

The Practical Origins of Epistemic Contextualism

Michael Hannon

Received: 17 August 2011 / Accepted: 5 December 2012 / Published online: 15 December 2012
© Springer Science+Business Media Dordrecht 2012

Abstract This paper explores how the purpose of the concept of knowledge affects knowledge ascriptions in natural language. I appeal to the idea that the role of the concept of knowledge is to flag reliable informants, and I use this idea to illuminate and support contextualism about ‘knows’. I argue that practical pressures that arise in an epistemic state of nature provide an explanatory basis for a brand of contextualism that I call ‘practical interests contextualism’. I also answer some questions that contextualism leaves open, particularly why the concept of knowledge is valuable, why the word ‘knows’ exhibits context-variability, and why this term enjoys such widespread use. Finally, I show how my contextualist framework accommodates plausible ideas from two rival views: subject-sensitive invariantism and insensitive invariantism. This provides new support for contextualism and develops this view in a way that improves our understanding of the concept of knowledge.

1 Introduction

One of the most popular views in recent epistemology is contextualism. The main thesis of contextualism is that the truth conditions of knowledge-ascribing and knowledge-denying sentences (sentences of the form “S knows that p ” and “S doesn’t know that p ”) vary in certain ways according to the context in which they are uttered (DeRose 2009: 2). In particular, what varies is the standard for how strong one’s epistemic position¹ with respect to p must be in order for one to count

¹ There are various ways of spelling-out the notion of the ‘strength of an epistemic position’. For our purposes, think of it as a matter of the range of alternatives to p with respect to which one’s evidence favors p .

M. Hannon (✉)
King’s College, University of Cambridge, Cambridge CB2 1ST, UK
e-mail: mh536@cam.ac.uk

as ‘knowing’² that *p*. Like many views in epistemology, contextualism has focused on knowledge attributions in order to answer questions about whether it is ever accurate to ascribe knowledge to people and, if so, when (DeRose 1995; Lewis 1996; Cohen 1988). As a semantic thesis, however, contextualism does not discuss why knowledge attributions are valuable, nor does it explain why ‘knows’ exhibits context-variability.

In his 1990 book *Knowledge and the State of Nature*, Edward Craig explains the value of the concept of knowledge by providing an account of its function. He also demonstrates why, given certain plausible assumptions about humans and their environments, the concept of knowledge has received its current shape.³ Craig proposes that the concept of knowledge is valuable because it helps us satisfy some very general needs in human life and thought; roughly, it enables us to flag reliable informants. Let’s call this ‘Craig’s hypothesis’.⁴

This paper appeals to Craig’s hypothesis in order to illuminate and support contextualism for knowledge claims. I argue that Craig’s ‘practical explication’ of knowledge provides a basis for the emergence of a particular version of contextualism that I call ‘practical interests contextualism’.⁵ In particular, the practical pressures that arise in an epistemic state of nature indicate why we should expect ‘knows’ to function in the ways that practical interests contextualism predicts. My central claim is as follows. Inquirers seek informants who are likely enough *in the context* to be right as to whether *p*. If the purpose of ‘knows’ is to flag informants that are reliable enough to satisfy our needs and interests, and if our needs and interests vary with context, then what it takes to be a reliable informant will change from context to context; thus, we should expect that ‘knows’ and its cognates also vary in order to serve their purpose.

1.1 Outline

Section 2 briefly outlines the thesis of epistemic contextualism and Sect. 3 discusses Craig’s practical explication of knowledge. Section 4 uses the method of practical explication to motivate practical interests contextualism and demonstrates how this view improves upon previous versions of contextualism. I also show how plausible

² I put quotes around ‘knowing’, ‘knowledge’, and ‘knows’ to indicate semantic ascent. My focus is not on the property of knowledge, but rather on our use of the word ‘knowledge’ (and its cognates) for epistemic evaluation.

³ Craig speaks of the role of the *concept* of knowledge rather than the *linguistic* role of ‘knows’. I shall sometimes speak of the concept of knowledge, but most of what I say can be adequately expressed with the linguistic framing. My concern in this paper is with our linguistic knowledge-attributing behavior. Also, if you agree with Frege that concepts are eternal, then substitute the notion of concept change with a change in what concept is expressed by ‘knows’ (and its cognates). Nothing should turn on this.

⁴ I do not assume that the *sole* purpose of the concept of knowledge is to serve as a marker for good informants. My view requires only that *a* central (common, important) purpose of the concept of knowledge is to certify such sources. For expository convenience, however, I shall speak of ‘the’ purpose of the concept of knowledge.

⁵ Fricker (2010), Henderson (2009, 2011) and Greco (2008, 2010) discuss the relevance of Craig’s approach to ‘proto-contextualism’, ‘gate-keeping contextualism’, and ‘attributor contextualism’ respectively. I have benefitted much from studying their views.

ideas from two views that rival contextualism can be accommodated within my contextualist framework: ‘subject-sensitive invariantism’ (Hawthorne 2004; Stanley 2005) and ‘insensitive invariantism’ (Brown 2005; Rysiew 2001). The view I defend will shed light on why knowledge attributions are valuable, why ‘knows’ exhibits context-variability, and why this term enjoys widespread use. This goes beyond standard formulations of contextualism which, as semantic theses, leave open these questions. Providing a practical explication will shed light on these issues, thereby improving our understanding of the concept of knowledge.

2 A Brief Outline of Contextualism

Contextualism is an umbrella term for a wide variety of semantic theories. My focus in this paper is on the most widely discussed form of contextualism, called *attributor* contextualism (henceforth ‘contextualism’).⁶ According to this view, the truth conditions of knowledge-attributing and knowledge-denying sentences vary depending on the attributor-sensitive epistemic standards that an agent, S, must meet in order for a statement of the form “S knows that *p*” to be true. As DeRose puts it,

In some contexts, “S knows that *p*” requires for its truth that S have a true belief that *p* and also be in a *very* strong epistemic position with respect to *p*, while in other contexts, an assertion of the very same sentence may require for its truth, in addition to S’s having a true belief that *p*, only that S meet some lower epistemic standard. (2009: 3)

What reasons are there for thinking that ‘knows’ is context-sensitive? Contextualism is typically motivated by appealing to *shifty judgments*. Shifty judgments occur in cases in which it seems true that an attributor, A, in a ‘low-standards’ context, can truly assert that he knows that *p*, whereas in a ‘high-standards’ context A says something false in asserting that he knows that *p*, even though the strength of A’s epistemic position has not changed across contexts.⁷ To illustrate this view, let’s briefly consider the following pair of bank cases made famous by DeRose (1992: 913).

Bank Case A. Keith and his wife, Rachel, are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Although Keith and Rachel generally like to deposit their paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so Keith suggests that they drive straight home and deposit their paychecks on Saturday

⁶ Cohen (1988), Lewis (1996), and DeRose (1995, 2009) are the chief defenders of this position.

⁷ Contextualism also allows for cases in which it seems true that an attributor, A, in a ‘low-standards’ context can say something true in asserting “S knows that *p*”, whereas another attributor, B, in a ‘high-standards’ context, can say something false in asserting “S knows that *p*”, even though A and B are talking about the same S and the same *p* at the same time *t*.

morning. Rachel says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.” Keith replies, “No, I know it’ll be open. I was just there two weeks ago on Saturday. It’s open until noon.”

Bank Case B. Keith and Rachel drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. He again suggests that they deposit their paychecks on Saturday morning, explaining that he was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, they have just written a very large and very important check. If their paychecks are not deposited into their checking account before Monday morning, the important check they wrote will bounce, leaving them in a very bad situation. And, of course, the bank is not open on Sunday. Rachel reminds Keith of these facts. She then says, “Banks do change their hours. Do you know the bank will be open tomorrow?” Remaining as confident as he was before that the bank will be open then, still, he replies, “Well, no. I’d better go in and make sure.”

Presume that in both cases the bank will be open on Saturday. Keith’s claim to know that the bank will be open in *Bank Case A* seems true; meanwhile, it also seems that he is saying something true in *Bank Case B* when he admits that he does not know that the bank will be open on Saturday. Contextualism diagnoses these intuitions (and similar intuitions about similar cases⁸) in a plausible way: in *Bank Case A* (a low-standards situation) the standards for ‘knows’ are relatively low, hence the knowledge attribution comes out true, whereas in *Bank Case B* (a high-standards situation) the standards for ‘knows’ are relatively high, hence the knowledge attribution comes out false.

The merits of this view are hotly debated in epistemology. This paper does not engage with arguments that challenge the linguistic data to which contextualism appeals (see Schaffer and Knobe 2012 for a survey). I am persuaded by recent work which suggests that both the salience of alternative possibilities and changes in practical stakes affect judgments about whether a subject ‘knows’ (*ibid*; Sripada and Stanley 2012). Much of the dispute in linguistics has shown that these data are compatible with both contextualism and its major rivals, particularly subject-sensitive invariantism and insensitive invariantism. My aim is to show that reflecting on the purpose of knowledge attributions lends support to contextualism over its rivals, thereby rendering the debate more tractable.

3 Craig’s Practical Explication of Knowledge

In *Knowledge and the State of Nature*, Craig provides a new way of investigating knowledge that he calls a ‘practical explication’. Instead of the traditional method of analyzing the concept of knowledge into necessary and sufficient conditions, Craig approaches epistemology by investigating the value of the concept of knowledge.

⁸ I have in mind cases like Jonathan Vogel’s ‘Stolen Car’ (1999: 161) and Stewart Cohen’s ‘Airport Cases’ (2000: 95). I will discuss Cohen’s airport cases below.

He asks us to consider, first, what the concept of knowledge does for us, then, second, what a concept having that role would look like, i.e., what conditions would govern its application (Craig 1990: 2). Craig explains the value of the concept of knowledge in terms of its function: it is valuable because it serves an important role in human life and thought. His hypothesis is that knowledge attributions serve as markers for reliable informants.⁹

Earlier statements of this view can be found in the writings of Sosa (1974: 118 and 1988: 152–153) and Williams (1973: 146). For example, Sosa says that knowledge has a “social aspect” and that certain

departures from the traditional account [of knowledge] may make better sense if we reflect that the honorific term “knowledgable” is to be applied only to those who are reliable sources of information, surely an important category for a language-using social species. (1974: 118)

While several people have endorsed this idea, I shall focus specifically on Craig’s view because he develops this idea in more detail.

This hypothesis emerges from the idea that the concept of knowledge functions to facilitate the survival and flourishing of human beings in communities. The idea is that humans need true beliefs about their environment in order to successfully guide their actions. For this reason they require sources of information that will lead to true beliefs (Craig 1990: 11).¹⁰ Often the easiest and most efficient way to acquire a true belief as to whether p is to ask someone reliable whether p (Schmitt 1992: 556); and since on any issue some informants will be more likely than others to provide a true belief, any community may be presumed to have an interest in evaluating sources of information. The concept of knowledge emerges in connection with this interest; roughly, the concept of knowledge has developed to flag reliable informants, which is an important role in the general economy of our concepts. Grounding this hypothesis in the purpose the concept is supposed to serve, and the needs it is supposed to answer, is what connects Craig’s approach to various ‘state of nature’ theories (hence his book’s title). State of nature explanations proceed by identifying certain basic human needs and argue that our conceptual and linguistic practices are necessary (or at least highly appropriate) responses to them (Craig 1990: 89).¹¹

⁹ Two points of clarification: first, Craig’s initial hypothesis is that “the concept of knowledge is used to flag approved sources of information” (*ibid*: 11), however he later settles on the more specific hypothesis I’ve mentioned concerning reliable *informants* (*ibid*: 35); second, Craig describes these informants as ‘good’ rather than ‘reliable’, but I shall use these terms interchangeably.

¹⁰ This approach assumes that true belief is conceptually prior to knowledge. For a contrary view, see Williamson (2000).

¹¹ State of nature theories only require us to *think* of concepts as having started out in a state of nature without actually presupposing that they began this way at some point in our ancestors’ history. Craig is not committed to the idea that people actually went through the stages he describes (1990: 84). We can disentangle the method of practical explication from Craig’s genealogical (or state of nature) account. A practical explication of a concept seeks to explain how, given some set of facts, and some interests or aims, we have a need that the concept we are explicating satisfies (see Kappel 2010: 72–73). A genealogical account, in contrast, is concerned with the origin of a concept. This paper is concerned with the essentially ahistorical method of practical explication (see Rysiew 2012: 274).

For the sake of brevity, we can skip many of the details of Craig's proposal in order to focus on the main features that a reliable informant will possess. Craig enumerates the following general properties of a 'reliable informant' (*ibid*: 85):

- (a) The informant tells the inquirer the truth on the question.¹²
- (b) The informant is detectable by the inquirer (via some property *X*)¹³ as likely to be right about *p*.
- (c) The informant is accessible to the inquirer here and now.
- (d) The informant is as likely to be right about *p* as the inquirer's concerns require.
- (e) The channels of communication between the inquirer and the informant are open.

The goal of Craig's inquiry is not to enumerate the necessary and sufficient conditions for knowledge, but rather to shed "light on the nature and origins of the present practice" (*ibid*: 8). We should therefore interpret the aforementioned general properties as describing a 'prototypical' reliable informant. Craig is concerned with identifying the common core of the concept of a reliable informant, which allows him to specify conditions that hold only for the most part, but not always.

A significant problem arises at this stage of Craig's account. It seems to be no part of our familiar concept of knowledge that a knower has a recognizable property *X* indicative of true belief, nor must he be willing to convey his information. Rather, it seems perfectly conceivable to describe many cases in which an agent either knows whether *p* even though inquirers cannot detect any property that correlates with having a true belief as to whether *p*, or knows whether *p* even though the agent is unwilling to tell anyone what he knows. In other words, our familiar concept of knowledge does not feature anything like conditions (b), (c), or (e). This creates a large gap between the concept of a reliable informant and the concept of a knower, which seems to cause trouble for Craig's proposal. His aim is to construct a concept that not only functions in the manner suggested by his hypothesis, but also fits our intuitions about knowledge (or otherwise explains why our intuitions diverge) (Schmitt 1992: 555). As Craig himself admits, his investigation is anchored by the everyday concept of knowledge that he is looking to explicate. Thus,

should it reach a result quite different from the intuitive extension, then, barring some special and especially plausible explanation of the mismatch, the original hypothesis about the role that the concept plays in our life would be the first casualty. (1990: 2)

¹² The truth condition is essential because knowledge and 'knows' are factive (see Hannon 2013). However, speaking of a 'reliable' informant while including a truth condition is slightly misleading because reliable informants may sometimes be mistaken. I think there is a sense of 'reliable' that captures the truth condition—i.e., "Tim is usually reliable, but he wasn't today" (he was wrong)—but you might prefer to follow Craig and use the term 'good informant' to avoid this confusion.

¹³ These 'indicator properties' bestow what Williams (2002: 42–43) calls "purely positional advantage", an example of which is having been looking in the right direction at the right time.

Given the gap that seems to have opened up between the concept we have arrived at by considering the practical situation of the inquirer and our natural ascriptions of knowledge, shouldn't we conclude that Craig's hypothesis is inadequate?

This takes us to the second part of Craig's story. In order to connect the concept of knowledge to the concept of a reliable informant, Craig appeals to a process he calls *objectivisation*. Jonathan Dancy pithily characterizes this process in the following way: "A concept is objectivised if it becomes progressively less tied to the particular concerns of the user" (1992: 394). How exactly does this process work and why does it come about?

The situation Craig has thus far described involves a number of individual inquirers with a shared problem of having to determine the truth as to whether *p*. However these individuals form a community and are in some degree responsive to the needs of others.¹⁴ Craig now asks us to imagine "a more complex situation in which agents in the community have diverse concerns and different cognitive capacities but can also collaborate in order to achieve their goals" (Kelp 2011: 56). On many occasions we hope that others will recommend informants to us (including when someone recommends himself), for others are often 'better placed' than we are to detect reliable informants. We also have a need to store information without knowing when, or for what purposes, it might be used. Facts such as these put pressure towards the formation of objective concepts, since we need reliable informants in an objective sense—i.e., someone who is a reliable informant for others, not just me; someone detectable by others, not just me.

According to Craig, when I identify someone as a reliable informant on behalf of others, presumably I have to concern myself with the reliability of the informant to meet their standards, not my own. The problem is that we have no idea of the practical needs or interests of many people to whom this informant may be reliable. As a result,

a practice develops of setting the standard very high, so that whatever turns, for others, on getting the truth about *p*, we need not fear reproach if they follow our recommendation... In recommending an informant to you I am indeed implying that the likelihood of his being right is as great as your concerns require.¹⁵ (1990: 94)

When saying that someone 'knows' whether *p* we are certifying him as an informant who satisfies this standard. With respect to the general properties of the concept of a reliable informant, objectivisation will abstract away from the speaker-relative features of these properties. In particular, it will leave (a) unaffected, it will relax

¹⁴ Even if I am not so inclined, I still need to appreciate their point of view if I am to be successful at getting them to help me (1990: 88).

¹⁵ Craig does not discuss whether it is *only* purposes that drive the standards up or other aspects of the epistemic context as well. I briefly discuss this point in the final paragraph of Sect. 4.3.

conditions (b), (c), and (e),¹⁶ and it will turn (d) into a strong reliability condition of the form

(d*) The informant is highly likely to be right about p ¹⁷

in order to eliminate the problematic speaker-relativity that creates a gap between Craig's initial hypothesis and the concept of knowledge.¹⁸ Thus, identifying reliable informants becomes what we might call 'socially-directed'; we are not evaluating whether someone is reliable enough for our own individual purposes, but rather whether that person is reliable enough for an indeterminate range of purposes.

This is a crucial feature of Craig's account. In order to avoid the mismatch between our intuitions about the extension of 'knowledge' and his initial hypothesis, Craig argues that our familiar concept of knowledge is what remains *after objectivisation*; thus, in saying that someone knows whether p , we are certifying him as an informant who satisfies the objectivised standard. In his own words, "[k]nowledge lies at the end of the road of objectivisation" (*ibid*: 91). The endpoint is the idea of "someone with a very high degree of reliability, someone who is very likely to be right—for he must be acceptable even to a very demanding inquirer" (*ibid*).

4 The Practical Origins of Contextualism

At the beginning of this paper I claimed that a practical explication of the concept of knowledge could be extended to motivate contextualism about 'knows'. In what follows, I will support this claim by defending a version of contextualism that is suggested by the practical explication discussed above.¹⁹ My goal is to show that

¹⁶ Whether or not the objectivised concept will feature descendants of (b), (c), and (e) that are relaxed enough to be satisfied by virtually anyone to whom we would intuitively attribute knowledge is a contentious issue. Some commentators, such as Kappel (2010: 86), claim that these conditions (or at least some of them) drop out of the picture altogether as a result of objectivisation. I simply leave this matter open, since my focus is on condition (d*).

¹⁷ The skeptic will claim that there isn't any reason to think that someone is more likely to be right about p than not- p . Craig discusses skepticism in chapter 12 of his book, and I will discuss it in Sect. 4.6.

¹⁸ I borrow this way of framing the issue from Kelp (2011).

¹⁹ Henderson (2009, 2011) uses Craigean insights to motivate his own version of contextualism, which he calls 'gate-keeping contextualism'. While there is much in Henderson's work with which I agree, we develop our views in different ways. A detailed discussion of gate-keeping contextualism will take us too far afield, but I will mention one essential difference. Henderson's account rests on a distinction between two broad types of communities, 'general source communities' and 'applied communities' (2009: 120). I find this distinction troublesome because, as Henderson recognizes, there will be much membership overlap: people will be members of several groups of both kinds of communities. This makes it difficult to explain how these groups can operate with distinct sets of epistemic standards. Given the inevitable overlap between these groups, it seems that the standards associated with one of these communities will influence the other, leaving no way to distinguish them. In contrast, my own view will show that an intersubjectively determined epistemic standard (the 'objectivised' standard) will set reasonable constraints on the context-variability of 'knows' without posing problematic groups or communities.

putative facts about the role (or at least one major role) of knowledge attributions provide evidence for contextualism about ‘knows’.

4.1 The Invariantist Interpretation of Craig

Recall that the process of objectivisation will strengthen rather than relax condition (d)—the reliability condition. According to Craig, this will

edge us towards the idea of someone who is a good informant as to whether *p* whatever the particular circumstances of the inquirer, whatever rewards and penalties hang over him and whatever his attitude to them. That means someone with a very high degree of reliability, someone who is very likely to be right—for he must be acceptable even to a very demanding inquirer. (*ibid*: 91)

At first glance, this characterization seems to support ‘high-standards invariantism’, a rival to contextualism. According to this view, the standard for ‘knows’ is set fairly high in all contexts. Kelp (2011) rightly thinks that recent debates over the semantics of knowledge attributions have revealed that people’s concerns can sometimes be quite pressing (as DeRose’s bank cases illustrate nicely). If ‘knows’ is used to identify informants who, among other things, are as likely to be right as almost anyone’s interests or purposes require, then “Craig’s proposal will arguably lead him to a so-called ‘high-standards’ version of classical invariantism” (Kelp 2011: 65).

Martin Kusch thinks Craig offers little in support of this view: “Why not hold instead, say, that objectivisation leads merely to demand that the knower will be reliable so as to satisfy *ordinary* concerns?” (2011: 11). Kusch is presumably recommending a moderate version of invariantism. I will not enter into a debate about how high the invariant standard will be set, especially since I intend to argue against invariantism. For our purposes, let’s presume that the standard will be fairly high, but not so high as to entail skepticism. Kelp has sympathy for this view (2011: 65).²⁰

The invariantist interpretation has *prima facie* plausibility. Admittedly, there seems to be a tension between the idea that the concept of knowledge is progressively less tied to the particular concerns of the user as a result of objectivisation, on the one hand, and the contextualist’s thesis that the meaning of ‘knows’ is dependent on features of the ascriber’s context, on the other hand. This lends some support to the idea that Craig’s account is a version of invariantism. The challenge, then, is to explain why this view is mistaken.²¹

²⁰ I do not consider versions of invariantism that set the standard fairly low because Craig is clear that the standard will be quite high as a result of objectivisation.

²¹ Whether or not Craig actually endorses some version of contextualism is not my concern. My aim is to show that Craig’s hypothesis can motivate contextualism; it is not to show that Craig would endorse this view. In other words, I shall argue that the reasons which seem to motivate the invariantist interpretation do not force someone attracted to Craig’s general idea to be an invariantist.

Notice that the non-objectivised concept of a reliable informant presents us with a contextualist-friendly picture. According to Craig, a reliable informant must be reliable enough to suit the inquirer's own needs and interests; however inquirers will encounter a range of contexts in which they will have different demands depending on their needs and interests: in some contexts the inquirer will require very little by way of an informant's reliability, whereas in other contexts the inquirer will demand much more (Fricker 2010: 64). As a result, what it takes for an informant to count as reliable will change from context to context.

But if the concept of knowledge arose from our need to flag informants reliable enough to satisfy our needs and interests, and if our needs and interest vary depending on the context, then we should expect 'knows' to vary in order to serve that purpose. Why, then, might we shift from the contextualist-friendly scenario described in the hypothesized state of nature to the invariantist-friendly concept that is the product of objectivisation?

It will help to recall Craig's motives for thinking that our familiar concept of knowledge is the product of objectivisation: first, we need to account for undetectable informants and informants with whom the channels of communication are not open (1990: 88–89); second, because we often do not know the practical needs or interests of agents to whom we might recommend an informant, the standards rise to a point where the informant is someone reliable enough to suit almost all practical purposes (*ibid*: 94). Craig provides two examples of this: (a) I may have an interest in collecting information while it is available, without knowing when, or why, or under what pressures it may be needed; (b) I want others to make assessments of informants because that may turn out to be useful to me, even though they do not have my particular concerns in mind when making such assessments (*ibid*: 91). The concept becomes objectivised when it is socially-directed; i.e., I come to think about the (possible or actual) needs and interests of others, and my own attitude towards the expected utility of being right and the perceived risk of being wrong is replaced by a more universal standard (Feldman 1997: 210).

These are reasonable motives for thinking that our familiar concept of knowledge is a result of objectivisation, since any constructed concept that fails to accommodate the aforementioned concerns will be very unlike the everyday concept of knowledge. Nevertheless, I will now argue that our concern with the needs and interests of others does not provide sufficient reason to think that the standard for 'knows' will be invariant. Rather, I will suggest a way to accommodate Craig's motives for objectivising the concept of knowledge without endorsing invariantism.

4.2 Practical Interests Contextualism

This section will sketch a view I call 'practical interests contextualism'. Much of what I say here is compatible with versions of subject-sensitive invariantism and insensitive invariantism. For this reason, I will further develop my view in ways that

favor contextualism over subject-sensitive invariantism and insensitive invariantism in Sects. 4.3 and 4.4.²²

I shall argue that our need to objectivise the concept of knowledge does not preclude the ability to tailor this concept to an individual's needs or interests.²³ Rather, the concept of knowledge objectively construed operates as the *default* or *ordinary* standard for 'knows', which nevertheless permits the standards to change in response to contextual shifts. This allows for cases where the person recommending an informant may tailor his or her recommendation to the needs or interests of a specific inquirer.

Contextualism explains our willingness to ascribe knowledge to Keith in *Bank Case A*, despite the fact that he does not seem to meet the standard to satisfy Rachel's practical concerns in *Bank Case B*, on the grounds that we are aware that nothing significant turns on whether the bank is open in *Bank Case A*. Thus, in attributing knowledge to Keith in *Bank Case A* we are implying that the likelihood of his being right is as great as Rachel's concerns require. In *Bank Case B*, however, we recognize that, given what is at stake, the likelihood that Keith is right is not high enough to satisfy Rachel's practical concerns (despite the fact that the likelihood that Keith is right has not changed); hence, we refrain from attributing knowledge.

This idea is in keeping with something Craig says: "In recommending an informant to you I am implying that the likelihood of his being right is as great as your concerns require" (94). What we should add is the following idea: if I do not know the particular needs or interests of the person to whom I am recommending an informant, as is often the case, the standard for counting as a reliable informant will be set to a point where one must satisfy the needs or interests of almost any inquirer. By operating with this standard—what I call the 'default' or 'ordinary' standard—a recommender need not fear reproach if others follow his recommendation. For example, if I recommend to someone an informant whom she should ask about the times of connecting flights in Chicago, my inquirer should not expect me to take into account how important it is for *her* to get the right answer (what is at stake if she misses her flight connection), unless of course her particular needs or interests are somehow salient to me. Ordinarily I use 'knows' to flag informants who are reliable enough to satisfy most practical concerns, and therefore presumably reliable enough to satisfy my inquirer's concerns.

I call this view practical interests contextualism because the epistemic standards are partly set by the practical needs and interests of a given inquirer, but they are also constrained by more general practical interests of humans *qua* information gatherers. On this account, we require the objectivised concept of knowledge for the reasons Craig mentions (e.g., we often do not know the particular needs or interests of many people to whom an informant may be reliable), but we can also tailor our knowledge ascriptions to individual inquirers in cases where an inquirer's concerns are particularly pressing. This offers a plausible diagnosis of the intuitions involved in various high- and low-standards cases. Let's set aside the bank cases for a

²² I leave aside truth relativism about knowledge (see MacFarlane 2005).

²³ 'Tailoring' works somewhat like Lewis's (1996) Rule of Accommodation. We are expected to be accommodating or sensitive to contextual shifts and this is presumed to be semantic.

moment and try out this hypothesis on another pair of cases. Consider the following scenario, which I borrow from Cohen (2000: 95):

Airport Case A. Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask if anyone knows whether the flight makes any stops. A passenger Smith replies, “I do. I just looked at my flight itinerary and there is a stop in Chicago.”

Presuming that the flight has a layover in Chicago, the intuitive verdict is that Smith ‘knows’ that the flight has a layover in Chicago. Now consider the following high-standards case (also from Cohen):

Airport Case B. The set-up is the same as the aforementioned scenario except now Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, “How reliable is that itinerary, anyway. It could contain a misprint. They could have changed the schedule since it was printed, etc.” Mary and John agree that Smith doesn’t really know that the plane will stop in Chicago on the basis of the itinerary. They decide to check with the airline agent.

In *Airport Case B* Smith no longer counts as ‘knowing’ given the standard introduced by Mary and John’s practical interests. Presumably one wants to say that in *Airport Case A* Smith is right, whereas in *Airport Case B* Mary and John are right. Plausibly, the different interests governing Mary and John’s travel plans make different judgments appropriate. Given their interests, Mary and John appropriately do not take Smith to ‘know’ because what counts as being reliable enough to ‘know’ is relative to different contexts. We are aware of Mary and John’s practical interests and we recognize that Smith is not reliable enough to satisfy those interests; thus, we refrain from attributing ‘knowledge’.²⁴

Practical interests contextualism illustrates how context-dependence is a key aspect of the practical explication of the concept of knowledge. According to this view, an inquirer is looking for someone who is likely enough *in the context* to be right as to whether *p*, which suggests that what it takes to be a reliable informant will change from context to context (Fricker 2010: 63). In contexts where the inquirer’s needs or interests are salient, the recommender will tailor her use of ‘knows’ to accommodate the relevant expectations (within reason); however, in cases where we do not tailor ‘knows’ to the particular needs or interests of a given inquirer, we can instead use the socially-directed application of ‘knows’, which flags informants that are reliable enough to suit the needs or interests of most inquirers (i.e. in ordinary contexts). In many contexts knowledge ascriptions pick out the socially-directed standard, but in some contexts knowledge ascriptions are ‘individually-directed’ rather than ‘socially-directed’.

²⁴ Whether it is the perceived stakes that matter or the actual stakes depends on the context. I agree with DeRose (2009: 55) that perceived stakes may take us from a low-standards context to a high-standards context; however the reverse is less plausible. If someone is in a context in which much is at stake but they fail to recognize this fact, the standards reflect the actual stakes.

My core claim is this: If the concept of knowledge arose from our general need to flag informants reliable enough to satisfy our (more specific) needs and interests, and if our (more specific) needs and interests vary depending on context, then we should expect ‘knows’ to vary in order to serve its purpose. Our general need to acquire information from a pool of reliable informants implies an interest-relativity for the ascription of reliability.

This view explains our need to objectivise the concept of knowledge without endorsing invariantism. This is a welcome result because practical purposes do not always require informants with a very high probability of being right, although we sometimes demand informants who satisfy very high standards. Craig’s example of a murder trial (1990: 94) is illuminating here because it suggests that although the objectivised standard of ‘knows’ is quite high, it is not high enough for all practical contexts. In the context of a murder trial, where the most serious consequences turn on it, “our inclination is to wind the standard up yet another notch” (*ibid*). Objectivisation thus explains why the default (or ordinary) standard for ‘knows’ is quite high without excluding the possibility that practical concerns can either push the standards up or pull them down.²⁵

Putting all this together, we get the result that knowledge attributions are sensitive to context because they involve different expectations of reliability, and expectations of reliability are sensitive to context. Knowledge attributions inherit the context-sensitivity of what counts as reliable enough.²⁶

Contextualism faces two major rivals: subject-sensitive invariantism and insensitive invariantism. Each position offers a competing account for our intuitive judgments about whether an agent ‘knows’. In the next two sections, I will show that my version of contextualism has several advantages over these views.

4.3 Against Subject-Sensitive Invariantism

The concept of knowledge that we have been constructing from our hypothesis bears some relation to Stanley’s (2005) proposal that facts about practical rationality must be incorporated into the truth conditions of knowledge ascriptions. Hawthorne (2004) defends a similar view that he calls ‘subject-sensitive invariantism’. I shall henceforth refer to these views collectively as ‘SSI’. While my view bears some similarities to SSI, this section will demonstrate that my proposal actually puts considerable pressure on SSI as a viable alternative to practical interests contextualism.

Both contextualism and SSI suggest ways in which whether one ‘knows’ depends on context, however SSI holds that the epistemic standards increase as the value of what is at stake for the *subject* increases (the person to whom knowledge is attributed), whereas standard versions of contextualism claim that the truth conditions for knowledge ascriptions are relative to the context of the *attributor*

²⁵ Admittedly, there is a limit to how low the standards can go—i.e., it is doubtful that they can go so low as to include cases in which S truly counts as ‘knowing’ that *p* even though S has an unjustified belief that *p*.

²⁶ This is to be expected since ‘reliable’ is a gradable adjective and as such context-sensitive, just like ‘tall’ or ‘flat’.

of the knowledge claim. Thus, the dispute isn't over whether practical interests play a role in determining whether a subject 'knows', but rather over *whose* practical interests play this role.

Stanley argues that there are certain cases which contextualism has trouble accommodating. Consider an amended version of *Bank Case B* in which Keith claims to 'know' that the bank is open on Saturday even though a lot depends on his being right (for him). Also suppose that nothing much depends on Keith's being right for *us*. Intuitively, Keith's claim is false. According to Stanley, this intuitive verdict spells trouble for contextualism because it is the interests of the subject that seem to govern the standards for 'knowledge', rather than the interests of the attributor.

This case (and similar cases) will cause trouble for contextualism only if the attributor's context cannot be sensitive to the needs or interests operative in the subject's context. However there is no reason to think that contextualism is committed to this claim (DeRose 2009: Chap. 7; Greco 2008: 424–425; Lewis 1996: 561). As Greco writes,

If we are considering whether *we* should go to the bank sooner rather than later, then it makes sense to consider what *S* knows relative to *our* needs and interests. But if we are considering whether *S* should go to the bank sooner rather than later, then it makes sense to consider whether *S* knows relative to *his* needs and interests. Such a position would remain a version of attributor contextualism, since the truth-values for knowledge attributions would continue to vary across attributor contexts. Nevertheless, some attributor contexts would be partly defined by subject interests. (2008: 425)

This suggests that the evidential standards for knowledge ascribing sentences may vary relative to different practical environments. It is therefore mistaken to think that the practical situation of the subject of a knowledge ascription is irrelevant by the lights of contextualism. Contextualism can allow the standards for the satisfaction of 'knows' to be influenced by the practical context of the subject of the attribution.

There are many reasons why ascribers might consider the practical situation of the subjects of knowledge ascriptions even though, ultimately, it is the ascriber's situation that calls the shots. For example, the ascriber might be trying to help the subject decide how to proceed on an issue that is of the utmost importance to the subject but not to the ascriber. In such a context the ascriber will be affected by the subject's interests, which call for more demanding epistemic standards. The ascriber might also be certifying someone else as an informant who is reliable enough for the subject's purposes.²⁷

My version of contextualism does not suggest that the needs or interests of some individual alone (i.e. the attributor or the subject) constitute the relevant practical environment. Rather, what determines the relevant epistemic standard in a given context is a complex function of the knowledge attributor, the needs or interests of the subject whose belief is epistemically assessed, and the epistemic community to

²⁷ Thanks to Robin McKenna here.

which the knowledge attributor and the subject in question belong. This is a virtue of my account because, as Greco writes, “It is more plausible that the truth-values of knowledge claims depend on the actual interests of some relevant group, than it is that they depend on the perceived interest of some individual” (2008: 431). If the point or purpose of knowledge attributions is to flag reliable informants so that a community can pool and share information for an indeterminate range of projects, it makes little sense to think that ‘knowledge’ should be tied to the interest of the subject of an attribution rather than a more general standard.²⁸

SSI also has difficulty explaining why ascribers in high-stakes contexts will use third-person knowledge attributions to deny ‘knowledge’ to subjects in far away low-stakes contexts (DeRose 2009). That speakers will apply to such far-away subjects the same demanding standards that they are applying to themselves is a problem for SSI, which predicts that features of the subject’s context will set the operative standard. Contextualism, on the other hand, can readily accommodate this judgment. This is another reason to think that variance in the truth conditions should be given a contextualist, rather than SSI, treatment.

My view ties the context-variability of ‘knows’ to the context-variability of reliability, and I have argued that who counts as reliable depends on one’s needs and interests. Do I thereby run into a familiar problem for SSI, namely, trouble explaining cases in which practical factors stay fixed (e.g. neither the subject’s nor the attributor’s needs and interests change and the stakes remain the same) and yet our willingness to ascribe or deny knowledge seems to shift? This would be a problem for my view if I maintained that practical factors are the *only* relevant contextual parameter that determines epistemic standards for ‘knowledge’; however I allow other features of context to affect the epistemic standards, such as whether or not an uneliminated alternative possibility has been made salient and is taken seriously in a context. I leave this matter open.

4.4 Against Insensitive Invariantism

According to insensitive invariantism, what counts as being in a sufficiently good epistemic position to qualify as knowing does not vary with practical facts about either the subject’s or attributor’s context. Whatever their practical interests may be, there is some good epistemic position in which an agent must stand with respect to any proposition in order for that agent to count as knowing it. In other words, the truth conditions for knowledge sentences do *not* vary with context. As a result, non-skeptical invariantists typically maintain that in *Bank Case A* Keith says something true, while in *Bank Case B* he says something false; similar considerations apply in Cohen’s airport cases and other examples.

Presuming that speakers really do make divergent knowledge attributions in different contexts, it seems uncontroversial that *in some sense* contextual shifts occur. What, then, can invariantists say about the contextualist ‘data’? While contextualists say that both of Keith’s utterances are appropriate because they are true, invariantists

²⁸ David Henderson (2009: 124) raises a similar objection to SSI.

maintain that both utterances are appropriate but *not* because they are true. Rather, what varies with context are the conditions of warranted assertability. Invariantists usually spell out the idea of warranted assertability in terms of Gricean rules of conversation (Grice 1989). According to this idea, speakers can be pragmatically warranted in asserting something false, which explains how different utterances in different contexts can be appropriate. Here I focus on non-skeptical versions of invariantism according to which our knowledge attributions in low-standards cases are true and appropriate, whereas our knowledge denials in high-standards cases are false but appropriate.

Invariantists often say that asserting the truth in a high-standards context would have a false implicature and thereby convey something false to the audience; e.g., that one's evidence is strong enough for current practical purposes. In contrast, falsely asserting "I don't know whether the bank will be open" has a true implicature and conveys that Keith's evidence is not strong enough for current practical purposes, given what is at stake (Baumann 2011: 158). Thus, the cases which seem to strongly support contextualism are actually neutral on this issue.

An important lesson from Grice (1989) is that facts about use do not by themselves reveal the semantics of 'knows'. The most we can discern from our use of sentences are the conditions for which speakers find it *reasonable* to make assertions, which is distinct from whether such assertions are true. Thus, determining whether a knowledge attribution is true is not a straightforward matter because linguistic behavior is not purely generated by the semantics of a term; it can also be influenced by conversational pragmatics and distorted by performance errors. While a complete refutation of invariantism cannot be accomplished here, I will provide two reasons to favor practical interests contextualism.

First, if knowledge attributions play anything like the role I have outlined, then it renders the invariantist account much less plausible and possibly inconsistent. Presuming that knowledge attributions serve as a marker for reliable informants, why would it be true (but inappropriate) to say that Keith 'knows' in *Bank Case B* if his evidence *isn't* strong enough for current practical purposes? Put differently, why would 'knows' function in a way that renders our knowledge attributions true when the subject of the attribution fails to qualify as a good (enough) informant? Invariantism seems to entail that the truth conditions for knowledge sentences will sometimes be inconsistent with the point or purpose of knowledge attributions. In contrast, my version of contextualism tells a consistent story about how our sense of the purpose of knowledge attributions will condition our knowledge-ascribing/denying practices.

Second, if I am right that practical purposes are relevant to determining the shape of our concept of knowledge, and if concepts function as (roughly) the meanings of words, then our hypothesis about the purpose of knowledge attributions suggests that 'knows' will exhibit the semantic context-sensitivity predicted by my account. If it is plausible that the concept of knowledge arose in connection with the practice of flagging good informants, and if that purpose is at least partly constitutive of the concept of knowledge (which is a semantic entity), then the semantics constitutive of

the concept turn on what makes for our success at flagging such informants (Henderson 2011: 86).²⁹

4.5 The Default Use of ‘Knows’

My view explains why epistemic standards may be affected by features of both the attributor’s and subject’s respective contexts (presuming these differ), as well as the more general needs and interests of humans *qua* information gatherers. The objectivised standard is the standard that, because of various facts about humans and their needs, often best serves the social role of knowledge ascriptions, which is flagging informants reliable enough for our purposes. If the purpose of knowledge attributions is to flag reliable informants, then it is natural to think that the interests of those in the community for which informants are thus regulated should condition what is required to ‘know’ (Henderson 2009: 123).

One might ask why ‘knows’ in the objectivised sense is the ‘default’ use of that term, rather than just one contingent application among many. I prefer the default view to the contingent view because it better explains some important features about our knowledge-ascribing practices. To illustrate, recall *Airport Case B*. This case highlights several considerations that favor my view.

First, inquirers tend to have certain *expectations* surrounding the ordinary usage of ‘knows’; for example, informants are typically expected to operate with a standard of ‘knows’ that is high enough to satisfy the needs or interests of most inquirers. However, it would be strange and perhaps even wrong for Mary or John to expect Smith (without telling him) to take into account how important it is for *them* to get the right answer, what is at stake for them if he is wrong. We can easily make sense of this idea if we presume the notion of a default standard. Smith has been an epistemically responsible informant if, when claiming to ‘know’, he is reliable enough to satisfy the needs or interests of most inquirers.

Second, even if Smith fails to count as reliable enough to satisfy Mary and John’s particular concerns, as in *Airport Case B*, it seems unlikely that Mary and John will *reproach* Smith for claiming to know that the flight stops in Chicago on the basis of the published itinerary.³⁰ This is likely because they recognize that Smith is conforming to a common and widely accepted epistemic practice, one according to which it is acceptable to base one’s beliefs about flight layovers on the information in published itineraries. This practice has likely emerged because the practical needs and interests of travelers are not such that they commonly require information that is more reliable than that which is published in flight itineraries. While more reliable information is surely better, we are not always willing to pay the “information costs” needed to acquire it.

Moreover, Mary and John would probably expect Smith to *retract* his claim to know if they convinced him that he is not reliable enough to suit their particular

²⁹ For a defense of the idea that evaluative concepts are semantic entities that are shaped by their roles or purposes, see Henderson (2011).

³⁰ This is compatible with Craig’s idea that “a practice develops of setting the standard very high, so that whatever turns, for others, on getting the truth about *p*, we need not fear reproach if they follow our recommendation” (Craig 1990: 94).

concerns. For example, if Mary tells Smith that it is absolutely vital that they land in Chicago and if she reminds him of the possibility that the flight itinerary may contain errors, it is likely that Smith will no longer consider himself reliable enough to satisfy their concerns; thus, he will withdraw his self-ascription of knowledge.

What if the recommender *doesn't* tailor his or her use of 'knows' to accommodate the relevant expectations? For example, what if Smith is aware of the high stakes present in *Airport Case B* but he nevertheless asserts that he 'knows' the flight has a layover? The intuitive response is that Smith is speaking inappropriately, and according to contextualism his attribution is inappropriate because it is false. Given Mary and John's pressing concerns, Smith doesn't count as a reliable enough informant for them and thus does not 'know'.

Finally, the idea of a default epistemic standard gains support when we reflect on the fact that having such a standard makes pooling and sharing information much easier. If the standards for 'knows' varied wildly across contexts, then our knowledge attributions wouldn't often pick-out informants on whom others could rely for a variety of purposes.³¹ Without a default standard that sets reasonable constraints on the context-variability of 'knows', knowledge reports could not "serve as a piece of common coin, usable by and useful to a number of people or groups not part of the immediate situation from which it emerges" (Rysiew 2012: 286). The notion of a default standard better accommodates the trans-contextual role of knowledge ascriptions.

These considerations render contextualism consistent with Craig's claim that our familiar concept of knowledge is what remains after objectivisation. The default use of 'knows' is the one that a competent user of that term will expect other competent users to employ when a specific inquirer's needs or interests are not taken into account. Even if the standards for 'knowledge' vary across contexts, there is a stable, central, default use of that term around which other uses center. This points to how the invariantist line of thought has some plausibility, but the connection with practical interests does not totally vanish. The objectivised use of 'knows', which is perhaps what contextualists refer to when they speak of a so-called "ordinary" sense of 'knows', reflects our epistemic community's inter-subjective needs and interests.

4.6 A Skeptical Worry Answered³²

The result of objectivisation is an epistemic position sufficiently strong to allow informants to fittingly serve as sources of information on which people can reasonably draw for a variety of purposes. This calls for a very high quality of epistemic position. However, why not heighten the standard to a level that would satisfy *everyone's* concerns?

The skeptical threat arises if we think that the role of knowledge attributions is to certify informants who are reliable enough to satisfy the demands of *any* inquirer whatsoever, including arbitrarily demanding inquirers. But the best conceptual

³¹ This objection has been made by Williamson (2005: 101) and Rysiew (2012: 286–290).

³² This section is heavily indebted to Henderson (2009).

strategy does not call for such a high standard. To insist upon this condition would exclude a vast number of cases where a particular informant is perfectly suitable for the purposes of many, perhaps most, inquirers. Our withholdings would suggest to others that these perfectly reliable informants do not have sufficiently good information on which we can base our beliefs and actions. But there is a clear benefit to treating such individuals as good informants if they are likely enough to be correct for our purposes: we do not always require our informants to have the highest quality of epistemic position.

According to our practical explication, the concept of knowledge is rooted in the social–epistemological need for pooling and sharing information, which generates a need to flag informants who can provide us with reliable information. The skeptical result runs against this approach because it would frustrate our communal epistemic practices. To deny knowledge to a potential informant would be to suggest that he is not in a strong enough epistemic position to appropriately contribute to the stock of information and beliefs on which others can draw for an indeterminate range of projects.

5 Concluding Remarks

The above considerations illustrate how a practical explication of our concept of knowledge provides support for contextualism about ‘knows’.³³ If the purpose of knowledge attributions is to flag informants that are reliable enough to satisfy our needs and interests, and if our needs and interests vary, then what it takes to be a reliable informant will change from context to context; thus we should expect that ‘knows’ and its cognates will also vary in order to serve that purpose.

This account is also profitable because it answers some important philosophical questions that contextualism *qua* semantic thesis leaves open, particularly why knowledge attributions are valuable, why ‘knows’ exhibits context-variability, and why this term enjoys such widespread use. Knowledge attributions are valuable and enjoy widespread use because they play a vital role in our cognitive economy, namely, they allow us to flag reliable informants. In addition, ‘knows’ displays context-variability because the standard required to count as an informant reliable enough to satisfy an inquirer’s needs or interests will depend on what those needs or interests are.

My view is also attractive because it honors some insights of SSI and explains the appeal of insensitive invariantism within one unified account that is motivated by a plausible story about the purpose of knowledge attributions.

Finally, nothing about my approach rules out the usual methodology on which contextualists and other epistemologists rely. It is a common methodological practice in epistemology to consult our intuitions about possible cases, as well as

³³ Rysiew (2012) raises doubts about the advisability of arguing directly from facts about the function of ‘knows’ to facts about its semantics. Rysiew’s paper was published after my paper was written, so I cannot provide a detailed reply here. However, much of what I have argued does go some way to answering Rysiew’s worry that “more needs to be said by [contextualists] about how their favored account of knowledge ascriptions enables them to play their broader, social role” (*ibid*: 291).

consult the linguistic data regarding actual language use. According to Greco, with whom I agree, “this sort of methodology has made for a number of interesting arguments and important insights” (2008: 423). It is therefore desirable that my approach can supplement this methodology.

Acknowledgments I am grateful to Hallvard Lillehammer, Ernest Sosa, Tim Crane, Arif Ahmed, Michael Blome-Tillmann, Mikkel Gerken, Berit Brogaard, Robin McKenna, Jess Kwong, members of the Cambridge Epistemology Group, and members of Ernest Sosa’s Epistemology Group at Rutgers University for helpful advice. This paper was written while I was generously funded by the Social Sciences and Humanities Research Council of Canada.

References

- Baumann, P. (2011). WAMs: Why worry? *Philosophical Papers*, 40(2), 155–177.
- Brown, J. (2005). Comparing contextualism and invariantism on the correctness of contextualist intuitions. *Grazer Philosophische Studien*, 69(1), 71–100.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2, 91–123.
- Cohen, S. (2000). Contextualism and skepticism. *Noûs*, 34(1), 94–107.
- Craig, E. (1990). *Knowledge and the state of nature*. Oxford: Oxford University Press.
- Dancy, J. (1992). Review of knowledge and the state of nature. *The Philosophical Quarterly*, 42(186), 393–395.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.
- DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104(1), 1–52.
- DeRose, K. (2009). *The case of contextualism*. Oxford: Oxford University Press.
- Feldman, R. (1997). Review essay: Knowledge and the stated of nature: An essay in conceptual synthesis. *Philosophy and Phenomenological Research*, 57(1), 205–221.
- Fricker, M. (2010). Scepticism and the genealogy of knowledge: Situating epistemology in time. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social epistemology*. Oxford: Oxford University Press.
- Greco, J. (2008). What’s wrong with contextualism? *The Philosophical Quarterly*, 58(232), 416–436.
- Greco, J. (2010). *Achieving knowledge*. Cambridge: Cambridge University Press.
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge: Harvard University Press.
- Hannon, M. (2013). ‘Knows’ entails truth. *Journal of Philosophical Research* (Forthcoming).
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Henderson, D. (2009). Motivated contextualism. *Philosophical Studies*, 142(1), 119–131.
- Henderson, D. (2011). Gate-keeping contextualism. *Episteme*, 8, 83–98.
- Kappel, K. (2010). On saying that someone knows: Themes from Craig. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social epistemology*. Oxford: Oxford University Press.
- 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. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 594–597.
- MacFarlane, J. (2005). The assessment sensitivity of knowledge attributions. *Oxford Studies in Epistemology*, 1, 197–233.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35(4), 477–514.
- Rysiew, P. (2012). Epistemic scorekeeping. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions*. Oxford: Oxford University Press.
- Schaffer, J., & Knobe, J. (2012). Contrastive knowledge surveyed. *Noûs*, 46(4), 675–708.
- Schmitt, F. (1992). Review of knowledge and the state of nature: An essay in conceptual synthesis. *Mind*, 101(403), 555–559.
- Sosa, E. (1974). How do you know? *American Philosophical Quarterly*, 11, 113–122.
- Sosa, E. (1988). Knowledge in context, skepticism in doubt. *Philosophical Perspectives*, 2, 139–157.
- Sripada, S., & Stanley, J. (2012). Empirical tests and 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. (1973). *Problems of the self*. Cambridge: Cambridge University Press.

-
- Williams, B. (2002). *Truth and truthfulness*. New Jersey: Princeton University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (2005). Knowledge, context, and the agent's point of view. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth*. Oxford: Oxford University Press.