

# What's the Point?

David Henderson

Draft 11/5/10 3:23 PM

i. Introduction .....	1
ii. Philosophical reflection: .....	2
iii. Some Data for Reflection on <i>Knowledge</i> , Illustrating Some Relevant Data-Types..	6
iv. Gate-keeping Contextualism: .....	15
v. Evaluative concepts and the regulation of practices .....	17
vi. Generalizing Beyond Evaluative Concepts? .....	22
vii. Other Cases and Other Apparent Uses .....	23
viii. Accommodating the diversity—an explanation. Is it the best? .....	28

## i. Introduction

Recently I began exploring what I believe to be a promising variant of contextualism about knowledge (Henderson 2009, 2011). It combines contextualism with a view about the point or purpose of knowledge attributions—specifically, that it is central to the concept of *knowledge* that attributions of knowledge function in a kind of epistemic gate-keeping for contextually salient communities. Greco (2008, 2010) has independently developed an account along similar lines. I will rehearse the main outlines of this gate-keeping contextualism, and discuss its virtues. However, a defense of gate-keeping contextualism is not my primary concern in this paper. Instead, the view and its virtues serve as a running illustration for a discussion of philosophical reflection. The main outlines of the account of philosophical reflection on offer have been presented elsewhere by Terry Horgan and myself (Henderson and Horgan, 2001, forthcoming). It holds that philosophical reflection is a multiply constrained abductive project—and that one

kind of data to be accommodated in reflection is one's sense for the purposes central to the concept on which one is reflecting. The case for gate-keeping contextualism is an instance of such philosophical reflection. Thus, in this paper, the case for gate-keeping contextualism is clarified while the character of philosophical reflection is exhibited. Both to the good. But, this paper has a yet more specific agenda. Both in my own discussions of gate-keeping contextualism, and in Henderson and Horgan's general discussions of how a sense for points and purposes of a concepts should inform one's analysis, talk of point and purpose has had a "seat of the pants" character. Much the same can be said of other accounts that seek to capitalize on the idea of concepts having central points or purposes. It is, for example, true of Craig's classic (1990), and this has prompted attempts at clarification (Kappel, 2010). I here seek to do what I can to clarify. Happily, there seems to be a kind of concept that provides relatively good traction for thinking about the point of concepts, and the concept of *knowledge* is an instance of this kind. Normatively-valenced, evaluative concepts constitute a broad class of concepts for which a socio-linguistic point or purpose may be readily sensed—and for which the intimate connection between point/purpose and the concept itself is readily understandable. This allows me to clarify my general thinking about the point of purpose of concepts. However, it must be admitted that, commonly at the inception of philosophical reflection one's sense for the point or purpose of a specific concept is seldom highly articulate. I argue that this does not prevent it from doing real work in philosophical reflection.

## ii. Philosophical reflection:

Henderson and Horgan (2011) advance an account of *a priori* knowledge and justification, particularly that characteristically arising from philosophical reflection. Our account provides for conceptual-competence-based reflective justification, and on this basis, we argue that it has good claim to providing an account of at least of an important species of *a priori* justification such justification. However, we argue that reflective justification, thus such *a priori* justification, itself must have an empirical dimension. We

David Henderson 11/3/10 9:35 AM

**Comment [1]:** This section needs streamlining. The general start is good, as is the ending. The central ideas are the two-stage process, the abductive character of the second stage, and the list of kinds of data. The parallel with the linguistics case is too long, and there is little need to make the case for the calling the whole low-grade a priori, or to say much about the empirical contamination. Perhaps a few sentences and a reference suffice.

mark this somewhat paradoxical result by writing of “low-grade” *a priori* justification—as opposed to the “high-grade” *a priori* justification of epistemic lore.

Our account takes its departure from the following three conjectures:

First, a prominent way of drawing upon one’s conceptual competence, conspicuously manifested in much well-received philosophical work, involves the generation of judgments about relatively specific and concrete scenarios. In response to such a scenario, conceptual competence spawns fittingly specific and concrete judgments *in a direct way*. It is not surprising that human conceptual competence is particularly suited to the generation of responses to such concrete specific scenarios. After all, in everyday contexts, where conceptual competence serves in a largely unnoticed way, it functions largely in the direct and automatic application of concepts.

Second, conceptual competence is much less steady and reliable when called upon to directly generate conceptually grounded *general* truths. Perhaps conceptual competence can generate general truths with reasonable success when they turn on certain concepts of relatively straightforward sorts. But particularly with respect to certain concepts that have commonly been significant for philosophers, one’s conceptual competence may be much better at directly delivering presumably veridical judgments regarding *specific* applications of those concepts than it is at directly generating veridical judgments regarding generalities.

Third, by drawing upon what conceptual competence *does* provide, one can manage to justifiably believe certain conceptually grounded generalities that are not themselves the direct deliverances of that competence. Consider the justly famous result of Putnam’s (and Kripke’s 1972) reflections:

In all possible worlds, water is stuff with the same micro-structure as the stuff in the salient samples.

This conceptually grounded general truth is justified in terms of more particular judgments about the referent of ‘water’ in various concrete scenarios. We ourselves believe that such general results are rightly taken to be *a priori*. The central question becomes just how one can get to such truths.

The above three observations strongly indicate that, at least as it has application to philosophical cases, an adequate model of conceptual-competence based *a priori* reflection will need to recognize at least *two stages*: one in which reflection on specific concrete scenarios generates correspondingly particular judgments (the direct deliverances of conceptual competence, mentioned in the first observation), and another in which one

reflectively draws upon these particular judgments to *inferentially* support a judgment whose content is abstract and general (the need for this is evident in the second and third observations above).

The methodology typically employed by linguists in constructing and evaluating theories of natural-language syntax presents a parallel two-stage epistemic process. The empirical data for syntactic theory includes certain judgments of competent language-users—in particular, judgments concerning the grammaticality or ungrammaticality of various sentence-like strings, and concerning grammatical ambiguity or nonambiguity of various sentences. Native speakers, after all, can be expected to have judgment dispositions about these matters that reflect a solid mastery of their own language (or their own regional dialect, at any rate). So, when native speakers are intersubjectively consistent and also uniformly confident about such syntactic judgments, then normally these judgments reflect the natives' *syntactic competence*, their mastery of the syntactic norms or syntactic structures underlying their language. And this psychological hypothesis, in turn, has a direct implication for linguistic theory—viz., that under an adequate theory of syntax for the natives' language (or dialect), those syntactic judgments will turn out mostly correct.

When the linguist proposes certain abstract general syntactic principles, claiming that they are the rules of grammar for a given language, these proposals are not themselves the direct deliverances of anyone's cognitive mechanisms of grammatical competence. To make the point vivid, suppose that the linguist is using himself or herself as an informant—perhaps seeking an account of his or her own idiolect. Still, in that event, the proposed syntactical principles will not be direct deliverances of the linguist's own grammatical competence. Rather, the proposed rules of syntax are *theoretical hypotheses* about the language. Thus, first, grammatical competence generates grammaticality judgments about *specific* sentence-like word-strings. Although these judgments presumably are generated in a manner consistent with the general rules of syntax—so that the competent speaker possesses an *implicit* mastery of those rules, whatever they are—grammatical competence does not generate explicit beliefs whose contents are the syntactic rules themselves.<sup>1</sup> Then, second, the linguist's own grammaticality judgments, with their default status of presumptive correctness, are *data vis-à-vis* these proposed rules. The evidential connection between the data and the theoretical hypotheses—here, as elsewhere in science—is *inference to the best explanation*.

Thus, a linguist who is seeking to produce an account of his or her own idiolect would begin with his or her own judgments, from the armchair—judgments that are

defeasibly presumed to be the relatively direct results of his or her grammatical competence.<sup>ii</sup> The judgments then become data for an overall theory that must provide an explanation of their status and their source. This theory will necessarily have dual dimensions—one regarding the rules of the grammar for the language, and the other regarding the extent of the informant's (linguist's) own grammatical competence. Accordingly, this second stage clearly involves an *empirical* abductive inference, rather than itself being the direct product of the linguist's grammatical competence.

Any actual competence is subject to interferences, distortions, and systematic errors. Accordingly, any satisfactory accounting for sets of judgments will need to sort out both the grammar of the language and the character of the individual's competence. Any account of the grammar will have implications for which judgments are competent judgments and which are performance errors. One's account of grammar thus makes demands on one's psychological theory of competence and of the ways that human competence falls short of ideal competence. Overall then, the linguist's proposed grammar will be informed by background and broadly empirical understandings of the finite and fallible character of people's actual (as opposed to ideal) grammatical competence; the account will then at least implicitly be committed to a rough understanding of this competence and its salient limitations; and this understanding will be rightly affected by, and answerable to, broadly empirical hypotheses about various human cognitive capacities and limitations. The results of such theorizing are clearly *not* the direct product of the linguist's grammatical competence. Nonetheless, syntactic theorizing can be effectively pursued by reflection in the armchair, drawing on the individual's own linguistic competence.

As in the case of linguistic investigation, philosophical reflection draws on the scattered concrete direct deliverances of a cognitive competence—*conceptual* competence, the analogue of grammatical competence. It seeks to arrive at conceptually grounded necessary truths, the analogues of the rules of grammar the linguist seeks to discover. Furthermore, the relevant conceptually grounded necessary truths are not reflectively generated in a direct way out of one's conceptual competence; so the ensuing reflection will need to employ the more particular concrete deliverances of conceptual competence as data for abductive inference. This reflective procedure has both psychological and conceptual faces, like investigations of grammar, and it is properly understood as having an empirical dimension.

Conceptual-analytic reflection is a philosophical enterprise continuous with relevant work in disciplines like cognitive science, linguistics, and sociolinguistics. Insofar as conceptual-analytic reflection relies primarily on armchair-accessible data, it can be effectively practiced by philosophers. One makes a case for a certain conceptual-analytic hypothesis—for instance, the contention that the meaning of natural-kind terms depends on the language-users' environment—by arguing that it does a better job, all things considered, of accommodating the relevant data than do any competing hypotheses. Such reasoning is broadly empirical: inference to the best explanation. Although it can be conducted from the armchair and it is aimed at discovering conceptually grounded necessary truths, it does not conform fully to the traditional understanding of the *a priori*—we think of it as a kind of low-grade *a priori*.

The kinds of armchair-obtainable data that are pertinent to philosophical conceptual-analytic inquiry appear to be fairly diverse (more so than in the linguistics case), with some kinds being more directly analogous to grammaticality judgments than others. The types of data that can figure in philosophical conceptual-analytic reflection include the following:

1. Intuitive judgments about what it is correct to say concerning various concrete scenarios, actual or hypothetical.
2. Facts about conflicting judgments or judgment-tendencies, concerning the correct use of certain concepts in various actual or hypothetical scenarios.
3. Facts about standardly employed warrant-criteria for the use of various concepts.
4. Facts about the key purposes served by various terms and concepts.
5. General background knowledge, including untendentious scientific knowledge.

Data of all these kinds can go into the hopper of wide reflective equilibrium whereby conceptual-analytic claims are defended in philosophy.

### **iii. Some Data for Reflection on *Knowledge*, Illustrating Some Relevant Data-Types.**

*Intuitive judgments about what it is correct to say concerning various concrete scenarios:*

Judgments provoked by scenarios have figured centrally in the case for contextualism about knowledge. DeRose (2009, p. 47) notes, “The best grounds for accepting contextualism come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, non-philosophical talk.” There are many scenarios and variants on scenarios to be found in the extant literature.<sup>iii</sup> In one very important class of such scenarios, agents claiming to know offer putative information to others, who then evaluate the epistemic standing of the agent or agents offering that information, all in a setting that makes salient the practical stakes faced in some decision to which that information would be relevant. Thus, DeRose’s bank cases:

*Bank Case A:* My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But, as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, ‘Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.’ I reply, ‘No, I know it’ll be open. I was just there two weeks ago on Saturday. It is open until noon (DeRose, 2009, p. 1).

*Bank Case B:* My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited in the bank before Monday morning, the important check we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, ‘Banks do change their hours. Do you know that the bank will be open tomorrow?’ Remaining as confident as I was before that the bank will be open then, still, I reply, ‘Well, no, I do not know. I’d better go in and make sure’ (DeRose, 2010, p. 2).

Many judge that the responses in these cases are correct—that in *Bank Case A* DeRose did know, while in *Case B* he did not. Of course, his epistemic situation with respect to the bank’s hours is identical in the two cases. What varies are the stakes faced—and this leads the reasonable agents to demand more on the part of one

serving as a source of information for purposes of the pairs joint decision making. The stakes and reasonable demands are paralleled in judgments about whether the agent knows.

It is worth noting that little changes when a third party is added. Perhaps a coworker who has little at stake himself is along for the ride—and suppose that this person seeks to provide information on the bank's hours, having himself exactly the background above accorded to DeRose. Apparently, what matters are the stakes faced by those who contextually are entertaining using the information, not the stakes faced by the agent who is said to know.

I do think that it is important to consider a rather different kind of case. Since the sort of case I have in mind has not been given much play in the discussion of contextualism, I must make up my own examples:

*Les Savants A:* I am a biologist, setting around the departmental office talking with some colleagues. We are discussing a rather surprising result published in a recent issue of some fancy-smancy biology journal. It purports to show that the biological process regulating some process is not DDT/EPA uptake and storage, as has been commonly thought, but STP/LSD inter-conversion. A colleague says, "Well, it is obviously pretty nice work, but can we say that the author of the study (and we) know on the basis of this study (that STP/LSD processes are responsible for the relevant regulation)?" Off we then go in conversation. We discuss the methodology. How various parameters were measured, the margins of error, the possible ways of going wrong in determining these—and decide that the data is solid (if gotten just as reported). We discuss what is known about various alternative processes, and how these might perhaps have produced what was here observed. We then look at the structure of the study reported, and whether it was structured to guard against such possibilities of spurious correlation. As we work thought such matters, we conclude that there is no (noncrazy) error possibility that the study does not guard against. Well, the author knows, and we now know this!

I (at least) judge the biologists here speak the truth—if indeed it is true that LSD/STP uptake regulates the process in question, then the author of the study, and now those scrutinizing the study, are each in a strong enough epistemic position relative to the claim in question that they know it.

David Henderson 11/3/10 10:35 AM

**Comment [2]:** Admittedly, we do not need both of these Savant eggs. I like the first because of my sense of humor. (Not a good reason to keep it, I guess.) I like the second because it is scientifically realistic. It should be kept, and refined as below—into a new A and B.



Here is a second case:

**Les Savants B:**

B<sup>1</sup>: Several of us biologists are discussing a recent study concerning the hibernation and winter survival of the arctic ground squirrel (*Spermophilus parryii*). For eight months of the year this squirrel curls up into a ball close to the ice of the permafrost, for it cannot burrow any deeper to escape seriously subzero temperatures. There, it allows its body temperature to drop to as low as -2°C, then maintains a body temperature at or below the freezing point of water. This itself is extraordinary. Further, it does not seem to have antifreeze in its veins—it is noted that extracted blood freezes at -0.6°C. Still, the little critters do not freeze solid while maintaining a subzero body temperature in a burrow that can easily hit -15°C. So, it seems likely that some form of supercooling is here in evidence. Risky, and not well understood. Still, we agree that the above facts about behaviors and temperatures are known (although the mechanisms in play are not yet established). The techniques, samples, structured investigations employed in by the investigators are agreed to be sufficient to tell strongly in favor of the above, versus error possibilities of standard sorts. These error possibilities are those that we and the authors (as biologists with familiarity with the field) have learned to systematically rule out in designing and implementing our studies. Of course, we note many open questions associated with the results of the study. While the survival value of such extreme hibernation is understandable in terms of energy calculations, we would love to understand how it is managed. Further, this new piece of work notes that these ground squirrels also have the ability to arouse spontaneously—to warm themselves all the way up to their body temperature when active, 37°C. Further, they do this every month of so, at a total seasonal cost of about half of their stored fat. Finally, once rewarmed, they squirrels show brain patterns associated with REM sleep. Again, we look at the methods and results reported, and we conclude that the authors (and now we) know this much.<sup>iv</sup>

B<sup>2</sup>: But, the authors of the study go farther, advancing the claim that the squirrels heat up “in order to (sleep and) dream.” We all understand this as a claim regarding the functional value of the intermissions, and about its basis being selected for due to this function. It is at this point that one of us objects: “Well, we now know that the squirrels heat up periodically, and that they then sleep and dream. But, lethargic things sleep, and such sleeping things dream. Whatever the exact processes by which they manage their extreme hibernation may be such as would somehow require periodic suspension, so that they are selected for only among those that are capable for such periodic self-arousal. Sleeping and dreaming may be mere by-products, and may have played no part in the explanation of why these critters have such periodic intermissions to their deep hibernation.” Immediately we see the point. There are clearly relevant alternative explanations that the authors have not

David Henderson 11/3/10 10:13 AM

**Comment [3]:** This might be usefully separated into two examples—as now suggested.

explored or ruled out, or shown to be less likely than the hypothesis they advance. We conclude: they do not know that the squirrels arouse *in order to* dream.

Again, I am inclined to judge that the biologists are correct in their attributions and denials of knowledge in these cases. By virtue of disciplinary training, they employ methods and judge methods in ways that are designed to rule out a range of alternatives (or at least to tell in favor of certain claims versus a very significant range of error possibilities). This range of alternatives is generally wider than that of concern to most who have practical stakes in the matter. At the same time, the range of relevant alternatives to which the scientists are responsive is certainly not unlimited. What is “known” about various instruments, sample properties, background processes, and the like are taken as informing what are the range of relevant error possibilities that must be guarded against. It is a wide range of alternatives that they seek to respond to, and they often manage. Here, the biologists readily note that, with respect to the claim about why the squirrels have periodic intermissions in their extreme torpor, there are relevant alternatives with which the authors (and their readers) are in no position to deal.

There is much that is striking about this case. One point of real importance for us in this paper is that nothing in the case depends on the conversational partners having anything of practical moment at stake. Perhaps these scientists work in very different areas of biology and feel little personally at stake (aside from their understanding of what counts as good contributions to biology). Perhaps they care a great deal for the little squirrel critters, perhaps little. What they implicitly demand before counting this as a case of knowledge does not, it seems, vary significantly with such stakes. This again seems fitting.

*Facts about conflicting judgments or judgment-tendencies:*

There is some tendency to judge that what is demanded in *Les Savants* cases is “real knowledge” and that “it does not vary with stakes.” This suggests one high standard—an epistemology “with backbone”—that such scientists comply with without wavering over stakes. Satisfying this standard would produce information

that would be enough for the purposes of any reasonable person, whether they had high stakes or low. One can then feel the temptation to repudiate the judgments made in connection with the bank cases. Such judgment tendencies then encourage classical or invariantist epistemology.

Still, the judgments made in the bank cases seem in order when considered on their own. Thus, when we juxtapose these cases we have an instance of conflicting judgment tendencies—a conflicting pull—mentioned above as characteristic of some of the data in philosophical reflection.

There will be other cases to consider below, but we can begin with those just presented—they are the sort of cases that are most directly suggestive of gate-keeping contextualism.

*Facts about standardly employed warrant-criteria:*

The above discussion already reflects something about the commonly employed warrant criteria deployed in attributing knowledge: one commonly looks to determine whether the agent is in an epistemic position to “rule out” error possibilities that one senses are relevant, or alive, or fitting. (If one is a reliabilist, one might understand being in a strong epistemic position with respect to some claim, that  $p$ , as a matter of having arrived at the belief that  $p$  in a manner that would reliably discriminate  $p$ 's obtaining versus some set of alternatives. If one is an evidentialist, one might understand this as a matter of having evidence that tells in favor of  $p$ , versus some range of alternatives or ways of going wrong.) One should also notice that the commonly employed warrant conditions may vary in certain systematic ways. The range of alternatives that are judged relevant seem commonly to vary in contexts in which one is engaged with folk who are engaged with some practical issue or decision. The range of alternatives that are judged relevant seem less variable when engaged with folk who are seeking to serve as sources of information to a (perhaps somewhat open-ended) set of folk who might come to have an interest of some sort in the matter in question.

*Standardly employed warrant-criteria and the key purposes served*

There is more to one's response to the above scenarios than a bare judgment that the attributions (or denials) of knowledge there are correct—or true. There is more than an evident dimension that is significant for warranting such attributions or denials (the range of error possibilities with which the agent in question is in a position to deal). There is also a strong sense that the warranting dimension is associated with the epistemic needs, or epistemic roles, of various communities of agents.

Consider first the kinds of variation in warranting criteria that is in evidence in connection with practically engaged communities. Discussions of practical rationality in economics and decision theory suggest a way of thinking about these matters (although those discussions may have sought an artificial level of precision). Roughly, as one's stakes go up, it seems reasonable to be willing to pay higher "information costs" in order to guard against a wider range of failures. Of course, one would be happy to draw on others who had paid those higher costs already. One seeks for the results of the more costly inquiry (on one's part or one's sources) insofar as it guards against error possibilities that could frustrate one's high-stakes projects. Thus, when facing the high stakes associated with having funds to cover the important check, DeRose would be unreasonable not to incur the inconvenience of at least pausing long and close enough to get a good look at the posted bank hours. Again, if someone else had already put themselves in a position to rule out the error possibilities (that it would be reasonable to be willing to rule out oneself), it is reasonable to draw on them. What one would not want, reasonably, is to draw one's information from someone that was not in a position to rule out error possibilities that would be reasonable to seek to rule out. Thus, DeRose and wife would jointly want to insist on a relatively strong epistemic position on their own part, or on the part of an informant. If, in attributing (or denying) knowledge they are regulating their sources in terms of such reasonable demands, then they would reasonably attribute knowledge to an informant in a low stakes case, and

reasonably deny knowledge of the same or similarly situated informant in a high stakes case.

They would they reasonably do so, were they, as a contextually practically engaged small community, concerned with sources of information. Further, my sense is that this is in effect what they would be doing in these scenarios. My sense is that, in these cases, some form of regulating of sources for communities is central. My sense is also that this is central to how we use the concept of knowledge. This serves as a kind of seat-of-the-pants take on the point of knowledge attributions. It is a point that is not lost on consuming communities—as they seem to automatically respond to such considerations in practice. Neither is it lost on those who contextually engage with those who seem in contextual need of actionable information that one might have. Were one asked by DeRose whether one knew whether the bank was open Saturday, one would be much more hesitant to say that one knew were one also to recognize that he was in the high stakes case. Of course, were one very well placed on the matter, one would know and be happy to say so. )

Thinking along these lines, it then becomes straightforward to make apparently parallel sense of the *Les Savants* cases. There are communities (communities of sources, of savants) whose socio-epistemic role is to serve as sources of information in a given domain to pretty much any contemporaries who might come to have need of such information. They as a community then seek to lay in a stock of information for which they are in a very strong epistemic position. How strong? Roughly, strong enough so that, were someone to need to draw on their stock of results, they would have already paid the informational costs that that agent could reasonably hope and expect. In the developed world, with its many divisions of labor, the sciences are commonly looked at as such source communities. Thus, when the biologists (in *Les Savants A* or *B*) employed their high standards in judging whether the author of the study knew, they were engaged in a kind of quality control for an epistemic community serving as a general source of information. Perhaps this result will answer a question of great moment to someone—perhaps

their life depends on intervening in the process regulated (in A) or perhaps they have staked their entire fortune on Arctic Ground Squirrel futures (in B). Are the savants up to the task? They want to be. It seems that they use attributions of knowledge to regulate what information is added to their stock of results, and what sources of information are treated as worthy contributors—and do so to be in a position to serve all-comers.

The kinds of regulation of sources of information just now sketched constitute the social-epistemic chore that I *have termed gate-keeping for various epistemic communities*. While it would be good to get better handles on this chore (see Henderson 2011), it is enough of a beginning to allow this sense for the point or purpose of the concept of *knowledge* to inform and constrain philosophical reflection.

In the above, one senses that attributions or denials of knowledge serve to mark out sources of good information—sources on which one can rely by virtue of their epistemic situation with respect to the matter in question. We do not suppose that this vague sense is of itself sufficient to bear of itself the weight of a philosophical account. Still, such a rough and preliminary sense for the point of the concept, how it functions to regulate practice, can roughly constrain one's account. It can condition what account one gives when so accommodating judgments about cases and the other forms of data mentioned above. At the same time, reflection on various cases can provide clues regarding what more precisely is the point or purpose central to the concept.<sup>v</sup> Accordingly, I do not advocate a cases-first approach, in which one would begin with judgments about cases alone.<sup>vi</sup> But, neither do I advocate a purpose-first approach. Instead, I insist that both judgments about cases and one's sense for the point of the concept each amount to data that should inform and constrain one's account of the concept. A proper account of the concept will be supported by what is ultimately a multiply constrained abductive inference.

Both judgments and one's sense for the point/purpose are defeasible—depending on how things sort out in the abduction. In the best overall account of the concept of *knowledge*, one might need to treat certain judgments as loose talk, or as mistakes. Similarly, one might need to accommodate certain facets of “what is there done” in attributing knowledge as secondary, derivative, or incidental. Thus, I argue (in sections vii and viii) that some cases in which one attributes knowledge at a juncture where no real gate-keeping is plausibly done, might yet be understood as cases of simulated gate-keeping.

#### iv. Gate-keeping Contextualism:

The suggestion in Henderson (2009, 2011) and here is that the concept of *knowledge* may be best understood as deeply conditioned by what would be fitting were one using the concept to keep epistemic gate for various contextually understood communities. In attributing knowledge, the attributor recommends a source to a contextually understood community, and the standards of knowledge should then be fitting in view of the cognitive and practical stakes of that community. Thus, I seek to accommodate the purposes that seem reflected in the above cases.

Contextualists posit a variation in the standards for knowing operative across the various “contexts.” They do so in order to accommodate common judgments of the sorts noted above. The working hypothesis here is that reflection on the epistemic needs of certain broad classes of communities would mesh well with, and provide a principled bases for, common judgments about when an agent knows. Contextualism gains significantly from being so alloyed, as the resulting contextualist account of knowledge affords a more principled treatment of contexts and of contextually fitting standards than is afforded by earlier versions of contextualism. Dialectically, the situation is broadly as follows. The contextualist faces two principal rivals—sensitive invariantists and insensitive invariantists. Each position seeks to account for judgments regarding when agents know. The subject-

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

This strength of gate-keeping contextualism becomes apparent when one draws on a broad taxonomy of communities for which one might be contextually keeping epistemic gate. Consider two broad classes of communities—each of which either needs sources of actionable information or looks to serve as a general source of 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 for



such communities, one finds that what counts as knowing in the context of such a community plausibly would be sensitive to the stakes faced by that community. In this way, a motivated contextualism can account for the apparent connections between knowledge and stakes and between knowledge and practical reasoning. When an agent is faced with a practical decision, and seeks to provide for himself or herself the necessary information, gate-keeping contextualism mirrors the verdicts on which the subject-sensitive invariantist relies. Thus, gate-keeping contextualism proves nimble and powerful in treating matters of concern to subject-sensitive invariantists (particularly the data of variation in knowledge-attributing judgments with salient stakes).

Further facets emerge when reflecting on the concerns of communities of inquirers who seek to generate a stock of information on which a somewhat indeterminate range of others can reasonably draw—call these *general-purpose source communities*. Think of general scientific disciplines as constituting paradigmatic examples. The contextual demands fitting for one keeping epistemic gate for such source communities would approximate those commonly advocated by classical, or insensitive, invariantists (again, this was apparently reflected in the judgments discussed earlier). Thus, motivated contextualism plausibly has the resources to accommodate the intuitive strengths of a range of approaches to knowledge, all within a single principled approach drawing on the plausible idea that a central point or purpose of the concept of *knowledge* is to serve in epistemic gate-keeping.<sup>vii</sup>

## v. Evaluative concepts and the regulation of practices

I take it that when thinking of “the” or “a key point or purpose of some concept,” one is not merely supposing that the concept is sometimes deployed in cases in which that purpose is significant. Rather, it is supposed that the point or purpose is intimately connected with the concept in question so as to deeply condition the how that concept works. One would like to be clearer on how some

some purposes can deeply condition and thus account for facets of an intimately associated concept. When reflecting on this general idea, one may find oneself less at sea when reflecting some classes of concepts than others. For certain classes of concepts one likely has a more determinate feel for how certain purposes condition the very semantics of the concept. Thinking about point or purpose in connection with such concepts provides a possible cognitive beachhead for thinking about the general phenomenon. It might serve as a model for saying something general about concepts and purposes. At the very least, it affords one perspective on the concepts in the beachhead.

In the case of evaluative concepts the idea of deep, conditioning, connections with purposes seems particularly clear. If we could articulate the idea of a key or central purpose here, perhaps it would provide some basis for generalization. Even were it not possible to generalize, it would afford one a better standpoint for thinking about the concept of *knowledge*, as it is itself an evaluative concept.

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

To posit such a deep association between the evaluative or normative concept and the associated practice is not to indulge in what is merely a plausible historical conjecture regarding the etiology of the concept. Rather, to say that a concept arose with “a constitutive eye to” the demands of successful practice (or that what makes for success there is central to its “core conception”) is to say that the semantics constitutive of the concept turn on what makes for such success, rather than on what is merely thought at some time to make for such success.<sup>ix</sup> Thus, to say that evaluative concepts arose with a constitutive eye to the demands of success in the regulated practice indicates something important about what is central to, even constitutive of, *the concept*, versus what is merely a matter of the correctable *conception* or *understanding* one might have in connection with the concept. This is to draw on the distinction between the semantic entity—the concept—and mere conceptions.

Plausibly, evaluative concepts arose with an eye to, or sensitivity to, what makes for success in some project. The main outlines of the concept would then be subservient to needs encountered in the pursuit of the project; such needs constitutively frame or structure the concept itself. Of course, individuals develop particular understandings—conceptions—in connection with the evaluative concept. But these conceptions are not constitutive of the concept. For example, the concept of a *good driver* presumably arose to regulate the operation of vehicles (paradigmatically on public roadways) and is sensitive to the associated needs for coordinating with others who are themselves operating such vehicles, for the mitigation of the risks to self and others posed by this technology, and for the expeditious transport of people and cargo. Those possessing this concept may have various more or less developed conceptions of the character traits, capacities, and skills that make for being a good driver. However, importantly, were one to find that some character trait one had thought was necessary for being a good driver, as a matter of fact, did not contribute to the project in question, perhaps even frustrated that project, then one would need to amend one’s conception of a good driver. In doing so, one would not have changed the subject—before and after one amended

one's conception, one would be using the same concept of a *good driver*. One merely would have refined one's understanding of what makes for good drivers, of what satisfies the concept—but what does satisfy the concept would not have itself changed.

These points provide a very powerful lesson: In analyzing or reflecting on an evaluative concept, it will be important to keep in mind the project in which the concept has its roots—and whatever sense one has for how the concept serves to regulate that ongoing project. This is to say that philosophical reflection should commonly be informed and constrained by one's sense for the point and purpose of the evaluative concept we seek to better understand. What is central to the concept, what is conceptually required rather than being 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 section iii, and Henderson and Horgan, 2001, 2011.)

The central ideas here are (a) that evaluative concepts arise in connection with practices, (b) that from their inception they are constitutively such as to be attuned to what furthers their associated practice, (c) that conceptions associated with a concept are, in contrast, attuned to understandings (or misunderstandings) about what furthers the associated practice. Further, it seems that the basis for the constitutive connection between the evaluative concept and what furthers the associated practice is that the point or purpose of the concept is to regulate elements of one's pursuit of the practice.<sup>x</sup>

Epistemic evaluative concepts such as *knowledge* and *justification*, like the concepts of a *good pass* or *well-executed turn*, cannot be well understood without keeping in view the needs faced in the practice that they serve to regulate. To say that someone knows that *p*, is to say that they have here put themselves in a sufficiently good epistemic position with respect to *p*. How good? Well what is suggested in the earlier cases, and what is proposed by gate-keeping contextualism is that the crux of the answer is that it is good enough so that it is reasonable for the

contextually salient ones to take it from the agent that *p*. Quality of epistemic position here amounts to something on the order of having arrive at the truth by virtue of an epistemically good process operating in suitable conditions so that they can fitting rule out a range of alternatives to *p*. The concept of *knowledge* then seems admirably suited to the point or purpose that we found strongly suggested by ones sense of what was done in the attributions of knowledge featured in the earlier scenarios. It regulates or certifies sources in a way that is suited to their epistemic needs.

We seem to edging towards this expanded gloss on the idea that the concept of knowledge has gate-keeping as its central or key purpose. It can be unpacked in three claims:

First, the evaluative concept of *knowledge* arose in connection with the practices of

- (i) laying in a stock of true beliefs on which oneself and others might reasonably draw in both epistemic and practical contexts, and
- (ii) certifying sources of information on which communities (epistemically engaged and practically engaged) can reasonably draw,

where (i) and (ii) can be seen as two “faces” of the one “coin”: keeping epistemic gate for various communities,

Second, from its inception the concept of *knowledge* is constitutively such as to be attuned to what furthers practice (i) and (ii).

Third, conceptions associated with a concept of *knowledge* are, in contrast, attuned to understandings (or misunderstandings) about what furthers the associated practice. When the conception is rooted in a misunderstanding of what furthers fitting gate-keeping, the conception is a misconception.

Gate-keeping contextualism could be mistaken or distorting in various ways. As I have indicated, the ultimate vindication of the position requires a nuanced comparison of the explanation afforded by the best worked out variants of gate-

keeping contextualism and its competitors. That is beyond the scope of the present paper, which seeks only to develop the virtues of the view and the general character of the case for it. Having now set out the central considerations favoring the view, we need to notice some reasons to worry about whether we have properly identified the key purpose of the concept of *knowledge*. It is certainly possible that the gate-keeping evinced in the earlier cases is not a central or key purpose or point of the concept—of course it is a use of the concept, but one might insist that this use is not *key* in the sense elaborated above. Further, while keeping epistemic gate might be *a* key point of the concept, it might not be *the* key point, as there might be additional coordinate key points. If epistemic gate-keeping is not at least *a* central or deeply controlling point of the concept of *knowledge*, then gate-keeping contextualism is wholly misguided. If epistemic gate-keeping is *one* central or deeply controlling point of the concept of *knowledge*, *one among several that each condition the concept*, then gate-keeping contextualism would be a partial view—and could easily be distorting were one not to recognize how the concept is conditioned by other points or purposes. We will soon (section vii) turn to what might be thought to be additional points or purposes of the concept of *knowledge* that are suggested by reflection on further cases—and (in section viii) to how these might be connected to gate-keeping.

## vi. Generalizing Beyond Evaluative Concepts?

**[In what ways can this be extended to some non-evaluative concepts?**

- **E.g., the concept of *free action*?**
- **E.g., natural kind substance concepts and commensurability?**
  - Here the practice is one of referring to, and manipulating, stuff of homogeneous sort. What makes it homogeneous can itself be discovered—microstructure, it turned out. If folk were to inquire into what these stuffs amounted to, it would be important that the concept

David Henderson 11/3/10 12:56 PM

**Comment [4]:** This should probably be dropped.

could not change with conceptions. The Greeks who wondered about water and thought it might be a simple, fundamental stuff, and the later folk who corrected them on this, all picked out the same stuff for investigation, ... and for irrigation, and for drinking. Both practical and theoretical engagement with elements of the world is furthered by the commensurability afforded by this principled indifference to mere conception. Here, the robustness of reference across differences in (mere) conception is important for referring to a homogenous stuff that can be both the subject of evolving investigation and ongoing manipulation.]

## **vii. Other Cases and Other Apparent Uses**

I have been unabashedly selective in the cases considered to this point, focusing on cases that suggest contextual variation in the standards for knowing, and cases that suggest invariance. I have noted that the contextual variation encountered in such cases would be fitting, were one regulating epistemic sources for various communities. Thus, the cases considered to this point give one strong prima facie reason to see epistemic gate-keeping as the (or a central) point or purpose of the concept of knowledge. However, there are cases that suggest other apparent purposes for talk of knowledge. In this section, I discuss two classes of cases. First, I discuss cases in which the evaluation concerns an episode that is so separated from the evaluator/attributor that it seemingly cannot really serve to keep epistemic gate for any epistemic community. Attributions of knowledge to historical figures provide illustrative cases in which one apparently has not the slightest intention of drawing on the historical figure's results. Second, I discuss judgments in what might be termed ethological contexts, or perhaps strategic contexts.

Galileo knew that a falling body is uniformly accelerated.<sup>xi</sup> One judges that Galileo did know this much. At the same time, it does not seem that, in so saying, one is really keeping epistemic gate for some community.

In many fields, both the understanding of relevant error possibilities and the understanding of how one needs to structure inquiry to deal with such possibilities, develop significantly over time. (Double-blind experiments with placebo or related controls provide one illustration of such developments.) Commonly, theoretical or experimental developments lead investigators to appreciate that some range of phenomena might perhaps give rise to hitherto unappreciated error possibilities. With regard to source communities associated with fields that are undergoing developments of the above sorts, it will be rare that one would even think to really depend on the results of investigators working significantly earlier. Prompted by recent developments in the field and the new issues that they suggest, more recent investigators typically would have re-examined the earlier results using more powerful tools. One tends to take one's certification from the up-to-date sources. Thus, while one may admire Galileo's work, honestly, one does not flatly rely on his results. Over time, communities of investigators have amassed vastly more nuanced data using significantly more sensitive devices or probes. It then seems strange to seek to draw on resources from communities working significantly earlier.<sup>xii</sup> If one were seeking to keep epistemic gate for this later community, in the sense of determining on whose results one in that community can flatly rely, and were someone then to urge one to draw on the results of the significantly earlier community, one would resist. This resistance would presumably be indicated by insisting that those earlier researchers did not know what they thought that they knew. But, at this significantly later stage, it seems that one would hardly think of adopting the beliefs in question on the basis of the earlier work. Typically, one does not look to the earlier community as a source on the matters of concern, and one does not feel the need to keep epistemic gate at this remove. 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, as one apparently is not keeping epistemic gate for some community, what is one doing when one then look back at this wonderful historical figure and judges that he then knew that falling bodies are uniformly accelerated? In a related discussion (Henderson and Horgan, forthcoming, chapter 5) we write of a *crediting use* in addition to a gate-keeping use of evaluative epistemic concepts. One feels that, even as one is not proposing to draw on Galileo's results, one is giving a kind of credit where due. He did wonderful work in his time and place. Although our contemporaries may be in yet a stronger epistemic position with respect to such matters than anyone could have been then, and although one will take one's information from these later (filtering) sources, one feels that Galileo should be credited for putting himself in a pretty fine epistemic position—an excellent one for his time and place.

This is to recognize a non-gate-keeping use of the concept of *knowledge*—and it suggests that that concept may have a non-gate-keeping point, in addition to the gate-keeping point mentioned above. However, I will soon explore how this might naturally be understood as a non-gate-keeping use of a concept which yet has gate-keeping as its key or central point.

Now to ethology:

*The Bear:* Suppose you are out hiking in Glacier National Park. As you emerge from a forested stretch of trail, a stream comes into view. Soon, you become aware of movement on your right. You look that direction, enjoying the breeze on your face. Then you locate the movement. It is not the wind over the small meadow, but a large bear, and one that is a little too close for comfort. Does it know that you are here? You begin to back out slowly and quietly. At first, the bear seemed to be intent on feeding on some sedges. But, now, it stops. It rises onto its hind legs, and nervously sniffs the air (happily you are down wind of it). It scans its environment. It may know that some animal is about. But, does it know where you are? Suppose that it now comes

to focus intently in what is your general direction. Yes, it knows you are here. But does it know who or what you are? Word on the street (or trail) is that, at this point, you should engage in some classic human display. (No, not running.) You should talk in a distinct clear (and calm) voice. You want the bear to know that you are not one of ... well, one of its common prey species.

One should note that it is common to judge that the agent in the scenario is correct in judging that the bear knows that some manner of critter is over yonder (where the agent is). I judge that it is true that the bear knows this much, bracketing fine issues about how to characterize the content of the bear's cognitive content.<sup>xiii</sup>

It is pretty clear that when the agent in the scenario judges that the bear knows that he or she is there, that agent is not regulating some epistemic community. The agent is not judging that other bears can take it from that bear that yonder the agent stands. (The serviceability of the bear as a source of information to epistemic communities of bears—or of bears, wolves, and ravens?—is not central to one's concerns.)<sup>xiv</sup> But, what then is the agent doing in judging that the bear knows the agent's location (while perhaps not knowing what manner of critter is yonder)?<sup>xv</sup>

This use of the concept seems intimately connected with the kind of situation that decision theorist envision when they write of what an agent knows, and of what agents know of what other agents know, in some "game" or choice situation. In the context of strategic games, one typically accords to oneself and to other agent(s) some epistemic access to elements of the situation—one supposes that they are tracking certain aspects of the situation. You may suppose that the others are tracking those elements of the situation that you are tracking, or that they are tracking somewhat different (perhaps more limited) elements of the situation. You may, for example, suppose that the other actors are tracking what elements of the situation other actors are tracking—or what some of that the others are tracking, or little of what the others are tracking, or something different and additional to what the others are tracking. On such bases, one projects what are the reasonable choices

to make—given what are the reasonable choices of the other agents in the situation.<sup>xvi</sup> One need not suppose that these roughly understood others are even other humans. Rather, for millennia, humans in real engaged situations have managed to project the behavior of other humans and other critters—commonly with better accuracy than the decision theorists have managed—and have managed to make strategic decisions on this basis. The essential elements of this success has been the ability to get some fitting take on three matters: (a) the “various information-processing capacities and information-gathering abilities” that the other parties to the situation possess, (b) elements of the situation that would cue the working of those processes, and (c) some sense for how that information would then provoke behaviors. In attributing knowledge in cases like *The Bear*, one seems to be focused on the first two matters—on how the other’s cognitive processes have provided them limited veridical information about their situation. In bear scenario, the agent is registering that Mr. Bear has veridical information about the agent’s location, and has yet to be afforded information about just what manner of critter is there.

What seems important here is that the bear’s processes have been such as to discriminate the location of the salient critter in his environment (they tell in favor of there being some critter about where the agent is, versus other locale, and versus there merely being some wind moving some branches perhaps); they have not been such as to discriminate what manner of critter (in bear-terms) is yonder. In registering this, the agent appreciates (in a very visceral and intuitive way, perhaps) that he and the bear are in a choice situation characterized in terms of certain (asymmetrical) payoffs and certain (asymmetrical) risks, and certain (differing) uncertainties. One finds here a *strategic* use of the concept of *knowledge*. In contexts in which one deploys the concept of knowledge with this *strategic purpose*, one is keeping track of what oneself and other actors in the situation are getting right and the discriminative capabilities by virtue of which they are getting some things right.

In such contexts, it seems a little strained to say that one is evaluating the bear’s (or the other agents’) epistemic situation. It is not as though one necessarily

aspires to the bear being in a good epistemic situation. Instead, one merely wants to get a handle on what information about the situation the bear is getting, what misinformation it might be getting, what it is oblivious to—all in order to project its choices, largely in order to make one's own strategic choices.

How does this strategic use of the (or a) concept of compare with the gate-keeping use?<sup>xvii</sup> In the gate-keeping use, one is concerned to evaluate the epistemic situation of potential sources of information—attributions of knowledge marks sources of information, and that is a good. In the strategic use, one is concerned to keep track of the extent and ways in which other agents are themselves privy to elements of the situation—attributions of knowledge may mark a good or bad state of affairs from one's strategic point of view. The gate-keeping concern is a more straightforwardly evaluative than this. The strategic concern is a more descriptive matter.

In any case, it is likely that this strategic use represents a very venerable purpose, of a sort that would from early have rooted or motivated the use of knowledge-attributions. If it is a use of one and the same concept of *knowledge*, which seems *prima facie* plausible, then it might also plausibly be so deeply associated with the concept to function as a key point or purpose of the concept.

### **viii. Accommodating the diversity—an explanation. Is it the best?**

Of course, the results in the above section provide some reason to think that the concept of knowledge might have multiple points or purposes. In some sense, it surely does. But, does it have multiple key points or purposes in the sense discussed in section ii? It will be instructive to consider how the multiple uses just canvassed might be accommodated naturally in keeping with the working hypothesis that gate-keeping is the central point or purpose of the concept. The ultimate issue is whether gate-keeping contextualism provides the best explanation for the range of data. For

now we can address the more limited question: does it provide at least a reasonably good explanation for the data?

I have admitted that one can truly say that Galileo knew that a falling body is uniformly accelerated—and that, in so saying, one is not really keeping epistemic gate. One is not signaling that one can “take it from” Galileo that falling bodies behave as described. One is not signaling to any extant source community that Galileo’s results are of themselves ready for prime time. But, can a gate-keeping contextualist consistently insist that the point or purpose of the concept of *knowledge* is epistemic gate-keeping while acknowledging that there are attributions of knowledge, uses of the concept, in which one not actually keeping epistemic gate? For one exploring gate-keeping contextualism, this requires that one make sense of the idea that, while epistemic gate-keeping may be a central point and purpose of the concept of *knowledge*, and while this may deeply condition the character and working of the concept, the concept can yet be used in ways that do not function to keep epistemic gate for a community. It is not at all hard to understand how this might be. Attributions of knowledge to historical figures might be thought of as a kind of *simulated gate-keeping*—and simulated gate-keeping is not actual 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 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 commonly automatically employed in understanding others.<sup>xviii</sup>

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 bit of simulated or pretend gate-keeping. In the course of this, we may “signal” that our subject (for example, Galileo) was a first-rate source on which other savants, and others practically engaged (such as artillerymen of the day) could draw without reservation. 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, are dead and gone, and thus cannot actually be regulated. Thus, we do not regulate any actual present community in saying that the historical agent knew, and we do not regulate any actual past community in so saying. Rather, we understand them in their setting.

One should recognize two uses of the concept of *knowledge* in pretense. First, there is the use of the concept in simulation so as to settle on a plausible explanation of what some agent thought. Galileo, for example claimed to know that a falling body is uniformly accelerated, signaling that he was epistemically well placed to serve as a source to his theoretically and practically engaged contemporaries. One might assemble his considerations, as reflected in the historical record, and use these as input into one’s cognitive processes. One might then conclude, “I (a.k.a...Galileo) know this.” Insofar as one has succeeded in retracing Galileo’s basis for confidence in his epistemic position, one has explained why he thought that he knew here. Second, one might use the concept of *knowledge* in a form of pretense in which one does not seek to retrace some agent’s thinking, but instead one seeks to evaluate for oneself the agent’s conclusions. Famously, some of Galileo’s contemporaries claimed that he did not know that the earth revolved around the sun. One can take the public information, available to these historical agents, as input to one’s own cognitive processes. This is not undertaken to retrace their thinking, but to ascertain if and where their thinking accorded or diverged from what would be fitting. I suppose that one should here judge that Galileo knew—that his epistemic situation did

indeed tell in favor of his conclusion, versus the relevant alternatives. In so judging, one is committed to thinking that his gainsaying contemporaries were mistaken in thinking that he did not know. They should have accorded him a place—a preeminent place—in the source community of savants. Your judgment has no impact on this historical community, and is not intended to have a gate-keeping effect there. Further, you are not signaling that Galileo’s epistemic position qualifies him as an unfiltered source in one’s contemporary community. In pretense gate-keeping for this historical community, one is evaluating their own gate-keeping.

Thus, as one uses evaluative concepts in pretense, one may do so to simulate and explain others evaluation, or one may do so in order to evaluate their evaluations. In the first case, one’s pretense must plausibly retrace the thoughts and actions of those we simulate. In the second case, it need not. In none of these cases is one actually regulating any community.

It is plausible that, such non-gate-keeping uses of the concept of *knowledge* are derivative. That is, it is plausible that the concept of *knowledge* centrally serves a gate-keeping purpose and works along the lines suggested in earlier sections because it is to serve this purpose. Then, in some applications it could readily be employed in a pretense of epistemic gate-keeping—notably when thinking about agents who one is *not* evaluating as a source for oneself or one’s contemporaries. Pretend gate-keeping is not really gate-keeping, but this does not preclude the concepts used in the pretense of gate-keeping to be themselves rooted in actual gate-keeping. Pretense gate-keeping seems wholly in keeping with gate-keeping being a central point or purpose served by the concept. The concept so employed plausibly yet has its roots in epistemic gate-keeping.

But, what of *The Bear*? The strategic use of the concept of *knowledge* that this scenario illustrates is clearly pretty common, and doubtless serving such strategic purposes by attributions of knowledge, is venerable—likely as venerable as the gate-keeping use of the concept. The limits of what I can presently manage to say regarding this matter explains why I here only claim to be exploring the case for gate-keeping contextualism. The limits of what I can say on this matter also highlights what is needed for a mature abductive conclusion to the question, “What

is (are) the central point(s) or purpose(s) of the concept of knowledge?”

What I can say is this: one can make a plausible case for the idea that, *were* gate-keeping to be the central point of the concept of knowledge, then the concept would yet also be serviceable in strategic contexts for strategic uses. The case for this again uses the idea that the concept might be employed in a kind of simulation. One special case of keeping epistemic gate is that in which one is keeping gate for a community of one—oneself or another. Typically, this is a matter of evaluating sources for an individual engaged in some practical decision. In simulation, one might gauge whether the individual agent’s own epistemic situation is sufficient to the agent’s own stakes as the situation unfolds. Of course, there need be no communication between oneself and the agent under consideration—so that one would not then actually be keeping epistemic gate for that agent (recommending, or cautioning against, that agent’s relying on the agent’s own information). This application to a community of one practically engaged agent is the sort of application in which gate-keeping contextualism most perfectly and naturally mimics SSI. Now, suppose that the practical situation with which the agent is engaged is itself a strategic situation. Then, it would be natural to engage in a set of parallel simulations for each of the classes of actors in the strategic situation.<sup>xix</sup> One might do so to evaluate the agents’ choices in their strategic situations, or one might do so to explain or predict (or retrodict) their choices. Thus, again, there is no conflict between the idea that the concept of knowledge centrally serves a gate-keeping purpose, and yet is used in a non-gate-keeping fashion.

One *cannot* conclude at this point that gate-keeping is indeed *the* central point of the concept of *knowledge*. A sober and mature inference to such a conclusion requires that one be in a position to show that attributing such a point to the concept of knowledge affords a superior explanation of the full range of data than is afforded by alternative accounts. At present, I sought to explore how the gate-keeping point might provide an account of the range of relevant data, and I have found that it provides a promising explanation of much of the data. This does not show that it provides a better account than alternatives. My remarks suggest that a concept for which the gate-keeping point was central could readily be



deployed with this strategic use in view—but this is far from making it plausible that the gate-keeping purpose is *the* central point of the concept, rather than a compatible and coordinate point. To be honest, I ultimately am attracted to an account that would understand the concept as responsive to multiple key points or purposes. Keeping epistemic gate would be a central, but not necessarily *the* preeminent, purpose. This point of the concept could readily condition the central outlines of the concept while being among a set of coordinate key purposes. The kind of crediting of others that is in evidence in the historical cases is certainly a use of the concept of *knowledge*. Recognizing this much is in no tension with gate-keeping contextualism, as argued above. The strategic use of the concept may constitute a key point or purpose of the concept that is as central as the gate-keeping point of the concept.

I have argued that an adequate account of the concept of knowledge must provide the best overall explanation of a range of data. The data includes judgments about cases, apparently conflicting judgment tendencies, a familiarity with commonly employed warrant criteria, and some sense for the what is done in applying the concept in cases—together with a seat-of-the-pants sense for what might be the point of this evaluative concept. I have sought to illustrate how this initial sense for the point and purpose of the concept might be developed, but the development itself is to be judged in the judging which account is best supported. A range of background information will be significant in developing any account of the concept—as would be expected in abduction. Among the background information that has been significant for us has been information about the human capacity for cognitive simulation both in evaluation and explanation. I hope to have clarified the epistemological character of the case for gate-keeping contextualism, and for the associated understanding of the point or purpose of the concept of *knowledge*. This case makes it highly plausible that gate-keeping is *at least a central point* of the concept—although not necessarily *the* central point, as it may be one of a set of points.

Works Cited:

- Barnes, B and Ritter, D. 1993. Patterns of body temperature change in hibernating arctic ground squirrels. In *Life in the Cold: Ecological, Physiological, and Molecular Mechanisms*, edited by C. Carey et al. Boulder, Colo.: Westview Press. Pp. 119–130.
- Barnes, B. and Ritter, D. 1989. Freeze avoidance in a mammal: body temperatures below 0°C in an Arctic Hibernator. *Science* 244: 1593-1595.
- 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.
- Fantle, J. and McGrath, M. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Greco, J. 2008. "What is wrong with contextualism." *Philosophical Quarterly*, 58, 416–436.
- Greco, J. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Grimm, S. 2011. "A Rising Tide Lifts All Boats." This volume.
- Goldman, A. 2006. *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford: Oxford University Press.
- Grimm S., ms. "Knowledge, Practical Interests, and Rising Tides." manuscript.
- Hawthorne, J. 2004. "Knowledge and lotteries." Oxford: Oxford University Press.
- Heinrich, B. 2009: *Winter World: The Ingenuity of Animal Survival*. Harper Perennial.
- 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. forthcoming-a. "Gate-Keeping Contextualism," *Episteme*.
- Henderson, D. forthcoming-b. "Let's be Flexible." *Journal of the Philosophy of History*.
- 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. 2010. "On saying that someone knows: Themes from Craig." In A. Haddock, A. Millar, and D. Pritchard, eds. *Social Epistemology*. Oxford: Oxford University Press.
- Knobe, J. and Shaffer, J., forthcoming. "Contrastivism Surveyed." *Nous*.
- Nichols, S. and Stich. S. 2003. *Mindreading: An Integrated Account of Pretense, Self-Awareness, and Understanding Other Minds*. Oxford: Clarendon Press.

- Sosa, E. 2009. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford University 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>i</sup> This leaves it open whether what we are here calling "implicit mastery" is a matter of the subconscious deployment of *explicit* representations of general syntactic rules by human language-processing mechanisms, or instead occurs without the deployment of rule-representations even at the subconscious level. Either way, rule-mastery is implicit in the sense that language processing conforms to the rules without deploying *conscious* explicit representations of them.

<sup>ii</sup> It is not implausible to suggest that these judgments, when they are the products of properly exercised grammatical competence, are themselves (high-grade) *a priori*, being the product of the linguist's implicit mastery of the concept *grammatical in L* (where L is the linguist's own language). But this suggestion is not essential to the use of armchair syntactic theorizing in linguistics as a model for armchair philosophical inquiry in pursuit of conceptually grounded necessary truths.

<sup>iii</sup> See, for example, DeRose (1992, 1995, 2009), Cohen (1999, 2005), Fantl and McGrath (2009), Hawthorne (2004), Stanley (2005).

<sup>iv</sup> These examples are inspired by Bernd Heinrich (2009), chapter 7. That discussion draws on Barnes and Ritter (1989, 1993).

<sup>v</sup> I begin here with a somewhat vague conception of epistemic gate-keeping, something along the following lines: In our epistemic life together, we rely on each other. One relies on different individuals and groups for information about different matters. One insists that they use processes and achieve results that put them in a good epistemic position with respect to matters that concern one. If it is a matter of practical concern one may hold out for one who is in a fairly good, or a truly wonderful, epistemic position with respect to the matter of moment—depending on the stakes one faces. So one looks for sources that are up to the corresponding demands—or one seeks to satisfy those demands oneself. Further, one may take on certain roles in which one, as a member of some group, seeks to provide information for others. Here, one must insure that both oneself and the others in the source group are uniformly up to this role—as the group is little better than its weakest link. Such are the crude ideas I associate with epistemic gate-keeping. I intend the reflection found in this paper, in Henderson (2009) and in Henderson and Horgan (forthcoming, chapter 5) as a first installment on clarification.

<sup>vi</sup> Thus, we are not recommending Craig's avowed practice, which supposedly starts with a conception of the purpose and explicates the concept largely on that basis. However, it is worth noting that even in Craig (1990) the explicated concept is also tested against judgments about cases. The proof is said to be "in the pudding"—so that Craig's actual practice is not naively dominated by his sense for the purpose of knowledge attributions.

<sup>vii</sup> It would be good to consider how there can be communities of an intermediate sort, and how expectations for the social flow of information enter into the demands in various contexts. See Henderson (2009) and Grimm (2011).

<sup>viii</sup> I will not here worry over nuances in the differences between concepts of *good X-er* and *competent X-er*—nor even between those evaluative concepts that treat of agents and those that treat of acts. Sosa (2009), Greco (2008, 2010), and others have explored such matters, and it is beyond the scope of this paper to contribute to their results. The present point has to do with something general to concepts of each sort: there is an intimate conceptual connection between the needs of a practice that is regulated using a given evaluative concept and what is semantically central to the evaluative concept itself.

<sup>ix</sup> Compare: natural kind substance concepts arose with a constitutive eye to what was deeply shared in the preponderance of samples—supposing that they did have some deep, shared, basis. This conception of a common shared basis shared across samples, constituting a stuff to which folk may refer—despite their disparate and changing conceptions of what exactly this shared basis might be—is a core conception in the case of natural kind substance concepts.

<sup>x</sup> The points just made regarding evaluative concepts, themselves could stand for much clarification and elaboration. One might, for example, wonder how one should identify and evaluate practices, and what this might suggest about how concepts should be individuated. For our purposes here, we need little more than provided already. But, it is perhaps worth discussing the open questions just a little. The examples we have used, that of a *good driver* and that of a *good archer* are associated with practices that have undergone significant change over time. At one point, archery involved a technology war and hunting, it later came to be largely a kind of contest. Arguably, what contributes to success in the practice (or practices?) is not the same early and late. The character traits and skills that contributed to success in war and hunting would include the trait of having "nerves of steel"—so that one's performance did not degrade in settings with "all hell breaking loose." In Now that archery is a modern contest, competitors are typically provided a setting conducive to composing oneself. Thus, it arguably required more to count as a good archer when one was putting together a war or hunting party than it does when one is putting together a team. Is this one practice, evolving over time, or two? Are there one or two (admittedly related) concepts here. Similar remarks and questions apply to the practice of driving and the concept of a good driver—as one moves from operating teams of animals on a public roadway to operating motorized vehicles.

One reasonable answer would be to suggest that in each case there is one generic concept—a *good archer*, a *good driver*—and that there are also more specific concepts—conceptini—such as a *good driver of animal teams*, and a *good driver of motor vehicles*. We need not insist on this answer here. What one should be confident of is that early and late the concept (and conceptinis, if any) is (are) directed to regulating the practice in which it is (they are) rooted.

<sup>xi</sup> That is, he knew that an equal increment of speed is added in each successive interval of time. In judging this to be known, I am also judging that it is true—and I must then be

understanding it as implicitly limited or qualified in various ways that are in keeping with applications near the earth, and bracketing the effects of the atmosphere.

<sup>xii</sup> Elsewhere Henderson (2011) suggests that what counts as “significantly earlier” and “significantly later” may be understood in terms of expectations for results in the field to become dated by methodological developments or by changes in the richness of data.

<sup>xiii</sup> Clearly there are several issues here. When we say that it is common to judge it correct to say that the bear know that the agent is yonder, we report our own immediate response to the case, and what we believe is the very common immediate response. Of course, many philosophers and other academics have long been taught to be cautious of belief attributions to animals—and this will produce a second thought among many. We do not share this aversion to attributing beliefs to animals—although we find the characterization of their mental contents a very delicate matter.

<sup>xiv</sup> One may believe that male bears over the age of 2 are relatively solitary creatures that take part in few significant epistemic communities. Perhaps bears use other bears as indicators, but do not engage in much ursine testimony.

<sup>xv</sup> Kornblith would suggest that this scenario reflects a use of the concept of *knowledge* that is common in cognitive ethology:

The knowledge that members of a species embody is the locus of a homeostatic cluster of properties: true beliefs that are reliably produced, that are instrumental in the production of behavior successful in meeting biological needs and thereby implicated in the Darwinian explanation of the selective retention of traits. The various information- processing capacities and information-gathering abilities that animals possess are attuned to the animals environment by natural selection, and it is thus that the category of beliefs that manifest such attunement—cases of knowledge—are rightly seen as a natural category, a natural kind (Kornblith, 2002, pp. 62-3).

However, to speak of an ethological/explanatory use with the features Kornblith lists is at least a little over-intellectualized in connection with our bear scenario. Focusing on what is done in cases of the venerable sort illustrated in the bear scenario, one does best to think in more strategically engaged terms.

<sup>xvi</sup> But, one need not suppose that the agents are highly rational. One certainly need not use any articulate mathematical theory of maximization.

<sup>xvii</sup> Biologists do use knowledge attributions loosely. They write of animals knowing when the time has come to migrate, to end hibernation, and the like. They then ask how they know these things. They even write of trees knowing when it is time to flower or to leaf out. (Several common northern trees form leaf and flower buds in midsummer. They then halt the development of these structures. Come spring, they then initiate development of these structures on differing schedules. Wind pollinated trees commonly flower before they leaf out.) One might see such talk as loose talk. Or one might think that the talk involves a separate knowledge concept. We do not think that the strategic use evinced in *The Bear* reflects either loose talk of a distinct concept from that familiar to epistemologists.

<sup>xviii</sup> See Goldman (2006), Steuber (2006), and Nichols and Stich (2003) for recent overview and discussions. Goldman has been a consistent proponent who maintains a balanced and empirically informed perspective. Steuber provides a useful perspective, making connections with wide literature in the philosophy of social science. Nichols and Stich provide a grudging and caution perspective, informed by a reading of much empirical psychology.

---

<sup>xix</sup> In practice, one commonly employs a mix of pretense and general descriptive information (Henderson 1996, forthcoming-b.)