

# Motivated Contextualism

David Henderson

Draft 12/8/2007 3:54 PM

Published *Philosophical Studies* (2009) 142:119-131.

## I. Introduction

Judgments about scenarios provide important data for philosophical reflection. For one who has achieved a competence with a concept, such judgments take fallible soundings of the extension of the relevant concept. But, this is not all that one has to do on in philosophical reflection. One typically has some sense for the point and purpose of the concept—why one would have developed and deployed the concept, the kind of sociolinguistic work the concept does for its users (Henderson and Horgan 2001). The philosophical literature on knowledge is certainly rich with judgment regarding scenarios. However, with a few exceptions, such as Craig (1990), few discussions of *knowledge* have drawn significantly on a sense for the sociolinguistic point and purpose of the concept.

I here evaluate contextualism and invariantism in terms of what I think should be an uncontroversial point regarding the point of the concept of *knowledge*: it is used to certify epistemic agents as good sources for an understood audience. Attributions of (and denials of) knowledge are used in a kind of epistemic gate keeping for communities with which the attributor and interlocutors are associated. (See also Craig, 1990.) To say that *S* knows about some matter is to say that *S* is epistemically positioned with respect to this matter so as to be a good source or storehouse of beliefs on the matter. To say that *S* knows about guanaco physiology indicates that *S* is a good source of true beliefs about guanaco physiology. To say that *S* knows that a major earthquake is *n* times more likely to hit Memphis than Lincoln indicates that *S* is so positioned with respect to that proposition that one can take it from *S*. This may not be the whole story regarding the point and purpose of the concept of *knowledge*, but it is central to the story. The claim here is not that knowledge-attributions *say*, or are semantically equivalent to the claim that, the agent is a good source. Rather, the idea is that this is central to the use of the concept, and that the social epistemic need for gate keeping allows us to appreciate why one would come to develop the concept of *knowledge* to pick out a kind of good

epistemic standing with respect to propositions—one that qualified agents to serve as sources.

Beginning with this simple observation regarding the point and purpose of the concept of *knowledge*, and reflecting on the epistemic gate keeping needs of various situated groups or communities, we can uncover rich implications.<sup>1</sup> In this paper, I highlight three implications and an open question.

One implication is that this central gate keeping point to the concept of *knowledge* gives one reason to prefer contextualism over various forms of sensitive invariantism. Contextualists treat attributions of knowledge as understandable in terms of the contextual concerns of attributor and interlocutors. In contrast, sensitive invariantists understand attributions of knowledge in terms of the interests and projects of the agent (to whom knowledge is ascribed). It makes a good deal of sense that gate keeping would be attuned to the interests of those of the community for whom the gate keeping is done (the attributor and interlocutors then being representative of, or responsive to, this community). It makes a good deal less sense that the gate keeping would be attuned to the interests of an epistemic agent who might have interests and projects rather at odds with the community for whom gate keeping is managed in the talk of knowledge.

Second, I sketch the outlines of a contextualist view—*gate-keeping contextualism*—that derives from combining the gate-keeping point of the concept of knowledge with a crude taxonomy of groups for which epistemic gate might be kept. I begin with some simple reflections on a spectrum groups and their needs for information and sources. At one end of the spectrum are cases in which the attributor and interlocutors may be engaged in gate keeping for a group that is focused on some practical project. (The group may be designing a machine or system, trying to install a program, trying to repair an engine or restore an ecosystem, trying to determine criminal responsibility, or trying to reach a lasting political compromise on a matter of dispute.) The members of such a group are concerned with sources of actionable information on which to proceed in their project. Call this *gate keeping for an applied community*. At the other end of the spectrum, there are cases in which the attributor and interlocutors may be *gate keeping for a source community*—a community devoted to developing a body of results on which folk in various other communities might confidently draw. Gate keeping

for such a community, one is including or excluding folk from a select community of experts or authorities, and doing so with a view to putting together a community up to producing actionable information for an indeterminate range of interested groups.

Some recent contextualist work highlights how interests or stakes shared by parties to a knowledge-attribution seem to condition contexts and the conditions operative in context. This idea is respected in gate-keeping contextualism. Particularly when gate keeping for applied communities, there would be significant variation in contextual requirements.

While contextualist themes seem most plausible in connection with applied communities, significant limits on contextualist themes emerge in connection with source communities that are *not* closely associated with some delimited range of practical projects. Think of the vast range of practical projects that have recourse to the results of general molecular biology—or some other general discipline. Such disciplines provide a body of results on which people engaged in an indeterminate range of practical projects might draw. (Of course, such disciplines are focused on subject matters that may have been associated with certain practical interests, but the range of practical projects that may come to draw on a discipline (think of particle physics, for example, or molecular biology) is really quite diverse. When *gate keeping for a general source community*, the concern is with *general purpose actionable information*. In contexts where the operative interests have to do with *general purpose information*, no concrete limited purposes are likely to provide a simple decisive understanding of just what counts as actionable information. Accordingly, while thinking about gate keeping for applied communities may encourage contextualist themes, thinking about gate keeping for general source communities mutes some of these tendencies. Here, one gets results that look rather like those with which insensitive invariantists such as Williamson (2000, 2005a) would be comfortable. This unexpected vindication of insensitive invariantist themes within gate keeping contextualism is the third central point pursued here.

I am convinced that something along the lines of what I here term gate-keeping contextualism is the most promising way in which to understand contextualism about knowledge. However, I do not have a settled view about just how to best combine the three lines of thought developed here. They might find happy development into a

thorough-going gate-keeping contextualism in which certain insensitive invariantist judgments and themes find limited vindication while being contained within a generally contextualist picture. I am attracted to this possibility, and highlight it in this paper. Still, in the final section I will also mention reasons for thinking that gate-keeping contextualism is ultimately an unstable position, containing the seeds of its own demise.

## *II. Stage Setting: the Contested Positions<sup>2</sup>*

### *a. Contextualism*

Contextualists hold that the semantics or truth conditions of knowledge attributions are keyed to the context of the attributor, so that a given sentence attributing knowledge to a given agent at a time can have different truth-conditions in the different contexts.<sup>3</sup> Holding fixed the agent, *S*, the proposition to be known, *p*, and the time, *t*, the attribution, “*S* knows that *p* at *t*,” can yet express different propositions when uttered in different attributor contexts. Aspects of the *context of attribution* condition what it takes to qualify as knowing that *p*. To know that *p*, the agent needs to be epistemically well positioned enough to be able to discriminate that *p* in contrast to the contextually relevant alternatives. Sometimes, the point is put by saying that the agent needs to be able to “rule out” the contextually relevant alternatives to *p*.<sup>4</sup>

What it takes to be epistemically well situated enough to discriminate one proposition’s obtaining from certain alternative possibilities obtaining, can be understood by importing one’s favored understanding of what processes and informational states make for objective epistemic justification. A reliabilist variant would require that by virtue of reliable processes the agent be able to reliably discriminate that *p* obtains in contrast to the alternatives. An evidentialist variant would require the agent to be situated so as to have a telling evidential basis for thinking that *p* obtains in contrast to the relevant alternatives.

Contexts have been variously understood. Commonly they have been understood as conversational settings in which moves by participants would raise or lower the standards for knowing (see, Lewis [1979, 1996] and DeRose [1995].) Standards are raised as moves in the conversation make salient various alternative possibilities. Particularly in Lewis’s (1996) discussion, one could easily get the impression that contexts are strangely hostage

to whatever “moves” are made by conversational participants in a “language game”—so that the mere mention of skeptical scenarios would raise standards, even when this would frustrate the practical interests of all participants. (Mentioning the epistemic possibility of an evil demon in the course of some practical project or some inquiry would be the cognitive equivalent of a suicide bombing.) With their bank and plane cases, DeRose (1992) and Cohen (1999) make it clear that the practical interests of the attributor and interlocutors significantly condition the contextually relevant alternatives. This suggestion seems to me a step forward. Whatever exactly is a context of a knowledge attribution, it is associated with the attribution, the attributor and the attributor’s interlocutors, and, commonly, with a wider “audience” represented by those parties—moreover, it is associated with their salient concerns and interests. The agent who is attributed or denied knowledge may have different concerns and interests—and these will only affect the context and the contextually relevant alternatives only to the extent that awareness of them affects the concerns and interests salient to the attributor and interlocutors.

Consider DeRose’s bank cases. In the first, I and my wife are driving by the bank with paychecks in hand on Friday after work. We generally like to deposit our checks fairly quickly, but not much hangs on this. Seeing long lines, I suggest depositing the check Saturday. My wife notes that some banks are closed on Saturdays. I respond that I know that our bank is open on Saturdays, having been there on a recent Saturday. (The bank is indeed open Saturdays.) Reasonably, she accepts my knowledge self-attribution. In the second, something significant hangs on getting our paychecks in the bank before Monday morning—perhaps a large and important payment will be posted against our account Monday. My wife then responds to my claim to know by mentioning a further possibility: that banks do change their hours, and our’s may have done so. (It has not.) I then reasonably retract my claim to know. It is natural to think that, in the first case, with the lower practical stakes there confronting my wife and I, fewer possibilities inconsistent with the bank being open Saturday are relevant. It is natural to think that in order to there count as knowing I would need to be in an epistemic position to “eliminate” only the most obvious and salient alternatives to the proposition in question. In contrast, it is natural to think that in the second case, with the higher practical stakes confronting my

wife and I, the possibility raised by my wife is relevant—so that to know that the bank is open Saturday, I would need to be in a position to “rule out” that alternative. According to many contextualists, the higher practical interests we face—here, I as attributor and my wife as interlocutor—makes a wider range of possibilities relevant, so that more is required in this conversational context to count as knowing that the bank will be open Saturday.

In the above bank cases, my wife and I are parties to a conversation in which we share stakes (either high or low). Further, the agent to whom knowledge is attributed or denied is a party to that same conversation—and shares those stakes—me. An important kind of case, differentiating contextualism from sensitive invariantism, is that in which the agent who is a candidate for knowing faces different stakes from those conditioning the context. So, imagine that my friends, Violet and her partner Scarlett, is trying to decide whether to stop at the bank to deposit her paycheck on her way home—and that they expects a large charge to be posted against her account sometime (perhaps early) Monday. Scarlett, calls me and asks whether the bank is open Saturday. I again say that I know that it is, for the reasons indicated. Scarlett relates the information to Violet. Suppose Violet then mentions the possibility that the bank may have changed its business hours, and insists that I do not know that the bank is open Saturday. According to a plausible contextualist account, when speaking to my wife in the low stakes case, I was correct to say that I knew that the bank would be open,<sup>5</sup> and Violet was also correct when she denied that I knew that the bank would be open Saturday. In the first two bank cases discussed, both the attributor and agent faced the same stakes—either they were both in a low stakes case, or they were both in a high stakes case. But, in this third case, the agent (myself) faces low stakes while the attributor (Violet) and interlocutor (Scarlett) face high stakes. Most people judge that Violet is correct to say that I do not know that the bank will be open. Contextualism readily accommodates the judgments made in these cases, holding that Violet speaks the truth. Such cases pose more of a challenge to sensitive invariantists.<sup>6</sup>

*b. Sensitive Invariantism*

Invariantists hold that the semantics of knowledge attributions are not variable across attributor contexts. Holding fixed the agent, *S*, the proposition to be known, *p*, and

the time,  $t$ , the attribution, “ $S$  knows that  $p$  at  $t$ ,” expresses a single invariant proposition; the semantic conditions for its truth do not vary when uttered by different attributors in different conversational contexts. Still, sensitive invariantists (such as Hawthorne [2004] and Stanley [2005]) hold that what is required for an agent to know a proposition can vary—as *facts about the agent,  $S$* , at a time,  $t$ , may determine just how  $S$  must be situated with respect to  $p$  and various not- $p$  alternatives in order to count as knowing that  $p$  at  $t$ . The central thing to notice is this: for the contextualist, knowledge attributions are keyed decisively to the attributor’s context (to features of the exchange between the attributor and the attributor’s interlocutors), while for the sensitive invariantist, knowledge attributions are keyed to facts about the agent at the time in question. The pivotal features of the agent are the agent’s interests—the stakes faced by the agent. Thus, the contrast with contextualism can be put in terms of whose interests condition what is required for an agent to satisfy a knowledge attribution: for the contextualist, knowledge attributions are keyed to the attributor’s and the audience’s interests, while for the sensitive invariantist, knowledge attributions are keyed to the agent’s interests.

Sensitive invariantists insist that knowing does not turn merely on epistemic matters—it does not turn merely on how well an agent is epistemically situated with respect to a proposition at a time. The agent’s own stakes in the matter in question supposedly determines just how well situated the agent must be at that time to count as knowing. For any given agent,  $S$ , and proposition,  $p$ , that  $S$  entertains, if  $S$ ’s stakes in connection with  $p$  vary across time, then sometimes  $S$  must be better situated with respect to  $p$  to count as knowing that  $p$  than  $S$  must be at other times to count as knowing that  $p$ . An agent in a relatively low stakes situation counts as knowing that  $p$  if the agent is epistemically situated so as to be able to discriminate  $p$ ’s obtaining from the obtaining of any one of a narrow range of alternative possibilities that constitute the most obvious and likely or pressing alternatives to  $p$ . In contrast, to count as knowing that  $p$ , the agent in a high stakes situation must be epistemically situated so as to be able to discriminate  $p$ ’s obtaining from a significantly wider set of alternatives.

When the agent and the attributor (at the relevant times) have highly similar stakes in the matter in question, when both are in a high stakes or both are in a low stakes situation, then sensitive invariantism and contextualism yields parallel verdicts regarding

the truth of knowledge ascriptions or denials.<sup>7</sup> However, verdicts regarding the truth of knowledge attributions diverge with respect to cases in which the agent and the attributor (and attributor's audience) face markedly different stakes in connection with the matter in question. This obtained in the third bank case. There, I (the agent) am in a low stakes situation with respect to whether the bank is open on Saturdays. For the sensitive invariantist, I know that the bank is open on Saturdays. At the same time, Violet and Scarlett are in a high stakes situation on that matter. For the sensitive invariantist, this may make a difference for whether *they* can know on the basis of my testimony that the bank is open (apparently they cannot). But, Violet speaks falsely when she says that I do not know that the bank is open. So the claim that David Henderson knows (on Friday afternoon) that the bank is open Saturdays is semantically invariant—whether it is expressed by an attributor in a high stakes situation or one in a low stakes situation. For the sensitive invariantist, there is no attributor context effect.

*c. Insensitive Invariantism*

Insensitive invariantism is the dominant traditional account of knowledge. An *invariantist* holds that to attribute to a given agent, *S*, at a set time, *t*, knowledge of a given proposition, *p*, expresses an invariant claim—rather than expressing different claim in differing attributor contexts. Both sensitive and insensitive invariantists hold this much. The *insensitive* invariantists holds, additionally, what counts as being in a sufficiently good epistemic position to know some proposition does not vary with—is not sensitive to—facts about the agent's interests at the time in question. What is consistently required for satisfying the concept of *knowledge* might be understood as exceedingly demanding (skeptical insensitive invariantism) or it might be understood in some moderate fashion satisfiable by flesh and blood humans (nonskeptical insensitive invariantism). In either case, for any agent, whatever that agent's practical interests or concerns, and for any proposition, there is some good epistemic position in which an agent must stand with respect to any such proposition in order for that agent to count as knowing it.

III. The Central Misgiving Concerning Sensitive Invariantism

That practical interests condition what constitutes knowledge has become common coin between sensitive invariantists, and contextualists.<sup>8</sup> The real issue separating them is

*whose* practical interests play a role in determining the conditions for knowledge, and how. The idea that the (or a) central point and purpose of the concept of *knowledge* is its role in a kind of epistemic gate keeping has real bite in connection with this central disagreement. In attributing knowledge to epistemic agents, one is, in effect, certifying those agents as sources of information on the matter in question or in the domain indicated. This makes it natural to think that *the interests of those in the community for which sources or members are thus regulated* should condition what is required to know (insofar as any interests condition what is required). To certify an agent as a source on some matter is to certify the agent to an audience—in effect to a community. The community in question may be a community composed of just the attributor and one interlocutor (as in the simple bank cases) or it may be an extended practical or epistemic community with which the attributor and interlocutors are understood to be contextually related. The community may be wholly present at the attribution (again as in the simple bank cases) or contextually understood (as when one understands one's discussion to be evaluating the agent's work as a fitting contribution to the results of some wide source community of scientists in some general subject domain). In each case, in attributing knowledge to the agent, one certifies the agent's belief or beliefs to an understood community with whom both the attributor and the immediate interlocutors are contextually related. In receiving such certification, members of the relevant community understandably are not concerned with whether the agent is epistemically positioned so as to render information that is fitting *to that agent's own interests* in the matter—which may be minor in comparison to their own. Rather, they are interested in whether the agent is epistemically positioned to render information that *is fitting to their understood community's own interests*. The interests that they share as members of the relevant community may make for higher stakes than the agent's. Similarly, in providing such certification, it seems fitting that the attributor be sensitive to the audience or community's practical interests when certifying the source. This general line of thought suggests a distinctive social-epistemological species of contextualism.

Before I sketch the gate keeping form of contextualism that seems indicated, I want to state in the most general fashion the misgiving concerning sensitive invariantism that derives from the recognition of the central gate keeping point or purpose of the concept

of *knowledge*. The misgiving is this: while it makes a good deal of sense that gate keeping would be attuned to the interests of those of the community for whom the gate keeping is done, and while it thus makes much sense that the attributions of knowledge then be keyed to the interests of that community, it makes a good deal less sense that gate keeping, and the attributions of knowledge by which it is managed, would be keyed to the interests of the subject of the attribution.<sup>9</sup> It will be common for the epistemic agent in question to have interests and projects rather different from those of the community for whom gate keeping is managed in the attributions of knowledge. Thus, in identifying the interests pivotally conditioning attributions of knowledge as those of the knowing (or opining) agent, sensitive invariantists must think of the semantics of the concept of *knowledge* as turning on interests that are curiously disconnected from the interests that would seem significant given the central point and purpose of that concept. It is not plausible, that folk would have developed an evaluative concept so that its very semantics were so disassociated with its point and purpose.

Consider a variant on the venerable question: why should one care about whether or not the agent knows?<sup>10</sup> One wonders why one should care about the agent knowing, if knowing is what the sensitive invariantist says that it is. To consider the question, imagine that one is a party to a discussion in which *S*, the potential source of information is truly said to know that *p*. According to the sensitive invariantist this is to say that *S* is epistemically so positioned as to be able to discriminate whether *p* obtained rather than any member of a contrast class of alternatives—where this contrast class is fitting in view of *S*'s own interests. Why care about this? Why care about whether  $K_{si}[S,p,t]$ ?

Admittedly, one special case poses no problem. One can understand why, at *t*, one with one's practical interests would care about whether *oneself* satisfies  $K_{si}[_ ,p,t]$ —since that involves being sufficiently epistemically well placed to discriminate between the alternatives that matter to one at that time, and to discriminate in the degree fitting to one's own interests. More generally, for any agent, *S\**, time *t*, and proposition *p*, such that *S\** at *t* faces *the same stakes as does oneself*, one should care whether  $K_{si}[S^*,p,t]$ —since this amounts to *S\** being epistemically situated well enough at *t* to discriminate between the alternatives that are fitting given *one's own interests* at *t*. Doing so makes the other *S\** a good source.

It is more tenuous why one should generally care about whether someone,  $S$ , other than oneself satisfies  $K_{si}[\_,p,t]$ . Of course, one might be interested in  $S$  as a source of information about whether or not  $p$ , and  $p$  might have consequences for one's practical projects. Suppose that  $S$  has just testified at  $t$  that  $p$ . Does one then want  $K_{si}[S,p,t]$ ? Intuitively, one should care whether or not  $S$  knows that  $p$  at  $t$ . But, should one care about  $K_{si}[S,p,t]$ ? We have already determined that one should care, when  $S$  and oneself share the same interests in whether  $p$ . This leaves two classes of cases to consider: those in which  $S$  faces lesser stakes than oneself, and those in which  $S$  faces greater stakes.

Of course, because one cares about whether or not it is true that  $p$ , and because knowing that  $p$  is factive, one should care in some sense—one should care about whether some necessary condition for  $S$  knowing that  $p$  obtains. But, this does not entail that one cares about whether the quality of  $S$ 's epistemic situation with respect to  $p$  is as is required for knowing on the sensitive invariantist account. So let us focus on this. Why should one care about whether  $S$  is epistemically situated so as to be able discriminate between  $p$ 's obtaining versus the obtaining of those alternatives that are salient or relevant given  $S$ 's stakes.

*If  $S$  at  $t$  is in a higher stakes situation with respect to  $p$  than is oneself, one need not really be concerned with whether  $K_{si}[S,p,t]$ .* Plausibly, if  $K_{si}[S,p,t]$  then you can rely on  $S$  here.<sup>11</sup> Still, for low-stakes-you and high-stakes- $S$ , it can reasonably be of little concern to you that  $\sim K_{si}[S,p,t]$ . What presumably is of concern to you is merely that  $S$  is epistemically situated to deal with the (narrower set of) alternatives to  $p$  that are relevant to lower stakes you. Thus, it would seem reasonable for one not to much care whether  $K_{si}[S,p,t]$ , even when one care's about whether or not  $p$ , and even when one is interested in  $S$  as a potential source of information about  $p$ .

On the other hand, *if  $S$  is in a lower stakes situation regarding  $p$ , and one is in a higher stakes situation, then one does not really care whether  $K_{si}[S,p,t]$ . That is not enough, given one's interests.* Given  $S$ 's lower stakes,  $K_{si}[S,p,t]$  merely means that  $S$  is in a position to distinguish  $p$ 's obtaining versus some relatively narrow range of alternatives—a set of alternatives that are significantly narrower than the set of alternatives that are relevant for oneself, given one's own higher stakes. What one needs, it then seems, is for  $S$  to “super-know<sub>si</sub>”  $p$ —to go beyond what would constitute knowing

that *p* on the sensitive invariantist account—otherwise *S* will not be in a position to say whether *p* obtains *versus one of the relatively wider set of alternatives that are relevant to oneself*, given one's high stakes in the matter.<sup>12</sup>

We can generalize the lessons of the three kinds of cases. *Given the sensitive invariantist account*, one should not so much care whether *S* knows that *p*, but rather one should care whether *an agent who is epistemically situated with respect to p as is S*, but who *faced stakes in connection with p on the order of one's own*, would know that *p*. This is to formulate in sensitive invariantist terms the lesson that one should care about pretty much what the contextualist would understand by *knowledge*.

The bottom line: why would one, or a community, develop a concept devoted to standards that seem as ill-coordinated with the epistemic interests of the community members employing the concept at any given time as is the concept of *knowledge* as understood by the sensitive invariantist? It seems to me implausible that a community would, and thus implausible that the concept of *knowledge* works as the sensitive invariantist suggests.

Once I have said more about gate-keeping contextualism, it will turn out to be comparatively easy to understand why one would desire that oneself and others satisfied the concept of *knowledge*.<sup>13</sup> In outline, the picture is this. The attributor of knowledge, along with the interlocutors in context, are, in effect, gate keeping for a community. They might be members of the community, or they might have been called on or put themselves forward as experts on the epistemic credentials of possible sources regarding some matter. In either case, they should be speaking with the practical/epistemic interests of the relevant community in view—let us say that they are *contextually invested* with the practical/epistemic interests of the relevant or understood community. Accordingly, they and the wider community care, and should care, whether or not the agent who is a possible source is epistemically situated well enough so as to be able to discriminate which of the relevant alternatives obtained. This is just what the gate keeping contextualist understands as the conditions on true knowledge ascription. So, any party to a knowledge ascription (the attributor, the interlocutors, and the wider community, if any, that is the intended contextual audience to whom the agent is certified) should care. The general idea that contexts are sensitive to some composite of attributor and interlocutor or

intended audience practical interests, much in the manner suggested by the bank and train cases. (The case of general source communities will require special developments to come.) To satisfy contextualist standards for knowing that *p* is then to be in a strong enough epistemic position to serve as a source to those to whom the certification is addressed, and this is to be qualified in a manner that suites one's stakes in the matter. It is understandable why one would then care about whether one could then truly say of a given agent that that agent knows something.

#### IV. Contextualist Gate Keeping

If a good part of the point and purpose of the concept of *knowledge* has to do with certifying sources, or withholding certification, and thereby keeping gate for an epistemic and practical community, then a form of contextualism gets a kind of principled motivation: one in which the contexts themselves come to have a kind of principled rationality, being keyed to the practical and epistemic concerns of attributor and interlocutors. The concerns of attributor and interlocutor, and thus the contextual demands on knowledge will not swing freely with moves made in language games, or with contests of will between the partisans of higher and lower standards.<sup>14</sup> Instead, contextual demands on knowing will be rationally conditioned by the stakes within the communities for which the attributor and interlocutor are gate keeping.

This may strike the reader as promising. To begin to make good on the promise, it is necessary to survey the kinds interests an attributor and interlocutors would have in common forms of epistemic gate keeping—the kind of interests with which they might be contextually invested.

Sometimes the attributor and interlocutors will be members of a special purpose applied community. The attributor and interlocutors might jointly constitute this community (as when a working group devoted to some project meets as a whole and discusses sources on which to draw, or to gauge what they jointly know). They might instead be a proper part of the applied community (as when some of a working group's members evaluate a source for the group). Or the attributor might be an expert from outside the group or community who has been approached for advice regarding sources on a matter of concern. In all these cases, the attributor is engaged in gate keeping for a group that happens to be focused on some practical project. The members of such a group

are concerned with sources of actionable information on which to proceed with their project. In such cases of *gate keeping for an applied community*, one who certifies sources of knowledge to or for such a group should be concerned with sources of actionable information in view of the practical interests of the applied community.<sup>15</sup> One who attributes knowledge when engaged with a special purpose applied community should be contextually invested with such interests.

There are also cases in which the attributor and interlocutors may be *gate keeping for a general source community*. Then, by certifying an epistemic agent or agents as a source with respect to a claim or subject matter, one is including those agents in a select community of experts, and including the claims they there advance in an associated body of accepted results on which others might draw. The attributor here may be acting as a member of the source community speaking to other members of that group. Or the attributor may be speaking to members outside the source community—to vouch safe an agent as a qualified member of the source community to that other community.<sup>16</sup> The most straightforward and revealing cases for my purposes here, are those in which epistemic agents (with their associated information) are certified to members of a source community itself dealing with certain matters. Suppose, for example, you read an article. As you do so, you evaluate the author's methods and results. Then, in conversation with others working on the same general subject matter, you might attest that the author knows certain of the claims made in the article (and perhaps not others). Here, your interlocutors are members of a source community with community interests that are not closely associated with some kind of practical project. Paradigmatic here would be 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 be associated with any one project in a defining or limiting way.<sup>17</sup> Think of the indeterminate range of practical projects that have and can have recourse to the results of some contemporary general scientific discipline. Insofar as it is helpful to think of practical interests being connected with such a discipline, it seems that the relevant concerns revolve around providing *general-purpose actionable information*. Thus, when *keeping epistemic gate for a general source community*, it would seem that the operative concern is with general-purpose information, and that no concrete limited purposes are

likely to provide a simple decisive understanding of just what counts as actionable information. *One who certifies sources of knowledge to or for such a group should be concerned with sources of actionable information for an indeterminate range of applied communities and their various practical interests and high or low stakes.* One should be contextually invested with the general role of the source community.

The foregoing suggests two prominent genera of attributor/interlocutor settings that would make for significantly different ways in which practical interests might condition attributor/interlocutor contexts.

First, the attributor and interlocutors are concerned with a more or less concrete practical project. For them, there will be high, low, or intermediate stakes associated with the project and the proposition,  $p$ , (given how courses of action and ensuing degrees of success may depend on  $p$ 's truth). As the bank and plane cases suggest, these stakes may condition what is required of an agent in knowing that  $p$ .

Second, in context, the attributor and interlocutors may be little concerned with any concrete practical project, but focused instead with something like quality control in a source discipline with which they are associated. This discipline need not be a contemporary academic discipline. It might be a stock of folk belief to which various folk make contributions of more or less quality. In any case, the attributor and interlocutors are then engaged in distinguishing worthy contributions—one's that are fitting to serve as actionable information for a more or less indeterminate range of practical projects. To judge that  $S$  knows that  $p$  is to certify  $S$ 's contribution here. This kind of general certification, which I believe is common in the connection with expert knowledge in domains such as the sciences, would then seem to *divorce such attributor/interlocutor contexts from the concrete and particular extant concerns of the attributor and a narrow set of interlocutors.* *If anything, their concern is with a very high quality of epistemic position on the part of the epistemic agent's to be certified.* How high a quality is demanded? It seems it should be *sufficient to allow the belief as produced and held by the epistemic agent to be such as is fitting serve as a resource on which folk within other groups or communities can draw in their practical or epistemic projects.* *This sets the bar high: to qualify as knowing that  $p$  in relation to a general purpose source community, an agent must be situated with respect to  $p$  so as to be able to distinguish  $p$ 's obtaining from*

*all the alternatives that would be relevant across all the communities that might draw on it, including those for which the stakes are high.* We will soon need to explore this sort of context further, for it points to a way in which the contextualist line of thought can naturally come to resemble a kind of insensitive invariantism. That said, the connection with practical interests has not totally disappeared, and I doubt that it should.

This sketch of gate-keeping contextualism provides at best a good beginning. It has employed an overly stylized, dichotomous, understanding of the kinds of communities for which one might be engaged in gate keeping. Even if interests condition epistemic gate keeping in something like the way sketched above, and even if the concept of knowledge has been developed so as to latch onto a family of epistemic standings fitting to such gate-keeping purposes, the above crude understanding of communities doubtless yields an overly dichotomous understanding of the contexts of knowledge attribution. Between the paradigmatic general purpose source communities and the applied communities or groups devoted to a particular concrete problem or project, there would seem to be spectrum of communities for which one might be gate keeping in attributing or withholding attributions of knowledge. I must forego adding details here, for reasons of space.

#### V. General Purpose Knowledge, Gate Keeping Contextualism, and Something Like Insensitive Invariantism

I do not mean to be paradoxical, but simply to reflect on the logic of gate keeping for a general source community. This section is not intended as an internal critique of contextualism; the suggestion is not that contextualism undermines itself, leaving the field to insensitive invariantism. Instead I argue that gate-keeping contextualism can and should accommodate some of the sensibilities characteristically motivating insensitive invariantists. In accommodating these sensibilities, certain excesses of familiar forms of contextualism are avoided. These excesses are avoided in part because the pivotal notion of a context is provided a better motivation. Here, I think, gate-keeping contextualism provides a marked improvement of contextualism at a point on which many have thought it problematic. It has seemed in some writings as though the contextual standards for knowledge are wildly variable—ranging from lax to skeptical—and subject to the whims of parties to a conversation who might simply choose to focus on some alternatives or

not. In Lewis's (1979, 1996) influential discussion, and to an extent, in DeRose's (1995), merely raising a possible alternative ("You could be an envatted brain, don't you know?") made it relevant, thus raising contextual demands and destroying all knowledge of contingent claims). It was allowed that parties to the conversation might resist this, but it was never made clear on what basis, beyond a contest of wills, they would do so.

Against this background, to key relevant alternatives to practical interests seems not only to make sense of some prominent thought experiments, but also promises to advance our understanding of contexts by affording them some principled anchor. In the gate keeping contextualism, I have sought to realize some of the promise, while also drawing on a complementary understanding of the point and purpose of the concept of *knowledge*.

One might think that the place for practical interests in gate-keeping contextualism would preclude its accommodating insensitive invariantist sensibilities. The point of the present section is to show that, some of these sensibilities can be given a surprisingly satisfying response when focusing on those contexts in which one is keeping gate for a general source community. In those contexts, something remarkably like what the insensitive invariantist would envision seems in the cards.

If the above sketch of gate-keeping contextualism is on the right track, then source communities will be associated with high standards for knowing, and standards that are not volatile or significantly variable. To see why, consider a parallel. Consider the difference between designing a system to work in some delimited environment in contrast to designing one to work in an indeterminate range of environments. If one designs a system to work in a delimited environment, more or less may be required of the system, depending on the challenges faced in the intended environment. In some benign environments, relatively little will be called for. In others, the more demanding environments, much. This is analogous to the variable requirements on knowledge associated with the differing contextual needs of differing applied communities. Then, consider designing a system to work in a wide, even somewhat indeterminate, range of environments. To work across a wide range of environments, the system would need to be up to working in the most demanding of them. This finds a parallel in the less variable, more demanding, requirements for knowledge that seem fitting to a source community. The information provided must be from sources that are up to providing actionable

information in the most demanding of the applied communities that might draw on their work.

To a first approximation, one might understand what is required for an agent to know something in the context of a source community along these lines:

(In connection with a general source community)  $S$  knows that  $p$  at  $t$  in  $w$  IFF  $S$  is sufficiently epistemically well positioned with respect to  $p$  that  $S$  can discriminate  $p$ 's obtaining in  $w$  from the obtaining of any member of the set of relevant alternatives that is the union of all sets of relevant alternatives that are fitting to the various applied or source communities that might draw on  $S$ 's result.

That is,  $S$  knows that  $p$  at  $t$  in  $w$  if and only if  $S$  is sufficiently epistemically well positioned with respect to  $p$  that  $S$  can “rule out” all alternatives that would be relevant to any applied or source communities that might draw on  $S$ 's result. (One might formulate this requirement in terms of a notion of the superset of relevant alternatives to  $p$ —the union of all sets of relevant alternatives for all communities and contexts in which they might draw on the result that  $p$ .)

This seems a good start. However, much clarification is needed. How, for example, should we understand the quantified phrase, ‘all the applied or source communities that might draw on  $S$ 's result’? Should the quantifier take as its domain only extant communities? Or should it take as its domain both present and future communities? Or possible communities?<sup>18</sup>

Suppose that investigators working in some contemporary source community,  $A$ , have settled on a result, that  $p$ . They happily announce, “We now know that  $p$ .” Suppose that another community,  $B$ , entertains the implications of that result for their projects or inquiries. But, suppose also that when the  $B$ -folk entertain drawing on the  $A$ -folks’ result they find relevant certain alternatives with which source community  $A$  is not positioned to deal. That an applied or source community would find certain alternatives to  $p$  relevant, while the supporting work in source community  $A$  does not position them to “rule out” that alternative, shows that the  $A$ -folk did not know after all. One can imagine this exchange:

$A$ -ologist: We know that  $p$ .

*B*-ologist: Wow. If  $p$ , then we could make real headway on this project. But, it would be important that  $p$  is not false because of something on the order of  $q$ . Does your data really allow you to discriminate between  $p$  and  $q$ .

*A*-ologist: Hum, ... I guess not. So, I guess that we really do not know that  $p$ .

*B*-ologist: Yep, that was my worry. But, when you get further results that more finely discriminate between such alternatives, please do let me know.

This illustrates the main idea, and does so using two (supposedly) actual communities and the alternatives that they find relevant on reflection. If an alternative is relevant to a community entertaining drawing on the results of a general source community, then it is relevant for purposes of gauging whether or not the source community knows.

Two communities need not be contemporaries in order for alternatives relevant to the one to undermine knowledge attributions concerning the other (considered as a potential source community). Suppose that the *A*-folk claiming to know were cutting edge researchers at some earlier time, and that the *B*-folk are our contemporaries. The *B*-folk may look back at the work in *A* that had been thought to yield knowledge that  $p$ . Realizing that there is an alternative that is relevant to them that *A*-folk were in no epistemic position to “rule out,” *B*-folk would be right to insist:

*B*-ologist: If  $p$ , then we could make real headway on our project. But, it would be important that  $p$  is not false because of something on the order of  $q$ . The data that those old *A*-ologists generated really does not allow one to discriminate between  $p$  and  $q$ . They really did not know that  $p$ . I wonder whether some of our contemporary *A*-ologists would find this an interesting problem.

Contemporary *A*-ologists: Yes, one needs more data before one can know that  $p$ .

Apparently what we know is merely that either  $p$  or  $q$ .

This allows us to say that the set of alternatives relevant for knowing in the context of a general source community is at least the union of all sets of relevant alternatives that are fitting to all actual *present and future* applied or source communities that might draw on *S*'s result.

Do nonactual but possible communities that might seek to use the source community's result contribute to the set of alternatives relevant in connection with a source community knowing? The crude answer seems to be affirmative. It seems that

*some* nonactual possible communities should be included among those communities whose relevant alternatives constitute the superset. There may be epistemic communities that could and ought to exist but never do. For example, perhaps for reasons of a lack of imagination or industry, an actual community does not pursue a promising lead. Suppose also that, had they done so, a new discipline (B-ology) would have developed. This possible discipline constitutes an epistemic community that is regrettably never actualized, but should be. Suppose further that there is some result,  $p$ , of an actual community, the A-ologists, such that, were B-ology to have become actual, it would have then found the A-ological result that  $p$  significant, and would have entertained drawing on that result. There would be relevant alternatives to  $p$  fitting to B-ologists. My sense is that these should be included in the superset of relevant alternatives that would be pivotal in whether or not the A-ologist's result that  $p$  counted as knowledge on their part. What the possible community of B-ologists would find relevant as alternatives seems significant. This seems important for the appropriate modal depth of the phrase, "all applied or source communities that might draw on  $S$ 's result."

Still, thinking in terms of an unrestricted class of possible communities would have very strange results. There is a vast range of possible epistemically injudicious or undesirable communities, and surely some of these cannot be allowed to condition the superset of relevant alternatives without disastrous results. For that matter, there are arguably *actual* injudicious epistemic communities that should not be allowed to contribute to the superset. There is for example, the community of skeptical epistemologists who imagine that deceiving demons and envatted brains are relevant alternatives to each and every contingent proposition. These cannot be allowed to contribute such alternatives the superset of relevant alternatives relevant to general purpose source community knowledge. Obviously, these considerations point to qualifications in general formulations made earlier. The superset of relevant alternatives must be the union of the relevant alternatives of *judicious* communities. The set of alternatives relevant to source community knowing must be the union of the sets of relevant alternatives *fitting* or *appropriate* to *judicious* applied and source communities that might come to draw on the result in question—and there is some *normative* account to be given of what possible communities are relevant here. However, a more articulate

formulation of the restricted class of possible communities contributing to the superset is beyond what I can supply here.

In any case, it is clear that the results of source communities *are* commonly taken up by a very wide range of other communities (both applied and source communities). It is clear that the yet wider range of communities that *might* draw on the results of a given general source community means that (a) some of these using communities could well face high stakes, and would then make for correspondingly wide ranging relevant alternatives, thus suggesting high demands on source community knowing, and (b) in any case, the union of the sets of relevant alternatives for these consuming communities will commonly make for high demands on source community knowledge. This is as it should be.<sup>19</sup>

It plausible, in light of the above, that the demands on knowledge on the part of source communities is effectively invariant—these demands are high, and there is a fact of the matter to just how high. The point is already illustrated by the stylize discussion imagined above between A-folk and B-folk regarding the earlier purported results of the community of A-folk. There, a source community knowledge attribution that may have seemed in order turned out not to be as the later community finds relevant an alternative with which the earlier community was not positioned to deal. In that case, one judges that the earlier community or agent did not know the proposition in question—although we might readily grant that they were justified in holding the belief. Such judgments indicate that the concept of *knowledge* is not pegged to the agent's or the community's extant understanding of relevant alternatives. Nor is it pegged to some composite of extant communities' relevant alternatives, or to the then extant standards for knowledge. An agent in a source community can fully conform to the community's best sensibility of what it takes to know and yet fail to know. An agent in a source community can be so epistemically positioned as to be able to discriminate between what would count as relevant alternatives on the best then extant sensibility and yet fail to know. For example, earlier source communities devoted to medical treatments may not have understood certain possibilities to be relevant alternatives, but they were. The possibility that observed effects might be placebo effects was always relevant, even though some earlier communities of wonderful inquirers did not appreciate the point. Interestingly, that this

alternative is highly relevant was *a discovery*. We discovered it as we (our applied and source communities) discovered something about related phenomena and the resulting possibilities for getting matters wrong on the matter in question.

This idea that there is an invariant requirement for knowing, and that it is not sensitive to, or indexed to, the situation of the knower is the central claim of insensitive invariantism. I have not championed it for all contexts of inquiry, but I have suggested that it holds true of those most exemplary of knowing contexts—those of the source disciplines.

Exemplary source disciplines such as physics, biology, chemistry, even psychology and economics, including their numerous special branches, each develop standards or methods. Their methods can be understood as evolving and correctible ways of guarding against pressing ways of going wrong (considering too few cases, considering too similar a set of cases, not shielding the study from processes that would produce spurious effects, ...). Each discipline has benefited from learning about what processes could produce spurious results, and what kinds of results might result from statistical accidents. Each is informed by substantive understandings of what are the plausible alternative explanations of the data one might get, and each requires that investigators structure their inquiry so as to allow them to discern whether such alternatives remain plausible in light of their results. In all this, disciplines have evolved correctible understandings of what are the relevant alternatives to the claims that they advance—and have evolved ways of structuring their inquiry so as to allow them to discriminate which of these alternatives likely obtains.<sup>20</sup> But, to say that these are correctible understandings is to say that there can be relevant alternatives that they do not appreciate now—even if they have done an excellent job given what they now have to go on. This should be as much insensitive invariantism as anyone could want.

## VI. Inconclusive Conclusion

Can the pieces—gate-keeping contextualism and the contextual approximation of insensitive invariantism—happily coexist. Will the insensitive invariantist themes just accommodated remain comfortably contained within the gate-keeping contextualist framework?

My worry, the open question, is can be understood as internal to gate-keeping contextualism itself. It is a question of whether it contains the seeds of its own critique. There is a reasonable suspicion that there is something about the project of gate keeping for variously situated, but significantly cross-cutting and interacting groups, that makes the kinds of concerns and ranges of relevant alternatives associated with source community knowledge more pervasive and invasive than is suggested above. To say that considerations fitting to source communities may be invasive is to say that these concerns (together with the associated wide range of relevant alternatives) may come to condition the contexts in which one keeps epistemic gate for a community with a special purpose applied focus. To the extent that this obtains, such contexts come to take on more of the epistemic sensibility of source communities—and the kind of standards and alternatives fitting there become pervasive. To the extent this obtains, then, gate-keeping contextualism ultimately comes to take on yet more of the form of insensitive invariantism—at least in its bottom line. There are several reasons for thinking that the concerns associated with source communities might be invasive and pervasive.

First, there is the matter of overlapping membership and trained sensitivities. Consider a representative epistemic agent, A, who is a member of various epistemic communities, and who moves freely between them. Perhaps A is a wetlands ecologist, a member of several conservationist organizations where A regularly discusses projects and causes with people of diverse educational and professional backgrounds. As A moves between the relevant groups—colleagues, funding agencies, editors, discussion groups, political organizing committees, and decision makers in connection with land and water use, A is likely to confront his or her epistemic landscape with a (only moderately flexible) sensibility regarding what is acceptable epistemic positions with respect to claims. In formulating talking points for political action in cooperation with A's fellow conservationists, A will commonly favor qualified or hedged statements of what is known about threats and what can be done. A will be concerned that the group not overstate its position. In all this, A will likely present a developed epistemic sensibility, with moderate variations—A will likely then insure that the knowledge claims of the applied communities not be too facile, from the point of view of one engaged with a source community.

Thus, to the extent that an agent's epistemic sensibility is significantly a matter of training, the overlapping membership of many groups will result in standards associated with source communities coloring what transpires in applied communities—it will do so first because it will color what representatives of the source community will attest to in testimony to the applied community, and it will do so insofar as those with source-community-like epistemic sensibilities are members of the applied community.

Second, there is the matter of *the bilateral or multilateral exchange of information*. Treatment teams at research and other hospitals collect data for, and collaborate with, research groups—applied groups dealing with concrete problem cases provide information to (and commonly overlap in membership with) source communities focused on a related problem area domain). The research groups then produce results that are taken up and inform “more fundamental” work in yet more general source communities—researchers in a problem-area domain provide information to general-purpose source communities. Work in the general-purpose source community can then yield results informing the both the work of the problem area domain source community and various applied groups—including the treatment team's practices. (This quick gloss is doubtless overly schematic and hierarchical in its portrayal of information exchange.) When one looks, some multilateral informational exchange is the norm.<sup>21</sup> To the extent that there are such multilateral information flows, the standards on the production of results within the more applied groups cannot be too out of step with the demands characteristic of general purpose source communities. Such relationships between communities reinforce the common phenomena of overlapping memberships. Both would seem to restrict the space for contextual standards that depart from the contextual standards appropriate to general source communities—which mimic what is commonly envisioned by nonskeptical insensitive invariantists.

What I am questioning whether the real differences in the alternatives that are salient to differently situated applied communities can ultimately support the differences earlier envisioned between epistemic contexts. I am questioning whether the differences in *salient* alternatives make for durable differences in *relevant* alternatives. I am seeking to make the point within the framework of the gate-keeping contextualism that I sketched—and which I believe is the most promising form for contextualism to take. The

point can be put this way. If the flow of information across groups—from most practical communities to more source communities—is significantly multilateral in character, then the kinds of gate-keeping fitting to even applied communities typically cannot be conditioned solely by some immediately entertained application of some source’s putative information. Even applied communities do not merely use information; they also pool or store it, produce new information using old, *and pass information along*. To the extent that this is in the cards, then the considerations associated with source communities should be pervasive—and the conditions on what it takes to know become significantly less variable or sensitive to interests distinctive of a narrow contextual moment.

What I find most intriguing about the general kind of gate-keeping contextualism I have entertained here is the way in which it promises to draw on a point about the point of the concept of *knowledge*, thereby making principled sense out of contexts, and to allow one to then accommodate some of the apparent interest-keyed variation in judgments discussed in the literature on contextualism as well as the kind of themes that attract insensitive invariantist. Whether the resulting picture should ultimately vindicate significant contextual variation then depends on the question of how much of our epistemic life can be understood as exemplified by the kind of narrow practical moment featured in the bank case where what is said in the car is used and stays in the car, and how much is instead to be understood as exemplified by the source community production of a pool of information for others. If gate-keeping contextualism is on the right track, knowledge attribution is pivotally the certification of a source on a matter. One can think of the certification as analogous to the certification of a financial instrument—a currency, something that is also passed around. The question then becomes whether attribution of knowledge in the context of many applied communities is aptly understood as like the certification of the old Slovene tollar (which would not be tendered beyond very narrow borders) or instead like the certification of a Euro note—intended reasonably for wider currency. To the extent that the former sometimes obtains, then gate-keeping contextualism with significant variability in connection with applied groups is descriptive. To the extent that the latter typically obtains, even gate-keeping connectionism comes to approximate insensitive invariantism.

## References

- Barke, A. 2005: "Epistemic Contextualism." In E. Brendel and C Jäger (eds.), *Contextualism in Epistemology*. Dordrecht: Springer, pp. 211-32.
- Cohen, S. 1998: "Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery." *Australasian Journal of Philosophy* 76: 289-306.
- Cohen, S. 1999: "Contextualism, Skepticism, and the Structure of Reasons." *Philosophical Perspectives* 13: 57-89.
- Cohen, S. 2005; "Knowledge, Speaker and Subject," *The Philosophical Quarterly* 66: 199- 212.
- Craig, E. 1990: *Knowledge and the State of Nature*. Oxford: Oxford University Press.
- DeRose, K. 1992: "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52: 913-29.
- DeRose, K. 1995: "Solving the Skeptical Problem," *Philosophical Review* 104: 1-52. Reprinted in K. DeRose and T. Warfield (eds.), *Skepticism: A Contemporary Reader*. Oxford: Oxford University Press (1999).
- DeRose, K. 2002; "Assertion, Knowledge and Context." *Philosophy and Phenomenological Review* 111: 167-203.
- DeRose, K. 2004: "The Problem with Subject Sensitive Invariantism." *Philosophy and Phenomenological Review* 118: 346-50.
- DeRose, K. 2005: "The Ordinary Language Basis for Contextualism and the New Invariantism." *The Philosophical Quarterly* 55: 172-98.
- Golley, F. 1993: *A History of the Ecosystem Concept in Ecology*. New Haven: Yale University Press.
- Hawthorne, J. 2004: *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hempel, C. 1965: *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: The Free Press.
- Henderson, D and Horgan, T. 2001: "The A Priori Isn't All It Is Cracked Up To Be, But It Is Something." *Philosophical Topics* ??: 219-250.
- Henderson, D and Horgan, T. forthcoming: "The Ins and Outs of Transglobal Reliabilism" (with Terry Horgan), in Sanford Goldberg (ed.), *Internalism and Externalism in Semantics and Epistemology*, Oxford University Press
- Kitcher, P. 1993: *The Advancement of Science*. Oxford: Oxford University Press.
- Lewis, D. 1979: "Scorekeeping in a Language Game," *Journal of Symbolic Logic* 8: 339-59.
- Lewis, D. 1996: "Elusive Knowledge," *Australasian Journal of Philosophy* 74: 549-67. Reprinted in K. DeRose and T. Warfield (eds.), *Skepticism: A Contemporary Reader*. Oxford: Oxford University Press (1999).
- Quine, W. 1953: "Two Dogmas of Empiricism," in W. Quine, *From a Logical Point of View*. Cambridge: Harvard University Press (pp. 20-46).
- Quine, W. 1960: *Word and Object*. Cambridge: M.I.T. Press.
- Schaffer, J. 2004: "From Contextualism to Contrastivism," *Philosophical Studies* 119: 73-103.
- Stalnacker, R. 2004: "Comments on 'From Contextualism to Contrastivism'," *Philosophical Studies* 119: 105-17.

- Stanley, J. 2005: *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Williamson, T. 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. 2005a: "Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge," *Philosophical Quarterly* 55: 213-35.
- Williamson, T. 2005b: "Knowledge, Context and the Agent's Point of View," in G. Pryer and G. Peters (eds.), *Contextualism in Philosophy*. New York: Oxford University Press, pp.

---

<sup>1</sup> One plausible observation regarding the point or purpose of the concept is not developed here: that in addition to the gate-keeping purpose focused on here, there is a kind of crediting purpose. Suppose that one is comfortable in thinking that one knows that *p*. For example, suppose one is a member of a research community for which *p* is a matter of well-established doctrine. Perhaps *p* has to do with the effectiveness of aspirin in relieving certain forms of minor discomfort—that is certainly well established. Here, one appreciates that one is so well placed on this matter that folk should be able to take it from you that *p*—even when being so well placed is gauged using the demanding standards fitting to source communities (see the discussion in the text to follow). Still, one may want to say of historical communities, that they had knowledge of propositions closely related to *p*—for example, that indigenous communities knew that willow bark was effective in relieving certain forms of minor discomfort. Here, one is not recommending that these communities be included as sources (be treated as fully contributing members of one's contemporary source community). Rather, one is crediting them as having gotten to the much the same result, and in a way that made them fitting sources for *their* respective communities. Were we to fold this observation about a crediting point to the concept of *knowledge* into the points made in the course of the present paper, we would certainly complicate the picture emerging here. Perhaps we would need to recognize distinct gate-keeping and crediting uses of the concept. In any case, I here set this second point aside in order to explore what is suggested by the gate-keeping purpose taken alone.

<sup>2</sup> Williams (2005) provides a formally elegant exposition of the competing positions.

<sup>3</sup> It is sometimes said that contextualists hold that the term 'knows' is indexical. I resist saddling coherentists with this commitment—although *some* presentations may encourage it. My reservation reflects the caution exhibited in formulations such as Cohen's: "Now, according to contextualism, 'knows' (along with 'flat', 'tall', etc.) is context-sensitive in ways analogous to indexical terms" (Cohen 2005, p. 202). Compare also Schaffer (2004) and Stalnacker (2005).

<sup>4</sup> The misgivings one might have with the "rules out" locution is that it seems to suggest a kind of epistemic certitude. So, while the "rules out" locution makes for a pithy formulation, I prefer to write of being well positioned to discriminate which of these obtains (if any).

<sup>5</sup> A more delicate question is whether I was correct when I, in conversation with Scarlett, attributed knowledge to myself. Here, much hangs on whether Scarlett and Violet's high stakes should condition the conditions on knowledge. My own thinking is that they should. In attributing to myself knowledge here, I am certifying myself as a source of actionable information *for the practical community comprised of Scarlett and Violet*.

<sup>6</sup> It is just at such places that the sensitive invariant will tend to invoke pragmatics instead of semantics to account for our judgments—distinguishing between what is correct to say and what is true.

<sup>7</sup> Here I am thinking of contextualist verdicts in those cases where knowledge attribution would be keyed closely to interests in practical projects.

<sup>8</sup> It is also common coin that only some practical interests condition what counts as knowledge. The interests of knowledge consumers—those who put information to work in their projects are relevant. Other practical interests surely illicitly condition knowledge attributions—for example, when the claims are motivated by political interests to justify a policy or to justify a group's professional authority (see Dawes () on the epistemic pretensions of clinical psychologists and clinical social workers).

<sup>9</sup> Concepts such as that of *rationality*, and of *subjective justification*, would seem more attuned to the subject's interests—and whether the subject is epistemically situated to distinguish between the alternatives that salient and relevant to the subject. But see also note 1.

<sup>10</sup> The problem I point to here is closely associated with problems that sensitive invariantists have dealt with concerning the treatment of high attributor stakes/low subject stakes cases. At this point, sensitive invariantists have tended to invoke conversational pragmatics to explain why one would find it incorrect to attribute knowledge to the low stakes subject when talking with high stakes interlocutors, even though (they insist) it is strictly true that the low stakes subject has knowledge. The contextualist semantics and the sensitive invariantist semantics combined with conversational implicature may each accommodate the judgments elicited by these cases. (Compare DeRose [2004].) My present point may be put this way: it is implausible that the semantics of the concept should be so out of step with its apparent point and purpose as the sensitive invariantist would have it.

<sup>11</sup> Although even this may depend on how the alternatives that *S*'s high stakes situation with respect to *p* make relevant line up with those that your lower stakes situation make relevant.

<sup>12</sup> These points are related to a matter that has received more attention: the verdicts rendered regarding cases of low subject stakes and high attributor stakes cases. (Stanley, 2005; DeRose ...)

<sup>13</sup> At least this is true as long as one sets aside the early contextualist accounts with their susceptibility to skeptical contexts. The susceptibility to skeptical moves, and the failure to give an account of contexts that ties them to real concerns and stakes for real epistemic agent constitutes a prominent reason for resisting more familiar forms of contextualism.

<sup>14</sup> Contrast what is suggested in Lewis (1996) and compare the concerns are expressed in Barke (2005).

<sup>15</sup> The point needs qualification in light of the next note

<sup>16</sup> Source-community attestations to an applied community may be understood in two significantly different ways. First, the attributor may be certifying to those in the other community that the agent in question counts as knowing within the source community—with its high, general purpose, standards—soon to be discussed. Alternatively, the attributor may understand his or her role as that of translating the source community epistemic situation with respect to various propositions into knowledge claims that would be fitting given the epistemic gate keeping demands for the applied community audience. The question will be which is the salient purposes in context.

<sup>17</sup> Admittedly, certain classes of practical projects may have motivated work in a general subject discipline, and at historical junctures may have motivated much of the discipline. Perhaps military applications motivated work on projectile motion in early modern physics. Still, the general discipline transcended these practical interests.

<sup>18</sup> Some of the formal questions here raise issues that are beyond my intuitions, inclinations, or sensitivity. I find it hard to say how exactly one should individuate communities across possible worlds, for example. Without this settled, it can be difficult to know exactly what to make of the modal requirement regarding communities that “might” draw on some result. Even so, it seems to me to make very good sense to say that a working group or a community might make more or better use of results in some other domain—for example, that archeologists might make some, more, or better use of extant anthropological work on supposedly parallel peoples and their social organization, or that they might use certain recent results in basic chemistry to construct a new measure of the age of certain artifacts. They might. This is to talk of the possible epistemic action of an actual community—in this case a source community.

<sup>19</sup> Various discussants (Albert Casullo, Matthew McGrath, John Greco, among other) have urged that my account sets the standards for source community knowing too high—inadvertantly making for a form of skepticism. Andrew Newman worries that the present formulation would make for problems for the autonomy of the various sciences—for the kinds of sensible standards that disciplines individually develop. To the extent that these worries are well founded, and to the extent that I am right about the motivated nature of contextualism, this would indicate that my present formulations must be too crude.

<sup>20</sup> Disciplines have evolving understandings of live possibilities—and of methods for treating these. To date, none of these involve BIV or demonic infestation as live possibilities. They involve competing hypotheses, including the null hypothesis and hypotheses regarding alternative causes. The disciplines have developed various reasonable ways of discriminating between an hypothesis that is vindicated and the contextually live alternatives—these employ devices such as shielding from known or plausible alternative causal agents, control treatment groups, longitudinal scope of study, diversity of samples, size of data set and sensitivity to statistical properties, and so on. Becoming a certified practitioner of a discipline (earning a PhD in a branch of physics, chemistry, biology, for example) involves learning about these threats to one's coming to know—and learning how to structure one's inquiry to deal with them. What one learns is not final, and is subject to ongoing refinement. This sort of thing is reflected in Kitcher (1993). Kitcher

---

portrays the informed evolution of such standards in the discussion of stylized cycles of training and work within a discipline. Conformity with such standards, when the discipline is on its game, is closely associated with objective epistemic justification (see also Henderson and Horgan [forthcoming] on suitable modulational control).

<sup>21</sup> The present concern is closely related to those Williamson (2005b) raises regarding the demands contextualism would seem to make in connection with belief acquired by testimony. It seems that one would need to keep track in memory of the context in which you acquired a testimonial belief, otherwise it would be employed and passed along in contexts where different standards would be appropriate. This puts a awkward demand on memory.