

Gate-Keeping Contextualism

David Henderson

This article is forthcoming in *Episteme*

ABSTRACT. This paper explores a position that combines contextualism regarding knowledge with idea that the central point or purpose of the concept of *knowledge* is to feature in attributions that keep epistemic gate for contextually salient communities. After highlighting the main outlines and virtues of the suggested *gate-keeping contextualism*, two issues are pursued. First, the motivation for the view is clarified in a discussion of the relation between evaluative concepts and the purposes they serve. This clarifies why one's sense for the point of an evaluative concept ought to constrain and inform one's understanding of the concept. Second, the paper explores ways of avoiding a problem in the author's earlier development of gate-keeping contextualism. The initial development of the view opened the door to a form of skepticism that would hobble an important facet of our social-epistemic lives.

BIO. David Henderson is Robert R. Chambers Distinguished Professor of Philosophy at the University of Nebraska, Lincoln. He works in both epistemology and philosophy of the social sciences. His recent work includes a number of papers in epistemology co-authored with Terry Horgan, culminating in *The Epistemological Spectrum* (Oxford University Press, forthcoming).

i. Introduction

In "Motivated Contextualism" (Henderson 2009) I began exploring the promise of a position that combines contextualism regarding knowledge-attributions with an attractive idea about the purposes served by the concept of *knowledge* – namely, that the central point or purpose of that concept is to feature in attributions that keep epistemic gate for various communities by marking out fitting sources of information. After reiterating the main outlines and virtues of *gate-keeping contextualism*, I develop two refinements. First, I make some general observations regarding the relation between evaluative concepts and the purposes they serve. This clarifies why one's sense for the point of an evaluative concept ought to constrain and inform one's understanding of the concept. Second, I correct a mistake. My initial development of gate-keeping contextualism did not adequately limit the epistemic demands on what I term *general-purpose source communities*. As

a result, it opened the door to a form of skepticism that would hobble an important facet of our social-epistemic lives. Such skepticism has no place in the gate-keeping contextualist approach being explored. Thus a central concern of this paper is to refine our understanding of the role of, and aspirations appropriate to, general-purpose source communities.

Plausibly, a central point of the concept of knowledge is reflected in Craig's (1990) idea that the concept serves principally in certifying good sources of information. When one attributes knowledge to an agent, one indicates to one's conversational partners that the agent is so epistemically positioned that one can freely draw on their veridical result. The suggestion is that the concept of *knowledge* is best understood as deeply conditioned by what would be fitting were one using the concept to keep epistemic gate for various contextually understood communities. In attributing knowledge, the attributor recommends a source to a contextually understood community, and the standards of knowledge should be fitting in view of the cognitive and practical stakes of that community. Contextualism gains significantly from being so alloyed, as the resulting contextualist account of knowledge affords a more principled treatment of contexts and of contextually fitting standards than is afforded by earlier versions of contextualism.

Contextualists posit a variation in the standards for knowing operative across the various contexts. They do so in order to accommodate common judgments about the correctness of knowledge attributions in various cases. The working hypothesis in (Henderson 2009) and here is that reflection on the epistemic needs of certain broad classes of communities would mesh well with, and provide a principled basis for, common judgments about when an agent knows. Dialectically, the situation is as follows. Contextualist faces two principal competitors – sensitive invariantism and insensitive invariantism. Each position seeks to account for judgments regarding when agents know. The sensitive invariantist emphasizes cases in which agents (who are evaluated as knowing, or not) themselves face various practical stakes. The sensitive invariantist insists on a

particular intimate connection between knowledge and practical reason in such cases – the agent must be in an epistemic situation commensurate with *his or her own* need for actionable information in order to count as their knowing. On the other hand, the insensitive invariant has emphasized the suspicion some have that the apparent variation in standards may be so much loose talk. They seek an epistemology with backbone – with standards that are fitting to mark out an epistemically well-positioned agent, without qualification. One may feel pulled strongly in both ways – I do. And this is why I find gate-keeping contextualism so attractive: it can honor variants on both approaches within one unified account – one that is not at all *ad hoc*, and one that draws naturally on one’s sense for the epistemic needs of various communities. On the one hand, it can accommodate the range of judgments that seem to indicate a variation in standards with the practical stakes faced – and can insist on a tight connection between knowledge and practical reasoning. On the other hand, it can honor the sense in which one might rightly insist on standards that qualify one as a source of actionable information for any who might reasonably look for actionable information. This is a neat trick. Of course, as a contextualism, it must add that the two pulls are honored in somewhat differing contexts – but they are each honored, and in a principled and uniformly motivated way.

This strength of gate-keeping contextualism becomes apparent when one considers two broad classes of communities for which one might be contextually keeping epistemic gate. When one thinks about the epistemic needs of *applied practical communities*, and what would qualify an agent to serve as a source of actionable information for such communities, one finds that what counts as knowing in the context of such a community plausibly would be sensitive to the stakes faced by that community. In this way, a motivated contextualism can account for the apparent connections between knowledge and stakes and between knowledge and practical reasoning. Motivated contextualism thus proves nimble and powerful in treating matters of concern to subject-sensitive invariantists. Further facets emerge when reflecting on the concerns of communities of inquirers who seek to generate a

stock of information on which a somewhat indeterminate range of others can reasonably draw – call these *general-purpose source communities*. Think of general scientific disciplines as constituting paradigmatic examples. The contextual demands fitting for one keeping epistemic gate for such source communities would approximate those commonly advocated by classical, or insensitive, invariantists. Thus, motivated contextualism plausibly has the resources to accommodate the intuitive strengths of a range of approaches to knowledge, all within a single principled approach drawing on the plausible idea that a central point or purpose of the concept of *knowledge* is to serve in epistemic gate-keeping.

ii. A sense for the point or purposes served an evaluative concept can and should constrain one's philosophical account of the concept.

It is overwhelmingly plausible that evaluative concepts such as those of a *competent* (or *good*) *driver*, of a *good* (or *competent*) *basketball player*, and of a *good* (or *competent*) *archer* arose in connection with, and along with, a general sort of practice. The concept of a *good driver* and of a *well-executed turn* arose in connection with the practice and technology of operating (first animal-drawn, then motor-powered) vehicles on public roadways. The concept of a *good archer* arose in connection with the use of bow and arrow in both warfare and hunting. As that technology was largely replaced, the concept retained a role in a kind of contest employing the technology. The concepts of a *good basketball player* and of a *good pass* arose in a connection with a game. In each case, the concept would not have arisen without the associated practice. In each case, it seems that the concept arose with a *constitutive eye to the demands of successful practice*. Such concepts are from their inception, and in what might be termed their *core conception*, to be attuned to what facilitates the associated practice.

To say that a concept arose with “a constitutive eye to” the demands of successful practice, and that what makes for success there is central to the “core conception” of the concept, is to say that the semantics constitutive of the concept turn on what makes for such success, rather than on what is merely thought at some

time to make for such success. This is a thesis about what is central to, even constitutive of, the concept, versus what is merely a matter of the correctable conception or understanding one might have in connection with the concept. Here I am appealing to the distinction between the semantic entity – the concept – and mere conceptions. (This is a distinction that both semantic internalists and externalists must honor.)

Plausibly, evaluative concepts arose with an eye to, or sensitivity to, what is needed in, or conduces to, successful pursuit of some project. The main outlines of the concept would then be subservient to some needs arising in the pursuit of the project, and such needs constitutively frame or structure the concept itself. Individuals develop particular understandings – conceptions – in connection with the evaluative concept. But these conceptions are not constitutive of the concept. For example, the concept of a *good driver* presumably arose to regulate the operation of vehicles (paradigmatically on public roadways) and is sensitive to the associated needs for coordinating with others who are themselves operating such vehicles, for mitigation of the risks to self and others posed by this technology, for the expeditious transport of people and cargo, and the like. Now, those possessing this concept will have some more or less developed conception of the character traits, capacities, and skills that make for being a good driver. However, were one were to find that some character trait one had thought was necessary for being a good driver did not contribute to the project in question, then one would need to amend one's conception of a good driver. In doing so, one would not have changed the subject – before and after one amended one's conception, one would be using the same concept of a *good driver*. One merely would have refined one's understanding of what makes for good drivers, of what satisfies the concept – but what does satisfy the concept would not have itself changed.

In analyzing or reflecting on an evaluative concept, it will be important to keep in mind the project in which the concept has its roots – and whatever sense one has for how the concept serves to regulate that project. This is to say that philosophical reflection should commonly be informed and constrained by one's

sense for the point and purpose of the concept we seek to better understand (see Henderson and Horgan 2001, forthcoming).

Those who are competent users of an evaluative or normative concept commonly have some sense for how the concept functions in their lives to regulate some ongoing project. Of course, one's sense can be tentative and vague. This is the case with our concept of *knowledge*. I find it plausible that something on the order of epistemic gate-keeping for various communities (various practically-engaged communities or groups, and various epistemically-engaged source communities) is at least *a* central point or purpose of the concept (see also Greco 2008, 2009; Pritchard 2010; Kappel 2010; and Grimm 2009). One does not need to have a precise, fully worked-out understanding of the point of the concept in order to get philosophical work out of the sense that one does have. A rough sense for the point of the concept, how it functions to regulate practice, can usefully constrain what account one gives when accommodating judgments about cases. At the same time, reflection on cases can provide clues regarding what more precisely is the point or purpose central to the concept. Accordingly, I advocate neither a "cases-first approach" nor "a purpose-first approach." Instead, both judgments about cases and one's sense for the point of the concept each amount to data that should inform and constrain one's account of the concept. A proper account of the concept will be supported by what is ultimately a multiply constrained abductive inference (Henderson and Horgan forthcoming, ch. 2). Gate-keeping contextualism is here developed as a working hypothesis, ultimately to be compared with alternatives.

iii. Gate-keeping contextualism, practical communities, and principled variation of standards with "stakes."

Earlier, we distinguished two broad classes of communities for which one might be keeping epistemic gate. Sometimes, one keeps epistemic gate for communities engaged in some practical endeavor – *applied practical communities*. Sometimes one keeps gate for *general-purpose source communities* – communities of inquirers having a social role of producing information of such a high epistemic quality that a

somewhat indeterminate range of groups might freely draw on their results without hesitation. We also noted that the stakes faced in practical communities should condition knowledge attributions in those contexts in a way not paralleled when gate-keeping for source communities. According to gate-keeping contextualism, the folk of contextual concern are those to whom the agent is ultimately certified as a source.¹ The truth of the attribution turns on whether the agent's epistemic situation at the time in question was or is sufficient to the demands of the context of attribution. (For our purposes here, we can understand one's strength of epistemic position with respect to some proposition p to be a matter of the range of alternatives to p with respect to which one's evidence and processes favor p .)

Subject-sensitive invariantism provides a form of sensitivity to stakes that makes little general sense when one is concerned with gate-keeping for a practical community, whereas a contextualist can provide a form of sensitivity that is attuned to the needs of gate-keeping (Henderson 2009). In epistemic gate-keeping for engaged practical communities, in certifying sources adequate for their practical reasoning, one should be sensitive to their stakes. Discussions of practical rationality in economics and decision theory suggest a way of thinking about these matters. Focus on the reasonable concerns of those in some practical community. Roughly, as one's stakes go up, it seems reasonable to be willing to pay higher "information costs" in order to guard against a wider range of failures. Of course, one would be happy to draw on others who had paid those higher costs already. One seeks for the results of the more costly inquiry (on one's part or one's sources) only insofar as it guards against error possibilities that could frustrate one's high-stakes projects. Thus, when facing the high stakes, folk would be unreasonable not to incur the costs of guarding against a relatively wide range of error possibilities. Again, if someone else had already put themselves in a position to rule out the relevant range of error possibilities, it is reasonable to draw on them. What one would not want, reasonably, is to draw one's information from someone that was not in a position to rule out the relevant error. According to gate-keeping contextualism, attributions of

knowledge to practical communities serves to certify sources that are fitting in light of such reasonable concerns.

iv. Gate-keeping contextualism and general-purpose source communities.

When keeping epistemic gate for a general source community, no concrete limited purposes provide a simple decisive understanding of just what counts as being in a sufficiently good epistemic situation on the matters in question. A somewhat indeterminate range of communities might seek to rely on one's results and one who certifies sources of knowledge to or for such a community should be concerned with sources of information adequate to the purposes and stakes of a correspondingly wide range of applied practical communities. One should also be concerned with sources adequate to various other source communities that might also draw on one's results in their investigations.

Suppose then that one is a member of some general-purpose source community, and that one is charged with evaluating some putative results gotten by an aspiring member of one's source community. In such a context, the attributor (oneself) and interlocutors (one's fellow source community members) will typically not be focused on any concrete practical project.² In such a context, folk are focused instead with something like quality control for their source discipline. They are engaged in distinguishing worthy contributions – results on which a more or less indeterminate range of folk (who might come to have interests in the matter) might reasonably freely draw and rely. This kind of general certification is common in the connection with expert or scientific knowledge. This calls for a very high quality of epistemic position on the part of the epistemic agents to be certified – an epistemic position sufficiently strong to allow the epistemic agents in the source community to fittingly serve as a source on which folk within other groups or communities can reasonably readily draw in their various practical or epistemic projects.

Accordingly, we can say to a first approximation what it takes to qualify as knowing that p , in the context of a general-purpose source community:

(In a source community) Ks,p,t IFF

s is sufficiently epistemically well positioned with respect to p that s can

discriminate p 's obtaining 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 (the union set of relevant alternatives).

The contextual variation of relevant alternatives with stakes associated with practical communities is thus lost in the union of relevant alternatives fitting to the audience of the general-purpose source community. For such communities, the relative alternatives do not fluctuate, not at least at reasonable time scales in which later investigators might seek to draw on the results of earlier research. Thus, in connection with general-purpose source communities, gate-keeping contextualism resembles a kind of insensitive invariantism.

For any given general-purpose source community, there are a range of source and practical communities that might be interested in their results – these are roughly those communities that might want to draw on the sorts of information on which the source community specializes. One might think of this as the *reasonable audience* of the source community – the source community serves to supply information adequate to the purposes of all within this reasonable audience. To do this, it must be that as its members contribute some result p to its stock of results, the contributing member s (or set of members) must be (individually or jointly) epistemic positioned so as to be able to tell in favor of p versus any of the alternative possibilities that are relevant to any member of its reasonable audience. Refinements will be needed.

v. A new skeptical worry?

In keeping with the above understandings of the commitments of source communities, we get this principle:

When some agent s within some source community claims to know that p , and some epistemic agent s^* who is entertaining drawing on s 's putative result that p finds that s 's epistemic position is not sufficient to tell against some error possibility that reasonably seems significant for s^* 's own project or inquiry, s^* should judge that s does not know that p .

Many will find this judgment intuitively correct. With refinements to be developed in the remainder of this paper, it is correct.

The above indicates that gate-keeping for general-purpose source communities reasonably gives rise to high standards that are largely invariant. However, one may worry that the envisioned high standards are *too* high. Much depends on just how one understands the *reasonable audience* of a source community. In this section I develop the worry that the contextual standard, as delineated to this point, would engender a disastrous form of skepticism. I also explain why demands on knowing that would result in such skepticism indicate that some mistake has been made in developing gate-keeping contextualism. In the following section, I provide a refined understanding of fitting gate-keeping for general-purpose source communities.

In characterizing to a first approximation the social role and commitments of a source community, and what it takes to know in such a context, I used variations on the phrase, “the various applied or source communities that might draw on *s*’s result.” This must be understood so as to make for a significant modal depth in the gate-keeping contextualist account of knowledge (see Henderson 2009). However, unless care is taken with this phrase, the standard sketched above could easily give rise to a virulent form of skepticism. It will be particularly important to refine our thinking about communities that are separated by an interval in which there is much development of evidential and methodological resources.

Commonly, in the course of scientific practice, one learns of error possibilities that one certainly needs to structure one’s inquiry to guard against. As a case in point, think of the demand in much medical research for double-blinded experimental designs and placebo controls (or controls involving understood comparison treatments). Enquiry is so structured in order to yield an epistemic situation that can deal with particular error possibilities. These error possibilities are far from idle; they were made salient by empirical results regarding placebo effects and experimenter bias.

One purporting to know in the context of a source community is taking on the burden of dealing with such error possibilities as would be found significant for

any who reasonably seek to draw on their results – these include others in the same source community as well as those in other source and practical communities with an interest in such matters. Once error possibilities are newly recognized in light of either empirical or theoretical developments, one who is gate-keeping for the source will rightly judge that earlier experiments or investigations were not structured so as to deal with important error possibilities – error possibilities that were initially not evident, perhaps, but which nonetheless must be dealt with in order to properly contribute to the source community’s stock of established results. In the context in which one is keeping epistemic gate for a source community, one will then judge that the somewhat earlier investigators, who understandably did not guard against the relevant error possibilities, consequently did not produce knowledge. One cannot flatly draw on their results and the needs to be investigated further. Continuing the above example, one will judge that the earlier experimenters, who understandably did not structure their study so as to be in a position to rule out placebo effects and experimenter bias, do not know that such and such a treatment regimen is effective. This much seems reasonable, and is fully in keeping with the demands of gate-keeping contextualism as set out above. Insofar as the later source community is reasonably entertaining simply drawing on the results of the earlier community, incorporating those results in its corpus of accepted results, such evaluations seem exactly right.

It is common to look to the results of recent studies and results when structuring our inquiries. For example, one who is setting up an experiment in a physics laboratory will rely on much information regarding supposedly well-understood processes and materials when designing instruments and shielding their experimental system from confounds of various kinds. This information will have been produced over the preceding few decades, and typically will be collected in handbooks of the most up-to-date results (themselves temporally situated codifications of institutional gate-keeping). Such practice draws on one’s community’s best present sense of relevant error possibilities. But, again, this sense for error possibilities is not definitive. New error possibilities soon may become apparent, and the investigations that do not tell against these error possibilities will

be found to have been inadequate to yield knowledge.

While all this seems generally reasonable and proper, it has skeptical implications. As we look back upon earlier source communities, or earlier stages in our own source community, one is likely to judge that those at work there did not know – because they will not have so structured their investigations as to guard against error possibilities of which we have subsequently learned. But, if we are honest, we should expect that in the course of future inquiry, folk will discover error possibilities that we today have not even thought to guard against. Future investigators will judge of us that we have not structured our inquiry so as to put us in an epistemic position to tell against error possibilities that they find relevant and important. Just as we are right to judge that earlier investigators do not know, so these later investigators will apparently be right to judge of us that we do not know. We may seek to be in an epistemic position sufficient to tell in favor of our results, and against error possibilities relevant to all who might entertain drawing on our results – we might signal to others that we have succeeded when we claim to know – but we will likely turn out to have failed here. Apparently, it is very likely that one does not know. This is skepticism.

The skeptical threat just sketched would certainly frustrate our community epistemic practices, specifically those practices associated with general-purpose source communities. This is evident when we think of attributions (and denials) of knowledge as being used to regulate or keep epistemic gate for such communities. In such contexts, to say that some investigators' result did not amount to knowledge would signal that they were not epistemically positioned as fitting to contribute to a stock of beliefs on which others could freely draw. To say that no one knew anything in that way would be to give up on the project of providing such a general-purpose stock of beliefs about subjects of general interest and significance (such as biology, geology, ecology, and physics). All would be epistemically worse off were folk to take this seriously.

In keeping with what was said in section *ii*, the skeptical result should be taken to indicate that some mistake was made in developing the gate-keeping contextualist account. The skeptical result is completely antithetical to any approach

that sees the concept of *knowledge* as rooted in the social epistemological need for keeping epistemic gate. Evaluative and normative concepts grow up in response to the needs of regulating ourselves and others in the course of pursuing some project. What is required to satisfy evaluative concepts cannot diverge from what is reasonable in view of needs faced in the project in which they have their roots. Were one to find that one's conception of knowledge would frustrate the individual and community practices that the concept of *knowledge* grew up to regulate, one's conception would need to give way. Properly understood, the concept would facilitate that project – certainly not utterly frustrate it. My initial analysis of *knowledge* has led to a broad skeptical result with respect to general-purpose source communities. I must have taken a wrong turn somewhere along the way.³

vi. Avoiding skepticism within gate-keeping contextualism.

The skeptical threat arose when thinking about epistemic gate-keeping for general-purpose source communities. These were characterized as aspiring to produce results that could serve as a source on which anyone (any practical or source community) could readily draw. The threat of skepticism derived from the observation that communities significantly far in the future are likely to have an informed perspective from which they might find relevant error possibilities that cannot be appreciated by contemporary communities of inquirers. In many fields, both the understanding of relevant error possibilities and the understanding of how one needs to structure inquiry to deal with such possibilities develop significantly over time. Commonly, as related work leads investigators to appreciate that range of phenomena that *might* be in play in the processes one studied, this affords a basis for appreciating error possibilities. For example, today one would notice two forms of acceleration in play in Galileo's famous inclined-plane experiment: the acceleration of a falling body (which the apparatus was designed to slow down and make observable to Galileo) and the angular acceleration of a rolling ball. Galileo was arguably in no position to determine that his result would hold up were only the one and not the other form of acceleration in play. Similarly, earlier biologists may have noticed an association of traits in a population and investigated whether

one trait somehow caused the other; later, after coming to understand meiosis, investigators may have investigated whether the separate genetic bases of both traits just happened to lay close together on some chromosome (so that neither trait need have caused the other); while subsequent investigators, with an understanding of how further processes may regulate gene expression, may investigate whether some process triggering the expression of the gene for one trait also is involved in regulating the expression of the gene associated with the other trait (as such processes could make for spurious causal connections between the two traits). In some of these cases, earlier investigators may not be in a position to even formulate or frame the error possibility that concerns later investigators.

With regard to source communities associated with fields that are undergoing developments of the above sorts, it will be rare that one would even think to depend on the results of investigators working significantly earlier. Prompted by recent developments in the field and the new issues that they suggest, more recent investigators typically would have re-examined the earlier results using more powerful tools. One tends to take one's certification from these up-to-date sources. (Thus, while one may admire Galileo's work, honestly, one does not flatly rely on his results.) Over time, communities of investigators have amassed vastly more nuanced data using significantly more sensitive devices or probes. It then seems strange to seek to draw on resources from communities working significantly earlier. If one were seeking to keep epistemic gate for this later community, in the sense of determining on whose results one in that community can flatly rely, and were someone then to urge one to draw on the results of the significantly earlier community, one would resist – insisting that those earlier researchers did not there know. But, typically, one does not look to the earlier community as a source on the matters of concern, and one does not feel the need to keep epistemic gate at this remove.⁴ Instead, one automatically looks to the results of more recent source communities whose inquiry is structured so as to deal with what may well be a more nuanced set of error possibilities, and whose results possibly benefit from more sensitive measurement.

There would seem to be some limit to the range of communities that might

reasonably expect to be able to draw on a given source community's results? It then seems fitting that one's contextual certification would be understood in terms of the epistemic needs of this restricted range of communities – fitting standards would be appropriate to informational needs of this restricted audience.

One might try to capture the limits on the reasonable audience of a source community this way. Let us say that a source community G^{early} counts as *significantly earlier* than another source community G^{late} when in the interval there has been work by some intermediate source communities $G^{\text{mediate1}}, \dots, G^{\text{mediate2}}$ that has put them in an epistemic situation with respect to the matter in question that is palpably superior to the epistemic position of G^{early} on the matter. Intermediate source communities that are not rendered “significantly earlier” by yet other intermediates may be termed *up to date* source communities.⁵ We can then say that a source community would be unreasonable to flatly draw on the results of a significantly earlier source community. What is more, it would be unreasonable for source communities not to notice that there is some reasonable “shelf-life” to their results. That is, in any developing field, one should expect that there will come a point at which source communities will not reasonably seek to flatly draw on one's results – because those results will no longer be up to date. Communities of inquirers can reasonably draw only on what are for them more up to date communities. This points in the direction of a refined way of thinking about the role of source communities, and thus of the reasonable standards for contextually knowing there.

In using the concept of *knowledge* to certify sources, thereby regulating our epistemic community life, one should be sensitive to the plausible shelf-life of one's and others' results. Some results are known in the sense that they satisfy the contextual standards for a source community, then are superseded by those of significantly later source communities. When gate-keeping in the context of a source community, one may speak the truth in saying that s knows that p – thereby signaling that various groups can draw on the result that p . This would obtain because s 's epistemic position (or that of s 's epistemic community) is sufficient to tell in favor of p against the range of relevant error possibilities fitting to *the*

practical and source communities that are not significantly later than s. The contextual aspiration of a source community is to play this role. This is a more reasonable aspiration for a general-purpose source community, and the contextual application of the concept of *knowledge* should enable us to regulate our epistemic life together in a manner that is fitting to this aspiration.

The present suggestion represents a well-motivated refinement on the earlier treatment of contextual knowing in a source community. However, it does not fully deal with the skeptical threat. Suppose that G^{early} produces the result that p , and asserts that they know that p . Then suppose that there are subsequent developments in related fields that suggest error possibilities with which G^{early} is not epistemically positioned to deal. At this juncture, suppose that some source community $G^{\text{transition}}$ entertains drawing on G^{early} 's result, but notices that G^{early} is not in a position to deal with the new error possibilities. Suppose also that at this point, there is yet no intermediate community that is better placed on the matter than is G^{early} . On the account as developed to this point, $G^{\text{transition}}$ would be right to deny that G^{early} knows. Further, it would seem that there would need to be such cases at the margins of the shelf-life of all knowledge attributions in developing fields – otherwise there would seem to be no reason for communities to do the work that makes them better placed than the relevant G^{early} . It is the judgment that G^{early} does not know that reopens inquiry on the matter.

If there are always or often such cases at the margins of the shelf-life of source community results, and if they involve *corrections* of earlier knowledge claims or suppositions, then, again, one seldom knows anything in the context of a general-purpose source community. We have yet to dispense with the skeptical threat.

For the gate-keeping contextualist, the question should be whether or not the judgment by $G^{\text{transition}}$ that G^{early} does not know is best understood as *in the same or commensurable context* as G^{early} 's earlier claim to know. If the context is the same or commensurable, then this denial would seem to conflict with, and correct, that earlier attribution. Skepticism then results.

What we are looking for is a notion of commensurability of contexts that is

fitting to the social-epistemic role of general-purpose source communities, to their reasonable contextual aspirations and their reasonable audiences.

Suppose that a community G^1 claims to know that p . Suppose that some subsequent community G^2 notes that G^1 was in no epistemic position to deal with an error possibility that reasonably concerns them in light of subsequent developments. Now, if the G^2 is among the reasonable audience contextually supposed in the earlier self-attribution of knowledge within G^1 , then (a) the error possibility of reasonable concern to G^2 are among those that are contextually relevant to the semantics of the knowledge ascription in G^1 , and (b) the denial of knowledge within the context of G^2 at least counts as occurring in a commensurable context, and thus, (c) conflicts with and corrects the earlier ascription. Much then turns on the social-epistemic role and aspirations of source communities. The standards fitting to source-community knowledge ascriptions are to be understood accordingly. In keeping epistemic gate in the context of general-purpose source community, we do not want to demand so much that the community would never or very seldom produce results on which others may draw. While recognizing that there is a shelf-life to source community results, we want such communities to keep a stock of such results on hand. We depend on them, and that is an epistemic fact of life to which standards of knowing must be contextually accommodated.

Here is a simple and attractive idea: we should think of the reasonable aspirations, social-epistemic role, and fitting contextual gate-keeping for a source community as each conditioned on what are reasonable expectations for development in the field. One commonly has reasonable, if somewhat vague, expectations for when one's field and related fields will easily have developed so as to reveal error possibilities that could not be anticipated at (that) present. It also seems reasonable for those in a source community to aspire to produce results that are adequate to guard against the union of reasonable audience error possibilities *for roughly that expected duration*. The certification of results in attributing knowledge could be understood as commensurate to such expectations and aspirations. This is a good way to understand the idea of a reasonable audience comprised of those practical and source communities that might reasonably expect

to draw on one's results. To know that p in the context of a source community then requires one to be epistemically positioned with respect to p so that one can deal with such error possibilities as are relevant to contemporary communities and to future communities within the reasonably expected shelf-life of such results – where the reasonably expected shelf-life of such results is a function of the projected rate of development within the fields or communities in question.

The reasonable expectation and aspiration for the shelf-life of one's results in a source community is a function of informed expectations for developments in the field and related fields. It is a function of expectations for development that could easily indicate error possibility that cannot readily be envisioned at the time. What error possibilities are relevantly envisionable at the time is a matter of open lines of research and theoretical discussion in a field (and related fields) at the time. If an error possibility is envisionable at a time, then that error possibility is arguably one that one must be in a position to deal with to count as then knowing. At any given time, there will be a sense among practitioners for open questions, for what would be the possible significance of results within ongoing and envisioned research. Perhaps there is research underway into some matter, and were the research to answer particular questions certain ways, it would indicate error possibilities relevant to one's own work. (Just as the research demonstrating a significant placebo effect indicated an error possibility plaguing earlier research.) These *open* questions and issues – the ones suggesting error possibilities with respect to the one's own work – presumably amount to then readily anticipated error possibilities – although they are not error possibilities that one is then in a position to deal with (if they depend on open lines of research). This prevents one from knowing. In contrast, I suppose that (for example) there are lines of research underway in contemporary molecular biology that indicate error possibilities that simply could not be envisioned by biologists a half century earlier. These error possibilities could not be anticipated by the earlier investigators. When they became a motivated glimmer in the eyes of the later researchers, that marks a development that constitutes the actual end of the shelf-life of the earlier results (if it had not come before in some parallel fashion). In attributing knowledge in the context of a source

community, one is saying that the agent or agents who are the subjects of the attribution are epistemically positioned to deal with error possibilities within what one reasonably projects to be the shelf-life of such results. This projected or anticipated shelf-life sets the truth-conditions for the attribution, because it reflects the role of the source community for which one is gate-keeping.

I cannot pretend to precision here, and suspect that precision would be inappropriate. But the above seems a plausible, promising way to think about knowledge attribution in the context of general-purpose source communities. This allows attributions of knowledge to do important work in keeping epistemic gate for general-purpose source communities – work fitting to a reasonable understanding of the social-epistemic role of these communities.

ix. Conclusion

In this paper I have explored how the concept of *knowledge* would work if it were largely developed in the service of a central point or purpose – that of keeping epistemic gate for various practical and inquiring communities. I suggest that it would naturally have a place in a set of contextually sensitive attributions. One set of these should be sensitive to the epistemic needs of practical communities, and would look a lot like attributor contextualism with standards varying with stakes of the contextually salient practical communities. In another set of these contexts, one should be sensitive to the aspirations of and for general-purpose source communities. In such contexts, the judgments should naturally approximate the attributions envisioned by classical or insensitive invariantists. I have worked to understand the role of general-purpose source communities so as to achieve a more nuanced account of the aspirations and standards that are there contextually fitting.

My goal has been to clarify and develop one promising account, an account that ultimately may prove to be the superior explanation of the workings of the concept of *knowledge*. Ultimately, alternative accounts of the concept of *knowledge* must be judged abductively. This requires that they each be developed to put their best face forward.

References

- Craig, E.** 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. 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.
- Greco, J.** 2008. "What Is Wrong with Contextualism?" *Philosophical Quarterly* 58: 416–36.
- Greco, J.** 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Grimm S.** 2009. "Knowledge, Practical Interests, and Rising Tides." Unpublished manuscript.
- Henderson, D.** 2009. "Motivated Contextualism." *Philosophical Studies* 12: 119–31.
- Henderson, D.** "What's the Point?" Unpublished manuscript.
- Henderson, D. and T. Horgan.** 2001. "The A Priori Isn't All It Is Cracked Up to Be, But It Is Something." *Philosophical Topics* 29: 219–50.
- Henderson, D. and T. Horgan.** Forthcoming. *The Epistemological Spectrum*. Oxford: Oxford University Press.
- Kappel, K.** 2010. "On Saying That Someone Knows: Themes from Craig." In A. Haddock, A. Millar, and D. H. Pritchard (eds.), *Social Epistemology*. Oxford: Oxford University Press.

¹ The flexibility in contextual concern is strongly suggested in (Greco 2008) and (Henderson 2009). It is also reflected in passages in (DeRose 2009).

² There can be mixed cases. For example, an investigator from a general science (perhaps molecular biology) can be engaged with a research community seeking a treatment for some disease. A certain result may be of significance both for the general science of molecular biology and for the prospects of treating the condition.

³ Greco (2008, 2010) suggests what are, in effect, constraints on reasonable contextualist gate-keeping, constraints that derive from an understanding of the point or purpose of the concept of *knowledge* that is similar to my own. On his account as well as mine, a proper understanding of *knowledge* will reveal standards that are neither too lax nor too demanding for purposes of the epistemic projects regulated (2008, particularly p. 429).

⁴ I would suggest that in application to historical cases, one commonly uses the *knowledge* concept in a kind of *simulated* gate-keeping that facilitates historical reconstruction (see Henderson [manuscript]).

⁵ I have been writing of earlier and later communities, but under one familiar way of identifying and individuating such things, communities may have long existences, and may be significantly different at different times. One might, for example, think of

the community of medical researchers as one entity that has persisted and developed for over a century. One then notes that the relevant alternatives envisioned and dealt with by medical researchers changed significantly between 1930, for example, and 1980. For my purposes here, what are strictly speaking at issue are temporal stages of this one source community. Relative to the community of medical researchers in the 1980s, the community in 1930 is clearly a significantly earlier community (or stage of a community).