'Knowledge' ascriptions, social roles and semantics

Robin McKenna

University of Edinburgh

rbnmckenna@gmail.com

Abstract The idea that the concept 'knowledge' has a distinctive *function* or *social role* is increasingly influential within contemporary epistemology. Perhaps the best-known account of the function of 'knowledge' is that developed in Edward Craig's *Knowledge and the state of nature* (1990, OUP), on which (roughly) 'knowledge' has the function of identifying good informants. Craig's account of the function of 'knowledge' has been appealed to in support of a variety of views, and in this paper I'm concerned with the claim that it supports a sort of epistemic contextualism, which is (roughly) the view that the semantic contents and truth-conditions of 'knowledge' ascriptions - instances of 'S knows that p'- depend on and vary with the context of ascription (see, for instance, Greco, 'What's wrong with contextualism', *Philosophical Quarterly* [2008]). *Prima facie*, this claim should strike us as surprising. A number of concepts and linguistic items (words, sentences) serve functions that have little or nothing to do with semantics. However, I argue that, on the best interpretation of talk of the function of a concept such as 'knowledge', the function of 'knowledge' is relevant to semantics. Along the way I also suggest how to improve on what I call the 'usual argument' that Craig's account of the function of 'knowledge' supports epistemic contextualism.

0. Introductory remarks

Plausibly, a successful analysis of the concept 'knowledge' must match and predict our intuitions about a wide range of cases and explain why knowledge is valuable (Greco 2007; Kvanvig 2003; Zagzebski 1996). Call this the 'value turn' in epistemology (Riggs 2008). Perhaps equally plausibly, a successful analysis must also fit with an account of the distinctive *function* or *social role* that the concept plays in our community. The best-known account of the distinctive function of 'knowledge' is that developed in Edward Craig (1990), on which 'knowledge', in particular 'knowledge' ascriptions - instances of 'S knows that p' - has the function of identifying *good informants*. The idea is that, in saying that a subject S knows that p, I identify S as a good informant on the matter of p. Call this the 'functional turn' in epistemology.

The functional turn has been suggested in a number of places. For instance, Duncan Pritchard claims that it's a point in favour of what he calls 'anti-luck virtue epistemology' that it accommodates Craig's insight (see also Greco 2007, 2008; Kelp Forthcoming-a, Forthcoming-b):

"Many philosophers find [Craig's account] of the source of the concept of knowledge to be very persuasive … once we start to think about this account of the source of the concept of knowledge more detail, it becomes apparent that it actually lends greater support to anti-luck virtue epistemology than it does to virtue epistemology (2012: 276).

Craig's insight has been appealed to in support of a number of other views in epistemology, and in this paper I'm concerned with the claim that it supports a sort of 'epistemic contextualism', which is (roughly) the view that the semantic content of the word 'knowledge' and its cognates depends on and varies with the context of attribution.[1] Here's John Greco:

"[A]n important function of our concept of knowledge and our knowledge language, perhaps its primary function, is to flag information and sources of information for use in practical reasoning. Now suppose this is right. Does that speak in favour of attributor contextualism or interest-dependent invariantism? To my mind, it speaks in favour of attributor contextualism" (2008: 433).

That the function of the concept 'knowledge' might support a particular view about the *semantics* of the word 'knowledge' should, *prima facie*, strike us as surprising. A wide range of concepts and linguistic items (words, sentences) serve functions that, on the face of it, have little or nothing to do with semantics. Example: The function of certain tokens of 'Smith has good handwriting' is to ensure that Smith isn't considered for the job, but that tells us little or nothing about the semantic content of this or any other token of 'Smith has good handwriting'. Or: Perhaps at some stage in human history the function of the concept 'marriage' was to keep parents together for the purposes of raising children, but that tells us little or nothing about the semantic content of 'marriage'. The point isn't that one *can't* appeal to the function of a concept or linguistic item in support of a semantic view (I will argue that, in a certain sense of 'function', one can) but that, *prima facie*, we should be suspicious of doing so.

---

[1] For other appeals to Craig's insight see Chrisman (2010); Fricker (2012); Neta (2006).

My aims in this paper are twofold. First, I will propose a particular way of understanding talk of the 'function' of a concept (what I call the 'genealogical interpretation'). In doing so I will sketch, in rough outline, Craig's methodology (§1; §2.3). Second, I will argue that, understood along the lines of the genealogical interpretation, the function of 'knowledge' supports a certain sort of epistemic contextualism (§2.1-2.3).

In broad outline, the idea to be defended is that talk of the function of a concept concerns long-standing and extensive patterns of use within a linguistic community. But the way in which a concept (and the word it expresses) is used over a long period of time within a linguistic community shapes and informs the conditions under which the concept applies, and the meaning of the word it expresses. Consequently, the function of a concept is relevant to semantics.[2]

1. Craigian genealogy

Here's Craig describing his methodology:

"We take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application" (1990: 2).

The intended contrast is with the traditional 'analytical' approach, which starts with intuitions about whether subjects in a range of imagined cases know. This raises three questions, which I'll take in turn. First, what motivates doing things this way? (§1.1) Second, how does one 'explore what a concept with that function would look like'? (§1.2) Third, what's the end result? (§1.3)

1.1. Motivations

Craig adverts to three motivations (1990: 1-2). First, as demonstrated by the post-Gettier literature, providing a successful analysis of 'knowledge' isn't easy (see Ichikawa & Steup 2012). If an alternative is available, why not give it a go? Second, intuitions about whether subjects in imagined cases know provide the data for an analysis of 'knowledge'. Call these 'extensional intuitions' (intuitions about the extension of the concept 'knowledge'). But we also have 'intensional intuitions' (intuitions about the conditions that must be satisfied for the concept to

---

[2] Clarificatory note: I assume that words express concepts, and that concepts are the meanings of words. So when I talk about functions I'll often switch between talking about the function of the concept 'knowledge' and the function of the word 'knowledge'.

apply). Sceptics argue that the two don't mesh: the intuitive intension of 'knowledge' determines an extension that is far smaller than that determined by our extensional intuitions.[3] This suggests that, in the case of 'knowledge', providing an intuitive intension that fits its intuitive extension is going to be difficult. Third, if these difficulties can be overcome and an acceptable analysis of 'knowledge' found, some questions will remain unanswered; in particular, why is knowledge *valuable*? Here's Timothy Williamson:

"[K]nowing matters; the difference between knowing and not knowing is very important to us … This importance would be hard to understand if the concept knows were the more or less *ad hoc* sprawl that analyses have had to become; why should we care so much about *that*?" (Williamson's emphasis) (2000: 31).

In Craig's view, ultimately only the success of his approach can motivate adopting it (1990: 2-4, 3, 106). But what are the 'success criteria' here? That leads to our second and third questions.

1.2. Craig's hypothesis

Craig's hypothesis is that 'knowledge', in particular 'knowledge' ascriptions, has the function of identifying good informants. The motivation: We need information about the world around us but due to cognitive and practical limitations we can't get all of it 'first hand', so we have a further need for a way of identifying those who have information (good informants) and, because others often have the information that we need and are willing to share it, we have a concept with the function of identifying good informants.

The procedure here involves three steps. First, we identify something that humans, both now and at all points in their development, need. Humans, both now and at all points in their development, need information about the world, and because they can't get all of that information themselves they need a way of identifying those who have it.[4] That humans have

---

[3] Motivation: Sceptics provide scenarios in order to elicit some intuition as to what is generally required for one to know, and they hope to demonstrate that the intuitive conditions that are generally required for one to know are so strict that we know little, if anything. If correct, this conclusion is surprising because, intuitively, we know a lot.

[4] Why 'now and at all points in their development'? In some sense of 'function', certain concepts have the function of fulfilling certain needs at some times but not at others, e.g. at some point perhaps the concept 'marriage' had the function of keeping parents together for the purposes of raising children. But the features of human psychology that Craig appeals to are those that "may plausibly be supposed to be possessed by all humans" (1990: 5).

these needs is, in some sense, contingent. If we were omniscient, we'd have all the information that we need at our fingertips. Second, we ask what meets our need for a way of identifying those who have information, and Craig's hypothesis is that the concept 'knowledge' does. Third, given that 'knowledge' meets this need, we identify the conditions we would have to put on the concept in order for it to meet it. To identify someone as a 'knower' is to identify her as a good informant, so what is required to be a good informant?

The following captures the conditions one would have to meet to be a good informant for the rest of one's small group (cf. Craig 1990: 84-5):

GOOD INFORMANT: Roughly, a good informant about p is someone who:
i.   Tells the truth about p (i.e. either p and she believes that p, or not-p and she believes that not-p).
ii. Is as likely to be right about p as the interests and purposes of her group require.
iii. Is detectable by the others in her group as likely enough right.
iv. Is immediately available to the others in her group.

An example: Otto is sitting on a hill and scanning the horizon for a sight of our enemy. Otto has good eyesight, a well-known good track record as a lookout and can shout loud enough for those in his group to hear. Assuming Otto is telling the truth when he shouts that the enemy is coming, he's a good informant as to whether the enemy is coming.

However, we don't identify good informants in a social vacuum. When I identify a subject as a good informant, I don't just recommend that subject as an informant to the others in my group, I recommend her to the 'world at large'. An epistemic evaluation is a *public property*. Others who are not in my group, and who have interests and purposes I'm not even aware of, can and will draw on my epistemic evaluation. These sorts of social pressures mean that we need to remove the 'subjective' conditions (iii-iv), and we need to remove the reference to the interests and purposes of some particular group (ii). This gives us the following (cf. Craig 1990: 88-9):

GOOD INFORMANT*: Roughly, a good informant about p is someone who:
i.   Tells the truth about p.
ii. Is as likely to be right about p as the interests and purposes of her community require.

In Craig's view, GOOD INFORMANT corresponds to what was required to be a good informant at an earlier stage of the development of the concept 'knowledge' whereas GOOD INFORMANT*

corresponds to what is required at this later stage in its development (1990: Chapter 10). So, an exploration of 'what a concept with that function would look like' gives us a picture of how the conditions one must meet to be a good informant, and so to 'know', will have developed through social pressures.

1.3. The end result

Why does this tell us anything about *our* concept rather than some other concept? Here's Craig:

"[I]t is not the idea to construct an imaginary concept, but to illuminate the one we actually have, though it be vague or even inconsistent; and to illuminate it by showing that a concept with the hypothesised role would have characteristics closely resembling those that it exhibits itself. But should our intuitions prove indeterminate or elastic, this type of investigation might reveal constructive ways of stretching them, and the rationale behind the stretch" (1990: 2-3).

So Craig gives us a reconstruction of our concept 'knowledge', where the reconstruction gives conditions one must meet to be a good informant, and so to 'know', such that the conditions (1) ensure that 'knowledge' can perform its distinctive function and (2) don't diverge significantly from its intuitive intension and extension (see also 1990: 5, 8, 13-14, 17). Assuming Craig's hypothesis about the function of 'knowledge' is right, his reconstruction will achieve (1). (I set aside (2) here, but see 1990: Chapters 2-3, Chapter 10).

I'll refer to the procedure just outlined as 'Craigian genealogy'. As Craig puts it, he is investigating the 'genesis' of the concept 'knowledge' (1990: 102, 107, 109, 117). I now turn to the question of how to interpret talk of the 'function' of a concept, and why one might think that the function of 'knowledge' supports epistemic contextualism.

2. From Craigian genealogy to Craigian contextualism

I start by sketching what I'll call the 'usual argument' that Craigian genealogy supports epistemic contextualism, henceforth 'contextualism'. I then highlight two problems with that argument (§2.1). I dedicate §2.2 and §2.3 to dealing with those problems. I finish by explaining why I take the argument of §2.1-2.3 to show that Craigian genealogy supports a sort of contextualism rather than a certain sort of 'invariantist' view (§2.4), and by clarifying the sort of contextualist view that Craigian genealogy supports (§2.5).

2.1. The usual argument

Contextualists think that the semantic contents and truth-conditions of 'knowledge' ascriptions depend on and vary with the context of ascription. A 'knowledge' ascription can be true in one context yet false in another context, in much the same way that 'I am tired' can be true in one context (a context where I'm tired) yet false in another (one where I'm not tired). For convenience, I'll work with a version of the view on which 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 in which not-p that are relevant in C, where the range of relevant alternatives depends on and varies with the context (see Lewis 1996).

The 'usual argument': Craig claims that the function of 'knowledge' ascriptions is to identify good informants. But whether a subject is a good informant depends on and varies with the context. If I'm talking to a friend over lunch about Isla's whereabouts last night and I have good but not conclusive evidence that she was at the party I'll volunteer myself as an informant on her whereabouts. However, if I'm giving a statement to the police and I have the same evidence I'll not volunteer myself as an informant on her whereabouts. Consequently, given that what one will require of a good informant depends on and varies with the context, the truth-conditions of 'knowledge' ascriptions must depend on and vary with the context too. One can find something like this argument (albeit in a rather less compressed form) in Greco (2007, 2008), Michael Hannon (Forthcoming) and David Henderson (2009).

The first obvious problem: Whether someone meets the conditions on being a good informant specified in GOOD INFORMANT clearly varies with and depends on the context. S may be as likely right as my interests and purposes require, but not as likely right as your interests and purposes require. So - modulo issues about semantic relevance - it looks like, at an early stage in the development of our concept 'knowledge', some sort of contextualism was true. But whether someone meets the conditions on being a good informant specified in GOOD INFORMANT* doesn't depend on or vary with the context. If someone meets those conditions, they are as likely right as our community's interests and purposes require. So it looks like, at this later stage in the development of our concept 'knowledge', contextualism is false. I deal with this problem below (§2.2). I'm going to argue that a rather more sophisticated way of arguing from Craigian genealogy to (a sort of) contextualism is available.

The second obvious problem: *Prima facie*, one shouldn't think that the function of 'knowledge' and 'knowledge' ascriptions tells us much about the correct semantics for 'knowledge' and 'knowledge' ascriptions. Again: The point isn't that the function of a concept or word is *irrelevant*

7

to semantics, the point is that further investigation is needed to determine whether it's relevant. So, by itself, Craigian genealogy doesn't support contextualism. I also deal with this problem below (§2.3).

2.2. The usual argument, improved

Some preliminaries: First, on the account of Craigian genealogy given earlier the conditions on being a good informant, and so 'knowing', given in GOOD INFORMANT correspond to an early stage of development of the concept 'knowledge' whereas those given in GOOD INFORMANT* correspond to a later stage in its development. Call this the 'standard interpretation'. Second, on the standard interpretation the end result of Craigian genealogy is a sort of invariantism, on which whether one 'knows' that p depends on whether one can rule out a relatively stable range of alternatives in which not-p. As I'll put it below, the idea is that, because of social pressures, for any proposition p our community has settled on a range of alternatives in which not-p that are 'default relevant' (I take this phrase from Hannon Forthcoming). Third, I'll assume that the sorts of alternatives that are default relevant are the sorts of alternatives that subjects in ordinary situations (such as chatting to a friend over lunch) are often able to rule out whereas the subjects in more high stakes situations (such as testifying to the police) often aren't.[5]

On the standard interpretation, the interests and purposes of the whole community determine the relatively stable range of alternatives in which not-p that one generally has to be able to rule out to count as 'knowing' that p (the default relevant alternatives). This, as I'll put it, makes the default relevant alternatives 'globally interest-relative'. The range of alternatives that are default relevant depends on the interests and purposes of the members of our community (the sorts of practical projects they engage in, the sorts of practical situations they find themselves in). Call this GLOBAL INTEREST-RELATIVITY. A prominent contextualist idea is that the range of alternatives that are relevant in a context depends on *particular practical situations*, most often the practical situation of those in the context, but perhaps some other practical situation, such as that of someone who is being advised. This, as I'll put it, makes the range of alternatives that are relevant in a context 'locally interest-relative'. Call this LOCAL INTEREST-RELATIVITY. On the face of it, GLOBAL INTEREST-RELATIVITY is *incompatible with* or *rules out* LOCAL INTEREST-RELATIVITY. We don't offer epistemic evaluations in a social vacuum, so we need to require that subjects can rule out a relatively stable range of alternatives in which not-p to count as 'knowing'

---

[5] Kelp (2011) challenges this assumption. But I find it hard to understand why social pressures would lead us to put conditions on 'knowing' which are such that subjects rarely, if ever, meet them.

that p. Consequently, in any context, the range of relevant alternatives must include all and only the default relevant alternatives, some of which need not be appropriate in that particular context.

I think there are two reasons why one might think that GLOBAL INTEREST-RELATIVITY rules out LOCAL INTEREST-RELATIVITY.

First reason: There are many cases where GLOBAL INTEREST-RELATIVITY clearly rules out LOCAL INTEREST-RELATIVITY. One sort would be a case where a 'knowledge' ascriber (call her Ascriber) is certifying a subject as an authority on some matter. An authority is someone that a range of people in different situations can draw on, so Ascriber has to consider the default relevant alternatives rather than alternatives that are appropriate to any particular practical situation (including her own). Another sort would be a case where Ascriber is giving advice to a range of people in different situations. Again, such advice is something that many people may draw on, so Ascriber has to consider the default relevant alternatives rather than alternatives that are appropriate to any particular practical situation. But are there any cases where GLOBAL INTEREST-RELATIVITY doesn't rule out LOCAL INTEREST-RELATIVITY? If not, not only does Craigian genealogy not support contextualism, it's actually incompatible with contextualism.

Response: *Prima facie*, there are cases where GLOBAL INTEREST-RELATIVITY doesn't rule out LOCAL INTEREST-RELATIVITY. One sort would be a case where a 'knowledge' ascriber is in an unusually pressing practical situation. For instance, take a standard high stakes bank case (DeRose 2009; Stanley 2005). Hannah and Sarah are passing the bank on a Friday afternoon, and they have a cheque that they simply must cash by Monday. Even though there's a queue they go in and, by way of justification, Hannah says 'I've been in this bank on Saturday before, but that was a few weeks ago so I don't know that it's open on Saturdays.'[6] Call this HIGH. In HIGH, Hannah's epistemic self-evaluation isn't a public property (it's not something that a range of people in various situations can or will draw on). Rather, her evaluation is something that Hannah and Sarah, in their particular practical situation, will draw on. Consequently, Hannah should consider the full range of alternatives appropriate to her practical situation, which will be a wider range than are default relevant. Another sort would be a case where, while the 'knowledge' ascriber isn't in an unusually pressing practical situation, she's giving advice to someone who is. Imagine it isn't vitally important that Hannah and Sarah cash their cheque

---

[6] Assume that the bank is open on Saturdays.

before Monday, but it is vitally important that their friend Lisbeth does so. Again, Hannah's epistemic self-evaluation isn't something that a range of people in various situations can or will draw on. Rather, it's something that Lisbeth, in her particular practical situation, will draw on. Consequently, Hannah should consider the full range of alternatives appropriate to Lisbeth's practical situation, which again will be wider than the range that is default relevant.

On the standard interpretation, GLOBAL INTEREST-RELATIVITY rules out LOCAL INTEREST-RELATIVITY due to certain social pressures. These social pressures require that 'knowledge' ascribers only consider the default relevant alternatives when they are offering an epistemic evaluation of a subject that a wide range of people in different practical situations can or will draw on. However, they don't require that 'knowledge' ascribers only consider the default relevant alternatives when they are offering an epistemic evaluation that only they, or others in their group, can or will draw on. But, if 'knowledge' ascribers are often reflective enough to understand what is required of them - if they often realise when their epistemic evaluations are a 'public property', and when they aren't - there is no reason to think that GLOBAL INTEREST-RELATIVITY rules out LOCAL INTEREST-RELATIVITY.

Second reason: One can never be certain that one's epistemic evaluation isn't going to be drawn on by some unknown party. Say I'm in a very low stakes practical situation and I'm evaluating whether a subject with rather meagre evidence 'knows' that p. Because I'm in such a low stakes practical situation, I don't consider certain alternatives in which not-p that are default relevant. Further, say that I'm very confident that my epistemic evaluation isn't something that a range of people in various situations can or will draw on (I think I'm alone). But because it could be that, unbeknownst to me, there are people that can and will draw on my evaluation (maybe I'm not alone), I should still consider all of the default relevant alternatives rather than some narrower range. 'Knowledge' ascribers can never be certain that, in taking particular practical situations into account, they aren't providing epistemic evaluations of subjects that may mislead others who draw on those evaluations.[7]

Response: This gets things half right. It gives us good reason to think that GLOBAL INTEREST-RELATIVITY rules out a type of LOCAL INTEREST-RELATIVITY on which the range of alternatives

---

[7] The thought here is that if one goes around making 'knowledge' ascriptions based on one's particular practical situation then one is making a lot of potentially misleading epistemic evaluations. Perhaps most won't be used by anyone, but some will, with potentially damaging effects.

that are relevant in a particular context can be *narrower* than those that are default relevant. Even if a 'knowledge' ascriber is in a very low stakes practical situation, she can't ignore alternatives that are default relevant because, if she does so, she'll give potentially misleading epistemic evaluations. But reconsider the HIGH case. Hannah and Sarah are in an unusually pressing practical situation, and Hannah is evaluating whether a subject (herself) with good (but not conclusive) evidence 'knows' that the bank is open on Saturdays. Like in the very low stakes case, she's confident that her epistemic self-evaluation won't be drawn on by a range of people in different situations (again, she thinks she's alone) but, unlike in the very low stakes case, she can be certain that, in taking her particular practical situation into account, she's not providing a *misleading* epistemic self-evaluation. Why? Because there's an asymmetry between the very low stakes case and HIGH. In the former I ascribe 'knowledge', and thereby identify the subject as a good informant; in the latter Hannah denies that she 'knows' but she isn't thereby identified as a *bad* informant (untrustworthy etc). Rather, she's just not identified as a good informant. To provide a negative epistemic evaluation involves failing to provide some potentially useful information, and there's a difference between misleading someone and failing to provide them with something that could be of use to them. Consequently, Hannah should consider a wider range of alternatives than those that are default relevant, and so deny that she 'knows' that the bank is open on Saturdays. So we haven't been given any reason to think that GLOBAL INTEREST-RELATIVITY rules out another type of LOCAL INTEREST-RELATIVITY, on which the range of alternatives that are relevant in a particular context can be wider but not narrower than the range that is default relevant. In what follows I'll call this CONSTRAINED LOCAL INTEREST-RELATIVITY.

Objection: Why think that Craigian genealogy *supports* CONSTRAINED LOCAL INTEREST-RELATIVITY, rather than that Craigian genealogy doesn't rule out CONSTRAINED LOCAL INTEREST-RELATIVITY? Because the reason given for thinking that the subjective conditions on being a good informant given in GOOD INFORMANT correspond to an earlier stage in the development of the concept 'knowledge' whereas the conditions given in GOOD INFORMANT* correspond to a later stage in its development was that we don't offer epistemic evaluations in a social vacuum, and as we've just seen that reason doesn't tell against CONSTRAINED LOCAL INTEREST-RELATIVITY. So, contra the standard interpretation, we should see GOOD INFORMANT and GOOD INFORMANT* as both corresponding to the present day, with GOOD INFORMANT* capturing the conditions one must meet to be a good informant, and so to 'know', in one range of situations (when one is identifying a subject as an authority, when one is offering advice to a diverse group, etc) and GOOD INFORMANT capturing the conditions one must meet in another

range of situations (when one is in an unusually pressing practical situation, or offering advice to someone else who is in such a situation).

Taking all this together, I propose that Craigian genealogy supports what I call 'Craigian contextualism':

i. The truth-conditions of 'knowledge' ascriptions are context-sensitive. For convenience, put this in terms of contextually varying ranges of relevant alternatives, viz. 'S knows that p' is true in context C iff S's evidence is sufficient to rule out the alternatives in which not-p that are relevant in C.

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.

Two things to note: First, Craigian contextualists hold that the range of alternatives that are relevant in a context can be wider but not narrower than those that are default relevant (ii). Therefore, Craigian contextualists accept CONSTRAINED LOCAL INTEREST-RELATIVITY, and reject any other type of LOCAL INTEREST-RELATIVITY. Consequently, Craigian genealogy supports Craigian contextualism. Second, Craigian contextualism isn't so much a view as a template. To get a concrete view one has to give some account of what makes an alternative 'appropriate' in a given practical situation. I myself favour a view on which an alternative is appropriate just in case those in the context *have a reason* to consider it, and I've defended that view at length elsewhere (McKenna 2013). Here, however, I focus on this more general view (see also Greco 2007, 2008; Hannon Forthcoming; Henderson 2009).

Of course, the conclusion only follows if we disregard the first problem with the usual argument, which concerns the relevance of the function of the concept 'knowledge' and 'knowledge' ascriptions to semantics. It's to that problem which I now turn.

2.3. The genealogical interpretation and semantic relevance

I proceed as follows. First, I diagnose why many will find the claim that one can derive conclusions about the semantic contents of 'knowledge' and 'knowledge' ascriptions from the function of the concept 'knowledge' problematic (§2.3.1). Second, I introduce and defend a distinction between two ways of interpreting 'function talk' (talk of the function of a concept),

what I'll call the 'Gricean interpretation' and what I'll call the 'genealogical interpretation' (§2.3.2). Third, I argue that, understood along the lines of the genealogical interpretation, the function of a concept is relevant to semantics (§2.3.3).

2.3.1. The diagnosis

A common way of thinking about the function of a piece of language is in terms of what speakers of that language may use the relevant piece to *do* (Grice 1989). But, on this way of thinking about 'function claims', it's clear that the function of a piece of language is often (though not invariably) a matter for *pragmatics* rather than semantics. Again: When writing a job recommendation I may use the sentence 'Smith has good handwriting' to ensure that Smith doesn't get the job. But the function of this token of 'Smith has good handwriting' tells us nothing about its semantic content. Grice has taught us not to jump to semantic conclusions based on observations about the functions that pieces of language serve. But, just like the usual argument, the more sophisticated argument that Craigian genealogy supports (Craigian) contextualism offends against this basic lesson. That 'knowledge' and 'knowledge' ascriptions have a particular function is as likely a matter for pragmatics as semantics. To clarify: The point isn't that functions are relevant to pragmatics *rather than* semantics. The point is that, absent further investigation, we can't conclude which a particular function is relevant to. But, if that's right, Craig's account of the function of 'knowledge' doesn't, by itself, support contextualism.

I'll call this way of thinking of claims about the function of 'knowledge' the 'Gricean interpretation'. (The label is ugly, but it will be useful).[8]

2.3.2. Two interpretations

There are a number of features of the Gricean interpretation that don't seem to fit with Craigian genealogy. First, on the Gricean interpretation claims about the function of 'knowledge' ascriptions concern what particular speakers at particular times may do by using pieces of language. Second, the reasons why particular speakers may use particular pieces of language concern their communicative intentions, their beliefs about their situation and the like. In contrast, Craigian genealogy starts with the identification of a need (for information, and therefore for a way of identifying good informants) that is common to all humans at all stages in

---

[8] The label is perhaps liable to mislead. One can separate the point that speakers of a language often use pieces of that language to do particular things, not all of which are relevant to semantics, from the 'Gricean programme', and the attempt to analyse sentence and word meaning in terms of speaker meaning.

our development, and then proceeds to the hypothesis that the concept 'knowledge', in particular 'knowledge' ascriptions, meets that need. So, contra the first feature of the Gricean interpretation, Craig's hypothesis concerns how a community used, and continues to use, certain concepts and words in order to meet a common need. Further, contra the second feature, the focus is on the reasons why a community uses those concepts and words, not particular speakers and their communicative intentions.

On what I'll call the 'genealogical interpretation' (again, an ugly but useful label), Craigian genealogy is concerned with the way in which a community of speakers over a long period of time use and have used certain concepts and pieces of language (sentences of the form 'S knows that p' and 'S doesn't know that p') in order to meet a common need. While, like the Gricean interpretation, the genealogical interpretation appeals to use, it appeals to extensive and long-standing patterns of use within a whole community - where the community comprises those who use the concepts and language, hence a 'linguistic community' - rather than particular uses at particular times. And, for the reasons given, Craigian genealogy is best interpreted along the lines of the genealogical interpretation.

2.3.3. Functions, concepts and use

What is the connection between meaning and use? I distinguish two views: First, the (controversial) view that meaning just is use. On this view, facts about the meanings of our words are facts about how they are used (or, more plausibly, facts about meaning are constituted by some subset of facts about use as in Horwich 1998). Second, the (very plausible) view that the meanings of our words stand in some sort of complicated causal relationship to the way in which they have been used over a period of time. On this view, what a word means is partly a function of how it has been used, but we can't 'read off' facts about meaning from particular uses, and facts about meaning aren't facts about particular uses, use in general, or constituted by some subset of facts about use. Rather, the way in which a word is used over a long period of time will 'shape' the conditions under which the concept it expresses applies. In what follows I assume the second view is true, but I don't take this to be a controversial assumption.

On the Gricean interpretation, talk about the function of 'knowledge' ascriptions concerns particular uses by particular speakers at particular times. And, as I just pointed out, one can't read off conclusions about the meaning of 'knowledge' and 'knowledge' ascriptions from particular uses at particular times. Consequently, as one would expect, on the Gricean interpretation Craigian genealogy doesn't support Craigian contextualism. But, as I've just argued, Craigian genealogy is best interpreted along the lines of the genealogical interpretation

rather than the Gricean interpretation and, on that interpretation, talk about the function of 'knowledge' ascriptions concerns extensive and long-standing patterns of use within a linguistic community. And, as I also just pointed out, what a word means - in the case, 'knowledge' - is partly a function of the way in which it has been used over a period of time within a linguistic community. The way in which a word has been used shapes the conditions under which the concept it expresses applies, so the way in which we've used 'knowledge' shapes the conditions under which 'knowledge' applies. So, on the genealogical interpretation Craigian genealogy is relevant to the semantics as opposed to the pragmatics of 'knowledge' and 'knowledge' ascriptions. Consequently, Craigian genealogy supports Craigian contextualism.

2.4. Strict and sensitive invariantism

Those who are familiar with the intricacies of the contextualism/invariantism debate in epistemology will, at this point, be demanding answers to two questions. First, why does the argument of §2.1-2.3 tell in favour of Craigian contextualism rather than some variety of invariantism? Second why does it tell in favour of Craigian contextualism rather than some other variety of contextualism? (The second question perhaps raises another: What's the difference between Craigian contextualism and other varieties of contextualism?). I sketch an answer to the first question here and an answer to the second (and third) below (§2.5).

First, *invariantism* is the view that the semantic contents and truth-values of 'knowledge' and 'knowledge' ascriptions don't vary with the context of ascription. Second, we need to distinguish between *strict* and *sensitive* invariantism. Strict invariantists hold that only truth-conducive factors, such as the reliability of the subject's belief-forming mechanisms, determine whether the subject knows some true proposition that she believes. In contrast, sensitive invariantists hold that some combination of truth-conducive and non-truth-conducive factors, such as the stakes (for the subject) or what sorts of error-possibilities the subject is considering, determine whether the subject knows some true proposition that she believes (Fantl & McGrath 2009; Hawthorne 2004; Stanley 2005). Third, contextualists motivate their view by appeal to pairs of cases where in one case an ascriber with seeming propriety claims that a subject 'knows' that p based on certain grounds whereas in the other the same ascriber with seeming propriety denies that the same subject 'knows' that p based on the very same grounds, where the cases differ in the stakes for the ascriber, or the sorts of error-possibilities she is considering. Strict invariantists usually deal with these cases by positing contextual variation in the *assertability-conditions* (but not truth-conditions) of 'knowledge' ascriptions (Brown 2006; Rysiew 2001). Fourth, sensitive invariantists agree with contextualists that there are pairs of cases such that in one case a subject

knows that p based on certain grounds whereas in the other the same subject doesn't know that p based on the same grounds, but they deny that these are pairs that differ in the stakes for the ascriber, or the sorts of error-possibilities the ascriber is considering. Rather, they are pairs that differ in the stakes for the *subject*, or the sorts of error-possibilities the *subject* is considering. So one of the crucial differences between contextualism and sensitive invariantism concerns the following sort of case:

HIGH-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 cheque. Since they have an impending bill, and very little in their account, it's very important that they deposit their cheque 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'

Contextualists will say that Hannah speaks truly (her and Sarah are in a high stakes situation) whereas sensitive invariantists will say that Hannah speaks falsely (Bill is in a perfectly ordinary situation). While this is *prima facie* a cost for sensitive invariantists (DeRose 2009, Chapter 7), there are various responses available, the most promising of which is to appeal to what Stanley (2005: Chapter 5) calls the *projectivist strategy*. The idea is that in cases like HIGH-LOW the ascriber (here, Hannah) *projects* her situation onto the subject (here, Bill), and she does this because when we're inquiring into whether a subject knows we're using that subject as a potential informant. What Hannah wants to establish is whether, if Bill were in her situation, he would 'know' that p. Consequently, Hannah requires that Bill rule out alternatives appropriate to her situation, rather than to Bill's.

The projectivist strategy appeals to a particular feature of our epistemological practice - that, in evaluating whether a subject 'knows' that p, we're evaluating that subject as an informant - in order to explain why an ascriber such as Hannah might *mistakenly* project her situation onto the subject of her 'knowledge' ascription. But, as we've seen, it's plausible that this feature of our practice reveals the *function* of 'knowledge'. So the defender of the projectivist strategy has to

hold that the function or 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.[9]

I also take the argument of §2.1-2.3 to provide good grounds to prefer Craigian contextualism to strict invariantism. First, I mentioned above that strict invariantists appeal to contextual variation in the assertability-conditions of 'knowledge ascriptions. One long-standing problem with this appeal concerns the observation that in cases like HIGH (reminder: Hannah and Sarah are in a high stakes situation, and with seeming propriety Hannah says that she doesn't 'know' that the bank is open) it's not just improper for Hannah to say that she 'knows', it's proper for Hannah to deny that she 'knows'. While it's often inappropriate to assert something true, it's not so often appropriate to assert something false (see DeRose 2009: Chapter 3). Second, I've argued that the function of 'knowledge' ascriptions - talk of which reveals a long-standing pattern of use within our community - shapes the conditions under which the concept 'knowledge' applies. I don't take this to show that it *could not* be the case that what it takes to 'know' doesn't depend on the context whereas what it takes to be properly said to 'know' does. But I do take it to show that it's *unlikely* that what it takes to 'know' is a context-insensitive matter. Third, if the strict invariantist who appeals to a contextual variation in assertability-conditions is right, then when a 'knowledge' ascriber is in a very high stakes situation she could often, quite truthfully albeit improperly, say that some subject 'knows' that such-and-such. But, in doing so, she could not be identifying that subject as a good informant on the matter of such-and-such (because she wouldn't qualify). The strict invariantist needs to hold that the truth-conditions of 'knowledge' ascriptions can, as this sort of case demonstrates, be inconsistent with the function of 'knowledge' ascriptions. But, if that function shapes the meaning of 'knowledge', how could that be (see Hannon Forthcoming: §4.4; Henderson 2009: 123-5)?

2.5. Conversational contextualism

In brief: Craigian genealogy rules out any sort of contextualism that accepts what I'll call EXTENSIVE VARIABILITY, and what I'll call 'conversational contextualism' accepts (whereas Craigian contextualism rejects) EXTENSIVE VARIABILITY. Consequently, Craigian genealogy supports Craigian contextualism but rules out conversational contextualism.

---

[9] The argument in the main body of the text is admittedly brief, but an in-depth evaluation of the respective merits of contextualism and sensitive invariantism would take me too far away from my main purpose here. For a more detailed discussion see McKenna (2013).

First, conversational contextualists think of contexts as points along a continuum, with contexts where an extremely wide range of alternatives are relevant at one extreme, and contexts where an extremely narrow range of alternatives are relevant at the other. For any proposition p, the limit at one extreme is just all the alternatives in which not-p and the limit at the other is perhaps just any alternative in which not-p that actually obtains (of course, there may not be such an alternative). The context shifts, and the range of relevant alternatives expands or contracts, according to the conversational kinematics (the raising and taking seriously of error-possibilities etc) (DeRose 2009; Lewis 1996). Consequently, *conversational contextualism* has the following two features:

1. There are no constraints on how wide or narrow the range of relevant alternatives can get.
2. Expansions and contractions in the range of alternatives relevant in a context are tied to conversational kinematics.[10]

Call the combination of (1) and (2) EXTENSIVE VARIABILITY.

Second, Craigian contextualists think that the default relevant alternatives are relevant in *all* contexts, and that the range of relevant alternatives in a particular context can't get narrower than the range that are default relevant. Consequently, they reject (1). Further, Craigian contextualists think that expansions and retractions in the range of relevant alternatives are tied to the practical situations of 'knowledge' ascribers, not conversational kinematics. Consequently, they reject (2).

Third, Craigian genealogy supports CONSTRAINED LOCAL INTEREST-RELATIVITY, but it rules out a sort of LOCAL INTEREST-RELATIVITY on which the range of alternatives that are relevant in a context can be both narrower and wider than the range that is default relevant. Anyone who endorses (1) thinks that the range of alternatives that are relevant in a context can be narrower and wider than the range that is default relevant. Consequently, Craigian genealogy rules out (1), and therefore EXTENSIVE VARIABILITY.

---

[10] While I put (1-2) in terms of relevant alternatives, I could re-formulate them in terms of 'epistemic standards' in the sense of DeRose (2009), with (1) saying there are no constraints on how high or low the standards can get, and (2) saying that changes in the standards operative in a context are tied to conversational kinematics.

I conclude that Craigian genealogy doesn't support, and actually rules out, conversational contextualism.[11]

3. Concluding Remarks

I have argued that Craigian genealogy supports what I call Craigian contextualism. If one takes seriously the idea that 'knowledge' ascriptions serve the function of identifying good informants, then one has good reason to endorse a sort of contextualism. I have not argued that one should take that idea seriously, but I have provided an outline of what I think taking it seriously involves (§1; §2.3).[12]

Bibliography

Brown, J. 2006. 'Contextualism and warranted assertibility manoeuvres'. *Philosophical Studies*, 130 (3): 407-435.

Chrisman, M. 2010. 'From epistemic expressivism to epistemic inferentialism'. In A. Haddock, A. Millar & D. Pritchard (eds.), *Social epistemology*, pp. 112-129. Oxford: Oxford University Press.

Craig, E. 1990. *Knowledge and the state of nature*. Oxford: Oxford University Press.

DeRose, K. 1995. 'Solving the skeptical problem'. *Philosophical Review*, 104 (1): 1-52.

———. 2009. *The case for contextualism*. Oxford: Clarendon.

Fantl, J. & McGrath, M. 2009. *Knowledge in an uncertain world*. Oxford: Oxford University Press.

Fricker, M. 2012. 'Group testimony? The making of a collective good informant'. *Philosophy and Phenomenological Research*, 84 (2): 249-276.

---

[11] Objection: The standard contextualist response to scepticism, as outlined in DeRose (1995) and elsewhere, relies on EXTENSIVE VARIABILITY, therefore any version of contextualism that rejects EXTENSIVE VARIABILITY is unable to utilise the standard contextualist response to scepticism. The upshot is that, while conversational contextualism gains no support from Craigian genealogy, Craigian contextualism gains no support from scepticism. Response: First, the standard contextualist response to scepticism faces a range of objections (see, for example, Rysiew 2011). If - as I am inclined to think - these objections are sound, one might think that contextualism is left looking somewhat unmotivated. But, if the argument of §2.1-2.3 is right, contextualists (at least, Craigian contextualists) can appeal to something else, viz. the function of 'knowledge' ascriptions. Second, my aim here has been to explore an alternative motivation for a form of contextualism, not to argue that Craigian contextualism is, all things considered, preferable to conversational contextualism.

[12]

Greco, J. 2007. 'The nature of ability and the purpose of knowledge'. *Philosophical Issues*, 17 (1): 57–69.

———. 2008. 'What's wrong with contextualism?' *Philosophical Quarterly*, 58 (232): 416-436.

Grice, H. P. 1989. *Studies in the way of words*. Harvard University Press.

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.

Horwich, P. 1998. *Meaning*. Oxford: Oxford University Press

Ichikawa, J. J. & Steup, M. 2012. 'The analysis of knowledge'. In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Winter 2012 Edition)*. URL = <http://plato.stanford.edu/archives/win2012/entries/knowledge-analysis/>

Kelp, C. 2011. What's the point of "knowledge" anyway? *Episteme*, 8: 53-66.

———. Forthcoming-a. 'Knowledge: the safe-apt view'. *Australasian Journal of Philosophy*.

———. Forthcoming-b. 'How to motivate anti-luck virtue epistemology'. *Grazer Philosophische Studien*.

Kvanvig, J. 2003. *The value of knowledge and the pursuit of understanding*. Cambridge University Press.

McKenna, R. 2013. 'Epistemic contextualism: a normative approach'. *Pacific Philosophical Quarterly*, 94(1): 101-123.

Neta, R. 2006. 'Epistemology factualized: new contractarian foundations for epistemology'. *Synthese*, 150 (2): 247 - 280.

Pritchard, D. 2012. 'Anti-luck virtue epistemology'. *Journal of Philosophy*, 109: 247-79.

Riggs, W. 2008. 'The value turn in epistemology'. In V. Hendricks & D. Pritchard (eds.), *New waves in epistemology*, pp. 300-323. Palgrave Macmillan.

Rysiew, P. 2001. 'The context-sensitivity of knowledge attributions'. *Noûs*, 35 (4): 477–514.

———. 2011. 'Epistemic contextualism'. In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition)*. URL = <http://plato.stanford.edu/archives/win2011/entries/contextualism-epistemology/>

Stanley, J. 2005. *Knowledge and practical interests*. Oxford: Oxford University Press.

Williamson, T. 2000. *Knowledge and its limits*. Oxford: Oxford University Press.

Zagzebski, L. 1996. *Virtues of the mind*. Cambridge University Press.