

# MOTIVATED REFINEMENTS TO MOTIVATED CONTEXTUALISM

David Henderson

VERY PRELIMINARY DRAFT—DO NOT QUOTE

APOLOGIES FOR INELEGANT FORMULATIONS AT THIS STAGE.

## *i. Introduction*

In this paper I argue for two refinements in what I term motivated contextualism, a contextualist account of knowledge that affords a principled treatment of standards and contexts. The principled motivation of context derives from a sense for the point and purpose of the concept of KNOWLEDGE—attributions of knowledge serve in a kind of epistemic gate-keeping for various communities of inquirers and communities of practical reasoners. In attributing knowledge, the attributor recommends a source to a contextually understood community, and the standards of knowledge should then be fitting in view of the cognitive and practical stakes of that community. This sense for the point and purpose of the concept of KNOWLEDGE provides reasons to prefer a form of contextualism over subject sensitive invariantism.

In Henderson (2009) I considered the contextual concerns of two broad classes of communities—each of which need sources of actionable information (and sometimes look to serve as a source of actionable information). 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, one finds that what count as knowing in the context of such a community would be sensitive to the stakes faced by that community. In this way, a motivated contextualism can account for the connections between knowledge and stakes, between knowledge and practical reasoning. Motivated contextualism thus proves nimble and powerful in treating of matters of concern to subject sensitive invariantists. Further strengths emerge when reflecting on the concerns of communities of inquirers who seek to *possess and provide* information on which a somewhat indeterminate range of others can reasonably draw—call these *general-purpose source communities*. I think of many

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 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.

Here I refine my understanding of motivated contextualism in two ways. First, I correct a mistake. My initial development of the view did not adequately limit the epistemic demands on general-purpose source communities. As a result, it opened the door to a form of skepticism that had no place in the general motivated contextualist approach. Thus a central concern of this paper is to refine our understanding of general-purpose source communities. Second, it should be obvious that there are plenty of cases in which folk attribute knowledge without such attributions serving to keep epistemic gate in any plausible sense. I suggest a way of accommodating this observation within my motivated contextualism.

*ii. A sense for the point or purposes served by a concept (particularly an evaluative concept) can and should constrain one's philosophical account of the concept.*

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. Often, plausibly, the concept was developed with an eye to, or sensitivity to, the needs of that project. The main outlines of the concept would then be subservient to the certain needs arising in the pursuit of the project. Individuals might develop particular understandings—conceptions—in connection with the concept. But these would not be definitory of the concept. For example, the concept of a GOOD DRIVER presumably arose to regulate the operation of motor vehicles (paradigmatically on public roadways) and is sensitive to the needs for being predictable by others, for mitigation of risks to self and others posed by this technology, for expeditious arrivals of people and cargo, and the like. At the same

time, those possessing this concept have some more or less developed conception of the character traits, capacities, and skills (or perhaps of various alternative constellations of character traits and skills) that make for being a good driver. However, were one were to find that some character trait one had thought was needed in a good driver, as a matter of fact, did not contribute to the project in question, then one would presumably amend one's conception of a good driver. In doing so, one would not have changed the subject—rather one would have refined one's understanding of what makes for good drivers. Before and after the change in one's conception, one would be employing the same evaluative concept of a GOOD DRIVER. If this much is correct, it suggests a very useful idea for philosophers: In analyzing or reflecting an evaluative concept, it will be important to keep in mind one's sense for the project in which the concept has its roots—how the concept serves to regulate an ongoing project that its users care about. This is to say, our philosophical reflection should commonly be informed and constrained by our sense for the point and purpose of the concept we seek to better understand. What is central to the concept, what is conceptually required rather than merely an element of one's present conception, will plausibly be closely related to the regulative role of the concept in the human projects in which it is rooted. (See Henderson and Horgan, 2001, forthcoming.)

A good/competent Xer—embedded in a practice. One can be confident that the concept is not such as would actually be antithetical to the practice. ...

These points will be particularly significant in section ix, in which I argue that a skeptical result would be highly implausible in connection with any understanding of the concept of knowledge that sees that concept as rooted in the regulation of sources of information to satisfy the various epistemic of communities.

*iii. A central point of purpose of the concept of KNOWLEDGE is keeping epistemic gate for various communities who need sources of actionable information.*

In “Motivated Contextualism”, I sought to honor the first general idea in sketching an account of the concept of KNOWLEDGE.<sup>1</sup> In particular, I sought to draw on

Craig's central idea that the (or at least a very significant) point of the concept of KNOWLEDGE is to certify good sources of actionable information—information on which one can reasonably or rationally rely in the course of one's practical and theoretical reasoning. Examples are everywhere. In judging that my mechanic knows about maintenance schedules and expected repair cycles, I feel free to “take it from him” that my car's timing belt is reaching the point at which it can no longer be confidently relied on not to break—and I rely on this information when I decide to replace the belt. The auto company's engineers relied on information about the properties of various materials when deciding to use one rather than another material in making certain parts—and in doing so, they relied on various material scientists as sources. One relies on one's mechanic, and the engineer relies on the material scientist, as sources of information in practical reasoning—and counts them as being in the know. Each of us keeps track of who would know of various matters, and we regulate the sources on which we draw (in our various capacities as members of various practically engaged communities). An experimental physicist may use information about various materials and processes when designing a particle detector. This is to rely on information in reasoning that is at once both theoretical and practical.<sup>2</sup> The theoretician may rely on results to inform their theoretical choices, and this is pretty clearly a part of their theoretical reasoning. As the various agents (the auto owner, the engineer, the experimentalist, the theoretical physicists) engage in their various practical and theoretical projects, they signal or certify good sources of information by attributing (or withholding attributions of) knowledge.

Kappel **O** usefully puts the point in terms of the regulation of the inquiry: attributions of knowledge signal when inquiry has proceeded far enough, withholding such attributions signals when a potential source needs to inquiry further or when we need to look for other sources. One works to put oneself in a good epistemic situation with respect to certain matters—one engages in inquiry. When one undertakes the inquiry oneself, one needs to gauge when one can rely on the results in one's theoretical or practical endeavors—when one can rest content

with one's inquiry and deploy one's results in one's own reasoning. Thus, to judge that one now knows on the basis of one's own inquiry is to judge that one can call an end to that inquiry because one now has provisioned oneself with actionable information. Others may have use for that same information. To attest that one knows is to certify oneself as a source of such information to an audience (an audience that may itself be representative of some wider community). To say that someone else knows is to certify the other as a source—indicating that those to whom the source is certified can freely draw on that source on that matter.

*iv. Various applied practical communities, various source-communities, and their needs for epistemic gate-keeping.*

It is very helpful to think of the role of attributions of (or denials of) knowledge in the certification (or decertification) of sources as a kind of epistemic “gate-keeping” for variously engaged communities. This invites us to think of the needs of various kinds of communities. Sometimes, we keep epistemic gate for communities engaged in some practical endeavor—certifying to, or for, them certain epistemically well-placed folk as sources on which they can freely draw. (The practical communities may be ad hoc—as when an set of folk just happen to come to confront some practical problem. For example, some folk may assemble in response to seeing a person fall in some public space. The practical communities also may be enduring—as when some group is assembled and maintained to be ready to confront a class of practical problems. For example, the staff of an emergency room at a hospital stand ready to deal with the range of health problems presented to them.) I call these *applied practical communities*. Sometimes we keep gate for communities of inquirers who are engaged with no definite class of practical problems. Some communities of inquirers are understood as having a social role of producing information of such a high epistemic quality that a somewhat indeterminate range of groups might draw on their results without hesitation. I call these communities, *general-purpose source communities*. Obviously, we are all part of various applied practical communities, some enduring, some ad hoc and relatively ephemeral. Some of us are part of certain general-purpose source communities.

There is room for various intermediate types of communities (consider, for example, a group of doctors, technicians, and nurses, at a research hospital engaged with a particular patient's case, and with an experimental treatment regimen that itself tests recent theories in cell science). When one thinks of our epistemic life together, it is easy to see that we clearly need to regulate our epistemic communities, and that we clearly engage in a lot of gate-keeping—and it is easy to appreciate why such social-cognitive critters as ourselves would develop the concept of KNOWLEDGE in answer to such needs.

Anyone competent with the concept of KNOWLEDGE recognizes how one fittingly responds to the attributions of knowledge, and to their denials. One has some sense for what the community of reasoners ought to do in response to attributions of knowledge, and in response to the denial that someone knows. It is this that I have in mind when I write of folk having a sense for the point or purpose of the concept—we have some understand of how it functions in regulating our cognitive life individually and socially. This should help constrain the philosophical account that we give of the concept of KNOWLEDGE. In the preceding paragraphs, I have provided an organizing idea—that of epistemic gate-keeping—with which to think about the purposes served by the concept of KNOWLEDGE.

*v. Whether an epistemic agent knows that  $p$  may depend on more than the agent's strength of epistemic position with respect to  $p$ —it may depend on "stakes."*

Thinking along the lines indicated by these first three big ideas has significant implications for the debates between contextualists, subject sensitive invariants (SSI), and classical or insensitive invariantists. But, before getting to these, we should pause over a point on which SSI and contextualism has some rudimentary agreement.

Both contextualists and SSI insist that what is required for the truth of knowledge ascriptions can turn on more than the quality of epistemic position of the agent said to know.<sup>3</sup> Thus, suppose that we have an agent whose epistemic position with respect to some proposition,  $p$ , does not change in the interval  $t$  through  $t^*$ .

According to the adherent of SSI,<sup>4</sup> it is possible for the agent's practical situation to change across this interval, even though the agent's epistemic position with respect to  $p$  remains unchanged. With certain changes in the agent's practical situation, it may be that the agent can truly be said to know that  $p$  at one time and not at another. For example, at  $t$ , the agent may be in a situation (Low) where relatively little hangs on whether or not  $p$ —so that the agent's epistemic situation makes it reasonable to act on his or her information to the effect that  $p$ —while at a later time,  $t^*$ , the agent may be in a situation (High) in which relative much hangs on whether or not  $p$ , and it may not be reasonable to act on the same information. According to SSI, the agent knows in Low and not in High. Thus, according to SSI, the truth of the attribution of knowledge that  $p$  to an agent at  $t$  turns both on the strength of the agent's epistemic position at  $t$  with respect to  $p$  and on the stakes faced by the agent at  $t$ .

For the contextualist,<sup>5</sup> stakes also may be pivotal, but enter in a somewhat different way. Think of an agent having a particular epistemic position with respect to  $p$  at a given time,  $t$ . The contextualist insists that, in one context of attribution, that agent be truly said to know that  $p$  at  $t$ , while, in another context of attribution, that same agent may not be truly said to know  $p$  at  $t$ . It is common to insist that in contexts in which the attributor and the attributor's audience face a high stakes situation, they may rightly withhold certification. That is, they may rightly deny that the agent knows  $p$ . In contrast, when they face a low stakes situation, the agent's epistemic position with respect to  $p$  may be sufficiently strong to warrant the attribution of knowledge. For the contextualist, the content of an attribution of knowledge turns on context, which is at least partially constituted by the stakes facing the parties to a conversation—or, more precisely, by stakes facing folk who are contextually in focus as those to whom the agent is ultimately certified as a source.<sup>6</sup> 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.

Thus, SSI and contextualism agree that the truth of knowledge attributions may turn on the stakes faced—although the two positions diverge on just whose stakes matter. For SSI, it is the stakes faced by the agent at  $t$  that determine what it takes for the agent to know at  $t$ . For the contextualism, it is the stakes faced by the attributor and relevant conversational partners at the time of the attribution that determine how good an epistemic position the agent needs to be in to count as knowing.

*vi. Gate-keeping, engaged practical communities, and general support for contextualism over SSI*

A central claim in Henderson (2009) was that the gate-keeping point of knowledge attributions affords significant support for contextualism over SSI:

Contextualists treat attributions of knowledge as understandable in terms of the contextual concerns of attributor and interlocutors. 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 it would be attuned to the interests of an epistemic agent who might have interests rather different from those for whom epistemic gate is kept (Henderson 2009, p. 120).

The point is most telling when thinking about gate-keeping for applied practical communities—where both my gate-keeping contextualism and SSI find that is required for knowing can vary with the stakes confronted. However, SSI provides a form of sensitivity to stakes that makes little 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.

When one thinks about keeping epistemic gate on behalf of some practical community, when one seeks regulate epistemic sources so as to meet the

informational needs of various practical communities, there are a number of things to keep distinct. Schematic representation is helpful here. Consider an attribution of knowledge to agent  $s$ , of proposition  $p$ , at time  $t$ , thus  $Ks,p,t$ . The attribution is made by an attributor,  $a$ , in a context in which  $a$  is certifying  $s$  as a source to a practical community  $P^{\dots}$ . Thus we can represent the context as  $\langle a, P^{\dots} \rangle$ . When listing practical communities, I will use superscripts without parentheses,  $P^1, P^2, \dots$ . When referring to their respective stakes, I will use parentheses as just now, or as in  $P^{(1)}, P^{(2)}, \dots$ . The practical community to whom a source is certified may be have high stakes related to the matter (thus,  $P^H$  with  $P^{(H)}$ ) or low states ( $P^{(L)}$ ). Thus, for the gate-keeping contextualist, the attribution of knowledge to  $s$  amounts to a contextual certification of  $s$  as a source to a practical community,  $P^{\dots}$ , and this can be represented:

[P]  $\langle a, P^{\dots} \rangle: Ks,p,t$

The attribution to a high-stakes practical community may be represented as:

[P<sup>H</sup>]  $\langle a, P^H \rangle: Ks,p,t$

The attribution to a low-stakes practical community is represented as:

[P<sup>L</sup>]  $\langle a, P^L \rangle: Ks,p,t$

According to the motivated contextualist, the truth of the attribution will turn on both the strength of  $s$ 's epistemic position with respect to  $p$  ( $E^{s,p}$ ) and on the states associated with the context, which seem fittingly determined by where the stakes of relevant practical community falls with the continuum,  $P^{(L)} \dots P^{(H)}$ . Thus, for the contextualist:

Truth  $[\langle a, P^{\dots} \rangle: Ks,p,t]$  is a function of  $E^{s,p}$  and  $P^{(\dots)}$

For the proponent of SSI, the truth of the assertion has nothing to do with the context of attribution  $\langle a, \dots \rangle$ . Instead, it turns on  $E^{s,p}$  and  $s$ 's stakes ( $P^{(s)}$ ). But, if gate-keeping is central to the point or purpose served by the concept of KNOWLEDGE, if attributions of knowledge centrally serve to regulate epistemic communities, the SSI approach faces an embarrassing question: If one is keeping epistemic gate for an epistemic community,  $P^{\dots}$ , which may face significantly different stakes than those faced by  $s$ , why should why should one's certification/attribution be attuned to  $s$ 's stakes,  $P^{(s)}$ . While the qualification of  $s$  as a source for  $P^{\dots}$  presumably is directly connected to  $E^{s,p}$ , and seems reasonably connected with  $P^{(\dots)}$ , it is at best

contingently connected to  $s$ 's stakes (for, except in special cases  $P^{(\dots)}$  and  $P^{(s)}$  can easily diverge and are only contingently related).

Of course, as  $s$  reflects and regulates  $s$ 's own practical reasoning,  $s$  should care whether  $E^{s,p}$  is adequate to  $P^{(s)}$ . But, on this point the gate-keeping contextualist will surely agree. In such cases of self-attribution one certifies oneself as an adequate source for oneself as a solo practical reasoner. When so engaged,  $s$  is in effect keeping gate for the limit-case practical community of comprised of just  $s$  (call this  $P_s$ ). Given that  $P_s$  just is the community for which  $s$  is contextually keeping epistemic gate, it is no mystery why any given agent,  $s$ , engaged in gate-keeping *for his or her own purposes*, should care whether  $\langle a, P_s \rangle: Ks,p,t$ . The agent cares for the same reason that members of any practical community,  $P^{\dots}$ , care whether for a prospective source,  $s$ ,  $\langle a, P^{\dots} \rangle: Ks,p,t$ . Such a source is epistemically well placed to provide the relevant information to the practical community in question. Accordingly, the gate-keeping contextualist can make perfectly good sense of this case in which SSI seems at its strongest.

Now consider the sort of case with which SSI seem problematic, one in which

1.  $a$  is asked by interlocutors from a practical community,  $P^H$ , facing high stakes,  $P^{(H)}$ , whether  $s$  knows that  $p$  (at  $t$ ), and
2. where  $s$  (at  $t$ ) faces a low stakes situation:  $P_s^{(L)}$ . (It follows that  $s \notin P^H$ ).

Keep in mind the central gate-keeping idea: in saying that  $s$  knows  $p$  (at  $t$ )  $a$  is certifying  $s$  (at  $t$ ) as a source of actionable information to or for the contextually understood community. In the case just stipulated, the attributor,  $a$ , is contextually keeping gate for a community,  $P^{\dots}$ , with  $P^{(H)}$ .<sup>7</sup> To make the point vivid, we may suppose that  $E^{s,p}$  is just sufficient for  $s$ 's own purposes, given  $P_s^{(L)}$ .<sup>8</sup> Those in  $P^H$ , should not think that  $s$  is an adequate source with respect to  $p$ . Further, if one is contextually gate-keeping for  $P^H$ , one should be contextually attentive to their informational needs. Thus, in  $\langle a, P^H \rangle$ , the stakes of  $P^H$  are of concern. For any agent,  $s$ , at  $t$ , and any proposition relevant to the agent's deliberations, that agent rationally cares whether it is true that  $\langle a, P_s \rangle: Ks,p,t$ .

The subject sensitive invariant idea is that what it takes to know is

conditioned by the stakes of the subject who whom knowledge is attributed. Thus, the truth conditions envisioned in SSI are those that would obtain on the motivated contextualist account just when the attributor was keeping epistemic gate for a community comprised of just that agent at that time—namely:  $\langle a, P_s \rangle: Ks, p, t$ . But, this makes it mysterious why folk facing stakes very different from those faced by  $s$  at  $t$  should care whether  $s$  knows, as that is understood by the proponent of SSI. The gate-keeping contextualist idea that what it takes to know is conditioned by the stakes of those for whom one is keeping epistemic gate makes it easy to understand why any community to whom an agent is fittingly certified as knowing should care about the truth of that attribution (see also Henderson 2009, pp. 124-5).

The point should be clear. In epistemic gate-keeping for engaged practical communities, in certifying sources adequate for their practical reasoning, one should be sensitive to their stakes. Insofar as the stakes of the prospective source at a given time is only contingently related to those of the community of concern, one should attune one's certification to the stakes faced by the community, not to the stakes faced by the agent.

The general point is telling. However, there seems to be one very big reason to proceed cautiously. While the gate-keeping may be a central purpose served by the concept of KNOWLEDGE, and while this may provide significant motivation for a kind of contextualism, there are *prima facie* reasons to doubt that epistemic gate-keeping is *the* point or *the* purpose served by the concept—rather than *a central* point or purpose. There are certainly attributions of knowledge in which one is not keeping epistemic gate in any significant sense. Thus, one must admit that the concept might be functionally diverse. We may serve various purposes in deploying the concept of KNOWLEDGE, and the concept may be rooting in these multiple points or purposes. Our concept of KNOWLEDGE might be a *motivational motley*. This misgiving will be addressed in the final section of this paper. However, I do not there attempt to decisively put the misgiving to rest. Rather I seek to show that at least many cases in which one is not gate-keeping when deploying the concept of KNOWLEDGE can readily be understood as deploying a concept that is yet centrally rooted in keeping epistemic gate for communities.

*vii. Gate-keeping contextualism and general-purpose source communities.*

Sometimes, while not engaged in some practical project for which  $p$  might matter, one is yet interested in determining whether or not  $p$ . Sometimes, when claiming to know that  $p$  (or when claiming that someone else knows that  $p$ ), one is not certifying oneself (or that other) to those who are engaged as members of an applied practical community. Indeed, commonly, inquiry is undertaken not to answer the informational needs of any determinate practical community, but to lay in a stock of information that an indeterminate range of folk can draw on in their various practical or other endeavors. One ultimately seeks to be a source for “all comers.” Various scientific disciplines commonly undertake to serve this role. It is the role of a general-purpose source community. When one thinks of well established scientific results as paradigmatic instances of knowledge (and I do so think of them) one cannot help but feel the pull of classical, insensitive, invariantism. However, it is one of the great strengths of gate-keeping contextualism that it can readily and naturally account for the judgments one makes in these cases.

Suppose, paradigmatically, that you and your interlocutors are members of a contemporary general scientific discipline which provides a body of results on which people with an indeterminate range of practical projects might draw while the discipline itself is not be associated with any one practical project in a defining or limiting way. (Think of the indeterminate range of practical projects that have and can have recourse to the results of some contemporary general scientific discipline.) When one is so engaged, one’s concerns revolve around providing general-purpose actionable information. It is then fitting that, when keeping epistemic gate for a general source community, 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 adequate to the purposes and stakes of a somewhat indeterminate range of applied practical communities. One should also be concerned with sources adequate to other source communities that might draw on one’s results in their investigations. One should be contextually

invested with the general role of the source community.

Suppose then that you are associated with some general-purpose source community, and that you are charged with evaluating some putative results gotten by an aspiring member of your source community. (Perhaps you are evaluating an article announcing some potentially significant result.) As you sit with your colleagues, you are asked: “Well, does she know that  $p$ ? Can we say that we now know that  $p$ ?” You, or another colleague may assert, “On the basis of this work, she (and we) now know that  $p$ .” The attributor and interlocutors may be little concerned with any concrete practical project, but focused instead with something like quality control in our source discipline. The attributor and interlocutors are then engaged in distinguishing worthy contributions—ones that are fitting to serve as actionable information for a more or less indeterminate range of others who might seek to draw on them. This kind of general certification, common in the connection with expert or scientific knowledge, divorces such attributor/interlocutor contexts from the concrete and particular extant concerns of the attributor and any narrow set of interlocutors. The concern is then with a very high quality of epistemic position on the part of the epistemic agent’s to be certified—one sufficient to allow the belief as produced and held by the epistemic agent to fittingly serve as a resource on which folk within other groups or communities can readily draw in their various practical or epistemic projects.

Thus, to qualify as knowing that  $p$  (in 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. Thus, for purposes of such a source community, one should keep epistemic gate so that the information allowed into its stock of established results, the information that its certifies as knowledge to “all comers,” is sourced so that it can fittingly serve as actionable information even for the more demanding of the applied communities that might draw on their work. Of course, other source communities may also want to draw on their results. (Think, for example, of the range of disciplines on which a discipline such as geology or ecology commonly draw.)

To a first approximation:

(In a source community)  $K_{s,p,t}$  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).

Building on our earlier notation, let  $P^{\dots}$  be some practical community, and  $G^{\dots}$  be some general-purpose source community, we could represent the context of an attributor,  $a$ , who is contextually keeping epistemic gate for a source community,  $G^c$ , as  $\langle a, G^c \rangle$ . Then we might represent the idea that knowledge attribution in the context of such a source community amounts to a kind of certification implicitly to the range of groups that might draw on the results of such a source community by crudely mentioning the communities that are contextually relevant consumers:  $\langle a, G^c \{P^1, P^2, \dots, P^n, G^c, G^1, G^2, \dots, G^n\} \rangle$ . Here,  $G^a$  represents the source community for which  $a$  is contextually keeping epistemic gate. The various communities in the inner brackets are those for which  $G^a$  might plausibly serve as a source. The "n" of this list plausibly will typically be quite large and somewhat vague. An attribution of knowledge to a general-purpose source community then is represented:

[G]  $\langle a, G^c \{P^1, P^2, \dots, P^n, G^c, G^1, G^2, \dots, G^n\} \rangle : K_{s,p,t}$

We can then compare [G] with the earlier gate-keeping contextualist schema for practical communities [P]  $\langle a, P^{\dots} \rangle : K_{s,p,t}$ . According to gate-keeping contextualism, the truth conditions for an attribution of the character [P] depend on the stakes of the contextually relevant practical community,  $P^{\dots}$ , for whom epistemic gate is kept. For  $s$  to count as knowing,  $E^{s,p}$  must be sufficient to discriminate  $p$ 's obtaining from that of the range of alternatives that are fitting and significant in light of  $P^{\dots}$ . This range of alternatives will be greater when  $P^{(H)}$ . Such variation with stakes disappears when keeping gate for a source community, because one then contextually aspires to take care of all  $P^{\dots}$  that might reasonably draw on source community results. Thus, gate-keeping for the general-purpose source community provides for no such variation.<sup>9</sup> The schema [G], indicates that  $E^{s,p}$  must be adequate

to discriminate  $p$ 's obtaining rather than any of the alternatives that would be relevant *given any of the stakes,  $P^{(1)}, P^{(2)}, \dots, P^{(n)}, G^c, G^1, G^2, \dots, G^n$* , fitting to the various applied or source communities that might draw on  $s$ 's result that  $p$ . The variation with stakes is washed away in the union of relevant alternatives fitting to various stakes. Or at least that is a decent, schematic, somewhat idealized, first approximation.

This points to a way in which the gate-keeping contextualist line of thought can naturally come to resemble a kind of insensitive invariantism. At least with respect to these paradigmatic cases of knowledge, general-purpose source communities, fitting epistemic gate-keeping yields a largely invariant standard in play, a high standard.

Several points warrant special mention.

First, note that gate-keeping contextualism only approximates insensitive invariantism in those contexts where one is keeping gate for a general-purpose source community. For such communities, the stakes of relevance, and thus the relative alternatives do not fluctuate, at least at reasonable time scales in which later investigators might seek to draw on the results of earlier research. (See section xi.)

Secondly, there is a matter that requires care in any contextualist framework: the matter of when one attributing or denying knowledge in one context can be said to disagree with one in another context. We can best approach this matter after developing the idea of a *context type*.

Schemata [G] is motivated by the twin ideas, (a) that there are general-purpose source communities whose role is to serve as sources on which many other practical and source communities may reasonably and confidently draw, and (b) that in attributing knowledge to  $s$ , one is certifying  $s$  as a good source, and thus keeping epistemic gate for the relevant community. The implication is that, in keeping gate for a general-purpose source community, one is keeping gate for a community that is to serve as a source for any community that might be interested in the matters on which it focuses. So anyone attributing knowledge in such a context should be responsive to such demands. Anyone. Thus, the identity of the attributor does not matter.<sup>10</sup> What matters is the source community for which gate

is kept. Different attributors, or the same attributor at different times, will be attributing the same knowledge-ish property to the agent in question as long as they are keeping gate for the one and the same source community.<sup>11</sup> We might say that, for any given source community that is the audience for a knowledge attribution,  $G^c$ , we have a *context-type*,  $\langle \dots, G^c \rangle$ , and that this context type determines the truth conditions of the knowledge-attribution. It works like this: 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, and it serves to supply information adequate to the purposes of “all comers” within this reasonable audience. To do this, it must be that, as its members contribute come result,  $p$ , to its stock of results, the contributing member,  $s$ , must be in an epistemic position 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. If and only if  $s$  is so epistemically positioned does  $s$ 's result count as knowledge for purposes of the source community  $G^c$ . Thus it is that all (source community directed) attributions of a given context-type have the same truth conditions.<sup>12</sup> All attributions, by any attributor, that serve to keep gate for a given source community have should be attentive to the informational needs of the same reasonable audience.

Attributions of knowledge that  $p$ , in context-type  $\langle \dots, G^c \rangle$ , are commensurable. Thus, suppose that one attributor,  $a^1$ , at  $t^1$ , certifies in  $\langle \dots, G^c \rangle$  that  $s$  knows  $p$  at  $t$ . That is:

$$\langle a^1, G^c \{P^1, P^2, \dots, P^n, G^c, G^1, G^2, \dots, G^n\} \rangle : Ks, p, t$$

Suppose also that another attributor,  $a^2$ , at  $t^2$ , denies in  $\langle \dots, G^c \rangle$  that  $s$  knows  $p$  at  $t$ .

That is:

$$\langle a^2, G^c \{P^1, P^2, \dots, P^n, G^c, G^1, G^2, \dots, G^n\} \rangle : \sim Ks, p, t$$

Then, they disagree. Perhaps  $t^2$  is quite a bit later than  $t^1$ . It does not matter, so long as the source community, and associated reasonable audience, are the same. Thus, at least with contextually gate-keeping for general-purpose source communities,

gate-keeping contextualism entails the kind of invariance in truth conditions that the insensitive invariantist would want.

For the gate-keeping contextualist, there is a place for disagreement across some contexts or context types. A prominent example would turn on the relation between a general-purpose source community and any community that is within its reasonable audience. Suppose that  $P^k$  is among the reasonable audience for  $G^c$  and (as above)  $a^1$  asserts in  $\langle a^2, G^c \rangle$  that  $Ks, p, t$ . As explained, for this to be true,  $E^{s,p}$  must be sufficient to rule out all error possibilities within the union of the relevant error possibilities appropriate to its various reasonable audience. Now suppose that the folk in  $P^k$  come to be concerned with whether or not  $p$ . Given their stakes, they will reasonably be concerned with some range of relevant error possibilities. They may learn of  $s$ 's epistemic position with respect to  $p$ . They may then realize that  $E^{s,p}$  is simply not sufficient to eliminate some of the error possibilities that seem to them fittingly relevant. Suppose that one of them,  $a^3$ , then denies that  $s$  knows that  $p$ . Plausibly,  $a^3$  could reasonably be saying either:

$\langle a^3, P^k \rangle: \sim Ks, p, t$ ,

or

$\langle a^3, G^c \{P^1, P^2, \dots, P^n, G^c, G^1, G^2, \dots, G^n\} \rangle: \sim Ks, p, t$

depending on what is salient for  $a^3$ . If what is salient is that  $E^{s,p}$  is not adequate to the needs of  $P^k$ , then it seems right to think of  $a^3$  as keeping gate for  $P^k$ . If, instead, what is salient is that  $E^{s,p}$  is then not adequate for purposes of a source community on which those in  $P^k$  should be able to draw, then  $a^3$  is plausibly thought of as gate-keeping for  $G^c$ . There is here a somewhat limited form of commensurability. Clearly, a parallel, and perhaps somewhat stronger form of commensurability obtains when members of some source community seek to draw on (and thus evaluate) another source community.

*viii. A new skeptical worry?*

Responses of the sorts just characterized should seem familiar when thinking about general-purpose source communities such as scientific disciplines. When some expert,  $s$ , claims to know that  $p$ , and some epistemic agent 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 the agent's own project or inquiry, the agent judges that *s* does not know that *p*. Many of us find this judgment intuitively correct. This response is perfectly intelligible in terms of gate-keeping contextualism which, again, holds that in contexts in which the epistemic standing of agents within a source community is at issue, to say that *s* knows that *p* at *t* is to say that *s*'s epistemically position tells strongly in favor of *p*'s obtaining in contrast to any member of the set of relevant alternatives that is the union of the sets of relevant error possibilities that are fitting to the various applied or source communities that might draw on *s*'s result (the union set of relevant error possibilities).

The above indicates that gate-keeping for general purpose source communities reasonably gives rise to high standards that themselves are largely invariant—invariant within a context type that plausibly has many instances. Again, this strikes me as a good result, one showing that the contextualist can provide an epistemology “with real backbone” of the sort to which insensitive invariantists aspire. However, one may worry that the envisioned high standards, at least as delineated to this point, are *too* high. A lot depends on just how one understands the phrase, “the various applied or source communities that might draw on *s*'s result.” In this section I want to develop the worry that the demands of gate-keeping for general purpose source communities (as delineated above) would engender a disastrous form of skepticism. In the following section (ix), I explain why any such skepticism should be completely implausible within the framework of gate-keeping contextualism. Demands on knowing that would result in skepticism would be diagnostic of some mistake in one's understanding of proper epistemic gate-keeping. Then, in section (x) I provide a refined understanding of gate-keeping for general-purpose source communities.

Let's begin with some quick observations, intended to make clear how the skeptical threat might be thought to arise.

It should be evident that the set of communities that *might* draw on one's results is different from the set of communities that *actually do* seek to draw on the result. Otherwise a community that fancied itself a general source community on

some matter, could readily speak the truth in claiming to know on that matter so long as others ignored them or never sought to draw on their results. An ignored community would vacuously count as having evidence that handled all error possibilities of concern to those actually seeking to draw on its results. Thus, if the set of communities that might seek to draw on its results were just the set of communities that do draw on it, such an ignored community would vacuously be in an epistemic position adequate to handling all error possibilities relevant to all communities that might draw on it. This is obviously an undesirable result. Being totally ignored should not make it easier for folk posing as a source for others to satisfy the requirements on knowing.

The set of communities that *might* draw on one's results also is wider than the set of actual communities that might easily or feasibly draw on one's results. Perhaps there are actual communities that need to know whether or not  $p$ , but fail to entertain drawing on  $s$ 's putative result only because  $s$ 's result is not well publicized, or because there are significant limits on communication of sources. For example, censorship in certain nations can preclude the communication of one's results to communities within their boundaries, and thus precludes their citizens from entertaining drawing on one's results. It does not seem that this should make it easier to count as knowing (by restricting the set of communities whose relevant alternatives with respect to  $p$  contribute to the set of relevant alternatives relevant to source community knowing).

The set of communities that *might* draw on one's results is different from the set of *actual contemporaneous* communities that might draw on them. Perhaps as a result of the epistemic situation within some actual extant communities today,  $t^1$ , a new field of inquiry will soon form and be pursued by a disciplinary community,  $G^n$ . It will take shape over the next few years. Suppose that, at  $t^2$ , those working in this new field will entertain drawing on the result that  $p$ , which a contemporary source community,  $G^a$ , produced at  $t^1$ . If there are error possibilities that are reasonably relevant to those in  $G^n$ , error possibilities with which those in  $G^a$  are not in an epistemic position to deal, then it is right to say that those  $G^a$  do not know  $p$  at  $t^1$ .  $G^a$  sought to produce results adequate to "all comers," and  $G^n$  materializes as a relevant

consumer of results on the matter in question.

Finally, arguably, the set of communities that might draw on one's results includes some members that are not actual communities. To illustrate, consider a variant on the case just described. Suppose that, as a result of the epistemic situation within some actual extant communities today,  $t^1$ , there is the basis and need for a new field of inquiry. Were folk in the extant communities not too intellectually timid, professionally cowed, or otherwise inhibited, a new field of inquiry would form and be pursued by a disciplinary community,  $G^p$ . This new community would take shape over the next few years, but it does not. Suppose that, were this epistemic community to have formed, those working in this new field at  $t^2$  would reasonably wonder whether or not  $p$ , and would reasonably entertain drawing on the result that  $p$  which an actual source community,  $G^a$ , had produced at  $t^1$ . If there are error possibilities that reasonably would be relevant to those in  $G^p$ , error possibilities with which those in  $G^a$  are not an epistemic position to deal, then it is right to say that those in  $G^a$  do not know at  $t^1$ . If proper and reasonable inquiry would soon have lead to a community of inquirers,  $G^p$ , such as could reasonably have sought to determine from  $G^a$  whether  $p$ , then surely  $G^a$  in the role of a general purpose source community aspires to deal with question of whether  $p$  in a way that would be adequate to  $G^p$ 's concerns. That contemporaries of  $G^a$  were so intellectually timid, or cowed, or so obtuse as not to have recognized and responded to some compelling open question for which the result that  $p$  would be relevant, should not make it easier for one in  $G^a$  to know that  $p$ .

Thus, the phrase, "the various applied or source communities that might draw on  $s$ 's result," should be understood so as to make for a significant modal depth in the gate-keeping contextualist account of knowledge.

All this seems right, at least as a first-approximation and with respect to the near-term sorts of cases that were just mentioned. However, unless care is taken to limit the demands on knowledge in the context of gate-keeping for source communities, the points just made could easily give rise to a virulent form of skepticism. It will be particularly important (in section x) to refine our thinking about communities that are separated in time so that there might have been much

development of evidential and methodological resources in the interval between earlier source community and later communities that might or might not seek to draw on its results.

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 classes). Enquiry is so structured in order to yield an epistemic situation that can deal with certain error possibilities. These error possibilities are far from idle; they were made salient by empirical results regarding placebo effects and experimenter bias. Apparently, we learn to inquire better.

For one who is keeping epistemic gate for a source community, the new appreciation of live error possibilities changes things—notably, it prompts a refinement in their understanding—of conception—of what it takes for one to count as knowing when seeking to contribute to the established stock of source community results. But, as noted already, this does not change what it takes for one to know for purposes of a source community. One contributing to the source community result contextually seeks to be in an epistemic position adequate to deal with all error possibilities of concern to all who reasonably seek to draw on their results. Thus, before the new understanding emerged, one purporting to know in the context of such a source community was taking on the burden of dealing with such error possibilities as would be found significant for those who would soon be keeping gate for the same source community. Once the error possibilities are newly recognized 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 (results on which folk may draw without hesitation). In the context in which one is gate-keeping 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—as one cannot flatly draw on their results. The matter at issue in the putative result needs to be investigated further. For example, one will judge that the earlier experimenters, who are 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. Lacking better structured studies capable of ruling out the relevant error possibilities, one will insist that these experimenters do not yet know whether the treatment is effective. (Although one might allow that their study suggests that the treatment may well be effective, or that it warrants further investigation.) 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 older community, incorporating those results in its corpus of accepted results, such evaluations seem exactly right.

As noted, this denial of knowledge to earlier investigators within a source community makes perfect gate-keeping sense as long as one is entertaining drawing freely on their results. And, 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, a kind of collection of what we seem now to know). Such practice reflects our best present standards, our best understanding of what it takes for one to count as knowing for purposes of a given source community. But, best present standards are not definitive. As inquiry develops, new error possibilities may become apparent, and the earlier standards will be found to be inadequate.

While all this seems reasonable and proper, it seems we may be walking into a skeptical trap. 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, error possibilities that we insist one guard against before counting as knowing. But, if we are ourselves 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—signaling that one cannot freely rely on them in the matter at hand—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 result—we might signal to others that we have succeeded when we claim to know—but we will likely turn out to have failed here. At least investigators in significantly later communities of inquirers will very likely know better.

If the earlier researchers are rightly judged not to know, it seems likely that the same will rightly be said of oneself—in light of one’s own epistemic position, and the error possibilities reasonably relevant to later investigators, one does not know. This is skepticism.

*ix. Why skepticism has no place in gate-keeping contextualism*

This skeptical result signals a fundamental mistake in developing the gate-keeping contextualist account—although it may not be obvious where the mistake was made. 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. It shows that, were the concept to work as just suggested, as it were *without qualifications*, it would frustrate, not facilitate meeting, the epistemic needs that source-communities seek to address. Without qualification or refinement, the picture just developed would result in a kind of paralysis in the broad epistemic community (at least with respect to source communities). It is time to pause, take a step back, and see why the suggested result is completely unacceptable. Doing so will allow us to gain a better understanding of the pivotal phrase: “the various

applied or source communities that might draw on *s*'s result."

Evaluative and normative concepts, such as KNOWLEDGE, or JUSTIFIED BELIEF, grew up in response to the needs of regulating ourselves and others in the course of pursuing some project (one with both individual and communal aspects). In this respect, they are not different from evaluative concepts generally. The concepts of a GOOD PASS, and GOOD POINT GUARD, developed in connection with the needs faced in the pursuit of basketball success and excellence. The concept of a GOOD JOINT and a GOOD CARPENTER grew up in connection with the needs faced in the pursuit of a certain building technologies. Plausibly the activities that are regulated, the needs faced, and the evaluative concepts, grew up together.

This suggests that the central epistemic evaluative concepts must be very old, and calls to mind a passage in Craig (1990):

Human beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome. That being so, they need sources of information that will lead them to believe truths. They have 'on-board' sources, eyes and ears, powers of reasoning, which give them a primary stock of beliefs. It will be highly advantageous to them if they can also tap the primary stock of their fellows—the tiger that Fred can see and I can't may be after me and not Fred—that is to say, if they act as informants for each other (p. 11).

This need to regulate sources is at the root of several epistemic evaluative concepts. Regulating both oneself as a source, and one's community of sources, is then undertaken with a diversity of epistemic evaluative concepts. The concept of KNOWLEDGE is an important one, signaling that the given agent, commonly the agent drawing on yet wider community resources, has pursued inquiry sufficiently so as to put her in an epistemic position to serve as a source for the theoretical and practical needs of a contextually relevant community.

It is highly plausible that such evaluative concepts cannot diverge from what is reasonable in view of needs faced in the project in which they have their roots. (This seems to me a conceptual truth on the same order as the conceptual truth that natural kind substance terms, if they refer, refer to stuff of the sort that is prominent

in the samples with which a folk have interacted.) Of course, folk's conceptions of these matters can get significantly out of step with the project. But, when this happens, the folk are mistaken.

Consider an easy case, the concept of a GOOD POINT GUARD in basketball. To unpack this concept, one needs to understand the project of playing (with good prospect of winning) in basketball. One quickly appreciates that, given the realities of human physical diversity, different folk will be suited to somewhat specialized roles on a basketball team, and the point guard represents one such role. One then notes the demands of the role in the pursuit of the endeavor, and begins to reflect on what characteristics, abilities, and skills contribute to that endeavor in that role. Let us now suppose that someone thinking along these lines thought as follows (a wrong turn, I think): "Well, a point guard needs to be (relatively) quick—specifically, a good point guard needs both rapid acceleration and to be able to achieve a sustained speed that brings the ball down court quickly. The quicker the better, with ball control. Some apriorist basketball-ologist might reason that the speed of light might serve as the idealized standard for point-guard speed. "Anything less than this won't really be a good point guard." One might imagine someone thinking along such idealizing dichotomous lines, but it would have to be a pretty flat-footed thinker. For, to so think would be to lose sight of the connection between the project regulated, and the concept on which one is reflecting. One can put the point this way: Suppose that an evaluative concept grew up to regulate a project, and suppose that one has a conception of the matter associated with one's grasp of the concept. In the present example, there is the concept of a GOOD POINT GUARD and there is one's conception of a good point guard—one's understanding of, or standards for, what makes for a good point guard. Suppose that one's conception is the high ideal conception above. Suppose that one deploys it in assembling one's basketball team. One would find no good point guards, and leave them off the team. Parallel analyses would result in one's recruiting no centers, no forwards, ... and ultimately, no basketball players. No basketball would be played. Now, either one's conception provides a proper analysis of the concept, or it does not. If it does, one should simply abandon that concept altogether—it is a lousy concept for purposes of the

project in which it should have its roots. However, I trust that none of my readers think that this is what is going on with the silly analysis. One is not ready to abandon the concept of a GOOD POINT GUARD, or more generally, of a GOOD BASKETBALL PLAYER. Instead, one insists that the silly analysis gets those concepts wrong. One insists that a proper analysis would reveal desiderata that fruitfully regulate the project of playing basketball.

Generally, when deploying one's conception in regulating the project with which it is associated would shut down or deeply frustrate the very project the associated concept supposedly serves to regulate, this is very strong prima facie grounds for thinking that one's conception is mistaken—that, properly understood, the concept would facilitate the further that project—not utterly frustrate it. An analysis of KNOWLEDGE that leads to a skeptical result with respect to general-purpose source communities is like the silly analysis of the concept of a GOOD POINT GUARD in that, were it applied to govern the gate-keeping for such communities, it would shut down the project—or that part of the project—that it governs. Apparently, were such an analysis correct, one would need to conclude that, for purposes of these communities, no one knows anything—that no one should take anything they say to be actionable information. Were this analysis to capture the workings of the concept of KNOWLEDGE, we should not care about knowledge in connection with such communities, and this seems clearly the wrong answer. One should conclude that, in getting to the skeptical result, we have taken a wrong turn somewhere along the way.

Cf. Greco <sup>13</sup>

*x. 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. To properly do this, folk in a source community would need to produce a stock of beliefs meeting this standard:

For a result to be contributed to the established results of the general-purpose source community,  $G^g$ , those in  $G^g$  must be (at least cumulatively)

sufficiently epistemically well-positioned with respect to  $p$  that they 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).

This, of course, was presented earlier as the standard for the truth of attributions of knowledge when contextually keeping epistemic gate for a general-purpose source community. Skepticism threatened when it was realized that, strictly speaking, it is highly unlikely that folk commonly satisfy this standard. At least it is unlikely if it is understood to require that groups significantly in the future, groups that may be the beneficiaries of much intervening inquiry into related matters, are thought to be among the possible consumers of a source community's results.

But, one might wonder, is there not some limit to the range of communities that might *reasonably expect* to be able to draw on a given source community's results. Perhaps equally significantly, is there not some limit to the range of communities to which one is *reasonably certifying* results when one, in a source-community, asserts that someone knows?

The threat of skepticism derived from the observation that communities significantly far in the future are likely to have a perspective from which they might find relevant error possibilities that cannot be appreciated by contemporary communities of inquirers. That this is likely results from some simple observations. Over time, communities of investigators learn to inquire better. 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, can change significantly. Double-blind experiments with placebo or related controls was given as one illustration. 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, in Galileo's famous inclined-plane experiment, one finds two forms of acceleration in play—the acceleration of a falling body (which the apparatus was designed to slow down and make observable to Galileo) and angular acceleration of a rolling ball. Galileo was clearly in no position

to determine that his result would hold up when only the one and not the other form of acceleration was in play. Then there is the increasing ability to measure certain quantities. Suppose that work in some source community strongly supports some generalization. Of course, their abilities to measure may be much cruder than those available to significantly later communities, and their abilities to measure may be restricted to much narrower ranges of values of the relevant variables. Obviously, the later community may then be in a position to be concerned with error possibilities that are beyond those that the earlier community could address. For example, different error possibilities associated with generalizations involving materials and temperatures are relevant to seventeenth century investigators and to our contemporary source communities. Not incidentally, different error possibilities are relevant to seventeenth century artisans and to contemporary engineers as they design and construct systems for various practical purposes.

One might conclude that source communities rarely, if ever, know. But, as explained in the previous section, this would seem to deprive us of source communities on which others can freely draw. That seems far too hasty. We had better look more closely at our fruitful practice. Again, I think that this question is crucial: is there not some limit to the range of communities to which one is *reasonably certifying* results when one in a source-community asserts that someone there knows? And, in thinking about this question, it seems fitting that one would condition one's contextual certification in a way that accommodates this observation: there is some limit to the range of communities that might reasonably expect to be able to draw on one's results.

With regard to source communities associated with fields that are undergoing developments of the sorts just sketched, it will be rare that one would even think to depend on the results of investigators working significantly earlier. More recent investigators would have typically have re-examined the earlier results using more powerful tools. One tends to take one's certification from the up-to-date sources. (One may admire Galileo's work, but 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. Again, it might

then seem silly to seek to draw on resources from communities working significantly earlier.<sup>14</sup> If one was seeking to keep epistemic gate for this later community, in the sense of determining whose results one in that community can flatly rely on, and were someone then to recommend the results of the significantly earlier community, one would presumably insist that those earlier researchers did not know what they thought that they knew. But, at just this stage, it seems that one would hardly think of adopting the beliefs in question on the basis of the earlier work. Typically, one simply does not look to the earlier community as a source on the matters of concern. 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.

So, let us focus on the reasonable epistemic dependencies between source communities for the moment. 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,  $G^{mediate1}$ , ...,  $G^{mediate2}$ , 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. We can then say that a source-community would be unreasonable to flatly draw on the results of a significantly earlier source community. (This is the burden of several previous paragraphs.) What is more, it would be unreasonable for source communities not to notice that there is reasonable some “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. One should then not aspire to serve as an unfiltered source for those in significantly later communities.

Here we come to the big point: 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. Of course, no result that is wrong (false) is ever known. But, some results are known in the sense that they satisfy the contextual standards for a source-community, then 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 (in the technical sense) significantly later than s*. The contextual aspiration of a source community is to play this role. This is the 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.

This calls for an explicit refinement on the earlier treatment of what it takes to know in the context of gate-keeping for a source community:

(In a source community)  $K_{s,p,t}$  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 reasonably draw on s's result, where no source community that is significantly later than s's community can reasonably seek to draw on s's result*.

#### *xi. Contextualism and Purposes*

But, damn it, Galileo knew that a falling body in uniformly accelerated in the sense that an equal increment of speed is added in each successive interval of time! How can I say this—since I am not really signaling that one can simply take it from Galileo? It may seem that I should not be able to say this, given my claims:

- (a) In attributing knowledge, I am keeping gate for some epistemic community.<sup>15</sup>
- (b) An attribution of knowledge will then only be correct when it would be

fitting for the community in question to draw freely on the relevant agents' results.

- (c) I cannot really so signal that one can flatly take this result on Galileo's say so, since I think that one should only take it from more up to date communities that have guarded against far more error possibilities than Galileo could envision or guard against?

My sense is that I can say this because, in so saying, I am not engaged in keeping epistemic gate at all. I am not signaling a good source. Apparently, while epistemic gate-keeping may be a central point and purpose of the concept of KNOWLEDGE, it is not *the* (unique) point and purpose served by the concept. Apparently, (a) needs to be qualified significantly.

The attribution of knowledge to historical figures provides illustrative cases in which one has not the slightest intention of drawing on the historical figure's results. The above attribution of knowledge to Galileo seems a good example. We are rightfully pretty confident that we today have better sources on the matter — one's who are epistemically better placed than was Galileo. Further, the attribution of knowledge, or withholding of such attributions, to historical figures such as Galileo does not serve significantly to regulate any epistemic community.<sup>16</sup> So, what are we doing in such cases.

It does not seem far-fetched to say that such attributions of knowledge might be thought of as a kind of *simulated gate-keeping*. Let us explore this idea.

In the philosophy of the social sciences and psychology, some have argued that our own ability to explain what agent's think and do commonly runs far beyond the generalizations we could deploy in some manner of subsumptive explanation. Rather than rely on generalizations to reveal the dependencies between what an agent may have initially believed and desired and what the agent ultimately thought or did, we instead seek to simulate in our own cognitive processing the thinking or decision-making of that agent. In doing so, we put our cognitive processes to work on inputs (pretend-beliefs and pretend-desires) that we suppose parallel those beliefs and desires that were salient for the subject we seek to understand. Our own cognitive processes are then deployed "off-line," to simulate the other. To the extent

that our pretend choices and conclusions match those plausibly observed in the agent, we suppose that the dependencies evinced in our own cognition may plausibly be paralleled in the agent we simulate. It is plausible that such a process is a pervasive piece of the cognitive repertoire of normal human cognizers and that it is natural and commonly automatically employed in understanding others (see, Goldman [2006], Steuber [2006], Nichols and Stich [2003]).

If this story is roughly correct, and I think that it is, then it would seem only natural, when thinking about historical agents and their epistemic situation to find ourselves engaged in a kind of vicarious gate-keeping. In the course of our historical reconstruction, we find ourselves doing a bit of simulated or pretend gate-keeping. In the course of this, we may “signal” that our subject, perhaps Galileo, was a first-rate source on which other savants, and others practically engaged (such as artillerymen) could draw without reservation. We judge this while realizing that his results would have a shelf-life. Our judgment, in simulation, takes up the perspective of Galileo, Galileo’s contemporaries, and those epistemic communities that do not come significantly later. We “signal”—signal in pretense—to a community that, outside of the pretense, we know are dead and gone, and thus cannot be regulated.

Such uses of the concept of KNOWLEDGE may in be derivative in a sense. In these instances, the concept of KNOWLEDGE outlined above—the concept that largely serves a gate-keeping purpose and that works as just described—is deployed when thinking about agents who one is *not* seeking to draw on as a source for oneself or one’s contemporary epistemic communities. It is deployed in a context in which one seeks to understand their situation, perhaps also that of their contemporaries. Such uses of evaluative concepts are not hard to understand. Consider a parallel. The concept of a GOOD HITTER has its roots in the project of baseball—where it is deployed in the regulation of certain practices. It might, for example, be deployed in recruiting one’s team—as one evaluates prospective players. It might be deployed in deciding who should enter the game at a pivotal juncture. These are the baseball analogs to the gate-keeping purpose of KNOWLEDGE. Now suppose that you are trying to understand certain moves made by decision-makers for team—say the move to acquire a certain player for the team. In attempting to understand, one may deploy

the concept of a GOOD HITTER—when concluding that they recruited the player because he was a good hitter. In arriving at this explanation of their action, one may have judged that the player was indeed a good hitter, that the team needed more good hitters, that the player was among the best hitters available, and so on. In running through what are likely their considerations, one is using one's own cognitive processes to simulate what plausibly moved the decision-makers. But, in the process, one's own judgment does not itself serve to regulate any baseball practice—it does not, for example, keep talent-gate for any team. Rather, it serves the project of understanding the relevant episode of baseball-talent-oriented gate-keeping. One deployed the concept to shadow or simulate the gate-keeping done there.

Related to the kind of explanatory simulation just envisioned, there are cases in which one uses one's cognitive processes off-line, deploying one's concepts, while evaluating some course of action in which one has no hand.<sup>17</sup> For example, one may use various of one's evaluative concepts, including that of a GOOD HITTER, when thinking about a the decision facing a team. Perhaps they have lost a good hitter to injury, are entering an important part of their season, have salary considerations, and considerations of interpersonal dynamics to consider. Thinking through such considerations oneself, in a kind of pretense, one might weigh out what the decision makers *should* do. In so doing, one is not really regulating the team's practice—or trying to. One need have no illusion that the relevant decision-makers will heed, or even hear of, one's judgment. Rather, one just seeks to evaluate the decision for oneself. One might judge that a certain player is a good hitter, and a bargain at this stage in his career, and that the player has a reputation for contributing to the morale “in the clubhouse.” One then concludes that acquiring that player for the team would be a good move. Again, one is not really keeping gate for team talent. Instead one is evaluating the evaluations and decisions of those who are keeping gate.

I think that we are commonly doing something parallel when we consider a historical case of epistemic practice, and judge that the agent in question knew or failed to know. We may take up in pretense the evidence and theoretical situation of

a historical scientist, considering their epistemic situation from the agent's perspective or from that of a community of scientists evaluating the results of one of their contemporaries. Perhaps they are all dead and gone so that there is no sense in which our evaluation really keeps epistemic gate for them, and perhaps we today have no need or thought of drawing on the result produced in the historical episode. We are not gate-keeping. Still, we are using a concept fitting to gate-keeping. We are using it in a kind of pretense in which we attempt to gauge for ourselves what the historical community *should have judged*. If, in the course of our pretense, we judge that the agent in question knew the result he announced back there, we have evaluated the agent's self-evaluation—or we have evaluated the evaluation or his contemporaries. In so doing, we judge, he spoke truly in there claiming to know.

So some uses of the concept of KNOWLEDGE do not serve the gate-keeping purpose of signaling good sources to some actual community. But seems wholly in keeping with gate-keeping being a central point or purpose served by the concept. Indeed, the thrust of this last discussion is that, even when one find the concept of KNOWLEDGE is not deployed in a form of epistemic gate-keeping, the concept so employed plausibly yet has its roots in serving the purpose of epistemic gate-keeping.

In this paper I have explored how the concept of KNOWLEDGE would work if it were largely developed in the service of a single 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. Another set of these should be sensitive to the aspirations of and for general-purpose source communities, and would look somewhat like the attributions envisioned subject insensitive invariantists. In this final section I have sketched certain applications of the envisioned knowledge-concept (the concept centrally conditioned by the point of epistemic gate-keeping) that themselves resemble what might be envisioned by subject sensitive invariantists. In these contexts, one is not so much keeping gate for an epistemic community as simulating the keeping of

epistemic gate for various communities. Here, as one judges in the course of a simulation that the agent in question knows, one is either explaining the judgment of some agent or agents who were themselves keeping gate, or one is evaluating what would have been the fitting gate-keeping by members of the relevant community. This is clearest in historical cases, where there is no real prospect of keeping gate. But, it can happen in contemporary cases.

What I have provided here makes it plausible that the concept of knowledge *is* preponderantly conditioned by, and subservient to, the epistemic need for a kind of gate-keeping for epistemic communities. A rather surprising range of uses of the concept can be understood in these terms, including some in which the attributor is not really engaged in gate-keeping. However, I do not think that what I have managed here amounts to a demonstration that the concept of KNOWLEDGE is rooted in and subservient to one central point: gate-keeping. While that understanding of the concept may be a plausible working hypothesis in light of what I have said here, I am not ready to claim to know that *the* point and purpose of the concept is epistemic gate-keeping. I certainly have not attempted to consider all uses of the concept and engage in an abductive inference regarding which set of purposes best explains the workings of the concept evinced in judgments about a fully representative set of cases. That, I fear, is a very big project.

## References

- 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, 913–929.
- DeRose, K. (1995). "Solving the skeptical problem." *Philosophical Review*, 104: 1–52.
- DeRose, K. (2009). *The Case for Contextualism*. Oxford: Oxford University Press.
- DeRose, K. (MS). On X=Phi.
- Greco, J. (2008). "What is wrong with contextualism." *Philosophical Quarterly*, 58, 416–436.
- Goldman, A. 2006. *Simulating Minds: The Philosophy, Psychology, and Neuroscience of*

*Mindreading*. Oxford: Oxford University Press.

Grimm S. (). "Knowledge, Practical Interests, and Rising Tides."

Hawthorne, J. (2004). "Knowledge and lotteries." Oxford: Oxford University Press.

Henderson, D. 1996. "Simulation Theory vs. Simulation Theory: A Difference Without a Difference in Explanation," *Southern Journal of Philosophy*, Spindel Conference Supplement 34: 65-94

Henderson, D. 2009. "Motivated Contextualism," *Philosophical Studies* 12: 119-131.

Henderson, D., & Horgan, T. (2001). "The A priori isn't all it is cracked up to be, but it is something." *Philosophical Topics*, 29, 219–250.

Henderson, D., and Horgan, T. (forthcoming). *The Epistemological Spectrum*. Oxford: Oxford University Press.

Kappel, K. (). "On saying that someone knows: Themes from Craig."

Knobe, and Shaffer ().

Nichols, S. and Stich, S. 2003. *Mindreading: An Integrated Account of Pretense, Self-Awareness, and Understanding Other Minds*. Oxford: Clarendon Press.

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

Stueber, K. 2006. *Rediscovering Empathy*. Cambridge, MA: M.I.T. 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, 213–235.

Williamson, T. (2005b). "Knowledge, context and the agent's point of view." In G. Pryer & G. Peters (Eds.), *Contextualism in philosophy*. New York: Oxford University Press.

---

<sup>1</sup> I have plenty of good company. Craig (1990) provided a seminal and increasingly influential treatment of KNOWLEDGE drawing on a sense for its point and purpose. Others, notably Greco (2008), Pritchard (2010), Kappel (), Grimm (), ... have fruitfully explored Craig's suggestion.

<sup>2</sup> I rely on a rough distinction here and make no effort to draw a sharp line between practical projects and communities, on the one hand, and more theoretical/inquiring communities, on the other.

<sup>3</sup> The agreement is rooted in judgments elicited for philosophers that seem to indicate significant variation in verdicts in connection with the stakes stipulated in the scenarios. Of course, how ultimately to accommodate such judgments is a tricky question (see, for example, Williamson (2005a, 2005b)). However, there now seem to be reasons to worry about the degree to which philosophers' common judgments themselves reflect a conceptual competence. Some recent work in experimental

---

philosophy has suggested that the folk do not share philosophers' intuitions regarding the relevant cases—and, by extrapolation, that philosophers might be thought to be biasing themselves in some way (for example, see Knobe and Shaffer, MS). I need not address the issue in this paper, in part because DeRose (MS) has recently provided a good start in addressing the concern. While I am open to being set straight by work in experimental philosophy, I think that it is plausible to believe that experimental philosophy has yet to provide us with the tests that would be needed.

<sup>4</sup> For example, see Hawthorne (2004) and Stanley (2005).

<sup>5</sup> For example, see Cohen (1999, 2005) and DeRose (1992, 1995, 2009). Other examples include Greco (2008, and MS), Henderson (2009), and Grimm (0). Williamson (2005a, 2005b) provides a nice comparison of the various positions.

<sup>6</sup> The flexibility in contextual concern is strongly suggested in Greco (2008) and Henderson (2009). It is also reflected in passages in DeRose (2009).

<sup>7</sup> To keep things simple, I suppose that  $s \neq a$ . However, to my mind, nothing hangs on this.

<sup>8</sup> So, were one gate-keeping for  $P_s$ , one could truly assert that  $\langle a, P_s \rangle: Ks, p, t$ .

<sup>9</sup> My presentation here is unabashedly schematic. It is easy to think of epistemic communities that are hybrids of practical communities and source communities. Take a narrower discipline such as restoration ecology, or restoration wetland ecology. Such disciplines serve as a source community for various policy makers and for field workers engaged in particular projects (with varying stakes). At the same time, they seem close enough to an identifiable set of communities and projects that one might expect that one can see contextual variation over time.

<sup>10</sup> Although arguably the intention of the attributor to certify a subject to a given general-purpose source community might matter.

<sup>11</sup> I take it that a source community is not a set of inquirers—as the death, retirement, or admission of new members—does not make for a different source community. For present purposes, the identity of a source community may be understood as flexible in the way that scientific disciplines are flexible.

<sup>12</sup> Obviously, this has to do with all attributions of knowledge that  $p$ , for a given proposition  $p$ .

<sup>13</sup> Greco (2008) suggests 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:

For one, if knowledge is for the purpose of practical reasoning, this requires that knowledge be widely attainable, and this puts pressure on the standards for knowledge in a downwards direction. The argument here is straightforward: knowledge is required for practical reasoning, and practical reasoning is widely necessary. Therefore, knowledge must be widely available to serve its purpose. Therefore, the standards required for knowledge cannot be so high as to make knowledge less than widely available. That is, the standards cannot be so high if knowledge is to have the role that it does in our practical activities.

Likewise, the demands of practical reasoning require that knowledge be imported across various practical environments, and this puts pressure on the standards for knowledge in an upwards direction. Suppose that I need to know whether  $p$ , and that in a different context  $S$



claims, 'I know that  $p$ '. For  $S$  to be a useful informant, for her to be a source of knowledge to be used in my practical reasoning, I have to be confident that the standards by which  $S$  counts as knowing in her context are at least as high as my practical reasoning requires. Again, the information-sharing function of our knowledge language puts pressure on the standards for knowledge in an upwards direction. The standards for knowledge cannot be so low as to make knowledge widely unusable. (check quotation , perhaps Greco, 428-9)

<sup>14</sup> One can understand what counts as "significantly earlier" communities in terms of expectations for results in the field to become dated by methodological developments or by changes in the richness of data.

<sup>15</sup> And one might automatically suppose that one is interested in the matter as a gate-keeper for a general purpose source community. The general problematic carries over to keeping gate for most engaged practical communities.

<sup>16</sup> It may contribute to an extended and somewhat indirect kind of epistemic self-regulation. Honoring Galileo's work with such attributions may have some salutary effect on us today, perhaps as we retrace how he managed to achieve what it did.

<sup>17</sup> It is commonly suggested that one uses one's capacity for pretense when planning, as one considers likely results of possible courses of action. Here, one evaluates what is reasonable to do in light of the information possessed. One can do this when evaluating the courses of action or thought open to others.