with the original case.

In distinguishing two sorts of *Low Ascriber-High Subject* cases I'm going to draw on the idea, mentioned above, that 'knowledge' ascriptions perform the social role of tagging good informants. An informant is a subject with information that someone, or some group, needs. Sometimes when an ascriber tags a subject as a good informant she will do so because that subject has the information that she - the ascriber - needs. A case of this sort would be one in which the ascriber- call her Inquirer - is interested in whether a subject - call her Informer - has the information that Inquirer needs in order to pursue some project in which she - Inquirer - is engaged. In this case, Inquirer will take her practical situation into account rather than the practical situation of Informer. (An example: Imagine that Informer is outside the bank and trying to decide whether she should go in or come back on Saturday. She has no idea whether the bank is open on Saturdays so she calls Informer). Other times, when an ascriber tags a subject as a good informant she will do so because that subject has information that some other party needs. One case of this sort would be one in which an ascriber - call her Advisor - is interested in whether a subject - call her Advisee - 'knows' in order to offer advice to Advisee about some project in which Advisee is engaged. In this case, Advisor will take Advisee's practical situation into account rather than her own. (An example: Imagine that Advisor is asked whether the bank is open on Saturdays. Whether Advisor will offer herself as an informant on that question will depend on what she takes Advisee's practical situation to be). Another case of this sort would be one in which Advisor is interested in whether a third-party - call her Informant - 'knows' in order to offer advice to Advisee about some project in which Advisee is engaged. In this case, Advisor will take Advisee's practical situation into account rather than her own or Informant's. (An example: Imagine that Advisor is trying to find out whether the bank is open on Saturdays for Advisee. If Informant tells Advisor that it is, whether Advisor will tag Informant as a good source of information on the question of whether the bank is open on Saturdays will depend on what she takes Advisee's, rather than Informant's or her own, practical situation to be).

If, as I've argued, 'knowledge' ascriptions play the social role of tagging good informants, then we should expect there to be the two sorts of *Low Ascriber-High Subject* case just discussed. In the first sort, it's the ascriber's practical situation that's important. In the second sort, it's the practical situation of some other party that's important. This is entirely consistent with the interests contextualist view, on which the practical importance of the project in which those in the context are engaged determines what alternatives those in the context ought to consider. Sometimes an ascriber will be engaged in the project of offering advice, and when she does so it's the practical situation of the advisee that will be important.

So: what about the original *Low Ascriber-High Subject* case? If it's a case of the first sort, the interests contextualist gives the result that Jill speaks falsely. However, if it's a case of the second sort, Jill speaks truly. So, if Stanley had the first sort of case in mind, I agree that Jill speaks falsely. However, if he had the second sort in mind, I disagree. Stanley's presentation of the case might lead one to think it's of the second sort. First, it's specified that Jill isn't aware of Hannah's situation. That means that Jill isn't going to be much use as an advisor. Second, one gets the impression that Jill's only interest is in whether going to the bank is liable to lead to her meeting Hannah. But perhaps things aren't so clear-cut. After all, we do sometimes give advice even when we aren't aware of the details of the advisee's situation. My contention is just that, in both sorts of *Low Ascriber-High Subject* cases I've discussed, the interests contextualist gives the intuitive verdict. The interests contextualist says that what alternatives one ought to consider in a context depends upon the practical importance of the project one is engaged in. But, because one can engage in the project of evaluating the actions available to others, the interests contextualist can allow that, when one does so, the subject's practical situation has an effect on the truth-conditions of 'knowledge' ascriptions.

My primary aim in this sub-section has been defensive, not offensive. The usual story is that contextualists have problems with *Low Ascriber-High Subject*, and my aim here has been to argue that, once we distinguish between two sorts of *Low Ascriber-High Subject* cases, we see that the interests contextualist can deal with both sorts. However, if I'm right, then IRI gets the wrong result in a certain sort of *Low Ascriber-High Subject* case, viz. cases like the Inquirer case. Of course, in order to respond to my argument the defender of IRI could appeal to an error theory. That error theory would, like the error theory she

appeals to in dealing with *High Ascriber-Low Subject*, involve denying a plausible account of the social role of ‘knowledge’ ascriptions. As I’ll discuss in the next section, I think this point is central to the debate between interests contextualism and IRI.

5. Why Interests Contextualism?

As I’ve argued, interests contextualism is better supported by Stanley’s cases than either IRI or standard contextualism. This, I take it, shows that interests contextualism is a viable alternative to the usual positions in the contemporary debate, but it’s a long way from a decisive argument for interests contextualism. While it’s a point in favor of interests contextualism that it does well with respect to Stanley’s cases, a lot more has to be said. My aim in this final section is to say a little about the further advantages of interests contextualism. First, I’ll discuss two further advantages that interests contextualism has over standard contextualism. Second, I’ll say a little about the debate between interests contextualism and IRI.

Interests Contextualism vs. Standard Contextualism

Contextualists think that the truth-conditions of ‘knowledge’ ascriptions depend upon the context of utterance. But what features of the context are relevant? As I mentioned above (§3), the usual response to this question is, first, that the answer is complicated and, second, that there are a number of relevant features of context. In contrast, of course, the interests contextualist gives a simple response. For the interests contextualist, the relevant features of context are facts about what alternatives those in the context ought to consider. As I’ll argue in what follows, this simple response has at least two advantages.

A common complaint about contextualist accounts is that they don’t say enough about the mechanisms by which context-shifts - expansions/retractions of the set of relevant alternatives or upward/downward shifts in the epistemic standards - occur.²³ What, exactly, does it take for the standards to go up? When, exactly, do the standards return to ‘normal’ after being raised? The first advantage of interests contextualism is that it gives a simple answer to these sorts of questions. For the interests contextualist, context-shifts occur when there’s a change in the practical situation of those in the context. When the stakes get high, standards go up; when the stakes return to normal, the standards go back

²³ Cf. Pritchard, 2001 & 2010.

down. There will, of course, be difficult cases, i.e. cases where it isn't clear, or it is difficult to determine, whether the practical situation of some group has changed (I discussed these sorts of cases in §3). But there will also be straightforward cases, i.e. cases where it is clear that the practical situation of some group has changed and, with respect to these sorts of cases, the interests contextualist is at the stated advantage. For example, I take it that Stanley's cases are of this sort.

Another common complaint is that contextualist accounts are too 'skeptic-friendly'.²⁴ Consider the following case, which is an example of a conversation involving what I'll call an *uncooperative participant*:

Ted: Have you read the latest issue of the *New Scientist*? There's some really good evidence for climate change in there. That should really answer the skeptics.

Dougal: Why should I read it? You do know the whole thing is a conspiracy, right?

Ted: Don't be ridiculous. Scientists *know* climate change is happening. I *know* climate change is happening.

Dougal: I'm not reading it. It's all lies anyway.

I think it's quite clear that contextualists shouldn't say that Ted speaks falsely in ascribing 'knowledge' both to himself and others here. Uncooperative participants can't affect the truth-values of our 'knowledge' ascriptions. However, it isn't clear how standard contextualists can deal with this. If, as the standard contextualist holds, the conversational score determines the epistemic standards, how can Dougal's attempts to change the score by raising the standards not affect the score?

The standard contextualist has two options here, and, as I'll argue, neither will help. First, she could say that Dougal's attempt to raise the standards fails, so Ted speaks truly in ascribing 'knowledge' both to himself and others. Second, she could say that, because Ted and Dougal disagree as to what the standards should be, Ted doesn't speak truly or falsely.²⁵ (Of course, she could say that Dougal's attempt to raise the standards succeeds, so Ted speaks falsely. This third option is obviously a non-starter, so I'll disregard it).

²⁴ This, of course, is a common complaint about Lewis, 1996. In what follows I argue that a version of the complaint applies to standard contextualism more generally. For a similar point see Montminy, forthcoming.

²⁵ DeRose, 2009, Chapter 4 takes this line.

If the standard contextualist takes the first option, she has to explain what the difference between this sort of case - one involving an uncooperative participant - and the sort of case where *philosophical* (as opposed to climate change) skeptics raise brain in vat-style error possibilities is. Contextualists such as DeRose want to allow that, at least in some cases, when one is in conversation with a philosophical skeptic one can't truly ascribe 'knowledge' of everyday empirical propositions, either to oneself or others.²⁶ If that's right, then there have to be cases in which philosophical skeptics succeed in changing the conversational score. So, if she takes this option, the onus is on the standard contextualist to explain why philosophical skeptics can succeed in changing the score whereas an uncooperative participant such as Dougal can't. But if the standard contextualist takes the second option, she can't say that Ted speaks truly in ascribing 'knowledge'. True, on the second option he doesn't speak *falsely*, but I take it that this option isn't much less palatable than the third.

The second - and, I think, far more important - advantage is that the interests contextualist has no problem with cases involving uncooperative participants. Ted, presumably, isn't in a position where, given his practical interests, he ought to consider the alternative that all of the evidence for climate change is part of a conspiracy. And because, for the interests contextualist, Dougal raising that alternative doesn't affect what alternatives Ted ought to consider, he speaks truly in ascribing 'knowledge' to himself and others.

Interests Contextualism vs. IRI

I argued in §4 that interests contextualism does better than IRI with respect to Stanley's cases. As I've noted, this is a point in favor of interests contextualism, but it's hardly a decisive one. Every candidate account of the semantics of 'knowledge' ascriptions - standard contextualism, IRI, classical invariantism and, presumably, interests contextualism - will have problems accommodating some intuition or other, and to deal with the problem they will have to develop some sort of error-theoretic explanation. So one might wonder: why is my pointing out that the defender of IRI has to deny a plausible account of the social role of 'knowledge' ascriptions in order to deal with *High Ascriber-Low Subject* (and, indeed, certain sorts of *Low Ascriber-High Subject* cases) not just another example of a case where one account of the semantics of 'knowledge' ascriptions

²⁶ Cf. DeRose, 1995.

has to invoke an error-theory? Contextualists (including interests contextualists) will have to invoke a number of error-theoretic explanations too, and one can only decide the debate between IRI and contextualism by taking all of these things into account.²⁷

I completely agree that a whole host of considerations need to be taken into account when comparing IRI and interests contextualism (and, indeed, standard contextualism and interests contextualism), but my aim in this paper has not been to argue that interests contextualism does better than IRI with respect to all of the relevant considerations. My aim has been to show that interests contextualism does better than IRI with respect to certain cases. However, in doing so I have managed to identify what I take to be a central consideration in the debate between interests contextualism and IRI. Interests contextualism, as I've argued, is entirely consistent with a plausible view about the social role of 'knowledge' ascriptions. However, the defender of IRI has to say that, if 'knowledge' ascriptions play the social role of tagging good informants, then that they do so reflects an error in our practice. (I take it that someone who invokes the IRI projectivist strategy discussed above *does* think that 'knowledge' ascriptions play something like this social role). So the key questions are (1) do 'knowledge' ascriptions play this social role and (2) if they do, what grounds are there for thinking that this reflects an error in our practice? I've given some grounds for thinking that 'knowledge' ascriptions do play the social role of tagging good informants (although, of course, more would need to be said) but I haven't given any grounds for thinking that this doesn't reflect an error in our practice.

Are there grounds for thinking that the social role of 'knowledge' ascriptions reflects an error in our practice? Answering this question would be a task for another paper, but I think it's important to bear the following in mind. While, as I've argued, the hypothesis that the social role of 'knowledge' ascriptions is to tag good informants is consistent with contextualism, it's also entirely independent of contextualism. That is, whether contextualism is true or not doesn't affect the plausibility of the hypothesis about the social role of 'knowledge' ascriptions. In my earlier discussion I dismissed two reasons for rejecting Stanley's projectivist strategy (that there's a default assumption to the effect that speakers aren't mistaken in applying a particular epistemic standard, and that he has

²⁷ This picture of the dialectic is pretty common. Something much like it can be found in DeRose, 2009, Hawthorne, 2004 and Stanley, 2005a.

given us no reason to think that our projection of epistemic standards is mistaken). The problem with both was that, if there are good reasons to think that contextualism is worse off than IRI in other respects, neither carries much weight. However, because the hypothesis about the social role of ‘knowledge’ ascriptions is plausible independent of contextualism, whether there are good reasons to think that contextualism is worse off than IRI or not is irrelevant to whether the hypothesis about the social role of ‘knowledge’ ascriptions reflects an error in our practice. When assessing the viability of the projectivist strategy, considerations against contextualism and for IRI aren’t relevant. All that is relevant is the plausibility of the hypothesis about the social role.

6. Conclusion

My main aim in this paper has been to show that a version of contextualism that emphasizes the role played by the ascriber’s practical interests - what I’ve called interests contextualism - is a viable alternative to the usual positions in the contemporary debate. Interests contextualism is better supported by Stanley’s cases than IRI or standard contextualism. Unlike standard contextualism, interests contextualism can accommodate our intuitions about *Ignorant High*. Unlike IRI, which has problems with *High Ascriber-Normal Subject*, interests contextualism can accommodate our intuitions about both *High Ascriber-Normal Subject* and *Normal Ascriber-High Subject*. If the CSAs I’ve considered in this paper support any view about the semantics of ‘knowledge’ ascriptions, it’s interests contextualism.²⁸

Bibliography

- Baumann, P. (2011). “WAMs: Why Worry?”, *Philosophical Papers*, 40 (2), pp. 155 - 177
- Blome-Tillmann, M. (2008). “The indexicality of ‘knowledge’”, *Philosophical Studies*, 138, pp. 29-53
- (2009). “Knowledge and Presuppositions”, *Mind*, 118 (470), pp. 241 - 294
- Brown, J. (2006). “Contextualism and Warranted Assertibility Manoeuvres”, *Philosophical Studies*, 130 (3), pp. 407 - 435

²⁸ I’d like to thank Jessica Brown, Matthew Chrisman, Patrick Greenough, Michael Hannon, Allan Hazlett, Sebastian Köhler, Jared Peterson, Duncan Pritchard, and audiences at the 2011 University of Edinburgh Graduate Epistemology Conference, the 2011 Notre Dame/Northwestern Graduate Epistemology Conference and the University of Edinburgh postgraduate work in progress seminar. Special thanks are due to two anonymous reviewers for this journal for providing extremely helpful feedback on earlier versions of this paper. The research for this paper was funded by the Carnegie Trust.

- Cohen, S. (1999). "Contextualism, skepticism, and the structure of reasons", *Philosophical Perspectives*, 13, pp. 57–89
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Oxford University Press
- DeRose, K. (1995). "Solving the skeptical problem", *Philosophical Review*, 104 (1), pp. 1-52
(2009). *The Case for Contextualism*, Oxford: Clarendon
- (2011). "Contextualism, contrastivism, and X-Phi surveys", *Philosophical Studies*, 156 (1), pp. 81-110
- Hawthorne, J. (2004). *Knowledge and lotteries*, Oxford: Oxford University Press
- Hazlett, A. (2009). "Knowledge and Conversation", *Philosophy and Phenomenological Research*, 78 (3), pp. 591-620
- Henderson, D. (2009). "Motivated contextualism", *Philosophical Studies*, 142 (1), pp. 119-131
- Kelp, C. (2011). "What's the Point of 'Knowledge' Anyway?", *Episteme*, 8, pp. 53-66
- Lewis, D. (1996). "Elusive knowledge", *Australasian Journal of Philosophy*, 74, pp. 649–667
- May, J., Sinnott-Armstrong, W., Hull, J. G., & Zimmerman, A. (2010). "Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study", *Review of Philosophy and Psychology*, 1 (2), pp. 265–273
- McKenna, R. (2011). "Interests Contextualism", *Philosophia*, 39 (4), pp. 741-750
- Montminy, M. (Forthcoming). "The role of context in contextualism", *Synthese*
- Pritchard, D. (2001). "Contextualism, skepticism, and the problem of epistemic descent", *Dialectica*, 55 (4), pp. 327–349
- (2010). "Contextualism, Skepticism and Warranted Assertibility Maneuvers", in *Knowledge and Skepticism*, (Eds.) Keim Campbell, J., O'Rourke, M. & Silverstein, H., MIT Press, pp. 85-103
- Rysiew, P. (2001). "The context-sensitivity of knowledge attributions", *Noûs*, 35 (4), pp. 477–514
- Stanley, J. (2005a). *Knowledge and practical interests*, Oxford: Oxford University Press
- (2005b). "Semantics in Context", *Contextualism in Philosophy: Knowledge, Meaning and Truth*, (Eds.) Preyer, G. & Peter, G., pp. 221-255, Oxford: Oxford University Press
- Vogel, J. (1999). "The new relevant alternatives theory", *Philosophical Perspectives*, 13 (s13), pp. 155-180
- Williams, B. (1980). "Internal and External Reasons", reprinted with postscript in *Varieties of Practical Reasoning*, (Ed.) Millgram, E., 2001, MIT Press