

## Knowledge in Practice

1

### 1. Knowledge as a Deontic Concept.

The mainstream tradition in modern philosophy teaches that knowledge is to be understood in *deontological* terms. In broad terms:

(i) To know is to enjoy epistemic authority. Knowledge is *justified*

true belief: belief to which one is epistemically entitled.

(ii) Epistemic authority is acquired/maintained living up to one's epistemic obligations.

(iii) These obligations involve the ability to give (thus the possession and proper use) of reasons and evidence. There is an essential connection between being justified and being able to justify.

Knowledge is a deontic concept in the sense that knowing is a matter of rights and duties.

For some years now, this approach to understanding knowledge has been subjected to a sustained assault from various forms of epistemic reliabilism, according to which knowledge is true belief acquired and maintained by some truth-reliable cognitive process, or true belief that somehow co-varies with the facts (as in Nozick's "truth-tracking" account). But in my view, the traditional approach remains the option of choice. In trying to make good on this suggestion, my strategy will be to examine knowledge—or more precisely epistemic concepts—from a pragmatic

standpoint, one that takes seriously the question “Why do we have such concepts in the first place?” I will end with a few words about skepticism.

## 2. Pragmatism, Meaning and Use.

The pragmatism I have in mind is not the classical pragmatism of Pierce, James and Dewey but contemporary neo-pragmatism (which of course has affinities with its forebear). David Macarthur and Huw Price have characterized this form of pragmatism in terms of an equation:

$$PRAGMATISM = LINGUISTIC PRIORITY + ANTI- REPRESENTATIONALISM$$

Let me offer a few words in explanation and defence of the *pragmatic* stance.

Linguistic Priority counsels not to start by asking about (say) the nature of values or modality or meaning or knowledge, but rather to start by examining what is distinctive about evaluative, modal, semantic or epistemic vocabularies. To see the force of this suggestion, consider meta-ethical expressivism: if this is even a coherent theoretical option, we can't just assume that there are “moral facts” that need to be “placed” in the world of fact or detected by a special faculty: metaphysical and epistemological problems thus disappear. If epistemic vocabulary is deontic vocabulary, it may not be fundamentally descriptive either.

In approaching these problematic vocabularies—indeed in approaching meaning generally--anti-representationalism counsels us to eschew any fundamental explanatory appeal to semantic notions, particularly truth and reference. Meaning is a matter of *use*. However, use is a complex notion. There is use as *usage*: *how* a word is to be used. But there is also use as *function*: what a word

may be *used to do*, or what having in the language is *useful for*. An explanation of meaning in terms of use (an EMU) will address both aspects.

Each dimension of “use”, usage and function, can be divided into two components. With respect to use as usage, there are material-inferential proprieties: the inferential connections (perhaps not fully specifiable, if we are suspicious of a strict analytic-synthetic distinction) between sentences involving the target word and certain other sentences. But there is also the epistemic character of those inferential commitments: for example, are they “primitive” (*a priori*) or are they open to empirical revision. In the case of use as function, we should distinguish *expressive or performative function* from *utility*: what we are doing in deploying the word in question (typically in a non-embedded declarative sentence), and the ends we accomplish by doing that.

In an EMU, the two aspects of usage specify a word’s content-determining use-proprieties (CDPs), thereby determining its conceptual content: *what we are saying* in deploying the word in question. The two aspects function specify its functional characteristics (FCs). The two dimensions of use, usage and function, are mutually constraining. The CDPs must be such as to *enable* the word’s FCs. Its functional character will *rationalize* its CDPs.<sup>2</sup> We have a sketch of a meta-theoretical analysis of an EMU.

The analysis reveals the important point that *expressive function* does not preclude *descriptive content*. We can therefore resist the suggestion, central to familiar forms of meta-ethical expressivism, that there are two types of meaning (or content): expressive and descriptive, where the latter associated with a vocabulary’s

being “robustly” representational.<sup>3</sup> However, there is a way in which this pragmatic approach to meaning is “function first.” CDPs are ten-a-penny: so why do we retain a word governed by some particular set? The FCs answer this question. “Function-first” is especially attractive in the case of philosophically problematic vocabularies. Instead of asking “what are they about?” or “What do they refer to?”, we can start by asking “What is the point of them?” (Klemens Kappel calls this “practical explication.”) This approach will vindicate a deontological understanding epistemic vocabulary, while allowing us to take due notice of reliabilist insights.

### 3. Attributing Knowledge: What is the Point?.

As Edward Craig notes, it is obvious enough why, we have always talked about the sheep and the goats and the fields and the corn. But why did we ever come to talk about knowledge? And what does the answer this question tell us about the content of such talk? Craig takes a “genealogical” approach to these questions.<sup>4</sup> He invites us to imagine a group of humans in a “state of nature”: they lack complex institutions, and they lack epistemic concepts. Now for Craig, the original home of knowledge-attributions is a question-context: what I want to determine is *whether* p, and I can’t easily settle all such questions single-handed. So, too, the humans in the state of nature since: they have informational needs that exceed their individual capacities for information-gathering, and therefore have a need to tag and keep track of others who can help: good informants. Introducing knowledge-talk enables them to do this. The practice of knowledge-attribution facilitates information-sharing among epistemically limited subjects, something that is evidently useful. But

in identifying someone as a (potential) good informant, I express a readiness to defer to him with respect to some question of current interest: so the core expressive/performative function of unembedded knowledge-attributions is to grant epistemic authority. Craig's thought-experiment offers suggestions about both aspects of the functional characteristics of knowledge-talk. His method is function-first.

Close in spirit to Craig's genealogical account is Brandom's pragmatic recasting of the traditional "justified, true belief" (JTB) analysis of knowledge.

According to Brandom:

In taking someone to be a knower, one *attributes a commitment, attributes entitlement* to that commitment, and *acknowledges commitment* to that same content oneself...These correspond...to taking to believe, taking to be justified in that belief, and taking the belief to be true.<sup>5</sup>

Brandom concludes that knowledge is a "complex, hybrid deontic status": complex because it involves both commitment and entitlement; and hybrid because the *attributor* acknowledges commitment (and assumes entitlement) on the *claimant's* authority. However, where Craig stresses utility, Brandom highlights expressive/performative function.

There is a methodological issue that needs to be noted. *Linguistic Priority* is really *conceptual* priority: we are not interested in the word "know" but in the functional role (inferential and expressive/performative) that the word (and its foreign counterparts) plays in our practical-linguistic engagements. Now while Craig approaches the *concept* of knowledge by way of a story about the introduction

of knowledge-*talk*, it is surely the case that I can show deference without the aid of explicit epistemic vocabulary. To adapt an example from Craig, if I see that Fred is up a tree, with a good view of the surroundings, I may be disposed to heed his warning that lion is approaching. Thus it looks as though, in some way, I must already grasp the *concept* of a person's being knowledgeable. This looks like a problem for a genealogical account of development of the concept expressed by "know," in so far as Linguistic Priority entails that possessing the concept depends on having the word.

Brandom's non-genealogical analysis helpful here. For Brandom, introducing a new kind of talk often functions to make explicit commitments that we could otherwise only display in practice. Accordingly, the doings that feature in his pragmatic recasting of the JTB analysis—the doings involved in epistemic deference--can be understood as describing a practice that potentially antedates the introduction of explicit knowledge-talk, though not the development of talk as such. Epistemic vocabulary is a kind of secondary vocabulary, presupposing more fundamental linguistic-conceptual abilities. Craig's genealogical story does not require us to have concepts before we have them, only that prior forms of discourse exemplify relations of epistemic deference that we lack the resources to make explicit in claims. To borrow a word from Michael Oakeshott, the practical basis must contain "intimations" of the more sophisticated way of talking, which allows us to *say* what we could otherwise only *show*. This suggests an expansion of the functional account. Explicit epistemic vocabulary makes the practice of information

sharing far more flexible and powerful. With it, we can share information *about* informants and investigate their trustworthiness in a more self-conscious way.

#### 4. Authority: Reliability, Responsibility and Reasons.

The next question is: what usage proprieties would comport with knowledge-talk's functional characteristics? That is, what material-inferential commitments does knowledge-talk embody, and what is their epistemic character? More concretely, what makes for a good informant, and how do we recognize one? Call this latter question "the Recognition Problem."

Knowledge is justified, true belief. Or as I prefer to state the JTB account,

Knowledge → true belief + epistemic authority.

The truth and belief conditions are unproblematic. Obviously, a good informant, with respect to some particular question, must possess correct information. Also, as Craig points out, settled belief entails confidence. Since good informants need to be recognized, this is important, for we will hesitate to take someone's word when (we suspect) he lacks conviction himself.<sup>6</sup> Much more controversial is the character of epistemic authority. According to the view I shall defend,

Epistemic Authority → Truth-Reliability + Epistemic Responsibility.

Epistemic authority is thus appraisable in two dimensions, and it is vital to get a handle on the ways in which considerations belonging to these dimensions interact.

Knowledge requires truth-reliability, expressed by some modal condition.

There are various candidates and I won't try decide which, if any, captures *the* reliability condition implied by our ordinary conception of knowledge. Perhaps

there is no single condition: it may be that the reliability requirement is in various ways contextually variable. For present purposes, the important question is why are we interested in reliability at all? Why do we care about what an informant *would have* believed in other (non-actual) situations? Again, Craig supplies the answer, which is that we don't always know *exactly* what situation we are in. We therefore want an informant who can be counted on to provide correct information across a range of possible situations.<sup>7</sup>

Truth-reliability is a vital component of epistemic authority. However, because of the need to address the Recognition Problem, the function-first approach to knowledge points beyond an austere reliabilism. As Craig notes, *determining* an interlocutor's reliability with respect to whether-p is the key to singling him out as a potentially good informant. A useful informant must be *better placed* than I am (by situation, visual acuity, whatever) to settle the question in hand. But recognizing that someone is better placed than I am to determine whether-p requires *some grasp of reliability-relevant factors*—particularly, a grasp of one's limitations and those of others--is built into the practice of knowledge-attribution on the ground floor, by virtue of its functional rationale. Primitively, this grasp of reliability-relevant factors will be implicit and practical. But it will be an essential element in the underlying practical basis of the sophisticated procedures that come with explicit epistemic vocabulary.

These reflections point to a vital distinction, elided by austere reliabilism, between *informants* and mere *information sources*. The clock tells the time, we say. But not really. We can tell (= find out) what time it is by looking at the clock,

assuming it is reliable and correctly set. But the clock doesn't tell us anything, in the sense in which a person can tell us things. Only beings with beliefs and communicative intentions can do that, and clocks have neither. Clocks indicate the time, but not by virtue of knowing what time it is. Epistemic vocabulary concerns informants, not information sources, which can be mere indicators. Of course, it is possible to treat someone's actions and reactions (including his saying this or that) as *evidence* that such and such. But to treat someone this way is not to treat him as an informant.

Turning from reliability to responsibility, epistemic authority involves epistemic responsibility in three senses of "responsibility": accountability, due diligence, and liability.

We are accountable (epistemically responsible<sub>1</sub>) for our beliefs in much the way that we are accountable for our actions. The antonym of 'responsible<sub>1</sub>' (if it has one) is 'non-responsible'. Being an epistemic subject *just is* being accountable for what one believes, in the way that being an agent just is being accountable for what one does. Just as agents alone can enjoy rights and be subject to duties, only epistemic subjects can possess epistemic authority and be under epistemic obligations. No authority without accountability.

Accountability draws the line between informants and mere information sources. It is not an accidental feature of informants that they can be questioned, challenged, asked to explain themselves. Susceptibility to epistemic questioning is essentially involved in the game of "telling." Accountability implies some measure of epistemic self-consciousness. This may be rudimentary and more displayed in

practice than explicitly articulated. But it cannot be altogether lacking. Beings that cannot *account for themselves* cannot be held to account: human knowledge differs from animal cognition, just as human action differs from animal behavior. This is a further reason why human beings need some sense of the epistemic strengths and weaknesses of themselves and others, which we noticed in connection with the recognition problem.

Epistemic authority depends on fulfilling one's epistemic obligations. Put another way, it depends on duly diligent (epistemically responsible<sub>2</sub>) belief-management. The antonym of 'responsible<sub>2</sub>' is 'irresponsible.' Sometimes, we *find* ourselves in a position to determine whether p. But often we need to *put* ourselves in such a position; and even when in position, we need to exercise due care and attention (e.g. take a good look).

Due diligence furthers but does not guarantee (situational) reliability. It confers what is sometimes called "personal justification" but does not guarantee knowledge. I can form a belief in an unreliable way, despite my best efforts. I may collect and make proper use of what is to all intents and purposes conclusive evidence. But through no fault of mine, it may be misleading. Often enough, responsibility and reliability go together. But when they fall apart. When they do, and I end up with a true belief anyway, we have a Gettier case: true belief that is justified *in a way* but that does not amount to knowledge.

Finally, we must distinguish accountability (responsibility<sub>1</sub>) as a *general condition* from liability (responsibility<sub>3</sub>): that is, being subject to sanction *in a particular case*.

In discussing accountability, I likened accountability for one's beliefs to accountability for one's actions. However, when it comes to liability, we find a structural difference. In the case of action, the *liability structure* is default and query. We are presumed liable (thus praiseworthy or culpable) for (at least some of the consequences of) an action unless we have an excuse that calls in question our (degree of) responsibility. In the epistemic case, by contrast, liability is more like what in the law of torts is referred to as *strict*. We can violate the reliability constraint even though we have fulfilled all our epistemic obligations. Although we have not been derelict in our epistemic duties, we are still subject to epistemic sanction in the form of loss of authority. If I know, and you don't, that you are in Goldman's Barn Façade County, I will not defer to you in the matter of whether there is a (real) barn in Farmer Jones's field or (in all likelihood) anyone else's.

At the outset, I said that a deontological approach to knowledge will give a central role to what I have now labelled epistemic responsibility<sub>2</sub>: living up to one's epistemic obligations. I also claimed that these obligations essentially involve the ability to justify one's beliefs, even though (as reliabilists insist) knowledge often—even typically--arises from the unself-conscious exercise of basic cognitive capacities such as perception and memory. How can this be? Is it sometimes permissible to believe without reasons, and if so why and when? To answer these questions, is to delineate what I call (with a nod to Sellars) "the deontic structure of the space of reasons."

For questions of epistemic responsibility to arise, there must be things that we are (potentially) authoritative with respect to, thus accountable for: judgments,

beliefs, or (as Brandom says) commitments, all of which imply conceptual capacities. Pragmatism treats such capacities as bound up with the possession of linguistic abilities, and treats linguistic meaning as constituted by use; and it is clear from the work of philosophers who have taken this approach, particularly Davidson and Wittgenstein, that a person could not possess conceptual capacities without making judgments and holding beliefs that are not only mostly true but also epistemically unproblematic.<sup>8</sup> Unless I get routine matters mostly right, I will not manifest the stable, interpretable patterns in my speaking and thinking which, from the use-theoretic standpoint, *constitute* conceptual content.

Wittgenstein is particularly interesting in this regard. For many routine judgments and beliefs, the idea of *making a mistake* gets no purchase: a mistake. Wittgenstein surmises, must fit in with other things that a person judges *correctly*. Now to say that there are situations in which it makes no sense to think of someone as making a mistake is not to say that, in such a situation, a person could not say (or believe) something false. However, were he to do so, or perhaps even were he to seem seriously to entertain doubts, we should wonder what (if anything) he was thinking. It follows that such situations are necessarily exceptional, on pain of a person's ceasing to be an epistemic subject (or fully a person). For many matters, at least in normal contexts, an epistemic subject automatically and quite properly takes himself to possess and is granted epistemic authority.

We all know that this is so. It is an evident feature of our practice of knowledge-attribution that there are situations in which authority is properly granted automatically, and situations in which it isn't. The knowledge of when it is

proper to grant or withhold authority is implicit and practical in character, more know-how than know-that. It is not self-consciously formulated but manifested in the structure of our practice of granting and questioning epistemic authority. That structure is *default and query*.<sup>9</sup>

(DQS) In many contexts and with respect to a wide range of claims, claimants possess and are properly granted default epistemic assertional entitlement: entitlement that does not depend on the claimant's doing or having done any *specific evidential work* or possessing any *citable reasons* for his commitment. Default entitlement is *situationally* unearned and imposes no *standing* defence requirement.

Default authority is perhaps most evident in what we may call a *prototypical communicative context*. This very much the sort of question-context envisaged in Craig's state-of-nature account of the function of explicit-attributions. Such contexts involve the sharing of information that (relative to the mutually-recognized background knowledge and cognitive capacities of the interlocutors) is not notably exotic or surprising, either in itself or with respect to the claimant's access to it. In such a context, epistemic queries—How-do-you-know or Why-do-you-believe-that questions will not normally arise. For one thing, an inquirer, whose interest is in whether-p, is likely to have settled them in the very act of identifying an appropriate informant.

There are other kinds of context in which information is exchanged. In a context that is investigative, explanatory, or otherwise involves controversial claims, there may be standing objections that preclude default authority. But even in

a prototypical communicative context, epistemic queries can come up. Concerns about concerning an informant's reliability or due diligence may be motivated by *red flags*: I have asked you when the next train to the city leaves, taking you to be familiar with the journey, but then notice that the timetable on your desk is last year's. Or I may have *background knowledge*, that I am not sure that you share. I heard on the radio that repairs to the track are producing delays: were you aware of that when you responded to my question? The reasonableness of questioning authority may also be influenced by the costs of error (so-called "pragmatic encroachment"): I *really* need to catch that train, so I want to be sure that you are sure (and have the right to be).

The upshot is that although DQS imposes no standing obligation for a claimant to defend or explain his claim, it does entail a *Contextual Defence Obligation*.

(CDO) Authority, whether default or earned, is typically suspended by a *contextually appropriate query*. To regain authority, a claimant must respond adequately to the query, undermine its appropriateness, or otherwise neutralize it. If the claimant cannot deal with the query, his authority lapses. Because questions must be contextually appropriate, both claims (or beliefs) on the one hand, and questions (or doubts etc.) on the other are equally subject to normative constraint. Just as claims (or beliefs) can lack authority, questions or objections can be inappropriate, weak, irrelevant, and so on. There is no automatic right to question epistemic authority.

We now see how *being* justified is connected with the *ability to justify*. Epistemic authority is essentially connected with being able to establish one's authority: by providing evidence for what one has claimed, displaying one's credentials, or whatever. However, "being able to establish one's authority" means "being able to do so *when required*." Whether it is required is a contextually sensitive matter, not a universal requirement.

It follows that, because authority is often possessed by default, it is not always irresponsible to believe without evidence, at least if evidence is equated with usefully citable reasons. Authority often, even typically, arises from the unself-conscious exercise of ordinary cognitive capacities, from testimony, and from other sources. Reliability may be (responsibly) presupposed. Nevertheless, because knowledge requires both responsibility<sub>2</sub> and reliability, default authority entitlement carries forward-looking *reliability commitments*. Reason to suppose that reliability constraints have not been met is a standard route to a contextually appropriate query.

Following Brandom, I used to call this structure "default and challenge". I changed to "default and query" because "challenge" suggests that "How do you know?" is asked, as Austin remarks, *pointedly*, the implication being that perhaps you *don't* know. But "How do you know?" may simply be a request for an explanation, no suspicions implied.<sup>10</sup> I now think that this is its *usual* force. In asking "How do you know that p" I concede that you do know that p: I want to know how you know. I want an explanation. I am not challenging you to justify yourself. The more usual form of a pointed query is "Why do you believe ... ?" Here I concede that you do

believe, but imply that I am dubious, or at least agnostic, as to whether you *should*. I therefore now reserve “challenge” for hostile queries. I call non-hostile queries “epistemic questions.”

Whereas appropriate challenges have a strong tendency to *suspend* authority, epistemic questions do (or anyway need) not. For example, you might want to know how I know something so that you can find out similar things (or get further information) for yourself. Far from impugning my authority, your question depends on conceding it. The appropriateness conditions for challenges and epistemic questions, and the ways in which they themselves can be questioned, will therefore differ. In the case of a challenge, I may need to ask you how you think I may be going wrong, a clearly inappropriate response to an epistemic question. But authority can *lapse* by virtue of a failure to respond to queries of either type.

Let me now return briefly to the function of epistemic vocabulary. I have endorsed Craig’s suggestion that the core function of the concept of knowledge is to identify and keep track of good informants, and Brandom’s claim that the speech-act function of saying that someone knows that *p* is to attribute to him epistemic authority with respect to *p*. However, I have also claimed that strong intimations of such a practice can—indeed must—underlie the introduction or use of explicit epistemic vocabulary, a point suggesting that a more nuanced account of the functional significance of explicit epistemic vocabulary is needed.

The need for such an account is reinforced by a point so far unmentioned: that the use of such vocabulary is rarer and more contextually restricted than we might expect, given the functional rationale as presented so far. The default form of

expressing a claim to authority is not to say that one knows that p, but simply to claim unqualifiedly that p. If I speak so as to settle a question for you, while having doubts about whether I am in a position to settle it, I violate a fundamental norm of co-operation. If I have reservations about my epistemic standing, I should issue a disclaimer. If I don't, and you have no reasons to distrust me, you are entitled to take me to be claiming authority. In the same vein, the default way of expressing epistemic deference is not to make an explicit, self-conscious judgment to the effect that I know that p, but simply to take my word for it. *Explicit epistemic vocabulary is primarily deployed in contexts of epistemic questioning*: i.e. when the question at issue is not whether-p, but whether an interlocutor knows that p. Explicit epistemic vocabulary belongs, as Sellars says, "in a higher court": it is for raising and answering questions *about* epistemic authority, though of course (in the higher court) it can also be used to express claims to or attributions of such authority.

In his genealogy of the concept of knowledge, Craig makes it clear that the primitive context in which knowledge-talk is introduced is special. It is a situation in which I want to identify and keep track of a good informant with respect to a particular and rather narrow question in a highly specific question-context. In Craig's example, I want to settle the question of whether a lion is approaching, so I ask Fred who is up a tree to keep a lookout. But of course, as our needs and interests expand, we need good informants with respect to questions that are not tied to such particular subjects, informants and situations. Craig calls the process that generates a more expansive practice of knowledge-attribution "objectivisation." I suggest that a prime function of explicit epistemic vocabulary is further

objectivisation by enabling us to investigate epistemic authority in ways that are not tied to highly particularized question-contexts.

### 5. Skepticism: the Agrippan Problem.

Let me close with a few words about regress (or as I like to call it “Agrippan” skepticism. At the heart of the problem is an enticing but apparently fatal trilemma, Agrippa’s Trilemma. I make a claim. The skeptic asks “How do you know?” I explain. The skeptic re-iterates his question: “How do you know *that*?” Three possibilities emerge. I can:

- (i) keep responding by finding something new to say, in which case I embark on a vicious infinite regress;
- (ii) at some point, refuse to answer, in which case my “knowledge” rests on a mere assumption; or
- (iii) recur to something I have already said, in which case my reasoning is circular.

In no case do I explain how I know. Since there is no fourth option, the skeptic concludes that knowledge is impossible.

Many philosophers continue to take the Agrippan problem at face-value. So to defend the possibility of knowledge, they must put a better face one of the three horns. They can argue that the regress is not really vicious, that there are foundational beliefs, or a belief’s being justified depends on its belonging to a “coherent” system, coherence arising from logical and other relations among beliefs that are not in themselves relations of justification. I reject all three approaches. I

have argued that, while epistemic authority can arise “non-inferentially” through the unself-conscious exercise of basic cognitive capacities, it also requires some sense of one’s reliability. This can seem to push me towards the coherentist camp. On the other hand, I do not think that authority can arise from coherence alone: it depends on the possession of capacities that are actually truth-reliable. Since infinitism strikes me as a non-starter, I seem to be left with no response to the Agrippan problem. But this is not so.

The key to the problem is to *question the pattern of questioning* that generates the trilemma in the first place. Outside philosophy class, we don’t tolerate the persistent, regress-threatening questioning that the skeptic goes in for. Should we? Since we are approaching epistemic concepts from a deontological point of view, the question to ask is who or what gives the skeptic *the right* to behave like this? The answer will be that I do. If I merely express an opinion, the skeptic has nothing to say to me. But if I lay claim to epistemic authority, the skeptic is entitled to ask me how I came by it.

This is the Agrippan skeptic’s *Founding Intuition*. But although it may seem innocent enough, a little reflection shows that it imposes a severe claimant-querier asymmetry:

(CQA) If I represent myself as knowing that P, there is nothing you have to do, no way that things have to be, or no special position you have to be in, in order for you to have the right to raise an epistemic question (i.e. to ask me how I know).

From the standpoint of practical explication, CQA *cannot* be a correct account of the deontic structure of the space of reasons, at least with respect to everyday claims to/attributions of knowledge. Imposing an unrestricted right to question and correspondingly unrestricted obligation to respond, in the manner of CQA, would make the game of giving and asking for reasons unplayable, thus pointless. But if my argument has been correct, there is no reason to read such a self-defeating structure into our everyday practices of epistemic assessment. On the contrary, there is every reason not to.

From an anti-skeptical standpoint, the merit of DQS is that it replaces the extreme pro-query bias of CQA with a *reasonable distribution* of epistemic burdens among claimants and challengers/questioners. Reasonable isn't equal. DQS involves a moderate claimant-challenger/questioner asymmetry in the opposite direction to that presupposed by the Agrippan skeptic: i.e. in favor of claimants.

This pro-claimant tilt is a practically reasonable. Both challenges and epistemic questions interrupt inquiry/conversation/practical projects and thus impose opportunity costs. Pragmatism counsels us to take practical interests/finite resources seriously. We are (rationally) not ordinarily interested in investigating a belief's pedigree, when it is one of our own, unless we find reason to do so.<sup>11</sup> If we find ourselves with conflicting beliefs, revision may be called for. But even belief-revision does not produce an automatic interest in doxastic pedigrees: we can often decide *where* we must have gone wrong, without knowing *how*.

Knowledge attributions are rooted in a practice of information exchange and mutual correction. Practices of epistemic deference and evaluation are rooted in the

fact of positional advantages and disadvantages with respect to determining whether p. In primitive situations, the advantages and disadvantages are literally positional: who is/was where, when. Later, they reflect education and training. Not only are we normally not interested in questioning a belief's pedigree (believing that p has settled the question as to whether p), we are not generally well-placed to do so. Challenges and questions generally come from interlocutors with interests and information differing from our own.

From the standpoint of practical explication, practical considerations matter. But as our discussion of the conditions of accountability suggested, more than practicality is at stake. Widespread default authority is a condition of epistemic subjecthood.

In previous writings, I have defended an understanding of the practice of reason-giving in terms of a default and query structure on the grounds that it gives the best diagnosis of how Agrippan skepticism goes wrong. To bolster this defence, I have argued that my account is *phenomenologically* preferable to the Agrippan skeptic's implied position in that it stays closer to the surface character of ordinary uses of epistemic vocabulary. In this paper, I have attempted to provide a more general *theoretical* defence that can be developed independently of skeptical concerns. If I have succeeded, my diagnosis of skepticism is that much better.

Michael Williams

The Johns Hopkins University

September 2012

---

<sup>1</sup> EPM § 6, p. 131.

<sup>2</sup> Since there may be more than one way of achieving the same practical ends, the relation of CDPs to FC is potentially many-one, while that of FCs to CDPs is one-many.

<sup>3</sup> HPCBLE

<sup>4</sup> KSN. Such an approach may not yield a strict set of necessary and sufficient conditions for the truth of “S knows that P,” once the holy grail for many analytic

---

epistemologists. But it may do something more illuminating: it may explain why the concept of knowledge has generated the kinds of disputes about how to understand it that epistemology has made familiar.

<sup>5</sup> MIE

<sup>6</sup>.<sup>6</sup> As Craig notes, a good informant must be willing to share his information. There is a norm of *truthfulness*. However, this norm is an application to communication of a general norm of co-operativeness, with no special connection to epistemic affairs. A potentially good informant can keep information to himself without compromising his epistemic authority. Accordingly, truthfulness is not required for knowing.

<sup>7</sup> Reliability can be both general and situational: someone who is generally reliable (to be relied on) with respect to a certain kind of information may not be reliable in a particular situation where (perhaps unknown to him) the conditions are not propitious for the exercise of his generally reliable cognitive capacities.

<sup>8</sup> Sellars is more problematic because of his conceptual relativism.

<sup>9</sup> Austin and Wittgenstein.

<sup>10</sup> Austin, OM. Neta.

<sup>11</sup> Levi.