
STABILIZING KNOWLEDGE

BY

MICHAEL HANNON

Abstract: If epistemic contextualism is correct, then knowledge attributions do not have stable truth-conditions across different contexts. John Hawthorne, Timothy Williamson, and Patrick Rysiew argue that this unstable picture of knowledge attributions undermines the role that knowledge reports play in storing, retrieving, and transmitting useful information. Contrary to this view, I argue that the truth-conditions of knowledge attributions are more stable than critics have claimed, and that contextualism is compatible with the role knowledge attributions play in storing, retrieving, and transmitting information across contexts. In particular, I discuss a social dimension of ‘knowledge’ that limits contextual variability. This indicates a new way of characterizing contextualism.

1. Introduction

The word ‘know’ seems to find a comfortable equivalent in every human language, which suggests that the concept expressed by this word plays an important role in our lives.¹ An adequate theory of knowledge attributions should therefore fit plausible assumptions about the role (or roles) that the concept of knowledge serves and the needs it satisfies.² It would be puzzling if the correct analysis of knowledge attributions failed to shed light on why we have a concept of knowledge; thus, our strategy should be to use the role(s) of the concept of knowledge as an adequacy constraint on a theory of knowledge attributions. Whatever our theory tells us, it had better be able to underwrite the role(s) that our concept of knowledge plays in epistemic evaluation.

One recent philosophical debate has centered on the question whether knowledge attributions are context sensitive. The traditional view is that sentences attributing knowledge have stable truth-conditions across different contexts of utterance. Contextualists deny this. They argue that the truth-conditions for knowledge claims vary depending on the context of utterance. It has been a matter of controversy to determine precisely how

the standards for 'know' get set in a particular context; however, it is generally agreed that the standards are at least partly determined by some complicated function of speaker interests, listener expectations, presuppositions of the conversation, what is at stake, and salience relations (Cohen, 1999).

An important criticism of contextualism is that this view implies an unstable picture of epistemic standards that prevents knowledge reports from storing and transmitting useful information. Patrick Rysiew argues that without stable truth-conditions, 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, p. 286). John Hawthorne (2004) and Timothy Williamson (2005a) raise a similar objection. They worry that the more extensively 'know' varies in reference with context, the less plausible it is that people can acquire epistemically useful information from knowledge attributions made in other contexts. Let us call this the *trans-contextual role problem*.

This article will show that the trans-contextual role problem isn't really a problem for contextualism. Section 2 will outline contextualism and Section 3 will discuss the trans-contextual role problem. Section 4 will show that a similar problem applies to views that rival contextualism, particularly subject-sensitive invariantism (as defended by Hawthorne) and moderate insensitive invariantism (as defended by Rysiew and Williamson).³ Section 5 will demonstrate that there are several ways for contextualists to stabilize the truth-conditions for 'know', which allows knowledge reports to serve a trans-contextual role. I suggest plausible constraints on how the salience of error and practical stakes can affect the epistemic standards; however, my main task is to stabilize the contextual variation of 'know' by appealing to a social dimension of knowledge.⁴ More specifically, I appeal to the increasingly popular idea that knowledge attributions are used to flag reliable informants. My goal is to defend a version of contextualism that sets constraints on the amount of cross-contextual variability that 'know' will tolerate, given certain plausible assumptions about why humans have the concept of knowledge. This allows contextualism to escape the trans-contextual role problem.

2. *Contextualism and the variability of 'know'*

In the *Meno*, Plato tells us that knowledge is 'tethered' (70a) and thus cannot easily run away. Unlike untethered true belief, knowledge is stable in a way that prevents it from easily slipping away from our minds. This idea is widely accepted. To cite one example, Williamson (2000) claims that knowledge is less likely to be undermined by future evidence and

therefore has cross-temporal permanence that true belief does not. Whether or not one accepts Williamson's view, many agree that the state of knowing is fairly stable and non-fleeting.

However, contextualists have suggested that whether somebody truly qualifies as 'knowing' can change abruptly, even when there is no change in one's evidence or reliability. How precisely the epistemic standards get set in a context is controversial, but it is uncontroversial that contextualism posits some degree of contextual variation in the truth-conditions of knowledge ascriptions. The contextualist is led to this conclusion as a way to accommodate the pattern of 'shifty epistemic judgments'. Shifty epistemic judgments occur when people are willing to ascribe knowledge in so-called 'low-standards' cases but unwilling to ascribe knowledge in so-called 'high-standards' cases. A number of examples have been used to illustrate this phenomenon, such as the following pair of bank cases that were originally developed by Keith DeRose (1992) and retold by Jason Stanley (2005):

Low Standards. 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 their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning'.

High Standards. 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 their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But as Sarah points out, banks do change their hours. Hannah says, 'I guess you're right. I don't know that the bank will be open tomorrow' (Stanley, 2005, pp. 3–4).

Presuming that the bank will be open on Saturday in both cases, many people think that when Hannah claims to know that the bank will be open in *Low Standards*, she is saying something appropriate; but it also seems that Hannah is saying something appropriate in *High Standards* when she concedes that she *doesn't* know that the bank will be open on Saturday (DeRose, 1992, p. 914).

Many philosophers with various theoretical predilections have supposed that people will attribute knowledge in low-standards cases and yet

deny knowledge in high-standards cases.⁵ To explain the supposed fact that Hannah speaks appropriately in both cases, contextualists posit that ‘know’ varies in reference with context. Thus, Hannah’s knowledge claim is true in *Low Standards* and her denial of knowledge is true in *High Standards*. For now let us assume this argument is compelling.

3. *The trans-contextual role problem*

We often retrieve and transmit epistemically useful information from knowledge reports made in other contexts. For example, I might read the headline, ‘Scientists now know there is water on Mars’, which signals that I may take the question at issue to be settled. I may then recall this information months or even years later in a different situation, such as a dinner party at which this question is being discussed. Knowledge reports come to us from other contexts all the time. Cross-contextual uses of ‘know’ thus seem essential to the role of knowledge ascriptions, not a marginal phenomenon.

That knowledge attributions play this trans-contextual role in epistemic evaluation is undeniable (or I won’t deny it, at any rate). However, a number of epistemologists have alleged that contextualism implies an unstable picture of epistemic standards that undermines the trans-contextual role that knowledge attributions play in storing, retrieving, and transmitting epistemically useful information. This is a significant objection that, if successful, would constitute a major blow to the plausibility of contextualism. An adequate theory of knowledge attributions should accommodate the idea that the word ‘know’ allows us to store, retrieve, and transmit useful information coded in knowledge sentences made in other contexts. Thus, if contextualism cannot accommodate this trans-contextual role, we will have strong grounds to reject contextualism. This section will outline the trans-contextual role problem as articulated by Hawthorne, Williamson, and Rysiew.

In *Knowledge and Lotteries* (2004), Hawthorne asks us to imagine a speaker who notices it is raining and says ‘It is raining here’. We are also asked to suppose (for the sake of simplicity) that people have a ‘belief box’ that contains sentences of English. On this model, to believe that p requires that one have a sentence in one’s belief box which expresses the proposition that p . According to Hawthorne,

If the speaker wants to *preserve* this information in the belief box, it won’t do to encode ‘It is raining here’ as a permanent inhabitant of the belief box, for the sentence will change its content as the speaker moves through time and space. Better to encode the information by a sentence that is not shifty in this way. Perhaps ‘It was raining there’ (accompanied by a mental snapshot or file of the relevant spatiotemporal location). Or ‘It was raining by the Eiffel Tower at noon on May 4th’ (Hawthorne, 2004, p. 110).⁶

Similarly, we often want to preserve the information we get from knowledge reports. If contextualism is true, the speaker cannot preserve accurate information simply by encoding 'S knows that p ' as a permanent inhabitant of the belief box, for that sentence will change its content as the speaker moves from context to context. Unlike the case of 'It is raining here', however, Hawthorne doubts whether an 'easy fix' is available for knowledge sentences.

Suppose that you are in a context C_1 in which Helen says 'I know that p '. You accept what she says and you store the sentence 'Helen knows that p (plus a date index)' in your belief box. Several weeks later you are in a new context C_2 in which the operative epistemic standard is higher and, as a result, your belief that Helen knows that p will now come out false. (If you were in a new context in which the standards fell, you could potentially lose much true information, since many beliefs that deny knowledge will now come out false.) This would not be a significant problem if we could update the knowledge-attributing and knowledge-denying sentences in our belief box; however, contextualists acknowledge that we are partially semantically blind to the context-sensitivity of the word 'know'⁷: the semantic content of 'know' shifts, but our language does not provide us with a standards index with which to enrich knowledge ascriptions that are tokened in the belief box (Hawthorne, 2004, p. 110). It is therefore unclear how we could manage to accurately preserve information in memory, which is a problem for our belief set.

Williamson develops this point in his paper 'Knowledge, Context, and the Agent's Point of View' (2005a). The reference of paradigm indexicals like 'here' and 'now' varies with easily identified parameters of context, namely time and place. These indexicals also seem designed for 'immediate consumption' within a context (ibid., p. 101). In contrast, our epistemic vocabulary is not intended solely for immediate consumption; the word 'know' plays an important trans-contextual role, such as saying who knew what when. Furthermore, the reference of the term 'know' varies with a less easily identified parameter: epistemic standards. This creates a problem when we consider the trans-contextual role of knowledge ascriptions. If 'know' is to usefully preserve information in memory and transmit it via testimony, then speakers and listeners must be able to retrieve information that is stored or coded by sentences using that word. The context-variability of the epistemic standards for 'know' makes it difficult to retrieve such information.

This reiterates Hawthorne's worry that a knowledge sentence 'P' will express different propositions when uttered in different contexts; in some contexts 'P' will express a truth while in others it may not. Supposing that you acquire some useful information expressed by 'P' in a context C_1 , where 'P' is true, it is not obvious that you can easily retrieve such information and pass it on to someone else in a new context C_2 , where 'P' no

longer expresses a truth.⁸ Thus, we cannot easily consume knowledge attributions and denials made in other contexts.⁹

Rysiew echoes this worry in his paper ‘Epistemic Scorekeeping’ (2012). He writes,

. . . for a stored or received knowledge sentence to be usable by (/useful to) me, I need to know and keep track of what the operative standards were. Alternatively, I need to know and remember what the relevant standard-affecting/or standard-setting facts – e.g., the stakes, the salient counter-possibilities, etc. – were, and what precisely their impact was. Very often, though, we receive and make use of knowledge reports unaccompanied by any of that kind of information (ibid., pp. 286–7).

It would be too demanding to expect people to store information about various factors that are present in a context, such as the relevant practical interests, the perceived costs of being wrong, the various error-possibilities that were salient and taken seriously, and so on, in addition to the relevant knowledge reports themselves. There are a wide variety of individuals with different interests, concerns, assumptions, projects, etc., in many different situations and at different times. Retaining all this information is not just costly – it isn’t feasible. And yet we *do* very naturally understand and make use of knowledge reports that come to us from outside our immediate context.

In order to avoid the trans-contextual role problem, the kind of epistemic scorekeeping involved cannot be implausibly demanding; in particular, it cannot require us to keep track of the particular stakes, error-possibilities, interests, and contextually-operative epistemic standards in each context that a knowledge attribution is uttered. This task is especially demanding if ‘know’ varies extensively in reference with context.

My aim in the remainder of this article is to provide a solution to this problem. Section 4 will show that a similar objection applies to views that rival contextualism, particularly subject-sensitive invariantism (Hawthorne, 2004; Stanley, 2005) and moderate insensitive invariantism (Rysiew, 2001; Williamson, 2005b). Section 5 will argue that the extent of contextual variation is much narrower than has been thought. This allows ‘know’ to have contextualist semantics and still play its trans-contextual role.

4. A problem for subject-sensitive invariantism and insensitive invariantism

The trans-contextual role problem has been formulated as a difficulty for contextualism, but this worry isn’t restricted to views whereby the *contents* of knowledge sentences are affected by context-specific factors. As Rysiew

(2012, p. 288) points out, any view whereby the *truth-values* of knowledge attributions are affected by context-specific features is vulnerable to this challenge. According to ‘subject-sensitive invariantism’ (Hawthorne, 2004; Stanley, 2005), the truth-value of a knowledge report partly depends on facts about the subject’s practical interests, i.e. the cost of being wrong. In order to know the extent to which a subject’s knowledge is due to the strength of her evidence and other truth-relevant properties, on the one hand, and how much it is due to her interest in the proposition in question, on the other hand, we need to know what the subject’s interests are (i.e. how high the stakes, etc.). If we do not know that, then it is not clear what use we might make of a report that they know something that is relevant to my inquiry (Rysiew, 2012, p. 288). Thus, subject-sensitive invariantism also faces this challenge.

Accommodating the trans-contextual role of knowledge reports is not just a problem for non-traditional views like contextualism and subject-sensitive invariantism. Even moderate insensitive invariantism, which is defended by Rysiew (2001) and Williamson (2005b), encounters this challenge. According to this view, the epistemic standards for ‘know’ do not vary either with the context of the knowledge attributor, as with contextualism, or with the context of the subject of a knowledge attribution, as with subject-sensitive invariantism. Rather, the epistemic standards one must satisfy to truly count as ‘knowing’ are stable across contexts. In order to accommodate our shifty epistemic judgments, many insensitive invariantists argue that our intuitions reflect when it is *conversationally appropriate* to ascribe knowledge, not when knowledge is *truly* ascribed (Rysiew, 2001; Brown, 2006). The contextualist is charged with mistakenly taking the fact that the conversational propriety of a knowledge attribution depends on context to show that the truth-conditions of such ascriptions are context sensitive.

Would the stable, contextually insensitive truth-conditions posited by an insensitive invariantist avoid our problem? I argue that it would not.¹⁰ When retrieving information from knowledge attributions made in other contexts, we must be careful not to allow the semantic content to mislead us. For example, an insensitive invariantist might maintain that in *High Standards* when Hannah says ‘I don’t know that the bank will be open tomorrow’ she is saying something contextually appropriate but strictly speaking false. Now, assume that Sarah wants to preserve this information, so she stores the sentence ‘Hannah doesn’t know that the bank will be open on Saturday’ in her belief box. There is a worry that in some future low-standards context, this sentence may mislead people into thinking that Hannah lacks information that she in fact possesses. For instance, Tom might also want to know whether the bank will be open on Saturday, yet there is not much at stake for Tom and so the epistemic standard is such that someone who has recently been to the bank on Saturday (i.e. Hannah)

would qualify as ‘knowing’. However, Sarah tells Tom that Hannah ‘doesn’t know’ and thereby conveys misleading information. Thus, the insensitive invariantist who goes in for a pragmatic story will have to say that we store and recall complex information about the context in which the assertion was made, which is the same problem the contextualist allegedly faces. The trans-contextual role problem is simply transferred to the level of pragmatics.

This suggests that the epistemic scorekeeping required to retrieve epistemically useful information from knowledge attributions is not a consequence of accepting contextualism but rather a consequence of certain empirical claims about our ordinary knowledge attributing practices.¹¹ The empirical claim is that our practice of attributing and denying knowledge appears to be sensitive to contextual factors like the salience of error and/or stakes. The supposed fact that people exhibit shifty epistemic judgments is responsible for the need for epistemic scorekeeping rather than any particular semantic theory of ‘know’. Thus, this difficulty is not a ‘significant disadvantage’ for contextualism (Williamson, 2005a, p. 102).

However, I want to do more than just extend this problem to views that rival contextualism. My larger aim is to provide a principled motivation for contextualism that resolves this worry. To this task, I now turn.¹²

5. ‘Know’ is common coin

The more extensively the truth-conditions for ‘know’ vary with context, the worse the problems of using knowledge reports to preserve and transmit information will be. If the contextual demands on ‘knowledge’ swing freely and widely across contexts, then knowledge reports could not serve as a piece of common coin for people outside of the immediate situation from which it emerges. If, on the other hand, it can be shown that the truth-conditions for knowledge reports do not often change from one context to the next, then we can accurately consume knowledge claims made in a variety of contexts.¹³ The remainder of this article will argue that the epistemic standards will not swing freely and widely from context to context.

5.1. SALIENCE AND STAKES

The two most widely recognized context-shifting mechanisms are salience relations and practical stakes. These may be stated as two rules:

Rule of Salience: the epistemic standards may increase when an otherwise irrelevant alternative is made relevant by drawing a subject’s attention to that alternative.¹⁴

Rule of Stakes: the epistemic standards may increase/decrease depending on what is at stake (i.e. if error would be especially disastrous).

The basic idea is that the standards for ‘know’ are partly set by what alternatives or possibilities are salient in the context, as well as what is at stake in a given context. To what extent do these context-shifting mechanisms raise or lower the epistemic standards? Let us first look at the Rule of Salience.

Although salience relations have played a central role in discussions about whether an alternative is relevant, there is disagreement about what counts as a ‘salient’ alternative and whether a salient alternative is necessarily a relevant alternative. According to one view, merely mentioning an alternative is sufficient to make it relevant: no consciously recognized possibility may be properly ignored.¹⁵ This view is commonly attributed to David Lewis, whose Rule of Attention states: ‘a possibility not ignored at all is *ipso facto* not properly ignored’ (1996, p. 559).¹⁶ This rule applies ‘[n]o matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context’ (ibid.). Lewis thus seems to hold that merely mentioning an alternative suffices for it to be attended to, and thus relevant in that context. Let’s call this the *mere mention view*.

No contextualist should accept the mere mention view.¹⁷ Simply making S aware of a previously recognized defeater (a possibility in which some proposition *p* is false) is not always sufficient to effect a change of an epistemic context with respect to S knowing that *p*. To borrow an example from Lewis (1996, p. 556), a jury that must decide whether Jones shot Smith may still properly ignore the possibility raised in the following desperate proclamation: ‘Ladies and gentlemen, I must point out that the prosecutor has failed to rule out the possibility that it was not Jones but rather Smith’s dog, marvelously well-trained, that fired the fatal shot!’ Assuming that the jury has overwhelming evidence that Jones shot Smith, this example illustrates that even when the stakes are high some possibilities may still be properly ignored. In addition, unless the jury members are ignoring other alternatives more relevant than the one mentioned above, they might rightly be said to know that Jones is guilty. Simply mentioning an alternative is not enough to make a hitherto ignored possibility relevant and thereby create a context in which it is no longer true to ascribe knowledge.

A more plausible idea is that in order for an alternative to be relevant it must not only be mentioned in a context but also *taken seriously* by participants in that context. Let’s call this the *serious alternative view*. DeRose (1995) shows some sympathy for this view, stating that merely mentioning a bizarre alternative (e.g. a well-trained murderous dog) is not

sufficient to render that alternative relevant. Rather, 'mentioning an alternative makes that alternative relevant only if one's conversational partner gets away with making it relevant' (1995, p. 14, fn. 21), where a failure to make it relevant means, roughly, that one or more of your interlocutors reject the possibility as relevant.¹⁸

Assuming it is difficult to get one's conversational partners to take seriously alternatives that they were not already considering, the epistemic standards would not easily change due to the salience parameter. If there is a shared sense amongst our fellow humans about which counter-possibilities to take seriously and which to ignore, then the salience parameter will remain fairly stable across a vast range of contexts, thereby allowing 'know' to serve a trans-contextual role. This assumes that our judgments about which possibilities to take seriously will sufficiently overlap. I do not find this assumption dubious. Even invariantists like Rysiew are committed to there not being gross differences in people's psychologies (see Rysiew, 2001, p. 489).

This does not suggest that *any* possibility which is not taken seriously in a context is irrelevant to that context. If the jury members were to (somehow) ignore the possibility that Jones was the killer and instead focused their suspicion on Smith's dog, the possibility that Jones killed Smith would still be a relevant alternative. We may not ignore just any possibility we please, for that would make true knowledge ascriptions very cheap indeed. Certain alternatives are relevant to a context irrespective of whether they are mentioned or taken seriously in that context. Nevertheless, the jury members in the aforementioned court case are entitled to ignore the possibility that the dog murdered Smith because it is not the kind 'we (normal) humans *take to be* the likely counter possibilities to what the subject is said to know' (Rysiew, 2001, p. 488).

I do not intend to settle this issue here. The purpose of my discussion is to illustrate that any plausible reading of the Rule of Salience will not allow the epistemic standards for 'know' to swing freely or widely. Rather, there are certain constraints on the effects that a salient alternative may have on the context.¹⁹

There are also plausible constraints on the Rule of Stakes. For instance, this rule does not suggest that an agent with little evidence or justification might qualify as a knower simply because the stakes are low. I would not qualify as knowing that there are forty students in my lecture on the basis of a brief visual scan simply because it does not matter to me whether there are exactly forty students. The epistemic standards for 'know' are fairly high even if the stakes are low: someone who is barely reliable does not qualify as a knower, no matter what anyone's cost-benefit balance of being right comes to (Kelp, 2011, p. 56).²⁰

This point is supported by recent empirical evidence reported by Chandra Sripada and Jason Stanley (2012). Having presented people with

a variety of high- and low-stakes cases, they found that ‘the sizes of the stakes effects observed in this study were modest’ (p. 16). This result is not worrying for theorists who maintain that practical stakes affect the epistemic standards, since these advocates need not maintain that ‘manipulating the stakes can magically transform epistemic dirt into gold, and vice versa’ (ibid.). While stakes may affect the epistemic standards, they need not do so often or extensively.²¹

This discussion of the Rule of Salience and the Rule of Stakes shows that the reference of ‘know’ does not vary to the wide extent that many critics of contextualism have supposed. However, you might still worry that the range of appropriate standards is fairly wide, albeit with certain limitations, and thus contextualism must encounter the trans-contextual role problem. In response to this objection, the next section will discuss a social parameter that restricts the contextual variation of ‘know’.

5.2. SOCIAL PRESSURES AND THE ROLE OF KNOWLEDGE ASCRIPTIONS

We need a compelling story about how ‘know’ varies in meaning with context in a way that does not prevent it from serving a trans-contextual role. To stabilize the epistemic standards for ‘know’, I will appeal to a plausible hypothesis about the role of knowledge attributions that was first proposed by Edward Craig in his 1990 book *Knowledge and the State of Nature*. Craig hypothesizes that knowledge attributions are valuable because they help us satisfy some very general needs in human life and thought; in particular, they serve as markers for reliable informants.²² I believe that this hypothesis can support contextualism and demonstrate how ‘know’ serves a trans-contextual role.

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. Often the most efficient way to acquire a true belief as to whether *p* is to ask someone reliable whether *p*. Since on any issue some informants will be more likely than others to provide a true belief, we 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. Epistemic evaluations thus facilitate the survival and flourishing of human beings in communities.

We all have a need to find informants who are reliable enough for our own purposes: I want someone reliable enough for *me*, you want someone reliable enough for *you*, and so on. Let’s call this the *subjective stance*. This picture makes it very difficult to volunteer oneself or another as a good informant because we would be required to keep track of what each

individual's interests and purposes are; however, we usually do not have such information. The subjective stance is therefore inadequate as a general strategy for human epistemic practices. Our needs and practices are more complex than 'I want information solely for my purposes here and now'. There is a need to assess and identify the adequacy of informants for persons and purposes beyond our own immediate interests. We have a need to share information and identify good informants *in general*. Since we often do not know the practical needs or interests of persons to whom an informant may be useful, we need to ensure that informants are reliable enough for an indeterminate range of people and purposes.

Craig therefore appeals to a process he calls *objectivization* (1990, p. 82). Jonathan Dancy describes this process in the following way: 'A concept is objectivized if it becomes progressively less tied to the particular concerns of the user' (1992, p. 394). We abstract away from details of the particular circumstances surrounding one individual's inquiry to arrive at a more intersubjective, socially-directed concept. Let's call this the *objectivized stance*. The objectivized stance resembles our familiar concept of knowledge: when one attributes knowledge to someone, one certifies that the agent is epistemically positioned such that we can freely draw on his or her information.²³ Knowledge ascriptions certify an informant's belief or information to a shared epistemic community.

This idea is nicely illustrated by David Henderson, who asks us to imagine:

... that you and your interlocutors are members of a contemporary general scientific discipline which provides a body of results on which people with an indeterminate range of practical projects might draw while the discipline itself is not associated with one project in a defining or limiting way (Henderson, 2009, p. 126).

In this context, the relevant concerns revolve around providing information that is good enough to satisfy general-purposes, not some concrete limited purpose. Thus, one who certifies sources of information for such a group should be concerned with informants with actionable information for an indeterminate range of projects and various practical interests. Henderson's point is that there will be a *general level* of stringency for knowledge ascriptions. There may be contexts in which a subject S is a suitable informant for a given inquirer, but if S does not meet the general threshold, we would not say that S knows. Rather, we can fall back on more guarded expressions like 'is reliable enough'. This establishes a boundary for our use of the word 'know'. It is necessary for one to be a 'knower' that one satisfies this general stringency requirement.

Can this story be used to motivate contextualism about 'knows'? According to Craig, objectivization will:

... edge us towards 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., p. 91)

At first glance, this characterization seems to recommend ‘insensitive invariantism’.²⁴ 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 objectivization, on the one hand, and the contextualist’s thesis that the meaning of ‘know’ 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 Craig’s hypothesis actually supports contextualism.²⁵

5.3. STABLE CONTEXTUALISM

Notice that the non-objectivized notion 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, p. 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 ‘know’ to vary in order to serve that purpose. Why, then, might we shift from the contextualist-friendly scenario described by Craig to an invariant concept of knowledge as a result of objectivization?

The subjective stance is inadequate as a general strategy for our epistemic practices because we often do not know the practical needs or interests of agents to whom we might recommend an informant. Craig provides two examples of this: first, I may have an interest in collecting information while it is available, without knowing when, why, or under what pressures it may be needed; second, I want others to make assessments of informants because that may turn out to be useful to me, even though the assessors do not have my particular concerns in mind when making such assessments (1990, p. 91). The concept of knowledge becomes objectivized when 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, p. 210).

Reflection on the epistemic needs of a broad class of people provides a principled basis for our common judgments about when a subject ‘knows’. Inquirers want information on which to base their belief and actions. Following Henderson (2009), let’s call this ‘actionable information’. We want sources of actionable information to help satisfy our needs and interests. Basic practical concerns thus generate a standard that is fitting to serve as certifying actionable information for a more or less indeterminate range of projects. However, the epistemic standard may change in some contexts. Even if an informant meets the general stringency requirement outlined above, an inquirer should not rely on his information if the informant is not reliable enough for one’s specific concerns; for example, if the stakes are higher and so one’s expectations are more demanding. Likewise, we should not recommend informants to others if we are aware that the inquirer’s purposes are particularly pressing and so the informant will fall short of the inquirer’s heightened expectations.

DeRose’s bank cases nicely illustrate this point. Hannah’s claim to know that the bank will be open in *Low Standards* seems appropriate; meanwhile, it also seems that she is saying something appropriate in *High Standards* when she concedes that she does not know that the bank will be open on Saturday. In *Low Standards*, Hannah is reliable enough for most people to base their beliefs and to act on her information, so she counts as knowing that the bank will be open. In *High Standards*, however, when Sarah is entertaining drawing on Hannah’s information, she finds that Hannah’s epistemic position is not sufficient to tell against some error possibility that reasonably seems significant for Sarah’s own inquiry, so she judges that Hannah does not know that the bank will be open on Saturday. When we attribute knowledge to Hannah in *Low Standards* we are implying that the likelihood of her being right is as great as Sarah’s concerns require. In *High Standards*, however, we recognize that, given what is at stake, the likelihood that Hannah is right is not high enough to satisfy Sarah’s practical concerns (despite the fact that the likelihood that Hannah is right has not changed); hence, we refrain from attributing knowledge.

The intuitive judgments regarding the two bank cases suggest something like the following principle:

For any agent A who is deciding whether or not to draw on the information of an epistemic agent E, if A finds that E’s epistemic position is not sufficiently strong to rule out an error possibility that reasonably seems significant for A’s inquiry (i.e. given A’s practical interests or concerns), then A should judge that E does not know that *p*.

While knowledge attributions are typically used to certify informants who satisfy the objectivized standard, this does not suggest that satisfying the objectivized standard is sufficient to count as a knower. Satisfying this threshold does not guarantee that one ‘knows’ in every context.

On this picture, the epistemic standards are partly set by the practical needs and interests of a given inquirer, but they are also constrained by a much wider set of interests, those of a variety of human inquirers. The principle of objectivization is needed to set a boundary on the epistemic standards required for 'knowledge', but the standards can also rise to as high a level as needed to accommodate the concerns of an inquirer. This offers a plausible diagnosis of the intuitions involved in various high- and low-standards cases. Our need for an objectivized stance does not preclude the ability to tailor the concept of knowledge to an individual's needs or interests. Rather, the concept of knowledge objectively construed operates as the *default* 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.

If the concept of knowledge arose from our general need to flag informants reliable enough to satisfy our own needs and interests, and if our needs and interests vary depending on context, then we should expect 'know' 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. However, there are significant constraints on the variability of epistemic standards across contexts. The operative standard will often reflect the objectivized stance, but more may be required from an informant in circumstances where an inquirer's concerns are unusually pressing. Objectivization thus explains why the default standard for 'knows' is quite high without excluding the possibility that practical concerns can change the standards.

This version of contextualism posits fairly stable truth-conditions across a wide range of contexts, but it also allows the epistemic standards for 'know' to shift in order for that word to continue serving its function of flagging reliable informants. As a result, this view does not presuppose an implausibly demanding picture of epistemic scorekeeping. In any complex society there will of course be individuals with different interests, assumptions, projects, concerns, etc., and it would therefore be too demanding to expect people to keep track of all these factors when storing a knowledge report. However, my view does not tie the epistemic standards to each individual's interests so tightly. If 'knowledge' were tied only to an individual's practical interests, then the epistemic standards would change whenever that individual's interests changed. Surely this is implausible, for what it takes to 'know' does not change radically from case to case. Rather, the epistemic standards are strongly conditioned by the general needs and interests of the epistemic community to which each knowledge attributor belongs. The purpose of knowledge attributions is to identify reliable informants so that a community can pool and share information, so it makes little sense to think that 'know' would be so

closely tied to the interests of a particular individual. This allows 'know' to remain fairly stable across contexts.

It is for precisely this reason that my proposal puts considerable pressure on subject-sensitive invariantism. Subject-sensitive invariantists hold 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 contextualists claim that the truth-conditions for knowledge ascriptions are relative to the context of the *attributor* of the knowledge claim. One of the advantages of contextualism over subject-sensitive invariantism is that the attributor's context can be sensitive to the needs or interests operative in the subject's context (DeRose, 2009, ch. 7; Greco, 2008, pp. 424–5; Lewis, 1996, p. 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, p. 425).

It is mistaken to think that the practical situation of the subject of a knowledge ascription is irrelevant by the lights of contextualism. 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 which the knowledge attributor and the subject in question belong. If the purpose of knowledge attributions is to flag reliable informants so that a community can pool and share information, it makes little sense to think that 'knowledge' should be tied to the interests of the subject of an attribution rather than a more general standard.²⁶

Does my view have any advantages over insensitive invariantism? My aim in this article is not to provide a refutation of insensitive invariantism. Rather, I want to provide a principled motivation for contextualism that enables this view to escape the trans-contextual role problem, which would otherwise constitute a major problem for contextualism. I believe this task has now been accomplished. However, if knowledge attributions play anything like the role I have outlined, it might also render the invariantist account less plausible.

According to the insensitive invariantist, the truth-values for 'know' do not vary with practical facts about either the subject's context or the attributor's context. Rather, 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; the truth-conditions for knowledge sentences do *not* vary with context. As a result, non-skeptical invariantists

typically maintain that in *Low Standards* Hannah says something true, while in *High Standards* she says something false.

In order to explain our shift in epistemic judgments, invariantists often say that asserting the truth in *High Standards* would have a false implicature and thereby convey something false to the audience; for example, that Hannah'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 Hannah's evidence is not strong enough for current practical purposes, given what is at stake.

Presuming that knowledge attributions serve as a marker for reliable informants, why would it be true (but inappropriate) to say that Hannah 'knows' in *High Standards* if her evidence *isn't* strong enough for current practical purposes? Put differently, why would 'know' function in a way that renders our knowledge attributions true when the subject of the attribution fails to qualify as a good enough informant? The invariantist seems committed to saying that the truth-conditions for knowledge sentences will sometimes be at odds with the purpose of knowledge attributions. One of the main functions of knowledge attributions is to identify reliable informants, yet the insensitive invariantist claims that Hannah truly counts as 'knowing' even though she is not reliable enough given the circumstances. In contrast, my version of contextualism predicts that it is false to say of Hannah that she knows in *High Standards* precisely because she is not reliable enough. This tells a congruous story about how our sense of the purpose of knowledge attributions will condition our knowledge-ascribing/denying practices.

The invariantist might say that most of the relevant reports are true, but sometimes their serving the social function in question involves people saying things that are literally false. However, the invariantist would then owe us a story about why, if the point of knowledge attributions is to identify reliable informants, it would be true to say that an informant who is not reliable enough (in the context) does in fact 'know'. I do not deny that such a story is possible, but it would fit less well with the functional role of 'know' that I have described.

These considerations show how Craig's idea of objectivization can supplement contextualism. The default use of 'knows' is the one that a competent user of that term will expect other competent users to employ, unless an inquirer's concerns are unusually pressing. Even if the standards for 'know' vary across contexts, there is a stable, central, default use of that term on 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 objectivized use of 'know', which is perhaps what contextualists refer to when they speak of a so-called 'ordinary' sense of that word, reflects our epistemic community's intersubjective needs and interests.

5.4. SOME REMAINING WORRIES

Any version of contextualism will entail that the epistemic standard used in one context will not always be the operative standard in another context. Thus, no matter what uniform value is selected as the 'default', there are still going to be cases for which this standard will not be correct. This is especially obvious for two kinds of cases: skeptical cases and cases in which the stakes are unusually high. The increased stringency triggered by skeptical cases might raise the epistemic standards to a point much higher than the objectivized stance, thereby shifting the semantic content of 'know'. Furthermore, people's concerns can sometimes be quite pressing, as DeRose's bank cases illustrate nicely, which shows that the default standard is not suitable for all contexts. Does this spell doom for contextualism?

While skeptical pressure sometimes leads us to withdraw or weaken our knowledge claims, these fleeting moments of doubt should not be overemphasized. Our non-skeptical intuitions are both pervasive and persistent. Furthermore, the inability to accurately store and transmit information from skeptical contexts to non-skeptical contexts seems precisely what any reasonable theory of knowledge attributions should predict. Consider an analogous case in which we lead competent speakers to question their everyday ascriptions of 'flatness' by pointing out bumps that we ordinarily do not pay attention to. People might deny that almost any surface is 'flat' because they are led, via certain conversational moves, to believe that in order to be 'flat' a surface must be completely, or absolutely, flat, meaning having no bumps whatsoever (Unger, 1975; Cohen, 1999). Presuming that 'flat' is a context sensitive term (which is widely acknowledged), it seems correct to say that 'flatness' reports will encounter difficulties if we try to retrieve and transmit information from contexts in which the standards are unusually high to those in which they are more ordinary. Similarly, it would be peculiar for a theory of knowledge attributions to predict that we can smoothly and accurately transmit information encoded in knowledge reports from skeptical contexts to ordinary contexts. Surely it is a *virtue* of my account that it predicts no such thing.

What about cases in which the stakes are remarkably high? Again I want to stress that we should not overemphasize the number of contexts in which people are unusually demanding. Furthermore, nothing suggests that we cannot adduce extra-semantic factors in order to handle these different epistemic standards. In particular, we can rely on certain notions and locutions in understanding and making use of knowledge reports that come to us from contexts in which the standards are unusually high. For instance, Sarah might say 'Hannah denied that she knew the bank would be open, but she was being paranoid' or 'she was justifiably cautious given that was at stake'. There are conversational moves to indicate that the

standard used was a departure from the default.²⁷ Moreover, we might unconsciously adjust to contextual variation when preserving and transmitting information, even if we are not consciously aware of such variation. This is not a surprising or controversial view since much of native speakers' competence with natural language is 'not open to their view' (Williamson, 2005a, p. 102). That these cases are infrequent also significantly lessens the burden placed on memory when relying on such knowledge reports outside of the context in which it emerged.

Finally, one might wonder why the 'default' standard is set lower than the standard imposed by the skeptic. If objectivization is supposed to give us 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, why not heighten the standard to a level that would satisfy *everyone's* concerns, including radically skeptical 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 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 they can base their 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.²⁸

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 they are 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. We have no use for such a concept.

6. *Concluding remarks*

By reflecting on the role that knowledge attributions play in flagging reliable informants, we are led to a version of contextualism that escapes the trans-contextual role problem. The force of the trans-contextual role problem presumes that the epistemic standards will change easily and frequently; however, I have curtailed the excessive variability in standards

that has been associated with familiar forms of contextualism by connecting these standards with the interests of a community of inquirers.

This account also 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 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.

Any reasonable version of contextualism will put limits on the effect that context-shifting mechanisms may have on the truth-conditions of knowledge ascriptions. This is true not just in the case of ‘know’, but for any context sensitive term. Suppose that A says ‘France is hexagonal’ and B accepts this claim. It is plausible, via Lewisian rules of accommodation, that A’s claim is true. In contrast, suppose that A says ‘France is boot-shaped’ and B accepts this claim. Presumably, this claim does not count as true. There must be constraints on how the course of a conversation may affect the truth-conditions of certain utterances. If contextualism is the correct view about the semantics of ‘knows’, there must be constraints on when subjects can truly be said to know, just as there are constraints on what shape France can truly be said to be. We cannot simply employ any standards we please and still be speaking a common language.²⁹

Philosophy Department
Fordham University

NOTES

¹ I put quotation marks around ‘know’, ‘knowledge’, and ‘knowing’ to indicate semantic ascent. However, my discussion will often be stated in the object language for ease of expression.

² I assume that words express concepts and that concepts are the meanings of words.

³ I leave aside truth relativism about knowledge. For this view, see John MacFarlane (2005).

⁴ In fact, I am inclined to think that the social dimension of knowledge can help explain why there are plausible constraints on how salience and practical stakes can affect the epistemic standards, but I will not argue for this view here.

⁵ In addition to DeRose’s bank cases, other pairs of cases include Stewart Cohen’s airport cases (1999), Fred Dretske’s zebra cases (1970), and Jeremy Fantl and Matt McGrath’s train cases (2009). All of these examples have been taken to illustrate flexibility in our willingness to ascribe knowledge. This suggests that, in some sense, whether someone counts as knowing depends on context.

⁶ Likewise, we are good at figuring out what is meant by context-sensitive terms like ‘rich’ and ‘tall’. They are typically easily replaced by some fairly obvious and easily articulable rendering of the relevant information in invariantist terms – e.g. someone’s income or their actual height in feet and inches (Rysiew, 2012, p. 289).

⁷ Contextualists must posit semantic blindness because speakers do not usually recognize the context-sensitivity of ‘know’ (e.g., when faced with skeptical challenges). Hawthorne coined the term ‘semantic blindness’ in his book *Knowledge and Lotteries* (2004, p. 107), but this issue was first discussed by Stephen Schiffer (1996).

⁸ Williamson acknowledges that you might be able to construct a new sentence ‘P₂’ that as uttered in a new context C₂ expresses the very truth that ‘P’ expresses as uttered in C₁; however, this places a heavy burden on memory. This proposal is dubious because ‘[m]ost people cannot remember where they got much of the information on which they rely’ (Williamson, 2005a, p. 100). I will return to this idea in Section 5.4.

⁹ Kent Bach makes a similar objection in Bach, 2005, pp. 60–1.

¹⁰ This section is greatly indebted to Emil Moeller. During the 2012 Epistemology Network Meeting, he gave a talk in which he argued that the epistemic scorekeeping required to retrieve epistemically useful information from knowledge attributions is not a consequence of accepting contextualism but rather a consequence of certain empirical claims about our ordinary knowledge attributing practices. This section summarizes the argument Moeller gives in more detail in his paper, ‘Consuming Knowledge Claims Across Contexts’ (MS).

¹¹ Moeller makes this point in the paper mentioned in the previous endnote.

¹² I remain neutral about whether subject-sensitive invariantism or insensitive invariantism can provide a compelling reply to the trans-contextual role problem. However, I will discuss certain advantages that contextualism has over both of these views in Section 5.3.

¹³ One might worry that since any version of contextualism is committed to contextual variation, this view will never fully solve the trans-contextual role problem. I discuss this worry in Sections 5.3 and especially 5.4.

¹⁴ I will speak as if the standards for ‘know’ consist in a set of *relevant alternatives* that must be eliminated if the subject is to have ‘knowledge’. See Cohen, 1988 and Lewis, 1996 for this approach. However, this is inessential to my account and I invite others to substitute their own preferred framework.

¹⁵ ‘Mentioning’ is taken to include not just stating aloud to others but also a conscious awareness in silent thought. I shall focus only on cases in which a possibility is mentioned aloud, since this view seems more plausible than the idea that merely thinking of different alternatives affects the set of relevant alternatives in a conversational context. If the former view is shown to be implausible, I assume the later view is thereby shown to be implausible as well.

¹⁶ Although it is debatable whether Lewis fully endorses this view (1996, p. 560).

¹⁷ The mere mention view is rejected by DeRose (1995) and Cohen (1999).

¹⁸ Michael Blome-Tillmann (2009) defends a similar view, although he uses the framework of ‘conversational common ground’ made familiar by Robert Stalnaker (2002).

¹⁹ My own sympathies are with the normative view that the relevant alternatives are those that people in the context *ought* to consider, and what they ought to consider partly depends on their practical interests. I will not explore this topic here. Robin McKenna defends this view in his 2013. Blome-Tillmann also defends a normative requirement that he calls the Rule of Presumed Interests: if it is presupposed in context C that the subject S, given her presumed practical interests, ought not to ignore w, then w is not properly ignored in C. See Blome-Tillmann, 2012, p. 113.

²⁰ One might question the connection between the point that the epistemic standards are not terribly low and the claim that they are fairly high. I will argue for the latter point in Section 5.3.

²¹ A statistically significant difference in responses between low- and high-stakes cases was not found in the following three studies: Adam Feltz and Chris Zarpentine (2010); Joshua May *et al.* (2010); and Wesley Buckwalter (2010). However, a number of compelling objections have been raised against these studies. See DeRose, 2011; Ángel Pinillos, 2012; and Sripada and Stanley, 2012.

²² I leave open the possibility that knowledge attributions serve other roles. My view only assumes 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. One might worry that, although the role of identifying informants leaves room for contextual variation in the truth-conditions of knowledge attributions, we might end up with no room for such variation once we add all the other roles. I see no reason to think that other roles will provide additional constraints that ultimately rule out contextual variation, nor has any published argument been made to this effect.

²³ What about first-personal cases in which someone knows something but does not function as an informant for anybody? Does the fact that identifying good informants becomes socially-directed leave mysterious why we have cases of knowledge that do not involve any interpersonal relations? These cases are not problematic because there are reasons to apply the concept of an informant to oneself in the first-person case. In particular, it is important that we have a practice whereby people can declare themselves to be qualified informants, since they themselves will often be the only person in a position to tell whether they are qualified or not. See Edward Craig, 1990, p. 63. In addition, reliable information might be important not just for other inquirers, but also potentially for oneself in the future. This answers what Hallvard Lillehammer calls the ‘Robinson Crusoe’ problem; i.e. the idea that we might want to say that somebody knows things even if that person is not a member of any community.

²⁴ Two people who take Craig’s account to support insensitive invariantism are Chris Kelp (2011) and Martin Kusch (2011).

²⁵ I develop this view in my (2013) paper ‘The Practical Origins of Epistemic Contextualism’.

²⁶ David Henderson raises a similar objection to SSI in Henderson, 2009, p. 124.

²⁷ It has been suggested to me that an invariantist might not deny that there is a plausible contextualist story to tell about how we consumers of knowledge reports might correctly interpret and make use of those reports, but the invariantist might argue that to the extent that there is such a story it will seriously undermine the case for contextualism. I see no reason to accept this point, especially since my view provides a principled explanation for why ‘know’ would be context-sensitive in the way I have described. But whether or not the invariantist finds this story about contextualism plausible, the trans-contextual role problem is not a problem for contextualism. (And, if my argument in Section 4 is correct, it is an unresolved problem for invariantism.)

²⁸ This section draws on Henderson, 2009 and 2011. While there is much in Henderson’s work with which I agree, we develop our views in different ways. A detailed discussion of his view will take us too far afield, but I will mention two essential differences. First, Henderson’s account rests on a distinction between two broad types of communities, ‘general source communities’ and ‘applied communities’ (2009, p. 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 in principle way to distinguish them. In contrast, my own view shows that an intersubjectively

determined epistemic standard (the ‘objectivized’ standard) sets reasonable constraints on the context-variability of ‘knows’ without positing problematic groups or communities. A second advantage of my view over Henderson’s is that he allows the epistemic standards for ‘know’ to reflect the interests of ‘applied source communities’ that are concerned with some specific project. However, if there are only a few members of such a community and if they do not demand or require very reliable information (say, because nothing important is at stake), then it seems that the standards for ‘know’ will be set *too low*. The fact that a small group of people does not care about reliable information is not ‘knowledge’-making. The thresholds relevant to ‘knowledge’ do not fluctuate depending on whether certain small communities (that may be formed on an *ad hoc* basis) require a higher or lower level of reliability. If Henderson were right, then ‘know’ probably could not serve the trans-contextual role that it clearly does play.

²⁹ For helpful advice I am grateful to Hallvard Lillehammer, Chris Cowie, the participants of the 2013 Midsummer Philosophy Workshop and the Roles of Knowledge Workshop, and especially to Robin McKenna and an anonymous reviewer for this journal. This paper was written while I was funded by the Social Sciences and Humanities Research Council of Canada.

REFERENCES

- Bach, K. (2005). ‘The Emperor’s New “Knows”,’ in G. Preyer and G. Peter (eds) *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Oxford University Press.
- Blome-Tillmann, M. (2009). ‘Knowledge and Presuppositions,’ *Mind* 118(470), pp. 241–294.
- Blome-Tillmann, M. (2012). ‘Presuppositional Epistemic Contextualism and the Problem of Known Propositions,’ in J. Brown and M. Gerken (eds) *Knowledge Ascriptions*. Oxford: Oxford University Press.
- Brown, J. (2006). ‘Contextualism and Warranted Assertability Maneuvers,’ *Philosophical Studies* 130(3), pp. 407–435.
- Buckwalter, W. (2010). ‘Knowledge Isn’t Closed on Saturday: A Study in Ordinary Language,’ *Review of Philosophy and Psychology* 1(3), pp. 395–406.
- Cohen, S. (1988). ‘How to Be a Fallibilist,’ *Philosophical Perspectives* 2, pp. 91–123.
- Cohen, S. (1999). ‘Contextualism, Skepticism, and the Structure of Reasons,’ *Philosophical Perspectives* 13, pp. 57–89.
- Craig, E. (1990). *Knowledge and the State of Nature*. Oxford: Oxford University Press.
- Dancy, J. (1992). ‘Review of *Knowledge and the State of Nature*,’ *The Philosophical Quarterly* 42(186), pp. 393–395.
- DeRose, K. (1992). ‘Contextualism and Knowledge Attributions,’ *Philosophy and Phenomenological Research* 52(4), pp. 913–929.
- DeRose, K. (1995). ‘Solving the Skeptical Problem,’ *Philosophical Review* 104(1), pp. 1–52.
- DeRose, K. (2009). *The Case of Contextualism*. Oxford: Oxford University Press.
- DeRose, K. (2011). ‘Contextualism, Contrastivism, and X-Phi Surveys,’ *Philosophical Studies* 156(1), pp. 81–110.
- Dretske, F. (1970). ‘Epistemic Operators,’ *Journal of Philosophy* 67(24), pp. 1007–1023.
- Fantl, J. and McGrath, M. (2009). *Knowledge in an Uncertain World*. 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), pp. 205–221.
- Feltz, A. and Zarpentine, C. (2010). ‘Do You Know More When it Matters Less?’ *Philosophical Psychology* 23(5), pp. 683–706.

- Fricker, M. (2010). 'Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time,' in A. Haddock, A. Millar, and D. Pritchard (eds) *Social Epistemology*. Oxford: Oxford University Press, pp. 51–68.
- Greco, J. (2008). 'What's Wrong With Contextualism?' *The Philosophical Quarterly* 58(232), pp. 416–436.
- Hannon, M. (2013). 'The Practical Origins of Epistemic Contextualism,' *Erkenntnis* 78(4), pp. 899–919.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Henderson, D. (2009). 'Motivated Contextualism,' *Philosophical Studies* 142(1), pp. 119–131.
- Henderson, D. (2011). 'Gate-Keeping Contextualism,' *Episteme* 8(1), pp. 83–98.
- Kelp, C. (2011). 'What's the Point of "Knowledge" Anyway?' *Episteme* 8, pp. 53–66.
- Kusch, M. (2011). 'Knowledge and Certainties in the Epistemic State of Nature,' *Episteme* 8, pp. 6–23.
- Lewis, D. (1996). 'Elusive Knowledge,' *Australasian Journal of Philosophy* 74(4), pp. 594–567.
- MacFarlane, J. (2005). 'The Assessment Sensitivity of Knowledge Attributions,' *Oxford Studies in Epistemology* 1, pp. 197–233.
- May, J., Sinnott-Armstrong, W., Hull, J. and Zimmerman, A. (2010). 'Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study,' *Review of Philosophy and Psychology* 1(2), pp. 265–273.
- McKenna, R. (2013). 'Epistemic Contextualism: A Normative Approach,' *Pacific Philosophical Quarterly* 94(1), pp. 101–123.
- Pinillos, Á. (2012). 'Knowledge, Experiments and Practical Interests,' in J. Brown and M. Gerken (eds) *Knowledge Ascriptions*. Oxford: Oxford University Press.
- Rysiew, P. (2001). 'The Context-Sensitivity of Knowledge Attributions,' *Noûs* 35(4), pp. 477–514.
- Rysiew, P. (2012). 'Epistemic Scorekeeping,' in J. Brown and M. Gerken (eds) *Knowledge Ascriptions*. Oxford: Oxford University Press.
- Schiffer, S. (1996). 'Contextualist Solutions to Scepticism,' *Proceedings of the Aristotelian Society* 96, pp. 317–333.
- Stalnaker, R. (2002). 'Common Ground,' *Linguistics and Philosophy* 25(5–6), pp. 701–721.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Sripada, S. and Stanley, J. (2012). 'Empirical Tests and Interest-Relative Invariantism,' *Episteme* 9(1), pp. 3–26.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2005a). 'Knowledge, Context, and the Agent's Point of View,' in G. Preyer and G. Peter (eds) *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Oxford University Press.
- Williamson, T. (2005b). 'Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge,' *Philosophical Quarterly* 55(219), pp. 213–235.