Who Knows What?

Epistemic Dependence, Inquiry, and Function-First Epistemology

Joshua DiPaolo | California State University Fullerton | jdipaolo@fullerton.edu

Abstract: Function-first epistemologists analyze epistemic concepts, norms, and practices by investigating their *functions*. According to the most prominent function-first account, the primary function of our concept of knowledge is identifying reliable informants. In this paper, I take for granted the function-first methodology to achieve three main goals: First, I argue against this prominent account: studying practices of knowledge attribution and denial related to epistemic dependence, coordination, and competition reveals that the primary function of our concept of knowledge is not identifying reliable informants. Second, I recommend that function-firsters accept an alternative account: the primary function of our concept of knowledge is identifying those who have eliminated all relevant alternatives. Finally, I diagnose the error that has led function-firsters to their account: they base their investigation into the nature of knowledge on the situation of inquirers seeking informants, a situation that does not reflect the complexity of our epistemic condition.

Keywords: Epistemic Dependence; Coordination; Competition; Inquiry; Function-First Epistemology; Relevant Alternatives; Edward Craig; Michael Hannon

1. Introduction

Why do we think and speak about knowledge? Because we are needy and limited. We need information that we can't get all by ourselves. Our concept of knowledge has emerged to help us satisfy this interest: the purpose of this concept is to identify reliable informants. By signaling whose information we can trust, knowledge attributions save us from needing to gather information first-hand thereby enabling us to overcome our limitations. Or so argue the function-first epistemologists Edward Craig (in *Knowledge and the State of Nature*) and Michael Hannon (in *What's the Point of Knowledge?*).

Three elements of this reasoning must be distinguished. First, the method: function-first epistemologists analyze epistemic concepts, norms, and practices by interrogating the roles they play in human life, hence the question why do we think and speak about knowledge? Second, the implementation: Craig and Hannon both focus on relatively basic interests and epistemic dependence, hence the premise about individuals needing information they cannot practically obtain themselves. Finally, the substantive proposal: Craig and Hannon both conclude that the primary

function of our concept of knowledge – knowledge has many functions, but the one that explains the others is its *primary* function – is to identify reliable informants. This proposal is not merely the associative claim that there is a correlation between being knowers and being reliable informants – a claim every epistemologist should readily grant – but the explanatory hypothesis that the concept of knowledge is *for* identifying reliable informants.

Epistemologists can part ways with Craig and Hannon at each step: reject the function-first methodology, implement it by focusing on different needs, desires, or epistemic predicaments, or reach a different conclusion about the function(s) of our concept of knowledge. While I have many questions about the function-first method, I will not argue against it here Craig and Hannon have developed a fascinating method well worth exploring. Instead, I question both their implementation of the method and their substantive proposal.

Epistemic dependence is a central part of our epistemic predicament that has largely been neglected but must be accounted for by any theory of knowledge. However, Craig's and Hannon's view is at once not social *enough* and also *too* social: it neglects important social *and* individualistic dimensions of knowledge. While most discussions of epistemic dependence – Craig's and Hannon's included – focus on *our* dependence on *others*, little is made of *others'* dependence on *us*. Investigating that side of the coin reveals interests in tracking knowledge arising from epistemic dependence unrelated to interests in identifying reliable informants. Moreover, we also need to track others' knowledge for the sake of coordination and competition (§2).[1] With this broader perspective in mind, I sketch an alternative to Craig's and Hannon's substantive proposal (§3). Finally, I diagnose where Craig and Hannon went wrong: roughly, their implementation of the function-first method focuses narrowly on the epistemic predicament of the inquirer searching for an informant (§4). This predicament is surely *one* important aspect of knowledge, but it is not the only (or main) one. Subjects acquire much of their knowledge passively rather than through active inquiry, and many of our interests in tracking others' knowledge are unrelated to the inquirer's interest in finding reliable informants. It's no surprise that attempting to understand knowledge by narrowly casting the knower as the inquirer produces an account of knowledge that reflects the inquirer's interests. But this points to an important lesson: function-first accounts run the risk of being parochial. When

[1] I strongly agree with Elgin (2021: 104) that "our epistemic interdependence runs far deeper than our depending on one another for reliable information." Most of Elgin's examples illustrating this point involve agents relying on others to exercise their knowledge. To continue mapping the territory, I develop other examples epistemic interdependence.

interrogating the purpose of a concept, we must carefully consider the question: the purpose *for whom*?

2. Who Knows What?

Call our broad practice of determining who knows what – or attributing or denying knowledge and ignorance to subjects – our practice of *tracking knowledge*. We often track knowledge for purposes unrelated to and inexplicable by our need for identifying reliable informants. For illustration, we can consider examples related to *epistemic dependence*, *coordination*, and *competition*. When investigating what interests the concept of know has emerged to promote, the basic needs and interests represented in these examples deserve as much theoretical attention as the need to identify reliable informants.

As we proceed, we must keep in mind the distinction between two claims:

PROPOSAL: The primary purpose of the concept of knowledge is to identify reliable informants.

CORRELATION: Knowing that *p* is positively correlated with being a reliable *p*-informant.

PROPOSAL is the controversial explanatory claim defended by Craig and Hannon that I will question. It answers the question of why we have the concept of knowledge by saying that concept is primarily *for* identifying reliable informants. CORRELATION, on other hand, is the uncontroversial claim that knowing *p* and being a reliable *p*-informant go hand in hand. On PROPOSAL, CORRELATION is true because knowing *p* is explicable in terms of being a reliable *p*-informant. But CORRELATION itself says nothing about this explanatory relationship. It's compatible with contrary explanations that entail being a reliable informant is a "byproduct" of or merely "tags along" with knowing.[2] As you consider the following examples, you will notice that the characters to whom knowledge is attributed will appear to be reliable informants and those to whom it is denied will not appear to be reliable informants. But remember: that doesn't favor PROPOSAL over CORRELATION. Following the function-first methodology, the question you need to ask is whether the practices of tracking knowledge in these examples are *for* identifying reliable informants. Or do they answer other basic needs and

[2] Hannon (2019: 123) uses these nice expressions to articulate how he sees the relationship between *assurance* and the primary purpose of the concept of knowledge.

interests? I submit that none of these unexceptional examples of knowledge tracking suggest that the purpose of these practices is identifying reliable informants.

The idea that knowledge is for identifying reliable informants is not the only idea about knowledge derivable from considerations of our needs, limitations, and epistemic dependence. We may also want to track who knows what in order to know who needs to be informed and who doesn't. Consider:

> TIGER: Up in a tree looking out onto the field, Trina sees a tiger. Noticing George is about to walk into the field because his view is blocked, Trina thinks to herself: "George doesn't know there's a tiger out there. If I don't tell him, he's going to be tiger lunch! I need to let him know."[3]

Trina tracks whether George knows not because she wonders whether he's a reliable informant about the field's tiger density, but for the sake of tracking what she needs to tell him to keep him safe. Of course, Trina knows *both* that George doesn't know there's a tiger in the field and that George is not a reliable informant on this matter. But she isn't tracking his knowledge *for the purpose of* tracking whether he's a reliable informant.

This is an instance of a general phenomenon: when others depend on us for information, we need to track what they already know. Just as finding information for ourselves is costly, so is informing and failing to inform others. Humans have a general interest in tracking who knows what in order (i) to inform those who need to know and (ii) to avoid wasting resources informing those who already know. While tracking others' knowledge, agents may conclude that they must inform others because they don't know – as in TIGER – or that they need not inform others because they already know—as in DIRECTIONS:

> DIRECTIONS: Jeff and Gina invited Tara to their house for a party. "Do I need to tell Tara our address?" Jeff asks. "No," Gina replies, "she already knows where we live."

Jeff tracks Tara's knowledge to decide whether she needs to be informed, not to determine whether she's a reliable informant.

[3] This example is inspired by Craig (1990: 11).

A more pervasive example of the need to track epistemic dependents' knowledge is education.[4] Children aren't born knowing what fire does to hands that touch it, where food is, or which berries are poisonous. Until they do know these things, responsible caregivers will avoid leaving them unattended near a flame or sending them off on their own to gather food. Before advancing children to reading or algebra, we need to confirm that they know their ABCs or arithmetic. Before letting engineering students build bridges, we must be confident that they know the fundamentals. And so on. We do sometimes track students' knowledge to determine whether they are reliable informants (e.g., when older siblings are responsible for teaching younger siblings). But knowledge tracking in education is not always – indeed, maybe rarely – for this purpose.[5] Groups want their members to possess certain knowledge. Knowledge tracking allows people to assess the extent to which those interests are satisfied.

We need to track knowledge for the sake of coordination as well. Effective communication, for instance, requires knowing what others know.

> VENTING: Suppose Veronica has been venting to friends about trouble she's been having with co-workers. Filling in Natasha on the latest developments, Veronica begins, "I can't remember what you already know. What was the last thing I told you?"

Veronica needs to determine what Natasha knows before proceeding in order to avoid confusing her by presupposing information she doesn't yet know. Veronica is not tracking Natasha's knowledge to figure out whether she's a reliable informant; she probably would not want Natasha to share this information with others. She's tracking Natasha's knowledge to pave the way for smooth communication. Related examples include: (i) tracking what your doctor knows about your medical history as you share new symptoms and (ii) research teams ensuring their members all know the same findings, data, and arguments. Failing to track confederates' knowledge can be costly. (Just consider the possible consequences of mistaking what your doctor knows about your medical background.) When working with others to achieve shared aims – communicating, advancing

[4] Elgin (2021) also mentions instruction as a problem for PROPOSAL to make the (correct) claim that not all instruction is for information transfer, while I'm using education to illustrate our need to track others' knowledge.
[5] Though we often *test* knowledge by determining whether students are reliable informants, we do this not *for that purpose*, but because it is an indirect — correlated — measure of whether they know certain information.

research, improving health, etc. – we need to know who knows what independently of any interest in determining who is a reliable informant.

A bit more precision here will forestall an objection. We might be tempted to say that patients *do* track doctors' knowledge to determine whether they're reliable informants: after all, patients want accurate assessments of their health from their doctors. Considering only PROPOSAL, we might therefore think this coordination example poses no problem to the view. But this oversimplifies the view. Notice there are two propositions playing different roles in this example: (i) a proposition about the patient's medical history *m* and (ii) a proposition expressing the doctor's assessment of her health *a*. On Craig's and Hannon's view, the primary purpose of a knowledge attribution "S knows that *p*" is to identify S as a reliable informant about *p*. But that doesn't capture what's happening in this example. Patients may track doctors' knowledge about one proposition (*m*) to determine their reliability about another (*a*). The purpose of tracking the doctor's knowledge *about one's medical history*, however, is not usually to determine whether the doctor is a reliable informant *about one's medical history*. Not all coordination examples will even plausibly involve agents with interests in determining whether confederates are reliable informants. But as closer scrutiny of this example reveals, even when people engaged in joint ventures *do* want to determine this, their practices of tracking knowledge will not always be accounted for by Craig's and Hannon's model.

Finally, in competition, tracking your rivals' knowledge is often beneficial, not for the sake of determining whether they're reliable informants but for the sake of gaining strategic advantage.

> ALLIANCE: Disrespected too many times by their allies House Stark, House Frey has entered into a secret alliance with their erstwhile enemies, House Lannister, to defeat House Stark. Tonight is the night of the double-cross. The Freys will attack the Starks at dinner. Before dinner, Lannister leaders confirm with Frey leaders, "The Starks know nothing about what's coming." All night before the attack, cautious Frey leaders watch carefully for any sign that the Starks know. Convinced that they know nothing, the Freys make their move.

Knowing that your rivals don't know what you know often gives you an advantage. Knowing that your rivals *do* know what you know can also be beneficial; it might give you an opportunity to revise your strategy. Since knowing what our rivals do and don't know affects strategic behavior, we have strong practical interests in tracking their knowledge.

Interests related to epistemic dependence, coordination, and competition are basic human interests. We need to know who knows what (i) to determine who needs to be informed or educated, (ii) to communicate smoothly and achieve shared goals, and (iii) to compete successfully with rivals. If PROPOSAL is true, the purpose of these knowledge tracking practices is either that they identify reliable informants or are explicable in terms of that purpose. But the function of these practices does not appear to *be* identifying reliable informants. Reliable informants are in fact identified when we attribute knowledge. Again though, by function-first standards, moving beyond correlations to explanation requires establishing that identifying reliable informants is the *purpose* of these knowledge tracking practices. But identifying reliable informants seems to be the purpose of none of them. Preventing harm, determining what information must be shared, communicating smoothly, achieving shared goals, and defeating one's rivals are the goals served by practices of knowledge tracking in these examples. The point, however, isn't simply that people use knowledge-talk to achieve a variety of aims (cf. Hannon 2019: 23). Rather, the point is that humans have many other basic needs and interests, in addition to the need to identify reliable informants, that the concept of knowledge may have arisen to serve. Why prioritize the need to identify reliable informants over these other needs? Why not think these needs, in addition to the need to identify reliable informants, determine the purpose of our concept of knowledge?

Hannon might respond to these examples by appealing to his claim that the function of identifying reliable informants and the function of closing inquiry are two sides of the same coin (2019: 109). This is because, he argues, you can reasonably terminate inquiry when you identify a sufficiently reliable informant.[6] Attributing knowledge to someone is a way of expressing the attitude that their epistemic position with respect to a given proposition is good enough to stop further inquiry. In turn, having an epistemic position good enough to stop further inquiry is what makes an informant reliable enough to accept what they say (if they say it). Appealing to these alleged relations to defend PROPOSAL, Hannon might claim, for example, that agents who can legitimately conclude inquiry need not be informed. In DIRECTIONS, when Gina says Tara already knows the address, she indicates that Tara's epistemic position is good enough for her to stop further inquiry. Since Jeff is trying to figure out whether he needs to tell Tara the address, once he is aware that Tara's epistemic

[6] When Hannon (2019: 109) first discusses this idea, he writes that "*the* way to reasonably terminate inquiry is by identifying a sufficiently reliable informant" [my emphasis], but whenever else he discusses it, he only endorses the weaker (and more plausible) claim that this is *a* way. I assume his view is the weaker view.

position is good enough for *her* to conclude inquiry, telling her the address would violate a pragmatic rule against giving unnecessary information (Grice's (1975) maxim of quantity). So, Hannon might conclude that when combined with some basic pragmatic rules his account captures the practice of tracking knowledge to determine who needs to be informed.

This response shouldn't satisfy us. First, just as a correlation exists between knowing *p* and being a reliable *p*-informant, plausibly a correlation exists between knowers who do not need to be informed and those who can close inquiry. But it's not obvious that the former should be explained in terms of the latter. Is the purpose of Trina's denying knowledge to George to signal that George cannot legitimately close inquiry? Not obviously. Second, even if *some* of our examples can be analyzed in this way, not all can. We do not track rivals' knowledge to determine whether they can legitimately close inquiry. Third, should we really think of these two functions – signaling the proper close of inquiry and identifying reliable informants – as two sides of the same coin? Finding a reliable informant is *one* way to reasonably terminate inquiry. But you can reasonably terminate inquiry on your own without finding a reliable informant. Acquiring knowledge this way leads you to *become* a reliable informant, but inquiry is not terminated *because* you found a reliable informant. Since finding a reliable informant and properly terminating inquiry do not appear to be extensionally equivalent, it's odd to refer to them as two sides of the same coin.

More can be said here, but I'd like to step back. I have described ordinary cases of knowledge attribution and denial that cast doubt on one version of the idea that the primary function of the concept of knowledge is identifying reliable informants. These cases were selected to show that we can adopt Craig's and Hannon's methodology and initial focus on basic and universal human interests without arriving at PROPOSAL. I have also cautioned against quickly accepting the explanatory hypothesis PROPOSAL on the basis of what may be mere correlations between knowing, being a reliable informant, and legitimately concluding inquiry.

3. Generalizing Hannon's Accounts

What function does knowledge tracking play if not (primarily) signaling reliable informants or the legitimate conclusion of inquiry? Though we may doubt this question's presuppositions, I will stick with function-firsters to see whether we can sketch an account that accommodates all of our observations. Indeed, a plausible answer can be derived from another of Hannon's proposals.

Hannon develops an interesting answer to the "threshold problem": the problem of stating the threshold above which a subject's epistemic position with respect to *p* is good enough to count as being in a position to know *p*. This problem emerges when we deny that knowledge requires conclusive justification. If conclusive justification is unnecessary, how good must an epistemic position be? Hannon (2019: 68) defends a "community-based" relevant alternatives answer:

> **The Reliable Informant Standard for Knowledge (RIS):** To know that *p*, an agent must be in a strong enough epistemic position with respect to *p* to eliminate all of the not-*p* possibilities that are relevant alternatives to members of the epistemic community that might draw on the agent's information.

If the purpose of the concept of knowledge is to identify reliable informants, you are in a position good enough to count as knowing *p* when you've eliminated all the not-*p* possibilities that count as relevant alternatives to epistemic community members who might use your knowledge. Who constitutes the epistemic community? Hannon adopts a "reasonable person" standard:

> a knower must have evidence sufficient to eliminate the alternatives a reasonable person would want eliminated… After all, the reasonable person standard presumably reflects whatever epistemic standard is taken to be reasonable by the relevant community's judgement. (2021: 125)

Though this certainly raises many questions, it's not trivial. You see that your cat is on the counter. Do you know the cat is on the counter? Your evidence doesn't eliminate incompatible brain-in-a-vat possibilities. But a reasonable person doesn't ordinarily care if these possibilities are eliminated. If that's correct, RIS entails that knowing the cat is on the counter is compatible with not ruling out such skeptical possibilities. Thus, RIS combines a relevant alternatives framework with a reasonable person account of which alternatives are relevant.

Because I'm skeptical of PROPOSAL, I will suggest a generalized alternative to RIS. To see what I have in mind, consider one of the education cases: You want your daughter to gather berries, but you need to know whether she can distinguish poisonous from safe berries. In other words, you want her to know of the particular berries she encounters "this is safe" (*s*) or "this is poisonous" (*p*). So,

you show her a poisonous berry and ask: is this poisonous or safe? If she answers "safe," then she doesn't know *p* and she requires further instruction. If she answers "poisonous," you'll probably try again to ensure it wasn't a fluke, that she actually knows *p* and didn't just get lucky. Again, this is an utterly unexceptional case of knowledge tracking that casts doubt on PROPOSAL. You're not interested in whether she knows because you want to determine whether she's a reliable informant; you care whether she knows because you want her to safely gather berries and she's unlikely to succeed if she doesn't.

On Hannon's view, we want to determine whether *others* know propositions in order to determine whether *we* can act on those propositions: "The reason we have a concept of knowledge at all…is to mark out people on whom we can rely for actionable information" (2019: 123; see also pp. 13, 67, 81). I suggest cutting out the middle man. Rather than tracking people on whom we can rely for actionable information, we're tracking who *has* actionable information, sometimes for our own uses, sometimes not. Hannon and Craig are absolutely right that we often want to know whether other people know *p* in order to know whether they're reliable *p*-informants. But that's because we want to know whether *p* is sufficiently well-supported to be acted on. When your daughter asserts *s* when you know *p* is true, you worry that *she* will act on *s*—she'll eat a poisonous berry because she thinks it's safe. More generally, we track people's knowledge to determine whether they have actionable information.

We can generalize RIS by looking beyond the action of drawing on agents' information to any uses to which agents, including the potential knowers themselves, might put the information.

> **Actionable Information Standard for Knowledge (AIS):** To know that *p*, an agent must be in a strong enough epistemic position with respect to *p* to eliminate all of the not-*p* possibilities that are relevant alternatives to reasonable members of the epistemic community.[7]

[7] Sometimes Hannon's comments about the purpose of knowledge reflect the generality AIS attempts to capture. For instance: "the core of knowledge is sufficiently reliable true belief, where 'sufficiently reliable' is unpacked roughly as "reliable enough to serve the purposes of members in the epistemic community.'" (2019: 150) This refers only to the "purposes" of community members saying nothing of reliable informants. It seems much closer to the truth for that reason.

Which alternatives would a reasonable member of the epistemic community consider relevant? Ones that bear on the actionability of *p*. Brain-in-a-vat possibilities? Not relevant: they don't bear on actionability. The possibility that the berry is poisonous when you believe it's safe? Relevant: it makes a difference to whether you should act on *s*.

We can also state a more general proposal about the purpose of the concept of knowledge.

> ALTERNATIVE PROPOSAL: The primary purpose of the concept of knowledge is to identify those who have eliminated all relevant alternatives.[8]

This proposal captures all of our cases. Epistemic dependence: We track others' knowledge to know which possibilities must still be or do not need to be eliminated. Coordination and Competition: if I'm interested in determining whether you know *p* so we can achieve our shared goals or so I can outmaneuver you, I'm interested in whether your evidence has eliminated the relevant not-*p* possibilities. In the epistemic dependence and coordination cases, if you haven't eliminated the relevant not-*p* possibilities, I'll typically help you eliminate them. In the competition case, when what you know affects my plans, my strategy will depend on which possibilities you've eliminated.

ALTERNATIVE PROPOSAL also captures the reliable informant purpose of the concept of knowledge. When I'm looking for a reliable informant about *p*, I'm looking for someone who has ruled out all relevant not-*p* possibilities. Suppose I've lost my way to Yankee Stadium; I need to know whether to go north, south, east, or west. Stepping into a convenience store to ask the clerk, I'm searching for someone who knows which way to go. Why? Because I need to rule out possibilities I can't eliminate on my own.

ALTERNATIVE PROPOSAL captures the concluding inquiry purpose as well. Inquiry advances by whittling down the field of possible answers to a question *Q*, legitimately terminating when all of the relevant alternatives to an acceptable answer have been ruled out. Knowledge attributions signal the legitimate end of inquiry because knowledge attributions function to identify those who have ruled out all relevant alternatives to certain propositions—in this case, possible answers to *Q*. These points also allow us to account for the fact that finding a reliable informant is typically a means to

[8] Or those with a reliabilist bent might prefer Schmitt (1992:557): "I would suggest an alternative hypothesis about the concept of knowledge…: it serves not only to pick informants as to whether *p* for our own use and that of others, but to pick reliable believers as to whether *p*—individuals on whom we can rely to arrive at a true belief as to whether *p*."

concluding inquiry: if I know you're a reliable *p*-informant, then I know you probably know *p*, which means you've probably ruled out the relevant not-*p* possibilities; so, I know you can conclude inquiry. That gives me good reason to conclude my own *p*-inquiry.

The final advantage ALTERNATIVE PROPOSAL has over PROPOSAL is that it is old school…in a good way. Function-first epistemology is partly motivated by the stagnation of the research program dedicated to providing a conceptual analysis of knowledge. Rather than stating necessary and sufficient conditions for knowledge, function-first epistemologists aim to illuminate its "core," freeing themselves from the demand of accounting for every logically possible counterexample to their views. I strongly sympathize with this motivation. And yet, as PROPOSAL departs from the traditional method, it also unnecessarily abandons key insights uncovered within that research program. I would have thought that once freed from the constraints of conceptual analysis, work done in the traditional paradigm would be fertile ground for function-first discovery.

More concretely: no one working in the traditional paradigm has ever argued that knowledge should be analyzed and explained in terms of being a reliable informant. It's not that no one ever saw a connection here—Hannon cites Sosa (1974). But this idea has never been the centerpiece of a traditional analysis. So, it's surprising that Hannon and Craig argue that the *primary purpose* of the concept of knowledge is to identify reliable informants. In contrast, the idea that knowledge should be analyzed in terms of eliminating relevant alternatives pervades the traditional literature (Goldman 1976; Stine 1976; Dretske 1981). ALTERNATIVE PROPOSAL, unlike PROPOSAL, is favorably old school because it retains a promising insight about the nature of knowledge from the traditional literature and then applies that insight using the function-first method to reveal the purpose of the concept of knowledge. Obviously, Hannon has not altogether abandoned the relevant alternatives insight; the point is that traditional insights are strikingly absent from his and Craig's account of knowledge's *function*. In short, Hannon and Craig recommend not only divorce from the traditional method; they want a messy divorce where the parties aren't on speaking terms. Instead, I think function-first epistemologists should try a more amicable break-up: things are no longer working, but let's keep talking!

4. Examiner Situation and Inquiry

What went wrong? Function-first epistemology and PROPOSAL are not inevitably intertwined. Function-first epistemologists can deny that the primary function of the concept of knowledge is

identifying reliable informants. The concept of knowledge might have arisen in connection with needs other than the need to identify reliable informants. It's striking, then, that the two most prominent practitioners of function-first epistemology both end up endorsing PROPOSAL. It's also striking that needs related to familiar practices of knowledge tracking were not among their explananda. These striking facts call for an explanation.

The problem is that Craig and Hannon almost exclusively analyze knowledge from the perspective of the inquirer. Hannon writes (2019: 37, 75):

> The central focus of epistemic evaluation is the activity of inquiry.

Craig (1990: 11-12) writes:

> Our investigation ought to start from the position in which we as yet have no belief about *p*, want a true belief about it one way or the other, and seek to get it from someone else. … Consider then the position of someone seeking information on the point whether or not *p*. *What does he want?* In the first place, he wants an informant who will tell him the truth on that question. [My emphasis]

Thus, Hannon notes there are two key players in his and Craig's framework:

> *The Inquirer*: someone who is trying to find out whether or not *p* is true.
> *The Reliable Informant*: someone who provides the information the inquirer seeks.

The inquirer needs a way to distinguish unreliable from reliable informants. From this need emerges the primary purpose of the concept of knowledge, according to Hannon and Craig.

Again, it's no surprise that an account of knowledge reflecting the inquirer's interests emerges from an investigation of knowledge that begins by narrowly casting the knower as the inquirer. But why should we accept this starting point? And why think the central focus of epistemic evaluation is the activity of inquiry?

Craig (1990: 11-12) defends starting with the inquirer's perspective in this way:

> I shall not for the moment be concerned with the evaluation of what I have called 'on-board' sources. In the ordinary way we simply take it that the beliefs they mediate are true. To find oneself in possession of a belief on the question whether *p* pre-empts inquiry; to take a self-conscious look at one's own apparatus with the doubt in mind that it may have delivered a falsehood calls for a considerable degree of sophistication. Our investigation ought to start from the position in which we as yet have no belief about *p*, want a true belief about it one way or the other, and seek to get it from someone else. … Our interest in our own faculties as sound sources of information has a part to play, since under certain circumstances that interest becomes acute, for very good practical reasons; but it would not be good method to begin with it.

With '*on-board sources*', Craig (1990: 11) refers to "eyes and ears, powers of reasoning, which give them a primary stock of beliefs." There are several claims in this passage. Due to space limitations, I focus on:

A. We largely take for granted that beliefs mediated by on-board sources are true.

Craig's argument that we should begin our investigation from the inquirer's perspective seems to go like this.

To explicate knowledge, we must identify the basic needs and interests the concept of knowledge answers. Because we want information that others have and because people vary significantly in their reliability, we should expect a concept to emerge out of inquiry-driven interests that helps us evaluate information sources for their reliability. On the other hand, it would be surprising if the concept of knowledge emerged out of interests in evaluating on-board sources: because on-board sources, in contrast to informants, are taken to be reliable (A) and vary little in their reliability, there's little need to evaluate on-board sources. Therefore, we should begin by investigating inquiry-driven interests.

The primary flaw in this argument is that it reduces (i) evaluating knowledge not driven by inquiry to (ii) evaluating the reliability of on-board sources. Even if we agree that there's little need to evaluate the general reliability of on-board sources, we're not forced to conclude that the only remaining significant interests from which to start investigating knowledge are the inquirer's. None

of our examples involve people evaluating the general reliability of their subjects' faculties. But neither are they searching as inquirers for reliable informants. It matters to us – for a variety of reasons, some inquiry-driven and some not – what others know. Thus, while we may have reason to expect a concept for evaluating informants to arise out of inquiry-driven interests, these are not the only basic interests from which we would expect such a concept to arise.

Hannon provides a different argument for focusing on the inquirer's perspective. Following Williams (1973), Hannon distinguishes between the "examiner situation" and what we can call the "inquiry situation." The examiner situation is:

> the situation in which I know that *p* is true, this other man has asserted that *p* is true, and I ask the question of whether this other man really knows it, or merely believes it. (Williams 1973: 146)

Williams discusses the examiner situation in the context of defending the possibility of knowing *p* without believing *p*. The contrary view that knowledge is belief-plus-something-else, Williams argues, is encouraged by concentrating on the examiner situation. But Williams thinks concentrating on this situation is a mistake because we normally occupy the distinct role of inquirer rather than examiner:

> [O]ur standard situation with regard to knowledge (in relation to other persons) is rather that of trying to find somebody who knows what we don't know; that is, to find somebody who is a source of reliable information about something. (*ibid.*)

Hannon (2019: 4) argues that we must adopt a "deeply social" account of knowledge that places our reliance on others at center stage; a very small part of our knowledge comes to us from our immediate experience, he tells us. Hannon agrees with Williams about our standard role:

> Whereas the examiner situation is concerned with whether some potential knower really qualifies as such, the actual business of inquiry involves an inquirer who does not know whether *p* but wants to. (2019: 37)

Thus, we should eschew the examiner situation in epistemic theorizing in favor of the inquiry situation: the situation in which an agent doesn't know whether *p*, but wants to. Again, the central focus of epistemic evaluation, Hannon writes, is the activity of inquiry.

I think this argument is also flawed. First, plenty of our concern for others' knowledge *is* captured by the examiner situation. As we've seen, in formal and informal education settings – by no means non-standard settings – teachers, parents, and other caretakers regularly try to determine whether subjects know what they, the "examiners," already know. Second, many of our standard interests in others' knowledge are represented by neither the examiner situation nor the inquiry situation. Epistemic dependence, coordination, and competition all involve tracking others' knowledge outside of the inquiry situation. Many of these cases are also unrelated to the examiner situation. We often track what others know not to determine whether they know what they claim to know, but for the other reasons we've discussed.

Williams and Hannon correctly argue that our situation with regard to others' knowledge is not *limited* to the examiner situation, but neither is it adequately represented by the inquiry situation. We *should* develop a deeply social account of knowledge if this means an account that reflects people's reliance on others for information as well as the variety of other basic interests people have in tracking each other's knowledge. But the inquiry situation represents a mere segment of the complete social dimension of knowledge. An appropriately social account of knowledge must reflect a broader range of epistemic predicaments.

To summarize, we should be convinced by neither of these arguments for prioritizing inquiry over other dimensions of knowledge in our theorizing. Both demand a false choice. Craig insists that we must choose between an investigation of knowledge motivated either by a concern for evaluating the general reliability of our on-board sources or by a concern for evaluating the reliability of informants; Hannon seems to suggest that our standard situation is captured either by the examiner situation or the inquiry situation. But we have reasons to track others' knowledge unrelated to interests in identifying reliable informants even if we take for granted the general reliability of their on-board faculties, and much of our social interest in knowledge is captured by the examiner situation in addition to the inquiry situation and other situations as well.

Finally, focusing on the inquiry situation ignores not only social considerations, but individualistic ones as well. Because social considerations have largely been neglected by epistemologists, recently epistemologists have rightly foregrounded social dimensions in their

theorizing. But we shouldn't make the mistake of going too far in the other direction. A comprehensive account of knowledge must reflect the social *and* individualistic aspects of our epistemic situation. Though few epistemologists run the risk of going too social, Craig's and Hannon's implementation of the function-first approach *does* dismiss important individualistic dimensions of knowledge.

So much of what we count as know simply *hits us*, without us seeking that knowledge on our own or from others. That is, we often passively and involuntarily receive knowledge through lone interactions with our environment, rather than acquiring it through active inquiry. Here's a small sampling of the knowledge I accumulated today on my car ride home:

> Construction is occurring on Jamboree (a local road); the police pulled over a red car; I haven't gotten a ticket in years; our local office supply store is going out of busines; the stock market is fluctuating; landscapers are working in our community; my neighbor is walking his dog; music is playing in the office; my partner is in the office; the light is on in the kitchen; we're running low on milk…

That is a snippet of the information that I passively received on my trip. I didn't inquire into any of it. I don't care about most of it. But I know it.[9] Though dwelling on this kind of knowledge may be unfashionable, the fact that epistemologists *have* concentrated so much on passive individual knowledge based in perception, memory, and inference to the exclusion of social considerations casts further doubt on the claim that the central focus of epistemic evaluation is inquiry.

5. Conclusion

Too much of epistemology has focused exclusively on individualistic aspects of knowledge. But investigating the nature of knowledge by treating knowledge surrounding inquiry – even more narrowly, the knowledge inquirers seek from others – as the paradigm case improperly restricts the scope of this investigation. We assess people's epistemic positions for purposes unrelated to our interests as inquirers, and much of what we know comes to us not through inquiry but through

[9] Although some knowledge gained through social interactions is knowledge we've inquired into, not all of it is. When interacting with others we learn much information we have not inquired into. So, again, a commitment to focusing on social aspects of knowledge does not require focusing on the inquiry situation.

passive interactions with our environments. A comprehensive view of knowledge must reflect the complexity of our epistemic condition.

Our discussion reveals how the function-first method runs the risk of being parochial. PROPOSAL emerges as the narrow function-first answer to questions about the purpose of knowledge when we focus on what purpose knowledge serves for the *inquirer*. Function-first epistemologists must think carefully about *whose* purposes are prioritized in their theorizing.

Finally, going forward, function-first epistemologists need to refine their method in order to settle disputes about *which* function goes *first*. We should all be pluralists about the purpose of our concept of knowledge; but if Craig and Hannon are right that knowledge has a *primary* purpose, then fruitfully applying this methodology requires knowing how to determine which is which. Correlations between epistemic credentials won't settle the matter. If the *primary* function is whichever explains the others, that function may be identifying those who have ruled out all relevant alternatives. More importantly, our main general conclusion is that deciding which function is first requires attending to all of our needs related to practices of determining who knows what.

Acknowledgments: I want to thank Michael Hannon for inviting me to participate in this symposium. I have learned a ton from engaging with his deep and rich book. I also want to thank Daniel Immerman, Hilary Kornblith, Luis Oliveira, and Robert Simpson for helpful discussion, advice, and feedback on this project. Work on this project was supported by a research scholarship, for which I'm grateful, from the College of Humanities and Social Sciences at California State University Fullerton.

References

Craig, E. (1990) *Knowledge and the State of Nature*. Oxford: Clarendon Press.
Dretske, F. (1981) The Pragmatic Dimension of Knowledge. *Philosophical Studies* 40: 363-378.
Elgin, C.Z. (2021) The Function of Knowledge. *Analysis Reviews* 81: 100-107.
Goldman, A. (1976) Discrimination and Perceptual Knowledge. *Journal of Philosophy* 73: 771-791.
Grice, P. (1989) Logic and Conversation. In *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
Hannon, M. (2019) *What's the Point of Knowledge?* Oxford: Oxford University Press.
-- (2021) Replies to Henderson, Elgin and Lawlor. *Analysis Reviews* 81: 114-129.
Lawlor, K. (2013) *Assurance*. Oxford: Oxford University Press.
Schmitt, F. (1992) Review of Knowledge and the State of Nature: An Essay in Conceptual Synthesis. *Mind* 101: 555-559.
Sosa, E. (1974) How Do You Know? *American Philosophical Quarterly* 11: 113-122.
Stine, G. (1976) Skepticism, Relevant Alternatives, and Deductive Closure. *Philosophical Studies* 29: 249-261.
Williams, B. (1973) Deciding to Believe. In *Problems of Self*. Cambridge: Cambridge University Press. 136-151.