

Normative scorekeeping

Robin McKenna

University of Edinburgh

rbnmckenna@gmail.com

Penultimate draft. Final version forthcoming in *Synthese*.

Abstract Epistemic contextualists think that the truth-conditions of ‘knowledge’ ascriptions depend in part on the context in which they are uttered. But what features of context play a role in determining truth-conditions? The idea that the making salient of error-possibilities is a central part of the story has often been attributed to contextualists, and a number of contextualists seem to endorse it (see Cohen 1999, 61; Hawthorne 2004, 63-6). In this paper I’m going to argue that the focus on salience relations is a mistake. On the view that I defend in this paper, the relevant features of context are facts about what error-possibilities and alternatives those in the context have a reason to consider, not facts about what error-possibilities and alternatives those in the context actually consider. As I’ll argue, this view has certain advantages over the standard view.

1. Introduction

Imagine that Sophie and Janine are discussing the shapes of European countries. Sophie says that France is hexagonal and Janine readily accepts this claim. Although France isn’t exactly hexagonal it’s plausible that, via Lewisian rules of accommodation, Sophie’s claim is true (see Lewis 1979). But what if Sophie were to say that France is boot-shaped and, for whatever reason, Janine readily accepted that claim? It isn’t at all plausible that this claim is true. While it’s plausible that conversational kinematics - in more prosaic terms, the course of the conversation - can play a role in determining the truth-values or truth-conditions of certain claims, there must be constraints on the effect of conversational kinematics on truth-values or truth-conditions.

Epistemic contextualism - henceforth, *contextualism* - is the view that the truth-conditions of ‘knowledge’ ascriptions and denials depend on and vary with the context in which they are uttered.¹ The idea is that the truth-conditions vary because sentences of the form ‘S knows that p’ express different propositions in different contexts of utterance.

¹ Prominent defenders include Cohen 1999, DeRose 2009 and Lewis 1996. To avoid any tricky issues with use and mention, I’ve followed the convention of always putting quotation marks around ‘knows’ and its cognates.

On this view, certain features of the conversational context play a role in determining the truth-conditions of ‘knowledge’ ascriptions. But what features? Stewart Cohen claims that “[i]n the case of knowledge ascriptions, salience relations play a central role” (Cohen 1999, 61). The idea that the making salient of error-possibilities and alternatives is an important part of the story has often been attributed to contextualists, and contextualists other than Cohen seem to endorse it (see DeRose 2009, 142; Hawthorne 2004, 63-6; Lewis 1996, 559-60; Stanley 2005, 23-7). Those that endorse the idea want to say that conversational kinematics are an important part of the explanation of how the truth-conditions of ‘knowledge’ ascriptions are determined by context.

I think that the contextualist focus on conversational kinematics and salience is a mistake. I’m going to argue that, just as there are constraints on what shape France can truly be said to be, there are constraints on the effect of conversational kinematics on the truth-conditions of ‘knowledge’ ascriptions. As I’ll also argue, a contextualist account of the features of context that determine the truth-conditions of ‘knowledge’ ascriptions should say something about the *nature* of these constraints. I’ll argue that there’s a view that meets this desideratum on which the features of context that determine the truth-conditions of ‘knowledge’ ascriptions are facts about what alternatives those in the context *have a reason* to consider. On this view - which I call *interests contextualism* - conversational kinematics aren’t a central part of the explanation of how the truth-conditions of ‘knowledge’ ascriptions are determined by context. Consequently, the focus on conversational kinematics is a mistake.²

I proceed as follows. First, I articulate four desiderata for a contextualist account of the features of context that determine the truth-conditions of ‘knowledge’ ascriptions. One of these desiderata - the one mentioned above - is that the account should say something about the nature of the constraints on the effect of conversational kinematics on the

² DeRose (2009, 141-2) considers what he calls a ‘pure reasonableness view’, on which the features of context that determine the truth-conditions of ‘knowledge’ ascriptions are facts about what epistemic standard it would be *most reasonable* for those in the context to use. This view differs from interests contextualism in that it’s put in terms of contextually varying epistemic standards rather than relevant alternatives (on this point see fn. 3), but otherwise it’s very much in the same spirit. However, DeRose doesn’t develop the view in any detail, and his main reason for rejecting it is that it’s hard to see what makes an epistemic standard reasonable. One of my aims in this paper is to explain what makes it the case that those in the context have reasons to consider certain alternatives.

truth-conditions of 'knowledge' ascriptions (§2). Second, I develop a version of contextualism - what I call interests contextualism - that meets the four desiderata (§3). Finally, I say a little about why interests contextualism does better than its competitors on this score (§4). While I'm not going to present any novel arguments for contextualism, I hope that, in showing that there's a way of developing the basic contextualist idea that focuses on what alternatives those in the context have a reason to consider, rather than on what alternatives are salient, or taken seriously, by those in the context, I'll succeed in making the basic contextualist idea more attractive to non-contextualists.

2. What makes an alternative relevant?

Contextualists think that the truth-conditions of 'knowledge' ascriptions vary with the context of utterance. I'm interested in the question of what features of the conversational context are supposed to explain this contextual variation. It will be useful to adopt a contextualist framework on which an utterance of a sentence of the form 'S knows that p' in context C expresses a proposition that's true iff S possesses evidence that rules out all the not-p alternatives that are relevant in C.³ On this framework, the features of context that determine the truth-conditions of 'knowledge' ascriptions are facts about what alternatives are relevant. So, in terms of this framework, my question becomes *what makes an alternative relevant?* In what follows I'll argue that there are four desiderata for a contextualist account of the relevance of alternatives.

The first desideratum: paradigm cases

Consider this pair of cases:

Normal: Ted and Dougal are driving to the bank on a Friday afternoon to deposit a cheque. They have no bills due, so they don't need to deposit the cheque immediately.

³ The framework is from Lewis (1996). I'll often talk as if the truth of a 'knowledge' ascription requires the ruling out of relevant alternatives but nothing of substance relies upon this. I could just as well talk of epistemic standards, and at a few points it will be convenient to talk in these terms. A translation scheme will be useful. Instead of saying, as I do here, that an utterance of a sentence of the form 'S knows that p' in context C expresses a proposition that's true iff S possesses evidence that rules out the not-p alternatives that are relevant in C one could say that an utterance of a sentence of the form 'S knows that p' in context C expresses a proposition that's true iff S's epistemic position with respect to p is strong enough to meet the epistemic standard operative in C (see DeRose 1995).

They notice a long queue, but Dougal remembers that he was in the bank on a recent Saturday, and he has no reason to think the opening hours have changed. Ted says 'Good, so you know that the bank will be open on Saturday. We can come back tomorrow.'

High: Jack is driving to the bank on a Friday afternoon to deposit a cheque. Jack needs to deposit the cheque by Saturday else he goes bankrupt. He notices the queue and calls Dougal, who tells him that he was in the bank on a previous Saturday, and that he has no reason to think the opening hours have changed. But Jack says 'I need to cash this cheque. How do you know that the bank hasn't just changed its opening hours?' Dougal agrees that he cannot rule that out and Jack says 'Right, so you don't know that the bank is open on Saturday' (I've taken these cases, with a few changes, from DeRose, 2009, Chapter 1).

Contextualists appeal to the supposed fact that we have the intuition that Ted speaks truly in *Normal* and Jack speaks truly in *High*. In this case, and others, our intuitions about the truth-values of 'knowledge' ascriptions supposedly shift with the context.⁴ By way of explanation, the contextualist posits contextual variation in the truth-conditions of 'knowledge' ascriptions. I take it that this sort of argument provides the basic motivation for adopting contextualism. As DeRose puts it: "The best grounds for accepting contextualism come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, nonphilosophical talk: What ordinary speakers will count as 'knowledge' in some non-philosophical contexts they will deny is such in others" (DeRose 2009, 47). So the first - and most obvious - desideratum for a contextualist account of the relevance of alternatives is that it should give the standard contextualist explanation of cases like *Normal* and *High*. Any answer that does this will, as I'll put it, *deal with the paradigm cases*.

The second desideratum: avoiding crazy results

Consider another case:

Conspiracy: Richard, Melvin and John are discussing their favourite topic: how blind everyone around them is to the various lies and myths perpetuated by the mainstream

⁴ For some contrary data, see Buckwalter (2010) and May et al. (2010). For responses, see DeRose (2011) and Sripada & Stanley (2012).

media.

Richard: Do you guys remember Dana? I was talking to her today and she still believes Neil Armstrong landed on the moon.

Melvin: Not that old myth. Didn't she watch that documentary you sent her? What about the clips where the flag is clearly fluttering?⁵

Richard: She said that she did, but that there was a straightforward explanation for everything.

John: Even the flag?

Richard: Even the flag. Something about how it appears to be fluttering because of the way it was stored. She told me I was crazy, that all the evidence points towards Armstrong landing on the moon, that she knows he landed on the moon.

John: How can she be so easily fooled? She doesn't know that. All of the scientists who try to explain away the evidence are just part of the conspiracy.

Melvin: Agreed.

Richard: And as for JFK...

The reader can, I'm sure, construct a number of cases along the lines of *Conspiracy*. Just take someone who holds what are often euphemistically termed 'non-mainstream' views and get them to deny that someone 'knows' all sorts of things that they, and everyone else, would ordinarily take themselves to 'know'. I take it that, in these sorts of cases, we don't think that these 'knowledge' denials are true. The second desideratum for a contextualist account of the relevance of alternatives is that it should avoid giving counter-intuitive results in cases like *Conspiracy*. Any answer that does this will, as I'll put it, *avoid crazy results*.

I'm going to offer two motivations for the second desideratum. First, I just think that it would be a black mark against any contextualist account if it gave counter-intuitive results in cases like *Conspiracy*. Second, I take it that, while the *Conspiracy* case is doubtless a little silly, it is representative of a feature of our epistemic discourse. Those who engage in the practice of ascribing and denying 'knowledge' to subjects need not always consider alternatives that we would regard as reasonable. The intuitive idea behind the second desideratum is that, when 'knowledge' ascribers consider unreasonable

⁵ For some claims made by moon landing conspiracy theorists (and their detailed rebuttals) see http://en.wikipedia.org/wiki/Moon_landing_conspiracy_theories-Photograph_and_film_oddities.

alternatives, their doing so should not have any effect on the truth-conditions of their 'knowledge' ascriptions and denials. (Similarly, that Sophie and Janine use a standard of precision on which France counts as boot-shaped should not have any effect on the truth of Sophie's claim that France is boot-shaped).

The third desideratum: the nature of constraints

One possible contextualist account of the relevance of alternatives is just that an alternative is relevant in a context iff it's taken seriously by those in the context. One way for an alternative to be taken seriously would be its being mentioned and not subsequently dismissed by the conversational participants. So, for example, in *High* the alternative that the bank has changed its opening hours is taken seriously because Jack mentions it, and both he and Dougal regard it as worthy of consideration. In *Normal* that alternative isn't taken seriously because it's not even mentioned. But there are alternatives that everyone in a context takes seriously that aren't mentioned in that context, perhaps because it's obvious to everyone that they are taken seriously. For example, imagine that Ted and Dougal are driving to the bank on Friday but need to deposit their cheque by Saturday else they go bankrupt. Dougal recalls being in the bank on a previous Saturday. Before Ted has a chance to respond, they hear a radio report about how banks in Greece frequently change their opening hours due to financial uncertainty. Once the report has finished, it may well become obvious to both Ted and Dougal that they are taking the possibility that the bank has changed its opening hours seriously, even if neither of them actually mentions it.

This answer satisfies the first desideratum because it deals with the paradigm cases. In *Normal* Ted's 'knowledge' ascription is true because Dougal's evidence is sufficient to rule out all of the relevant alternatives in which the bank isn't open on Saturdays (as above, the alternative that the bank has changed its opening hours isn't taken seriously). However, in *High* Jack's 'knowledge' denial is also true because Dougal's evidence isn't sufficient to rule out all of the relevant alternatives (as above, the alternative that the bank has changed its opening hours is taken seriously). Unfortunately, though, it doesn't satisfy the second desideratum because it doesn't avoid crazy results. In *Conspiracy* a scenario in which the moon landing was faked is taken seriously by all of the conversational participants. That makes it relevant, so John's 'knowledge' denial must be true. The problem with this answer is that it doesn't put any *constraints* on the effect of conversational kinematics on the relevance of alternatives. The third desideratum for a

contextualist account of the relevance of alternatives is that it should say something about the *nature* of these constraints (in particular, what the constraints are and why we have them).

I'll also offer two motivations for this desideratum. First, any contextualist account of the relevance of alternatives that meets the second desideratum is going to have to put constraints on the effect of conversational kinematics on the relevance of alternatives. In *Conspiracy* Richard, Melvin and John take the moon landing conspiracy scenario entirely seriously so, if it's not relevant, that must be because alternatives that are taken seriously need not be relevant. So meeting the second desideratum will already require at least saying what the constraints *are*. Second, and more importantly, if the proposed constraints are not to be ad hoc the contextualist owes us an explanation of *why* these are the constraints. Contextualists can't just posit constraints without explaining why we have them.

The fourth desideratum: tough questions

Here are some questions that contextualists need to answer:

1. What are the features of context that determine the truth-conditions of 'knowledge' ascriptions?
2. What are the mechanisms by which context shifts - expansions/retractions of the set of relevant alternatives, or upward/downward shifts in the epistemic standards - occur? (Pritchard 2001; 2010).
3. It's perhaps not that hard to see how the set of relevant alternatives might expand, or the epistemic standards rise. But how does the set *contract*, or the standard *fall*?

The first question is, of course, just the question I've been discussing in this section. The second and third questions are questions that contextualists have often found hard to answer. All else being equal, a version of contextualism that can answer these questions, and in particular a version that can give simple and informative answers to these questions, should be preferred over versions of contextualism that either give only relatively complicated or uninformative answers to these questions, or no answers at all. This is the fourth desideratum for a contextualist answer to the question what makes an alternative relevant.

I'm now going to argue that there's a version of contextualism - what I call interests contextualism - that meets all these desiderata.

3. Normative Scorekeeping

The basic idea behind interests contextualism is that an alternative is relevant in a context iff those in the context have a reason to consider it. I'm going to develop the view in two stages. First, I'll explain what I mean by 'having a reason to consider an alternative', and, drawing on that explanation, I'll present the version of interests contextualism I've defended in previous work (McKenna 2011; 2013). Second, I'll explain why that version needs to be modified, and I'll propose a modification that meets all of the desiderata on a contextualist account of the relevance of alternatives.

Interests contextualism: a first pass

Sophie has a glass in front of her and she's trying to decide whether she should drink from it.⁶ She wants a drink of gin but has no desire to drink something unpleasant such as petrol. The glass contains petrol but Sophie is completely unaware of this. Unfortunately, she believes that it contains gin. We can use this case to distinguish two sorts of practical reasons. In an objective sense of 'having a reason', Sophie hasn't got a reason to drink from the glass, but she mistakenly thinks that she does. In this same sense, Sophie has got a reason to consider the alternative that the glass contains an unpleasant liquid, but she mistakenly thinks that she doesn't. Sophie's practical interests - her desire to drink from the glass if it contains gin but not if it contains an unpleasant liquid - are such that she hasn't got a reason to drink from the glass and she has a reason to consider the alternative that the glass contains an unpleasant liquid. However, in a subjective sense of 'having a reason', Sophie has got a reason to drink from the glass (this is the sense of 'having a reason' relevant to explanations of an agent's actions). She thinks that it contains gin, and she wants some gin. In this same sense, Sophie hasn't got a reason to consider the alternative that the glass contains an unpleasant liquid. She's completely unaware that the glass contains petrol, and nothing she has seen suggests that it does. As far as Sophie is concerned, her practical interests are such that she has a reason to drink from the glass and she hasn't got a reason to consider the alternative that the glass contains an unpleasant liquid.

⁶ This example, and the discussion that follows it, draws heavily on Williams (1980).

On the interests contextualist view, the reasons that ‘knowledge’ ascribers have to consider alternatives are of the same type as the reason Sophie has to consider the alternative that the glass contains an unpleasant liquid. So the view I’m calling *interests contextualism* is the combination of two theses:

- i. ‘Knowledge’ ascriptions have context-sensitive truth-conditions. In terms of the framework I’m working with in this paper: A ‘knowledge’ ascription of the form ‘S knows that p’ is true in context C iff S’s evidence is sufficient to rule out all of the alternatives that are relevant in C.
- ii. An alternative is relevant in a context C iff those in C have a reason, in the objective practical sense, to consider it.

Thesis (i) is just the contextualist semantic thesis. Thesis (ii) is distinctive of interests contextualism. (Note that the two are connected, in that thesis (ii) explains why ‘knowledge’ ascriptions have context-sensitive truth conditions).

It might be helpful to illustrate the interests contextualist view by explaining how it deals with some of the cases I’ve discussed so far. Consider Jack’s situation in *High*. It’s in Jack’s practical interests to get the bill paid, and that won’t happen unless the cheque is cashed by Saturday. Given these facts about Jack, he has a reason to consider the alternative that the bank has changed its opening hours. His practical interests are such that he has a reason to consider it. This is comparable to the reason Sophie has to consider the alternative that the glass contains an unpleasant liquid. Compare this to Ted and Dougal’s situation in *Normal*. They don’t want to waste time unnecessarily, but otherwise it’s all the same to them whether they go to the bank on Friday, Saturday, or next week. Given these facts, Ted has no reason to consider the alternative. Ted and Dougal’s practical interests are such that they have no reason to consider the alternative. None of this would change if Jack were in the same practical situation but didn’t consider the alternative. He would still have a reason to consider the alternative that the bank has changed its opening hours. Similarly, if Ted were in the same practical situation but did take the alternative seriously he would still not have a reason to consider it.⁷

⁷ It’s worth emphasising that the view I’m presenting here differs from the view defended in Stanley (2005). First, Stanley’s view - which he calls ‘interest-relative

On the interests contextualist view, there are *normative* constraints on the effect of conversational kinematics on the relevance of alternatives. Alternatives that are actually taken seriously can still be irrelevant, provided those in the context haven't got a reason to consider them, and alternatives that are actually ignored can still be relevant, provided those in the context have got a reason to consider them. There are normative elements - facts about what alternatives those in the context have a reason to consider - of the conversational score. We can think of this as a sort of *normative scorekeeping*.

I'm now going to consider two objections to interests contextualism. I'll argue that the interests contextualist can deal with the first objection, but that the second objection shows that the view needs to be modified. It's the modified view that, as I'll argue, satisfies the four desiderata for a contextualist account of the relevance of alternatives.

The interests contextualist says that an alternative is relevant in a context iff those in the context have a reason to consider it, where by 'reason' is meant an objective practical reason. But what about subjective practical reasons? Are there not cases where, intuitively, certain alternatives are relevant yet those in the context only have a subjective, not objective, reason to consider them? An example might be Jonathan Vogel's parked-car case (see Vogel 1999). 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'm not particularly worried about it. Consequently, 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, as a result of her making these possibilities vivid, I deny that I 'know' where my car is parked. Clearly, her making

invariantism' - is a view about knowledge, whereas interests contextualism is a view about the truth-conditions of 'knowledge' ascriptions. Second, for the interests contextualist it's the practical interests of the 'knowledge' ascriber that matter, whereas on Stanley's view it's the practical interests of the subject that matter. One dimension along which these views can be evaluated and compared is with respect to the range of cases discussed in Stanley (2005), and I carry out this task in McKenna (2013). A brief summary: First, I show that, while interests contextualism and interests-relative invariantism give the same result in certain cases (cases like *Normal* and *High*, and variants of those cases where those in the context are ignorant of the stakes), they give different results in cases where the 'knowledge' ascriber and the subject of the 'knowledge' ascription are in different contexts. Second, I argue that, on balance, these cases are more problematic for the interests-relative invariantist than the interests contextualist.

these possibilities vivid might alter what alternatives I have a *subjective* practical reason to consider (the making salient of possibilities may well affect how I perceive my practical interests). But, so the objection would go, it's not clear that her doing so is going to alter what alternatives I have an *objective* practical reason to consider (it's not clear that the making salient of possibilities can affect my actual practical interests). But that means that the interests contextualist would have to say that I would be wrong to deny that I 'know' where my car is parked even though my friend has made possibilities in which my car has been stolen particularly vivid. This, so the objection would go, is counter-intuitive.

There are three responses that the interests contextualist can make here. First, she could just deny that interests contextualism says anything counter-intuitive about this case. I regard this option as a non-starter. Second, she could modify her view. Instead of saying that an alternative is relevant in a context C iff those in the context have an objective practical reason to consider it, she could say that an alternative is relevant in a context C iff those in the context have a practical reason, whether objective or subjective, to consider it. Third, she could argue that the making salient of possibilities can sometimes affect what alternatives one has a reason, in the objective practical sense, to consider.

I'm inclined to think that the interests contextualist can make the third response work, at least in the case currently under discussion. The making salient of possibilities - here, the possibility that my car might have been stolen - may well affect my actual practical interests, and so affect what alternatives I have a reason to consider. Prior to talking to my friend, I wasn't at all worried about the possibility that my car might be stolen. But perhaps her mentioning those possibilities made me realise that I *should* have been worried. In that case, I would have a reason, in the objective practical sense, to consider the possibility. (While the third response seems to work in the case currently under discussion, what if there are cases in which this response isn't convincing? If there are such cases, the interests contextualist should endorse the second response).

However, the *Conspiracy* case shows that, as it stands, interests contextualism is inadequate. In *Conspiracy* John, Richard and Melvin consider a rather unlikely error-possibility on which scientists are part of some grand moon-landing conspiracy. Presumably, given the importance those who believe in them usually attach to conspiracy theories, John, Richard and Melvin's practical interests are such that they actually have a

reason to consider the rather unlikely error-possibilities. If that's right, then the interests contextualist has to say that John speaks truly when he denies that Dana 'knows' Neil Armstrong landed on the moon. So interests contextualism doesn't satisfy the second desideratum for a contextualist account of the relevance of alternatives.

One might wonder why the interests contextualist can't just appeal to a substantive view of our practical interests. The idea would be that despite the fact that, if we attach a great deal of importance to conspiracy theories, it may be in our immediate interests to consider certain outlandish error-possibilities, our longer term interests are best served by believing the truth, or something like that. Insisting on alternatives in which the US government conducts elaborate hoaxes for no obvious reason is not a good way of achieving true beliefs. However, while the interests contextualist could say this, I don't think that she should. I agree that, looking at things from a certain perspective, there's something wrong with considering outlandish conspiracy theories. But I don't think that perspective is available to the interests contextualist. One might want to say that, while John, Richard and Melvin have a *practical* reason to consider outlandish conspiracy theories, they don't have an *epistemic* reason to consider them. Plausibly, epistemic reasons to consider alternatives, unlike practical reasons, must somehow be connected with the aim of believing truly. But interests contextualism is a view about the connection between the relevance of alternatives and practical reasons, not epistemic reasons.

I'm now going to argue that, suitably modified, interests contextualism can meet all of the desiderata for a contextualist account of the relevance of alternatives. I'll start with an explanation of why, if contextualism is true, we should expect there to be constraints on the effect of conversational kinematics on the relevance of alternatives.

Constraints and Craigian genealogy

My explanation draws on Edward Craig's account of the function of 'knowledge' ascriptions. On Craig's account, 'knowledge' ascriptions serve the function of identifying good informants.⁸ The idea is that, in ascribing knowledge that *p* to a subject *S*, I certify

⁸ Craig develops his account in Craig (1990). I'm only going to discuss Craig's account in very broad outline here, but I should emphasise that, while there may be problems with the details of Craig's account, all I need to support my argument is that something in the general ballpark be correct. For further discussion see Kappel (2010) and Kusch (2011).

or ‘tag’ S as a good informant on the matter of p. I’ll start by sketching the considerations that Craig uses to motivate adopting this idea, and I’ll then outline the sort of contextualist view that Craig’s account of the function of ‘knowledge’ ascriptions might support. Just to be clear: My aim here is not to show that Craig’s account favours a sort of contextualism over invariantism. Rather, my aim is to outline a sort of contextualist view that incorporates certain constraints on the effect of conversational kinematics on the relevance of alternatives.⁹

Craig motivates his account by appealing to the observation that humans have a need for information. We can obtain some of the information that we need under our own steam through perception, introspection and the like. However, due to our limitations - both cognitive and practical - we can’t obtain all of the information that we need ourselves. So, if I can’t obtain some information myself, I need to be able to find out who does. The practice of ‘knowledge’ ascription is what meets this need. To illustrate: Say a police officer is chasing a suspect. She has no way of tracking the suspect but if she were to call her boss he could dispatch the unit helicopter. Hopefully, the pilot would locate the suspect and tell her where he is. Of course, at this point it’d be very important to identify the right man. She might question the pilot, and, if satisfied, ascribe ‘knowledge’ of where the suspect is to the pilot. In doing so she’s identifying the pilot as a good informant as to the whereabouts of the suspect. The police officer, and anyone else who is trying to catch the suspect, can use the pilot’s information.

Why might Craig’s account support a form of contextualism? At first blush, things look fairly simple: If what I’m doing in ascribing ‘knowledge’ that p to a subject S is identifying S as a good informant on the matter of p, there will be contextual variation in the truth-conditions of ‘knowledge’ ascriptions because what it takes to be a good informant in one context may differ from what it takes to be a good informant in another context. When I’m in a rather pressing practical situation, it’s going to take a lot for someone to qualify as a good informant. Putting things in terms of relevant alternatives, someone is going to have to be able to rule out a wide range of alternatives in order to qualify. However, when I’m in a less pressing situation, it’s going to take a lot less for someone to qualify. Again putting things in terms of relevant alternatives,

⁹ While I don’t argue that Craig’s account favours a sort of contextualism over invariantism here, Fricker (2008), Greco (2009), Hannon (forthcoming) and Henderson (2009) do.

someone is going to have to be able to rule out a narrower range of alternatives to qualify. In the example just given, the police officer is in an unusually pressing practical situation. She can't risk arresting the wrong man and, in the process, losing the suspect, so she's going to require that the pilot can rule out a wide range of alternatives before she tags him as a good informant. Compare this with a case in which, as part of a training programme, the pilot has been dispatched to locate a man fitting a certain description. The same police officer might question the pilot and, if satisfied, ascribe 'knowledge' of where that man is to the pilot. In this example, though, she's not in a particularly pressing practical situation so it's not going to take as much for her to tag the pilot as a good informant. She'll require the pilot to be able to rule out a narrower range of alternatives.

On the Craigian picture as I've sketched it so far, it looks like whether a subject qualifies as a good informant depends entirely on the practical situation of the person ascribing (or denying) 'knowledge' to that subject because, for the reasons just given, people will differ in what they require of a good informant. But Craig emphasises that, often, people will actually *not* differ in what they require of a good informant. This is because we don't tag good informants, or make 'knowledge' ascriptions, in a social vacuum (see Craig 1990, Chapter 10 on 'objectivisation'). Often, when I tag a subject S as a good informant on the matter of p I'm not just identifying S as someone that those in the immediate vicinity can rely on for the information that p, I'm identifying S as someone who the community at large can rely on. This, argues Craig, leads to pressure for each member of a community of 'knowledge' ascribers to require that a subject can rule out a similar, or perhaps even the same, range of alternatives in which not-p when evaluating that subject as an informant on the matter of p, or ascribing 'knowledge' that p to the subject. It will be helpful to think of this in terms of a range of *default* relevant alternatives. The idea is that, for each proposition p, there is a range of alternatives in which not-p that are default relevant.¹⁰ This might lead one to think that the end result is something that

¹⁰ I borrow this way of putting things from Hannon (forthcoming). Can the defender of this Craigian story say anything more about the default relevant alternatives? The idea is that, for each proposition p, a community of 'knowledge' ascribers will converge on a relatively stable range of alternatives in which not-p that are relevant. It's important to note that someone who defends this part of the Craigian story doesn't have to specify how wide the range of default relevant alternatives is. After all, a 'traditional' or 'strict' invariantist might think that, for each proposition p, a stable range of alternatives in

approximates (but maybe isn't quite) a 'traditional' invariantist view, on which the range of alternatives one must be able to rule out in order to count as 'knowing' that p doesn't vary from context to context (for this interpretation see Kelp 2011).

However, I don't think this is a good reason for concluding that Craig's account leads to a sort of invariantism. A prominent contextualist idea - and an idea that I endorse - is that what alternatives are relevant in a context somehow depends on the practical interests and purposes of those in the context (see also Greco 2009 and Henderson 2009). I take it that, on the Craigian picture just sketched, the relatively stable range of alternatives in which not-p that one must be able to rule out to count as 'knowing' that p somehow depends on the practical interests and purposes of the *community at large*. We're a community of knowledge ascribers, and the default relevant alternatives will depend upon facts about *us*, in particular the sorts of practical projects that we engage in. I'm going to call this *global interest-relativity* ('global' because the default relevant alternatives depend on the whole community, 'interest-relative' because the default relevant alternatives depend on the practical interests and purposes of the whole community). But, just because the community of which we're all a part generally requires that a subject can rule out a similar, or perhaps even the same, range of alternatives in which not-p in order to count as 'knowing' that p doesn't mean that, in every case, that subject must be able to rule out that range of alternatives in order to count as 'knowing' that p. Instead, it's plausible that, in certain cases - cases where our practical situation is unusually pressing - alternatives that are usually irrelevant can be relevant. I'll call this *local* (as opposed to global) interest-relativity ('local' because the range of relevant alternatives depends on particular practical situations).

Examples of such cases include the sorts of cases that provide the basic motivation for contextualism (i.e. *Normal* and *High*). What cases like *High* - that is, cases where 'knowledge' ascribers are in unusually pressing practical situations - show is that, sometimes, alternatives that are usually irrelevant can be relevant. While in our community we don't generally require that a subject can rule out alternatives in which the bank has changed its opening hours in order to credit that subject with 'knowing' that the bank will be open on Saturday - those alternatives aren't default relevant - such

which not-p are relevant, but she doesn't need to specify how wide the range is. The Craigian story is a story about how the range came about - through social pressures - not about how wide it is.

alternatives can be relevant in cases where ‘knowledge’ ascribers are in unusually pressing practical situations.

Summarising the above, the sort of contextualism that Craig’s account of the function of ‘knowledge’ ascriptions might support consists of two theses:

- i. ‘Knowledge’ ascriptions have context-sensitive truth-conditions. For convenience, we can understand this in terms of contextually varying ranges of relevant alternatives.
- ii. The alternatives that are relevant in a context C are the alternatives that are appropriate given the practical situation of those in C. Often, the alternatives that are appropriate given the practical situation of those in C will just be the default relevant alternatives. However, when those in C are in an unusually pressing practical situation, a wider range of alternatives will be appropriate.

I’ll call the combination of these theses *Craigian contextualism*. I want to note two things about Craigian contextualism.

First, what I’m calling Craigian contextualism is really just a template. To get a concrete view, something has to be said about what makes an alternative appropriate in a particular practical situation. In the next sub-section I’ll argue that, suitably modified, interests contextualism gives us a good way of understanding what makes an alternative appropriate.

Second, the Craigian contextualist not only says something about the nature of the constraints on the effect of conversational kinematics, she can also explain why there are constraints in the first place. The Craigian contextualist holds that contextual variation in the truth-conditions of ‘knowledge’ ascriptions is going to be rather limited for two reasons. First, in a lot of contexts the alternatives that are appropriate will just be the default relevant ones. Second, contextual variation will be tied to changes in the practical situation of ‘knowledge’ ascribers, not to the making salient of error-possibilities, and conversational kinematics more generally. The Craigian contextualist explains both reasons by reference to the function of ‘knowledge’ ascriptions. So, on the Craigian contextualist view, the truth-conditions of ‘knowledge’ ascriptions don’t, to borrow a (borrowed) phrase, “swing freely with whatever moves happen to be made in language games” (Henderson 2009, 126). The view is one on which contextual variation is actually

rather limited, and is tied to shifts in the practical situation of 'knowledge' ascribers. For any proposition p , certain alternatives in which not- p are always relevant - the default relevant ones - but, when 'knowledge' ascribers are in unusually pressing practical situations, certain other alternatives may become relevant.

Interests contextualism: a second pass

On the Craigan story just sketched, 'knowledge' ascribers are required to consider certain alternatives - the default relevant ones - just in virtue of being members of an epistemic community. The idea was that, if I'm a member of a community in which someone must be able to rule out a certain range of alternatives in which not- p in order to count as 'knowing' that p , then I'm required to consider those alternatives when ascribing 'knowledge' that p to a subject. We can say, then, that 'knowledge' ascribers have reasons to consider default relevant alternatives in virtue of being members of epistemic communities.

On the version of interests contextualism that I discussed above, 'knowledge' ascribers have reasons to consider certain alternatives in virtue of their practical interests. The idea was that an alternative is relevant in a context iff those in the context have a reason, in the objective practical sense, to consider it. As we saw, the problem with this version of the view was that 'knowledge' ascribers might have reasons, in the objective practical sense, to consider alternatives that, intuitively, aren't relevant (for example, the reason John, Richard and Melvin have to consider moon landing conspiracy theories in *Conspiracy*).

A modified version of interests contextualism results from combining these two different ways of thinking about reasons to consider alternatives. The Craigan contextualist says that an alternative is relevant in a context C iff it's appropriate given the practical situation of those in C . I want to propose that we can understand the appropriateness of an alternative in terms of those in the context having a reason to consider it. As above, there are two sorts of reasons those in the context might have to consider an alternative. 'Knowledge' ascribers have reasons to consider alternatives just in virtue of being a member of an epistemic community. But, when 'knowledge' ascribers are in unusually pressing practical situations, they will also have reasons to consider alternatives in virtue of their practical interests. These sorts of reasons will be objective practical reasons.

This gives us a modified version of interests contextualism consisting of two theses:

- i. 'Knowledge' ascriptions have context-sensitive truth-conditions. For convenience, we can understand this in terms of contextually varying ranges of relevant alternatives.
- ii. An alternative is relevant in a context C iff those in the context have a reason to consider it. Those in the context will always have reasons to consider alternatives in virtue of being members of a certain epistemic community. However, when those in the context are in an unusually pressing practical situation, they will have reasons to consider alternatives in virtue of their practical interests. These will be objective practical reasons.

I'm going to argue that this version of interests contextualism satisfies all of the desiderata for a contextualist account of the relevance of alternatives. First, though, I want to make two comments.

First comment: How does this version of interests contextualism differ from the original version discussed earlier? The key difference is that, on the original view, those in the context just have reasons to consider alternatives in virtue of their practical interests whereas, on the modified view, those in the context also have reasons to consider alternatives in virtue of being members of a certain epistemic community. Again: Those in the context are required to consider certain alternatives - the default relevant ones - just in virtue of being members of an epistemic community. So, on the modified version of the view, certain alternatives - the default relevant ones - are relevant in all contexts. This is not so on the original version of the view.¹¹

Second comment: It's plausible that there are multiple epistemic communities rather than one very large one, and that different epistemic communities will have different practical interests and purposes. But, because the range of alternatives that are default relevant in a context depends on the interests and purposes of the community, different sets of alternatives are going to be default relevant in different contexts. So why do I claim that there is going to be 'a set' of default relevant alternatives, or that there's going to be 'fairly limited' contextual variation in the truth-conditions of 'knowledge' ascriptions? These are good questions, and they raise further questions about the structure of

¹¹ Thanks to a referee for this journal for pushing me to say more about the differences between the modified view and the original view.

epistemic communities, for example: How are communities individuated? What are the membership criteria?. Of course, my aim here is to outline a contextualist view, not to develop it in full detail. However, I do need to say something by way of answer to these questions.¹²

First, recognising that there are multiple epistemic communities means acknowledging that there are many sets of default relevant alternatives, one for each epistemic community. Examples of epistemic communities might include the scientific community, the legal community and the historical community. For each of these communities there may be a different set of default relevant alternatives, and in each case what that set is will depend on the interests and purposes of those in the community. Second, while recognising that there are multiple epistemic communities will add a great deal of complexity to the view just outlined, that complexity won't affect the crucial point that there are going to be certain limits on contextual variation in the truth-conditions of 'knowledge' ascriptions. On a fairly natural picture, any 'knowledge' ascriber is a member of a number of epistemic communities, and she will have reasons to consider certain alternatives in virtue of being a member of those communities. While, for any given alternative, she will have reason to consider it in virtue of being a member of some communities but not others, the alternatives that she has reason to consider will still partly depend on and be constrained by which epistemic communities she is a member of. So, even though a 'knowledge' ascriber is a member of a number of epistemic communities, it will still be the case that, for any proposition *p*, that 'knowledge' ascriber has reason to consider certain alternatives in which not-*p*, viz. those that are a member of one of the sets of default relevant alternatives. So, in any context, certain alternatives are always relevant.

I'll now argue that the modified version of interests contextualism satisfies all four of the desiderata for a contextualist account of the relevance of alternatives. I'll take each desideratum in turn.

First, the view deals with the paradigm cases. In *Normal* Ted and Dougal haven't got a reason to consider the alternative that the bank has changed its opening hours (it's not default relevant, and they aren't in an unusually pressing practical situation) so they speak

¹² Thanks to a referee for this journal for pressing me on these questions.

truly; in *High Jack* has got a reason to consider that alternative (it's not default relevant, but he is in an unusually pressing practical situation) so he also speaks truly. Second, the view avoids crazy results. In *Conspiracy John*, Richard and Melvin haven't got a reason to consider moon landing conspiracy theories because such alternatives aren't default relevant and, even though they have got an objective practical reason to consider the alternative, they aren't in an unusually pressing practical situation. Third, the view says what the constraints on the effect of conversational kinematics on the relevance of alternatives are, and it explains why we have them. As I put it above, these constraints are normative. There are normative elements - facts about what alternatives those in the context have reason to consider - of the conversational score. We have these constraints because of the function played by 'knowledge ascriptions'. Fourth, the view gives simple and informative answers to the three questions identified above, viz.

1. What are the features of context that determine the truth-conditions of 'knowledge' ascriptions?

Answer: Facts about what alternatives those in the context have reason to consider, where having a reason to consider an alternative is either cashed out in terms of belonging to an epistemic community, or objective practical reasons.

2. What are the mechanisms by which context shifts - expansions/retractions of the set of relevant alternatives, or upward/downward shifts in the epistemic standards - occur?

Answer: A relatively stable range of alternatives is always relevant (the default relevant alternatives). But, when ascribers are in unusually pressing practical situations, the range of relevant alternatives expands.

3. It's perhaps not that hard to see how the set of relevant alternatives might expand, or the epistemic standards rise. But how does the set *contract*, or the standard *fall*?

Answer: Both expansions and retractions in the set of relevant alternatives occur when the practical situation of those in the context changes. There is no reason why there should be a special mystery about how the set contracts.

In short, interests contextualism is a plausible view that satisfies all of the desiderata discussed in this paper. I'm going to finish by saying a little about why other versions of contextualism fall short on this score.

4. Rival accounts

The reader might worry that I've been too quick in dismissing contextualist views that appeal to the role played by salience relations. While it's common to distinguish views on which the alternatives that are taken seriously in a context play the central role in determining what alternatives are relevant (views that emphasise the role played by salience relations) from views on which the practical situation of those in the context plays the central role (such as interests contextualism) people who defend the former sort of view readily accept that salience relations can't be the whole story. Here's 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, 61).

The basic idea seems to be that specifying the features of context that determine the truth-conditions of ‘knowledge’ ascriptions is a difficult task, and while contextualists can pick out features that are particularly important - like salience relations - we shouldn't expect there to be an exhaustive list of the relevant features.¹³ If Cohen's right, that there are constraints on the effect of conversational kinematics on the relevance of alternatives isn't surprising. There are a number of features of conversational context that determine the truth-conditions of ‘knowledge’ ascriptions, and some of these features might provide constraints on conversational kinematics. It will be useful to have a name for this general line of thought so I'll call it ‘pluralism’.

I must admit that it's not entirely clear to me what pluralism amounts to. If the thought is just that because there are a number of features of context there might be features such that, in *Conspiracy* moon landing conspiracy theory scenarios aren't relevant, then it seems clear that pluralism doesn't do at all well with respect to the third and fourth

¹³ DeRose (2009) and Ichikawa (2011) both seem to endorse a similar view.

desiderata.¹⁴ On the third desideratum: That there might be constraints on the effect of conversational kinematics on the truth-conditions of ‘knowledge’ ascriptions doesn’t even tell us what the constraints are, never mind why we have them. On the fourth: if the pluralist isn’t able to say what exactly the features are, she isn’t going to be able to give simple and informative answers to the three questions discussed above. Clearly, then, pluralism must amount to more than the mere claim that there are a number of features of context that determine the truth-conditions of ‘knowledge’ ascriptions.

As far as I’m aware, the only real attempt to specify all of the relevant features of context is that of David Lewis. Lewis provides a number of rules governing the relevance of alternatives, some of which - for example, the rule of actuality (the rule that the alternative that actually obtains is always relevant) - can be considered as constraints on the effect of conversational kinematics on the relevance of alternatives (whether those in the context consider the actually obtaining alternative or not, it’s still relevant) (see Lewis 1996, 554-560). While there are no doubt problems with Lewis’s particular rules, it will be useful to take his view as an example of the sort of approach that the pluralist must take. In explaining why I think the interests contextualist does better with respect to the desiderata than Lewis, I’ll hopefully demonstrate the challenges that the pluralist will face in dealing with the desiderata.

As it stands, Lewis’s view has problems with *Conspiracy*. The problem is Lewis’s rule of attention, on which any alternative that those in the context attend to is automatically relevant (see Lewis 1996, 559-60). It seems that, no matter how one understands attending to an alternative, John, Richard and Melvin attend to conspiracy theory scenarios in *Conspiracy*. (Note that the best way of understanding attention that I’m aware of - on which an alternative is attended to iff it’s negation isn’t part of the conversational common ground in the sense of Stalnaker 1974 - doesn’t help here (see Blome-Tillman 2009). In *Conspiracy*, no matter how the common ground is understood, the negation of conspiracy theory scenarios isn’t going to be in it). However, with sufficient ingenuity, I imagine one could alter the view to deal with that sort of case (and,

¹⁴ As a referee for this journal pointed out to me, Lewis (1979) puts ‘within certain limits’ provisos on the rules of accommodation he proposes (see, for example, 340-1; 349). Lewis accepts that there are certain limits (or, constraints) on conversational kinematics, but the challenge for Lewis (and the pluralist) is to specify those constraints (or limits) in a non-ad-hoc and informative way. I discuss some of the constraints proposed in Lewis (1996) below.

indeed, any sort of problematic case). The advantages that I want to claim for interests contextualism concern the third and fourth desiderata.

As I understand it, the fourth desideratum for a contextualist account of the relevance of alternatives is that it should provide simple and informative answers to a number of questions. The interests contextualist does well on simplicity because, on her view, the relevance of an alternative is just understood in terms of the reasons that ‘knowledge’ ascribers have to consider alternatives. As I illustrated above, this framework gives us simple answers to the relevant questions. The interests contextualist also does well on informativeness because the notion that she appeals to - reasons to consider alternatives - can be elucidated, along the lines pursued in this paper. Compare this with Lewis’s list of rules governing the relevance of alternatives. In order to deal with problematic data - for example, cases like *Conspiracy* - new rules will need to be added, and the end result will be a complex, unwieldy and entirely ad hoc list. (In contrast, the interests contextualist gets by with just a single rule). Further, Lewis’s existing rules make use of notions, such as attending to an alternative, or an alternative saliently resembling actuality, that aren’t clear, and haven’t been elucidated.¹⁵ (In contrast, the interests contextualist makes use of a notion - having a reason to consider an alternative - that can be elucidated). Relatedly (and this is relevant to the charge of ad-hocery) Lewis doesn’t offer any sort of explanation *why* these are the rules governing the relevance of alternatives. As I argued above, a contextualist account of the relevance of alternatives should at least say something about the nature of the constraints on the effect of conversational kinematics on the relevance of alternatives (the third desideratum). What are the constraints? Why do we have them? Because the interests contextualist endorses the Craigian story told above, she can answer these questions, whereas it’s not clear whether a defender of Lewis’s version of contextualism can.

The moral I want to take from this is that any attempt to make good on the pluralist’s claim that there are a number of features of context that determine the truth-conditions of ‘knowledge’ ascriptions will face a number of challenges. First, can the pluralist avoid appealing to a very large number of rules governing the relevance of alternatives? Second, can she avoid ad hoc appeals to rules in order to deal with tough cases like

¹⁵ Another of Lewis’s rules - the rule of resemblance - has it that any alternative that saliently resembles actuality is relevant (see Lewis 1996, 556-8).

Conspiracy? Third, can she appeal to rules that make use of notions that are clear, or can at least be elucidated? Fourth, can she explain why we have these rules? As I've argued in this paper, the interests contextualist can meet all of these challenges. Whether the pluralist can is an open question.

5. Conclusion

In this paper I have argued that a contextualist account of the relevance of alternatives should satisfy four desiderata. First, it should deal with the paradigm cases like *Normal* and *High*. Second, it should avoid giving counter-intuitive results in cases like *Conspiracy*. Third, I've argued that what cases like *Conspiracy* show is that there are constraints on the effect of conversational kinematics on the relevance of alternatives. A contextualist account should say something about the nature of these constraints. Fourth, it should give simple and informative answers to certain questions. I then argued that there's a version of contextualism - interests contextualism - that satisfies all four desiderata.

As I mentioned in §1, the default assumption has often been that contextualists think that conversational kinematics and salience relations are an important part of the explanation of how the truth-conditions of 'knowledge' ascriptions are determined by context. On the interests contextualist view, this assumption is mistaken. For the interests contextualist, the features of context relevant to determining the truth-conditions of 'knowledge' ascriptions are not facts about what alternatives those in the context consider or take seriously; they are facts about what alternatives those in the context have a reason to consider. While I haven't said much about why we should think contextualism is true, I think that the view I have presented here is a more plausible way of developing the basic contextualist idea than the other versions of contextualism that I've discussed in this paper. I hope that this paper has at least succeeded in making contextualism more attractive to non-contextualists.¹⁶

¹⁶ Thanks to two anonymous referees for this journal for helpful feedback on an earlier version of this paper. Thanks also to Michael Blome-Tillmann, Jessica Brown, Matthew Chrisman, John Greco, Patrick Greenough, Michael Hannon, Allan Hazlett, Torfinn Huvenes, Dirk Kindermann, Sebastian Köhler, Jared Peterson, Duncan Pritchard, Amy Seymour, audiences at the Arché Epistemology Seminar, the 2012 European Epistemology Network Meeting in Bologna, the 2011 Notre Dame/Northwestern Graduate Epistemology Conference and the 2011 Edinburgh Graduate Epistemology Conference. The research for this paper was funded by the Carnegie Trust.

Bibliography

- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind* 118 (470): 241-294
- Buckwalter, W. (2010). Knowledge isn't closed on Saturday: a study in ordinary language. *Review of Philosophy and Psychology* 1 (3): 395-406
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives* 13: 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): 1-52
- _____(2009). *The case for contextualism*. Oxford: Clarendon
- _____(2011). Contextualism, contrastivism, and x-phi surveys. *Philosophical Studies* 156 (1): 81-110
- Fricker, M. (2008). Scepticism and the genealogy of knowledge: situating epistemology in time. *Philosophical Papers* 37 (1): 27-50
- Greco, J. (2009). Knowledge and success from ability. *Philosophical Studies* 142 (1): 17-26
- Hannon, M. (forthcoming). The practical origins of epistemic contextualism. *Erkenntnis*
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press
- Henderson, D. (2009). Motivated contextualism. *Philosophical Studies* 142 (1): 119-131
- Ichikawa, J. (2011). Quantifiers and epistemic contextualism. *Philosophical Studies* 155 (3): 383-398
- Kappel, K. (2010). On saying that someone knows: themes from Craig. In Haddock, A., Millar, A. & Pritchard, D. (eds.), *Social epistemology*. Oxford: Oxford University Press: 69-88
- Kelp, C. (2011). What's the point of 'knowledge' anyway? *Episteme* 8: 53-66
- Kusch, M. (2011). Knowledge and certainties in the epistemic state of nature. *Episteme* 8: 6-23
- Lewis, D. (1979) Scorekeeping in a language game. *Journal of Philosophical Logic* 8: 339–59
- _____(1996). Elusive knowledge. *Australasian Journal of Philosophy* 74: 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): 265–273
- McKenna, R. (2011). Interests contextualism. *Philosophia* 39 (4): 741-750
- _____(2013). Epistemic contextualism: a normative approach. *Pacific Philosophical Quarterly* 94 (1): 101-123

- Pritchard, D. (2001). Contextualism, skepticism, and the problem of epistemic descent. *Dialectica* 55 (4): 327–349
- _____(2010). Contextualism, skepticism and warranted assertibility maneuvers. In Keim Campbell, J., O'Rourke, M. & Silverstein, H. (eds.), *Knowledge and skepticism*. MIT Press: 85-103
- Stalnaker, R. (1974). Pragmatic presuppositions. Reprinted in his *Context and content* [2001]. Oxford: Oxford University Press: 47–62
- Sripada, S. & Stanley, J. (2012). Empirical tests of interest-relative invariantism. *Episteme* 9 (1); 3-26
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press
- Vogel, J. (1999). The new relevant alternatives theory. *Philosophical Perspectives* 13: 155-180
- Williams, B. (1980). Internal and external reasons. Reprinted with postscript in Millgram, E. (ed.), *Varieties of practical reasoning* [2001]. MIT Press: 77–97