



## Summary of *What's the Point of Knowledge?* Oxford University Press, 2019

Michael Hannon

Department of Philosophy, University of Nottingham, UK

### ABSTRACT

This article summarizes the key ideas in my book, 'What's the Point of Knowledge?'. First, I articulate the method of function first epistemology, which looks at our epistemology from a practical point of view. Second, I outline my core hypothesis about the purpose of the concept of knowledge, namely, we have this concept to identify reliable informants. Third, I explain how this account is used to resolve a number of epistemological issues. I also draw connections between pragmatism and my own theory of knowledge.

**ARTICLE HISTORY** Received 4 August 2022; Accepted 4 August 2022

**KEYWORDS** Knowledge; function-first epistemology; reliable informants; skepticism; Edward Craig; pragmatism

There is a word for 'know' in every human language. This might not seem remarkable, but it is. According to the evidence from cross-linguistic semantics, almost every word in the English lexicon does *not* have an equivalent in every other language. Even words that refer to common emotional states like 'sad' and 'angry,' as well as words for seemingly universal mental states and processes like 'believe' and 'remember,' are language and culture-specific. In contrast, linguists have isolated 'know' as one of a very small number of words that are allegedly culturally universal (Wierzbicka 2018). This suggests that it answers to highly general needs of human life. But what are those needs?

*What's the Point of Knowledge?* (Oxford University Press, 2019) is about why humans were driven to develop the idea of knowledge. The method, broadly speaking, is to look at our epistemology from a practical point of view. Instead of asking what knowledge *is*, we should ask what the concept of knowledge *does for us*. What role does this idea play in human life?

**CONTACT** Michael Hannon  michael.hannon@nottingham.ac.uk

© 2022 Informa UK Limited, trading as Taylor & Francis Group

To answer this question, I draw on foundational work by Edward Craig (1990) who appeals to general facts about the human condition to explain why the concept of knowledge would come about. According to Craig, we need to form accurate beliefs about the world to successfully navigate it. But there is only so much information we can gather by ourselves, so we need to pool and share information with others. Yet people vary in terms of their reliability, so we must be able to distinguish the people on whom we should rely from those we should not rely on. The concept of knowledge emerges in connection with these basic and universal needs. In particular, we speak of 'knowers' to mark out good sources of information. By anchoring the practical value of the concept of knowledge in basic and universal needs that humans can be assumed to have anyways, this approach demonstrates why the concept of knowledge was bound to emerge as a matter of practical necessity: it provides a solution to a problem that virtually all humans will face.

I call this method 'function-first epistemology'. A function-first epistemologist seeks to explain the nature and value of an epistemic concept, norm, or practice by reflecting on its function or purpose. A function-first epistemologist will ask questions like: Why do humans speak and think in terms of 'knowing,' 'understanding,' and 'rationality'? If we lacked these concepts, what else would we lose? Which epistemological distinctions and norms would best facilitate human survival, cooperation, and flourishing? Do our epistemic notions carry much weight in science, philosophy, or daily life?

Function-first epistemology is really just another name for pragmatism. As Huw Price writes,

Pragmatism begins ... with phenomena concerning the use of certain terms and concepts, rather than with things or properties of a non-linguistic nature. It begins with linguistic behavior and asks broadly anthropological questions: How are we to understand the roles and functions of the behavior in question, in the lives of the creatures concerned? What is its practical significance?. (2011, 231–232)

Although function-first epistemology is a type of pragmatism, I have given it a new branding for people who think they don't like pragmatism. Whatever we call it, the general strategy is to explain why creatures like us would go in for a particular kind of discourse and thought. We explore the nature and value of knowledge by reflecting on the role or purpose of our concept of knowledge. This puts the explanatory weight on the function

of a type of discourse, not on the existence of the entities referred to in that discourse (see also Queloz 2021, 35).

*What's the Point of Knowledge?* has nine chapters, but these can be meaningfully divided into three parts. First, I articulate the *method* of function-first epistemology, outline some of its benefits, and compare it with four alternative approaches: traditional conceptual analysis, knowledge-first epistemology, reverse engineering epistemic evaluations, and epistemological naturalism. Second, I outline and defend my core *hypothesis* about the purpose of the concept of knowledge: we have this concept to identify reliable informants. Third, I *apply* this account of knowledge to a number of epistemological issues, including: the nature and value of knowledge, the foundations of epistemic normativity, the semantics of 'know', the differences between knowing and understanding, the epistemology of testimony, pragmatic encroachment, the Gettier problem, epistemic relativism, and skepticism. I attempt to show that we can answer many interesting and difficult questions in epistemology by examining the concept of knowledge from a practical point of view.

Although my book traces the concept of knowledge to its practical origins, it is not an attempt to document the history of our conceptual practice. *What's the Point of Knowledge?* is neither a historical genealogy nor a speculative 'just-so' story about when this conceptual practice emerged. Rather, it attempts to reconstruct the point of the concept of knowledge by rooting it in human needs so basic and circumstances so familiar that they would be at work in any human society. Thus, my book is not concerned with what actually triggered the occurrence of our knowledge concept, or why some particular individual applies the term 'know' as they did, or how this concept may have taken different variations throughout human history. Instead of investigating the distant past, my book reconstructs the practical problems to which the concept of knowledge offers a solution.

What is the function of the concept of knowledge? There have been numerous answers to this question. It has been suggested, for example, that we speak of knowing to signal the appropriate end of inquiry, to provide assurance to others, and to encourage good testimony. This plurality of hypotheses should make us wonder whether the concept of knowledge has just one function. If we suppose it does have a single function, how do we decide between these various proposals? In my book, I acknowledge that we speak of knowing for a variety of purposes, but I also argue that the *primary* purpose of the concept of knowledge is to identify reliable informants. I call this the 'informant-flagging function'.

This function is explanatorily prior to others functions because it can explain those other functions, while those other functions cannot explain it.

On this approach, knowledge turns out to be a type of social phenomenon. It is not something we find in the world, like rocks or water, but rather something we impose on it: it reflects human needs and interests. In short, knowledge is a social or artifactual kind, not a natural kind. Yet my approach is still a form of philosophical naturalism; it aims to naturalize the potentially puzzling or mysterious idea of knowledge by showing that we can fully account for it by appealing to natural facts, without resorting to any metaphysical or extra explanatory material. In this way, my view aims to demystify the nature and value of knowledge. It explains why certain needs would naturally lead creatures like us to develop this conceptual practice. This allows us to rid ourselves of the need for metaphysical explanations and unnecessary ontological commitments.

The idea that knowledge is important to humans has been roundly criticized in recent epistemology. It has been argued that our concept of knowledge is arbitrary (Risberg 2022) or a relic of a bygone age (Papineau 2019), that knowledge is not distinctively valuable (Frise 2017), that we can fully explain rational behavior without appealing to knowledge (Kaplan 2003), that knowledge is an incoherent notion (Schiffer 1996) or a myth (BonJour 2010), and that we should instead focus on other epistemic states that are allegedly more important, such as understanding (Elgin 2017) or wisdom (Grimm 2015). Against this trend, I believe that knowledge is fundamental to our epistemic life and belongs at the heart of epistemological theory.

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## References

- BonJour, L. 2010. "The Myth of Knowledge." *Philosophical Perspectives* 24 (1): 57–83.
- Craig, E. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- Elgin, C. 2017. *True Enough*. Cambridge, MA: MIT Press.
- Frise, M. 2017. "No Need to Know." *Philosophical Studies* 174: 391–401.
- Grimm, S. 2015. "Wisdom." *Australasian Journal of Philosophy* 93 (1): 1–16.
- Kaplan, M. 2003. "Who Cares What You Know?" *Philosophical Quarterly* 53 (210): 105–116.

- Papineau, D. 2019. "The Disvalue of Knowledge." *Synthese* 198 (6): 5311–5332.
- Price, H. 2011. *Naturalism Without Mirrors*. Oxford: Oxford University Press.
- Queloz, M. 2021. *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering*. Oxford: Oxford University Press.
- Risberg, O. 2022. Meta-Skepticism. *Philosophy and Phenomenological Research*. Early view: doi:[10.1111/phpr.12871](https://doi.org/10.1111/phpr.12871).
- Schiffer, S. 1996. "Contextualist Solutions to Skepticism." *Proceedings of the Aristotelian Society* 93: 317–333.
- Wierzbicka, A. 2018. "I Know: A Human Universal." In *Epistemology for the Rest of the World*, edited by Masaharu Mizumoto, Stephen Stich, and Eric McCready, 215–250. Oxford: Oxford University Press.



## Pragmatism, skepticism, and over-compatibilism: on Michael Hannon's *What's the Point of Knowledge?*

Georgi Gardiner

University of Tennessee, Knoxville, TN, USA



### ABSTRACT

Function-first approaches illuminate phenomena by investigating their functional roles. I first describe virtues of this approach. By foregrounding normal instances of knowledge, for example, function-first theorising offers a much-needed corrective to epistemology's counterexample-driven momentum towards increasingly byzantine, marginal cases. And epistemic practices are shaped by human limitations, needs, vices, and power relations. These non-ideal, naturalistic forces of embodied sociality form the roots of function-first theorising, which creates a fecund foundation for social epistemology. Secondly, I consider an objection to function-first theorising. The objection holds that function-first approaches lack adjudicatory power. That is, function-first proposals are overly compatible with diverse claims about knowledge, which encourages "just so" speculation. In response to this concern, I advocate methodological pluralism in epistemological theorising. All methods have limitations; researchers must be mindful of those limits and fruitfully combine multiple methods. I illustrate with Hannon's function-first based arguments for "epistemic pragmatism", which denies that the meaning of "knows" is determined by truth conditions. Finally, I argue that function-first theorising motivates staunch anti-skepticism about knowledge. Practical forces cannot chisel a sharp threshold for how much evidence is required for knowledge. But thinking this supports radical skepticism about knowledge conflates fuzzy thresholds and high ones.

**ARTICLE HISTORY** Received 4 August 2022; Accepted 4 August 2022

**KEYWORDS** Skepticism; methodology; function-first epistemology; teleological approaches; semantics of knowledge ascriptions; epistemic pragmatism

Function-first approaches in epistemology aim to illuminate the nature and value of epistemic phenomena by asking what roles they serve. The guiding idea is that investigating what each phenomenon does – what function it fulfils – reveals contours of the phenomena itself. The method can be applied to concepts, norms, institutions, artifacts, and practices. It is typically applied to epistemic assessments and ascriptions,

**CONTACT** Georgi Gardiner  georgicloud9@gmail.com  Philosophy Department, University of Tennessee, 801 McClung Tower, Knoxville TN 37996-0480, USA

© 2022 Informa UK Limited, trading as Taylor & Francis Group

such as attributing knowledge or understanding. It is also sometimes applied to epistemic states – knowledge and understanding – themselves.<sup>1</sup>

This approach was first proposed, at least in its recent incarnation, in Edward Craig's ground-breaking *Knowledge and the State of Nature* (1990). To investigate what the concept of knowledge does, Craig imagines a society that lacks the concept and asks what they would be unable to do. Craig hypothesises that denizens of his imagined society would lack a way to tag good informants, and so proposes this is the primary function of knowledge ascriptions. We can call this a 'functional hypothesis' and compare it with rivals, such as that knowledge ascriptions mark when evidence suffices to end inquiry.<sup>2</sup> Craig uses his favoured functional hypothesis to plumb the nature and value of knowledge.

Michael Hannon, a leading scholar of function-first epistemology, provides the first book-length development of these ideas. Hannon retains key Craigian claims, including the central hypothesis about the function of knowledge attributions, but departs from Craig in significant ways. (See 52–52, for example.<sup>3</sup>) Hannon describes methodological and metaphysical foundations of the function-first approach (chapters one and two), and he illuminates skepticism (189–221), the Gettier problem (73–77), the lottery paradox (77–80), group knowledge (125–133), knowledge norms (103–116), justificatory thresholds for knowledge (81–102), epistemic injustice (133–136), understanding (222–255), and explanation (226–229; 237–242). The final chapter applies the method to the nature of understanding.

In the interests of space, I assume familiarity with the basic contours of the function-first approach.<sup>4</sup> I'll instead highlight some key virtues of the approach and one potential pitfall. These features provide the foundation for my three critical comments. My critical comments concern Hannon's 'epistemic pragmatism' about the semantics of knowledge ascriptions, his concessions towards skepticism, and his contention that his favoured functional hypothesis has explanatory priority over rivals.

---

<sup>1</sup>Gardiner (2021b) suggests Hannon (2019) conflates a function-first approach to X and the concept of X.

<sup>2</sup>Knowledge ascriptions can serve multiple functions and debates centre on which function is primary.

<sup>3</sup>Unless otherwise specified, all references are to Michael Hannon's (2019) *What's the Point of Knowledge?* OUP.

<sup>4</sup>Gardiner (2021b), by contrast, is written for audiences not already familiar with function-first epistemology.

## Virtues

Following Gettier, epistemology foregrounded the pursuit of reductive necessary and sufficient conditions on knowledge. On some interpretations of this endeavour, any counterinstance suffices to undermine a proposed theory. But this methodology can obscure the distinction between important counterinstances – ones that reveal deep flaws – and marginal, insignificant counterinstances. Indeed combining a case-based method with the pursuit of necessary and sufficient conditions generates inevitable momentum towards peripheral cases. This is because theorists already considered normal and central cases when developing their theory. If progress is driven by counterinstance, then detractors must devise increasingly bizarre, arcane cases to find ones the target theory didn't already account for. But attention on weird cases, and the resulting byzantine codicils and epicycles – raises questions about the value and viability of the enterprise. Dissatisfaction with the resulting aims, methods, and theories that became known – often derisively – as Gettierology are well-documented (17–18). If nothing else, undue attention to peripheral cases can be distorting, as theories and debates contort around tortuous cases, perhaps failing to properly weigh other theoretical and explanatory virtues. Spending time on peripheral cases exaggerates their importance.

Function-first epistemology offers a useful corrective. It foregrounds ordinary, non-deviant, central examples. This is because function-serving phenomena are shaped by normal conditions of operation.<sup>5</sup> According to function-first epistemology, non-conforming examples can be disregarded if they are sufficiently peripheral. They are not fatal to function-first hypotheses, since function-first theories typically don't aim for exceptionless necessary and sufficient conditions.<sup>6</sup> There is far more to say, of course, about what makes examples peripheral or disregardable. But here I'll note simply that, at least in comparison to the inevitable pressures towards convoluted, divisive cases characteristic of Gettierology, function-first's orientation towards central cases is a welcome corrective.

A second virtue is that the function-first approach is, or can be, methodologically naturalistic.<sup>7</sup> It grounds practices of epistemic evaluation in

---

<sup>5</sup>Conditions can be abnormal globally, but normal for the relevant functional phenomenon. A firehose is designed to extinguish uncontrolled fires, for example. Such fires are abnormal, but not abnormal conditions for a firehose.

<sup>6</sup>Function-first methods can be combined with these aims (Gardiner 2015; Hannon 2019, 19).

<sup>7</sup>Kornblith (2011) demurs; Gardiner (2017, 106–110) responds.



ordinary day-to-day needs, and so renders intelligible why knowledge – or the concept of knowledge – has the structure and value it does. Function-first epistemology starts with realities of the human epistemic condition, including that we are embodied, social, emotional, cognitively limited, error-prone, and susceptible to vices such as self-interest. Relatedly the method foregrounds social features of epistemic lives. Traditional epistemology focuses on the individual inquirer. At its extreme, this is embodied by the lone Cartesian thinker. (I could say ‘non-embodied’, since non-naturalistic, idealised epistemology sometimes overlooks the physical realities underpinning thought.) And social practices were tacked on, as a relative afterthought, beginning with testimonial resources from the individual inquirer’s perspective. Function-first epistemology, by contrast, again offers a useful corrective. It typically begins with social practices. Sociality is no mere afterthought; it is baked into the very conception.<sup>8</sup> For this reason, function-first epistemology is a powerful – albeit underexplored – approach for theorising epistemic injustice.<sup>9</sup> Finally, the method foregrounds how practical considerations shape epistemic phenomena. I have reservations about listing this as unambiguously a virtue, but Hannon certainly would, and this facet of the method informs my comments.

### ***Adjudicative tenuity***

One potential pitfall of function-first approaches concerns whether they are sufficiently powerful to adjudicate amongst rival claims. Critics charge that the method lacks discriminatory power and so risks indulging in just-so stories. I do not press this objection against function-first methods myself – I employ function-first approaches – but the concern is worth understanding and learning from. Judicious use of any method requires understanding and respecting its limitations.

To see this worry, consider the functional hypothesis that knowledge ascriptions serve to tag good informants. And consider case-based objections, including seemingly good informants who lack knowledge and knowers who are not good informants (43-44). The former includes, for example, a political spokesperson who reliably reports well-researched vaccine facts that she, as an anti-vaxxer, does not believe and so does

---

<sup>8</sup>One can employ function-first approaches with an emphasis on individual inquiry, but most functional hypotheses foreground social features.

<sup>9</sup>For a survey, see Gardiner (2021b). Mills (2007) insightfully describes the rational path from naturalistic epistemology to a thoroughly non-idealised, social epistemology that centres human embodiment, limitation, vice, and injustice.

not know. The latter includes isolated ‘Robinson Crusoe’ knowers and non-linguistic people.

If the hypothesised link between knowledge and good informants purports to be an exceptionless claim, reductive analysis, or necessary or sufficient condition, such counterexamples are ruinous. When confronting such cases, advocates of the target theory have few options and the dialectical moves are fairly clear. In response to such counterinstances, the theorist could dispute whether the person is a knower or good informant, or – failing that – either revise or abandon the target theory. There is scant latitude for dismissing counterexamples as unimportant or ignorable.

For a function-first epistemologist, by contrast, substantially more moves are available. They might claim a pertinent condition of a good informant, such as availability to the inquirer, was essential to a ‘proto-concept’ of knowledge but has since been shed as the concept ‘objectivised’, that is, evolved into something less user-centric and more universally applicable (Craig 1990, 96–99). Or they might contend the target phenomenon does not reflect recent social innovations. The concept of knowledge predates spokespersons, for example, and so may conflict with recent practices. Or they might claim the conditions are met in an extended sense. Hannon argues that non-linguistic people qualify as good explainers because they can provide explanations to themselves, for example (233). Or they might argue discrepancies derive from some secondary functional role. Or they can concede such cases do not accord well with the target proposal, but aver they are peripheral and not decisive objections because function-first epistemology need not countenance every case. Hence function-first theorising provides significant latitude in the face of case-based objections.

The corresponding worry arises: Is the method so accommodating that any observation can be massaged to fit any functional hypothesis, and that any functional hypothesis can be made to support various putative upshots about knowledge? Rival claims can appear omni-compatible. That is, they seem to absorb any putative objection and accordingly it is unclear which considerations carry weight. As noted above, I endorse function-first methods in philosophy and do not regard the concerns as lethal, but they are worth taking seriously as an edifying source of caution.

### **Epistemic pragmatism and the semantics of knowledge ascriptions**

As described above, function-first epistemology has a relatively low capacity to adjudicate amongst rival claims. Its strengths lay elsewhere,

such as illuminating the contours of epistemic practices, highlighting suggestive ideas, and (arguably) underwriting a naturalistic explanation of epistemic phenomena. My three critical comments all stem from this relative weakness of function-first theorising. That is, if my three critical comments identify errors, there is commonality in the source: The errors stem from the method's adjudicative tenuity, which allows theories to be massaged to fit various claims.

Hannon himself displays an admirable tendency towards compatibilism. He perceives the good in various views and synthesises them into a valuable whole. But the danger of 'just so' reasoning lurks and – especially given the concerns articulated above – his compatibilist orientation can veer towards over-compatibilism. Hannon's attraction towards compatibilism, or perhaps over-compatibilism, are in the background of my critical comments. To forestall confusion, I should clarify that I call it 'compatibilism' and 'over-compatibilism' because Hannon aims to capture the good in rival views, but his conclusions are far from the anodyne platitude that each theory says something correct. As will be apparent, some of Hannon's resulting claims are bold, revisionary, and controversial.

In two places, Hannon makes the following dialectical move: The function-first approach, or Hannon's specific functional hypothesis, does not favour or adjudicate amongst incompatible claims. Therefore, Hannon concludes, there is no reason to favour one over the other. Or – more extremely – there is no fact of the matter about which is true; the truth is not merely underdetermined by the method or total available evidence, but is instead indeterminate.

The most striking example of this move concerns the semantics of knowledge ascriptions. (The second example, discussed below, concerns his treatment of skeptical challenges.) Hannon investigates whether proposed functional hypotheses adjudicate amongst rival semantics for knowledge ascriptions (157–169) and concludes that 'putative facts about the function of knowledge ascriptions will not instruct us as to whether contextualism, sensitive invariantism, or insensitive invariantism is true' (158). This is because various semantic theories are wholly compatible with proposed central functions of knowledge ascriptions.

Hannon then uses the idea that proposed functional hypotheses are compatible with rival semantic theories to develop and 'tentatively endorse' a more radical claim: There is no fact of the matter about which semantic theory is true.<sup>10</sup> Hannon's suggestion, which he calls

---

<sup>10</sup>Hannon (2021, 118, 2019, 177) 'fully agree[s]' with a parallel claim about other parts of speech.

‘epistemic pragmatism’ is that the meaning of ‘knows’ should not be understood by appeal to truth conditions, but instead by its practical functions (158, 173, 180–188).

It will be helpful to articulate three claims:

**Not Dispositive.** Function-first theorising cannot determine which rival semantic theory is true.

‘Not Dispositive’ claims that, to use Hannon’s term, ‘we cannot *derive* the correct semantic theory of “knows” from putative facts about the function of knowledge ascriptions’ (158, emphasis added).

**No Support.** Function-first theorising is not probative about which rival semantic theory is true.

‘No Support’ is a stronger claim. It holds that functional hypotheses do not favour one semantic theory over another; they are neutral with respect to rival semantic theories.

**Epistemic Pragmatism.** Rival semantic theories like contextualism, sensitive invariantism, and insensitive invariantism are all incorrect because they rest on the mistaken foundational assumption that the meaning of epistemic terms is given by truth conditions. Instead their meaning is determined by function.

It is worth emphasising how radical epistemic pragmatism is. Hannon writes,

I make a more radical proposal: I suggest this entire debate about the semantics of ‘knows’ mistakenly presupposes that we should account for the meaning of epistemic claims by determining their truth conditions. (158) [...]

Epistemic pragmatism attacks a widely held picture of language (173) [...] When investigating our knowledge claims, [these] theorists jointly presuppose that the question ‘Did she say something true or something false?’ is always in order, in the sense that the agent must have spoken truthfully or falsely. I want to question this presupposition. (175–176)

Although epistemic pragmatism is a strikingly bold suggestion, it nonetheless exhibits Hannon’s compatibilist tendency because he presents the view as seeing the good in conflicting claims. He writes,

I find it plausible to say there is no fact of the matter as to whether contextualism or some type of invariantism is the correct account of [the] utterance. They are equally good semantic models for understanding [the utterance], but neither has priority. Rather, they simply offer different ways to carve up semantics. (177)

Tempering the perceived radicality of the proposal, Hannon writes,

While this is a fairly radical view, it is really only radical at the level of theory. It does not call for a revision of our ordinary practice of knowledge ascriptions; at most, we are required to rethink philosophical practice. (173)

Hannon's comment that ordinary practice is preserved is gravid. He emphasises that adherents of rival semantic theories can effectively communicate and agree about which utterances are appropriate. (See the airport and fridge illustrations (170; 177)). Hannon leverages this observation to motivate epistemic pragmatism. He writes,

Knowledge ascriptions can do the work that we need or might reasonably expect them to do [...] *without our ever needing to attend* to the question of whether their semantics should be understood along contextualist or invariantist lines. (171)

In response: The fact we can use language effectively without attending to its semantics is no surprise. Ordinarily people use indexicals like 'I' and 'now' without attending to Kaplanian character. But this does not indicate much about their semantics and certainly does not suggest current semantic theorising rests on a mistake.

Secondly, it is not surprising if function-first theorising is ill-equipped to adjudicate amongst rival semantic theories. We should expect this from reflecting on contours of the method. This is because function-first epistemology focuses on central, paradigm uses. It is relatively quiet about unusual and marginal cases, such as embedded uses or direct quotation. But marginal cases and other nuanced linguistic evidence adjudicate rival semantic claims. Central paradigm uses won't cut it. This is because any plausible semantic theory will have already taken central uses of the term into consideration. There will be very little, if any, tension between central paradigm uses of 'knows' and live rival semantic theories.

We thus should not expect a theory about the central function of knowledge ascriptions to have much force for deciding amongst semantic theories. This does not impugn function-first theorising; it is simply a boundary of the approach. We can combine function-based theorising with other methods to fully illuminate the nature of knowledge. The methods of semantic modelling – such as investigating language acquisition, error patterns, semantic blindness, indirect and embedded contexts, hedged assertions, crosslinguistic patterns, and comparisons to other parts of language, such as indubitably context-sensitive terms – are better equipped to adjudicate semantic claims.

I described the oft-aided objection that within function-first epistemology it can appear that any functional hypothesis can be massaged to support any putative upshot. This worry can be exacerbated by attempts to deploy function-first theorising outwith its area of power, such as to adjudicate rival semantic claims. We can garner support for function-first epistemology by acknowledging and respecting its limits.<sup>11</sup>

Hence NOT DISPOSITIVE and NO SUPPORT might well be true, but this doesn't motivate epistemic pragmatism. Instead function-first epistemology's relative lack of adjudicative power indicates the importance of employing a plurality of philosophical methods; function-first epistemology cannot answer every question about the nature of knowledge.<sup>12</sup>

Hannon's response will, I think, contend that I do not sufficiently appreciate the radical depths of pragmatism, because I assume the orthodox truth-conditional view (182). Hannon holds that NOT DISPOSITIVE and NO SUPPORT cast doubt on this assumption. His view is, I think, something like this: The fact that reflecting on the function of knowledge ascriptions does not adjudicate amongst rival semantic claims indicates there is no practical pressure to settle the question one way or the other; epistemic evaluations perform their function equally well, regardless of which semantic theory is correct. And if there is no practical pressure to settle the question one way or the other then, at least in the case of knowledge claims, there is no practical pressure for the fact to be settled at all. This is a bold proposal about the metaphysics of language; I encourage meta-semanticists to engage with these ideas.

### The skeptic's error

A second example of Hannon's tending towards compatibilism – or perhaps over-compatibilism – is his treatment of radical skepticism. Hannon denies epistemic contextualism and global skepticism. But he also rejects anti-skepticism because, he claims, skeptical challenges are sometimes appropriate (216). This exemplifies Hannon's compatibilist tendencies because he aims to vindicate both skeptical challenges and everyday knowledge claims. I argue this is too concessive towards the skeptic. Given his broader theory, Hannon should instead endorse anti-skepticism.

---

<sup>11</sup>Gardiner (2015) and Henderson (2021) discuss the place of function-based theorising within a methodologically plural reflective equilibrium.

<sup>12</sup>Lawlor (2021) raises similar objections.

Hannon contrasts two conceptions of knowledge. The infallibilist ‘Cartesian requirement’ holds that knowledge requires eliminating all error possibilities (198). The less demanding, fallibilist ‘Austinian requirement’ restricts this to the *relevant* error possibilities (202). The more demanding conception underwrites many skeptical challenges. Hannon investigates how we can adjudicate between the two conceptions.

Function-first epistemology examines the forces that shape practices of epistemic evaluation. Craig, Hannon, and other function-first epistemologists note that practical factors exert a downward pressure on justificatory thresholds for knowledge and other epistemic standards, such as warranted belief and assertion. The lower, fallibilist Austinian standard is more useful than the extremal Cartesian standard. One reason is simply that more beliefs meet it and so a concept marking the lower threshold is deployed more often. Like any term, concept, or meme, it gains traction from use. The infallibilist standard, by contrast, faces extinction – or simply never arises – from desuetude.

There are further reasons the function-first approach favours the lower standard. By hypothesis, no one can eliminate radical skeptical error possibilities. They are phenomenologically indistinguishable from current experiences. There can be no practical purposes for tagging informants who can distinguish radical skeptical hypotheses, because no such informants exist. That mark of epistemic evaluation would not take hold because no one meets the standard (Craig 1990, 116; Hannon 2019, 210–211). As Hannon notes, ‘expecting informants to satisfy the Cartesian requirement would frustrate our communal epistemic practices’ (211). Similar claims apply *mutandis mutatis* for rival functional hypotheses. There is no point marking a threshold of appropriately ending inquiry, for example, that no one can possibly reach.

Practical factors thus exert downward pressures towards a fallibilist conception of knowledge. They also exert upward pressures. But – crucially – not towards skepticism-inducing infallibilist heights (212–213).

Yet rather than simply endorsing anti-skepticism, Hannon’s compatibilist tendencies kick in (216–219). He argues that skeptical challenges are sometimes appropriate, specifically in contexts of ‘pure inquiry’ wherein Cartesian thinkers aim to avoid all falsehood and during philosophical investigation about human knowledge. These contexts, Hannon argues, render radical skeptical doubts about knowledge claims apt. We should concede these defeats to skepticism, Hannon says, because he sees no promising account of where the skeptic goes wrong (219).

Hannon concedes too much to the skeptic and his own resources warrant a more anti-skeptical stance than he credits them with. When Hannon characterises the enterprise of ‘pure inquiry’ and the method of Cartesian doubt he does not use the term ‘knows’. This is telling. The pure inquirer aims to avoid false belief and seeks rational certainty. But what ties these aims of inquiry to knowledge or knowledge ascriptions?

Pure inquiry that aims to avoid all falsehood and philosophical theorising about the nature of knowledge and are not contexts that honed the concept of knowledge. Practices of knowledge ascriptions gestated in prosaic contexts. And knowledge, whatever its ontology and ontogeny, must predate theorising about knowledge. Why (or when, or how), on Hannon’s account, would these pure and philosophical contexts have left their mark on the concept of knowledge?

Skeptical challenges get this much right: We value rational certainty and infallible levels of epistemic support. Our beliefs could be wrong; we cannot eliminate every possibility of error. And radical skeptical hypotheses are phenomenologically indistinguishable from our experiences. What skeptics have not shown is that these facts undermine knowledge or knowledge claims.<sup>13</sup> A legitimate skeptical challenge in almost any context – philosophical or mundane – notes we cannot rule out every possibility of error. This contention is true for almost every claim. (The cogito is a notable exception.) But this legitimate skeptical challenge does not involve knowledge or knowledge ascriptions.

My opening comments noted that function-first theorising centres on practical forces that shaped epistemic practices. And I registered, but didn’t endorse, the worry that function-first epistemology lacks the adjudicatory power to support some upshots over others. Tying these themes together, I suggest Hannon’s practical orientation has more anti-skeptical ammunition than he credits it with. Not only is the answer not underdetermined by his functional hypothesis, or left indeterminate, Hannon’s own resources tell us which view to endorse: The skeptic is wrong about knowledge.

Craig (1990, 116) notes there are not practical pressures to create precise boundaries marking the justificatory standards for knowledge. This is true. Craig then claims ‘the resultant area of indeterminacy hosts the controversy about scepticism’ (117). This is mistaken. Radical

---

<sup>13</sup>Gardiner (2015, 42–43) conceives of skeptical challenges as attempts to show we lack something of epistemic value that we antecedently thought we possessed. It is not sufficient to show we lack something that we didn’t think we had. Gardiner (2021a) employs this anti-skeptical stance against mundane ‘doubt-mongers’ about, for example, rape accusations and vaccine science.



skepticism lies at a radical extreme, not at the vague and indeterminate boundary of how much evidence is required for knowledge. To see this we must distinguish between whether practical forces chisel a sharp boundary on a scale and whether practical forces push towards an extreme of that scale.

Here is a vivid example. There is usually no practical reason to determine precisely how many decibels counts as talking too loudly in a shared office. There is also usually no practical reason to shout in a shared office. Saying the lack of a sharp boundary is the reason one shouts in a shared office conflates the two. Saying the lack of a sharp justificatory threshold gives rise to skepticism risks making a related, albeit more subtle, error. The vagueness of justificatory thresholds concerns, for example, whether merely recalling eating berries suffices to know they are safe or whether one must recall it vividly. This is nowhere near the scale's radical skeptical extremes.

To provide another analogy: There is a practical reason to create a sharp threshold for what qualifies as entering someone's land. And there are often good reasons to enter someone's land. To contrast, there is no practical reason to forge a sharp boundary for what qualifies as being too close to a poisonous gas cloud. And there is rarely practical reason to get close. In both cases, whether there is reason to create a sharp boundary differs from whether there is reason to approach or traverse that boundary. Such analogies illuminate the terrain of skeptical challenges, especially within a function-first framework.

## Plural and rival functions

Chapter five describes various functions of knowledge, knowledge ascriptions, and the concept of knowledge. This includes, for example, assuring others, countering doubts, signalling the appropriate end of inquiry, marking the degree of warrant normally sufficient for practical reasoning, and incentivising good epistemic conduct (103). If theorists claim these functions are central or primary, they constitute rival functional hypotheses and compete with Hannon's claim that the central function of knowledge ascriptions is tagging good informants.

Hannon concedes that knowledge ascriptions play some of these other roles – exhibiting his admirable orientation towards compatibilism – but contends his favoured function is more foundational. Tagging good informants, Hannon claims, is the primary role. It is explanatorily and temporally first and explains why knowledge ascriptions have the structure

they do.<sup>14</sup> Hannon provides a hammer analogy (24; 136-137). The central function of a carpentry hammer determines its shape and heft, which helps explain why such hammers exist and how they can serve derivative functions.

The worry, though, is that sometimes these claims feel unconvincing or 'just so'. When considering the rival functional role of appropriately terminating inquiring, for example, Hannon concedes that knowledge ascriptions serve this need. But, he says,

This idea is not incompatible with Craig's hypothesis about our need to identify reliable informants. Rather, I think these two functions are just different sides of the same coin. This is because *the* way to reasonably terminate inquiry is by identifying a sufficiently reliable informant. (109, emphasis added)

But testimony is not *the* way to reasonably terminate inquiry; it is one of many. Testimony is important and epistemologically rich, but other ways of ceasing inquiry, such as remembering, looking, and testing are far more common. In everyday life, we look to see whether the oven is off far more often than we ask someone else whether it is off. Indeed, at least relative to perception and memory, testimony is likely an information source that *fails* to end inquiry, because we doubt each other's word. And forgetting the question, being distracted, and ceasing to care are also commonplace, epistemically legitimate ways to end inquiry in some contexts.

Hannon continues,

The dual roles of identifying reliable informants and terminating inquiry can therefore be given a unified treatment [...] [B]oth hypotheses appear to be mutually supporting. In fact, the role of identifying reliable informants appears to be *more fundamental* than the inquiry-stopping function because the former is explanatorily prior. That we have found a reliable informant explains why we may terminate inquiry at a certain point. Without such an explanation, it is unclear when we should reasonably end inquiry. (109, emphasis in original)

But, in response, ending inquiry seems more foundational than identifying reliable informants. Good informants are typically those who have already reasonably ended inquiry. And it is because we want to end inquiry that practices of testimony emerge. Thus if either is explanatorily prior, or more basic, it is the need to cease inquiry. This seems especially plausible for the function of knowledge self-attributions and knowledge

---

<sup>14</sup>We should not oversimplify the temporal priority. Multiple functions can play symbiotic sculpting roles concurrently as the phenomenon evolves. But Hannon must claim the informant-tagging function played a substantial sculpting role early in the concept's ontogenetic development, since otherwise the function cannot principally explain its structure.

itself, as opposed to third-party knowledge ascriptions. As noted above, function-first epistemology admirably foregrounds social aspects of epistemic life, but here it risks inappropriately centring them. Social facets of cognitive life are important, but the need to end inquiry is plausibly more basic.<sup>15</sup>

In a similar vein, Hannon considers Mikkel Gerken's (2015) claim that knowledge ascriptions serve to mark the threshold that normally corresponds to appropriate assertion and practical reasoning (113–116). But, Hannon argues, Gerken's view does not explain what that threshold is. Hannon claims that his favoured functional hypothesis can explain this (115).

The point of the concept of knowledge, I claim, is to certify reliable informants to members of our epistemic community. To say that an agent knows that *p* is to say that she is epistemically positioned with respect to *p* so as to be a good source of *actionable information*, which means that we may take it from her that *p*.<sup>16</sup> (115, emphasis added)

Hannon claims his favoured function 'is more explanatorily fundamental' than Gerken's threshold-marking function (114), is 'the primary purpose' of the concept of knowledge, and 'provides the backbone for the other functions we have thus far considered' (114–116). Again this conclusion seems backwards. Appropriate action, including actions underwriting practices of testimony and 'certifying' one another, seems more foundational than certifying good informants to the community. Indeed, Hannon's gloss on a good informant includes 'actionable information', which suggests thresholds of proper action are more foundational than his favoured functional hypothesis.

There are advantages to foregrounding social aspects of epistemic life, such as tagging good informants. But plausibly these advantages accrue in benefits like theorising varieties of epistemic pathology, dependency, and injustice, rather than showing which roles and practices are more explanatorily fundamental. This is because plausibly proper action, inquiry cessation, and our desire for surety are more foundational than certifying good informants to our epistemic community. Thus Hannon's tendency towards compatibilism can veer into the tendentious. He absorbs rival claims about the central function of knowledge ascriptions

---

<sup>15</sup>Rysiew (2012) and Elgin (2021) also argue for this conclusion.

<sup>16</sup>Hannon's 'is to say' formulation is stronger than normal function-first hypotheses and seems stronger than Hannon needs or intends.

into his own view and claims his view explains others. One worries that claims are being massaged.

I offered three potential objections to Hannon's ideas. Although heterogeneous, each is rooted in the relative tenuity of the function-first method at adjudicating amongst rival claims and – in some sense – arise from Hannon's compatibilist tendencies.<sup>17</sup> I described Hannon's epistemic pragmatism about the semantics of knowledge ascriptions, and I instead suggest that concordance between plausible functional hypotheses and semantic theories is precisely what we should expect from reflecting on their respective methodologies. And so this harmony offers no reason to reject mainstream semantic theorising; it simply shows the need for multiple methodologies in reflective equilibrium. Secondly, I suggested Hannon is too concessive towards radical skepticism and that – rather than adjudicative debility – his theory offers substantial anti-skeptical resources. Finally, although function-first epistemology offers substantial latitude for hypothesising about which functions are fundamental, I doubt that tagging good informants – especially understood as certifying or marking informants for a broader epistemic community – is more foundational than the need to appropriately end inquiry or rely on information for action.

Thus my latter two critical comments suggest that, rather than lacking probative power, the function-first method exhibits adjudicative force, but its areas of probative power can be difficult to discern. To emphasise, I offer these critical comments within a backdrop of great respect for Hannon's research. His monograph is full of valuable insights, which advance both function-first epistemology and other areas of epistemological inquiry. I look forward to engaging with him further about these ideas.

## Acknowledgements

Many thanks to Jon Garthoff and Michael Hannon for invaluable discussion on these topics.

## Disclosure statement

No potential conflict of interest was reported by the author(s).

---

<sup>17</sup>Another instance of Hannon's compatibilist tendency is his aim of reconciling a shared justificatory standard for knowledge with the impurist claim that standards for knowledge are determined by an individual's practical context (82). As a result, his 'communal impurist' justificatory threshold for knowledge has a bipartite 'default and defeat' structure. See Lawlor (2021, 110–111) and Gardiner (2021b) for critical discussion; the expression 'default and defeat' is from Lawlor.

## Funding

This research was supported by an ACLS Fellowship from the American Council of Learned Societies.

## References

- Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- Elgin, Catherine. 2021. "The Function of Knowledge." *Analysis* 81 (1): 100–107.
- Gardiner, Georgi. 2015. "Teleologies and the Methodology of Epistemology." In *Epistemic Evaluation: Purposeful Epistemology*, edited by John Greco, and David Henderson, 31–45. Oxford: OUP.
- Gardiner, Georgi. 2017. "In Pursuit of Understanding." Doctoral thesis. Rutgers University.
- Gardiner, Georgi. 2021a. "Banal Skepticism and the Errors of Doubt: On Ephecticism About Rape Accusations." *Midwest Studies in Philosophy* 45: 393–421.
- Gardiner, Georgi. 2021b. "Book Review 'What's the Point of Knowledge? A Function-First Epistemology' by Hannon, Michael. Oxford University Press, 2019, ix + 275 pp. ISBN: 9780190914721. \$78.00." *European Journal of Philosophy* 29: 674–678.
- Gerken, Mikkel. 2015. "The Roles of Knowledge Ascriptions in Epistemic Assessment." *European Journal of Philosophy* 23 (1): 141–161.
- Hannon, Michael. 2019. *What's the Point of Knowledge?* Oxford: Oxford University Press.
- Hannon, Michael. 2021. "Replies to Henderson, Elgin and Lawlor." *Analysis* 81 (1): 114–129.
- Henderson, David. 2021. "What's the Point of Knowledge?" *Analysis* 81 (1): 87–100.
- Kornblith, Hilary. 2011. "Why Should We Care about the Concept of Knowledge?" *Episteme* 8 (1): 38–52.
- Lawlor, Krista. 2021. "Comments on *What is the Point of Knowledge?*" *Analysis* 81 (1): 107–114.
- Mills, Charles. 2007. "White Ignorance." In *Race and Epistemologies of Ignorance*, edited by S. Sullivan, and N. Tuana, 11–38. Albany: State University of New York Press.
- Rysiew, Patrick. 2012. "Epistemic Scorekeeping." In *Knowledge Ascriptions*, edited by J. Brown, and M. Gerken, 270–293. Oxford: Oxford University Press.



# Who knows what? Epistemic dependence, inquiry, and function-first epistemology

Joshua DiPaolo

California State University Fullerton, Fullerton, USA

## 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 and competition; 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,

**CONTACT** Joshua DiPaolo  [jdipaolo@fullerton.edu](mailto:jdipaolo@fullerton.edu)  California State University Fullerton

© 2022 Informa UK Limited, trading as Taylor & Francis Group

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).<sup>1</sup> 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

---

<sup>1</sup>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.

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 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 knowledge 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.<sup>2</sup> 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 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.'<sup>3</sup>

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.'

---

<sup>2</sup>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.

<sup>3</sup>This example is inspired by Craig (1990, 11).

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

A more pervasive example of the need to track epistemic dependents' knowledge is education.<sup>4</sup> 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.<sup>5</sup> 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

---

<sup>4</sup>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.

<sup>5</sup>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.

knows about your medical background.) When working with others to achieve shared aims – communicating, advancing 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.<sup>6</sup> 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

---

<sup>6</sup>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.

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 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.<sup>7</sup>

---

<sup>7</sup>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.<sup>8</sup>

This proposal captures all of our cases. Epistemic dependence: We track others' knowledge to know which possibilities must still be, or need not 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 cannot 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

---

<sup>8</sup>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  $p$ ."



typically a means to 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 knowledge 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 business; 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.<sup>9</sup> 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 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

---

<sup>9</sup>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.

credentials won't settle the matter. If the *primary* function is whichever explains the others, then, I have argued, 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.

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## 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* 81: 100–107.
- Goldman, A. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73: 771–791.
- Grice, P. 1975. "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.
- Hannon, M. 2021. "Replies to Henderson, Elgin and Lawlor." *Analysis* 81: 114–129.
- Schmitt, F. 1992. "Review of Knowledge and the State of Nature: An Essay in Conceptual Synthesis." *Mind; A Quarterly Review of Psychology and Philosophy* 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*, 136–151. Cambridge: Cambridge University Press.



## Reply to Gardiner and DiPaolo

Michael Hannon

Department of Philosophy, University of Nottingham, Nottingham, UK

### ABSTRACT

In this article, I reply to comments on my book by Georgi Gardiner and Joshua DiPaolo. I will first reply to Gardiner's comments, focusing primarily on her doubts about the adjudicative power of function-first epistemology. I will then reply to DiPaolo, who argues that I have misidentified that primary function of the concept of knowledge.

**ARTICLE HISTORY** Received 4 August 2022; Accepted 4 August 2022

**KEYWORDS** Knowledge; function-first epistemology; skepticism

In this article, I reply to comments on my book by Georgi Gardiner and Joshua DiPaolo. I will first reply to Gardiner's comments, focusing primarily on her doubts about the adjudicative power of function-first epistemology. I will then reply to DiPaolo, who argues that I have misidentified that primary function of the concept of knowledge.

### 1. Reply to Gardiner

Gardiner makes three critical remarks that are unified by a common theme. Her overarching concern is that my view is *underpowered*. She writes, 'One potential pitfall of function-first approaches concerns whether they are sufficiently powerful to adjudicate amongst rival claims.' Gardiner worries that my view lacks such discriminatory power; for example, she says, 'function-first epistemology has a relatively low capacity to adjudicate amongst rival claims'. Following Gardiner's terminology, let's call this *the problem of adjudicative tenuity*. Gardiner argues that my treatment of issues like pragmatism, skepticism, and epistemic pluralism all succumb to this problem. As she says, 'I offer three potential objections to Hannon's ideas. Although heterogeneous, each is rooted in the relative tenuity of the function-first method at adjudicating amongst

**CONTACT** Michael Hannon  michael.hannon@nottingham.ac.uk

© 2022 Informa UK Limited, trading as Taylor & Francis Group



rival claims.’ In what follows, I will first clarify the problem of adjudicative tenuity and then I will examine Gardiner’s three critical comments in more detail.<sup>1</sup>

We can interpret the problem of adjudicative tenuity in at least two ways. On the first interpretation, my view does not have the resources *to justify the selection of one functional hypothesis over alternative proposals*. In the book, I argue that the primary function of the concept of knowledge is to ‘identify reliable informants to members of our community’. I call this the ‘informant-flagging function’. Yet epistemologists have proposed other hypotheses about the point of our knowledge concept, such as ‘signaling the appropriate end of inquiry’ or ‘providing assurance to others’. Can I adjudicate amongst these rival views? Is there any reason to favor the informant-flagging hypothesis over these other functions? If I cannot explain why my favored hypothesis about the functional role of knowledge is more plausible than alternatives, then I face one version of the problem of adjudicative tenuity.

The problem of adjudicative tenuity can be interpreted another way. Instead of being unable to tell us which functional hypothesis is correct, my preferred functional hypothesis might be *unable to adjudicate between rival (non-functional) epistemological theses*. For example, we might wonder whether skepticism is true or false, or whether knowledge claims are contextual or invariant, or whether pragmatic factors encroach on the epistemic or not. In my book, I attempt to cast new light on these issues by reflecting on the purpose of epistemic evaluation. But one might doubt that my account provides the resources to settle these disputes. In other words, *even if we accept my hypothesis* about the social role of knowledge, my approach may be too underpowered to tell us which of these rival views (e.g. skepticism or anti-skepticism, contextualism or invariantism, etc.) is correct. This is the second interpretation of the problem of adjudicative tenuity.

Gardiner doesn’t distinguish these two interpretations, but she seems to have both in mind. In the section of her commentary titled ‘Epistemic Pragmatism and the Semantics of Knowledge Ascriptions’, Gardiner worries that my informant-flagging hypothesis is ‘over-compatible’ because it does not tell us whether contextualism or invariantism is the

---

<sup>1</sup>Gardiner’s commentary is ambivalent on the issue of ‘adjudicative tenuity’. On the one hand, Gardiner says she ‘does not press’ the objection about adjudicative tenuity. She claims to ‘register, but not endorse’ this worry. On the other hand, she acknowledges that all three of her objections are rooted in the adjudicative tenuity of function-first epistemology. I found this puzzling. If her objections to my view ‘stem from the method’s adjudicative tenuity’, then why does she claim to not press this objection?

correct semantic theory of 'knows'. This supports the second interpretation of the problem of adjudicative tenuity: my account does not adjudicate amongst incompatible (non-functional) theses about knowledge. At other points in her commentary, however, Gardiner seems to adopt the first interpretation: my view is unable to explain why one functional hypothesis is more foundational than another. For example, Gardiner asks, 'Is the method so accommodating that *any* observation can be massaged to fit *any* functional hypothesis, and that *any* functional hypothesis can be made to support the various putative upshots about knowledge?' (emphasis mine). She later asserts that my argument in favor of the informant-flagging hypothesis is 'tendentious' and she 'worries that claims are being massaged'.

Now, I doubt that Gardiner really believes *any* functional hypothesis can be made to fit the data. If I were to hypothesize that we have a concept of knowledge to identify ham sandwiches, then I could easily be refuted. So, not *any* observation can be made to fit *any* functional hypothesis. More charitably, Gardiner might mean there is no obvious reason to favor one functional hypothesis over other *live contenders*. But she cannot mean this, either. In the section titled 'Plural and Rival Functions', Gardiner argues that my own account gets things backwards: instead of thinking that my favored hypothesis 'is more explanatorily fundamental' than other plausible hypotheses (as I argue in the book), she says that *other hypotheses are more explanatorily fundamental than the informant-flagging function*. As far as I can tell, this is incompatible with the charge of adjudicative tenuity in the first sense. If Gardiner can show that some particular hypothesis about the function of knowledge is more fundamental than my preferred hypothesis, then the method of function-first epistemology is able to adjudicate between competing functional hypotheses, at least to some degree. Thus, her objection to my view must be that I've *misidentified* the primary role of our knowledge concept. I will dispute this claim below.

In the section of her commentary titled 'The Skeptic's Error', Gardiner argues that my treatment of radical skepticism goes awry in the following way: while I try to vindicate both skeptical challenges and everyday knowledge claims, Gardiner says I should simply 'endorse anti-skepticism'. In other words, she thinks I am too concessive to the skeptic. Indeed, Gardiner says my own view *provides the resources* to favor anti-skepticism over skepticism. I will question this claim below. For the moment, I again just want to emphasize that this very objection seems incompatible with the charge of adjudicative tenuity, but this time it is incompatible

with the second interpretation of this challenge. If reflecting on the purpose of knowledge leads us to reject skepticism in favor of anti-skepticism (as Gardiner claims), then it seems my account *can* adjudicate between at least some rival epistemological theses. Thus, Gardiner's objection must be that function-first epistemology is not sufficiently powerful to adjudicate amongst other rival claims in epistemology. I am happy to grant this general point, but I would like to know more about which debates she has in mind (and whether my view is at a comparative disadvantage relative to other approaches).

Let's now consider Gardiner's objections to my treatment of epistemic pragmatism, skepticism, and epistemic pluralism, respectively.

### 1.1. Epistemic pragmatism

In Chapter 7 of my book, I make two claims about the semantics of knowledge claims. First, I argue that appealing to functional considerations will not help us decide between contextualism, sensitive invariantism, and insensitive invariantism: these semantic theories are all compatible with the informant-flagging function of our knowledge concept. Gardiner calls this view 'NO SUPPORT'. Second, I argue that *none* of these semantic theories are correct. This entire debate about the semantics of 'knows' rests on an error: it presupposes that we should account for the meaning of epistemic claims by determining their truth conditions. A more natural way to approach the meaning of epistemic claims, I argue, is to ask what practical functions they serve us in communicating with each other. I call this view 'epistemic pragmatism'.

As Gardiner correctly points out, 'NO SUPPORT' does not motivate epistemic pragmatism. It would be illegitimate to move from the claim that (a) functional considerations are unable to decide between contextualism, sensitive invariantism, and insensitive invariantism to the claim that (b) the entire debate about the semantics of 'knows' is misguided. But this is not my argument. Here's what I say in the book:

Even if we assume that knowledge ascriptions can fulfill their role irrespective of whether their semantics is contextualist or invariantist, *it does not follow* that neither of these views is correct, or that the answer is indeterminate. One might argue that it is far more plausible that the semantics of knowledge ascriptions could have been one way or the other, at some point, but then the semantics became settled. . . The problem with this reply, however, is that it simply assumes the truth of the orthodox view about how language works, namely, that the semantics of 'knows' is determinate between contextualism and

some type of invariantism. What is under investigation is the way we insist, in the course of doing philosophy, on the truth (or falsity) of a sentence that is uttered. Why expect our knowledge ascriptions to come equipped with a semantics that is determinate between contextualism and invariantism? The fact that our knowledge ascriptions achieve their purpose whether or not a specific semantic theory is true of them should lead us to wonder why we must assume that one of these semantic accounts is somehow fundamental to our concept of knowledge. (Hannon 2019, 182, emphasis mine)

In other words, I argue that (a) the function of our epistemic concepts plausibly shapes the semantics of our epistemic terms, (b) our epistemic evaluations achieve their characteristic purpose whether we understand their appropriate use in terms of literal content or conversational pragmatics, and thus (c) there is no clear reason to assume the semantics must determinately settle either in favor of contextualism or some type of invariantism *precisely because* our practice of epistemic evaluation works perfectly well in the absence of such an answer. We should therefore question the motivation for thinking there *is* such an answer. If we assume that debates about the meaning of 'knows' must be settled by determining truth conditions, then we obscure from view accounts that explain the meaning of the target claims in terms other than truth conditions.

This is a controversial view. In the book, I admit a good deal of uncertainty about it. If it were to turn out that NO SUPPORT is true but epistemic pragmatism is false, I would not consider this a devastating blow to my theory. As Gardiner correctly observes, it would not impugn function-first theorizing if it were unable to decide between contextualism and invariantism; this would just be a boundary to what the approach can achieve. The point I want to emphasize, however, is that my argument does not attempt to motivate epistemic pragmatism by appealing to NO SUPPORT. Instead, I encourage us to rethink a common presupposition in philosophy about how best to account for the meanings of epistemic claims. If we take the purpose of epistemic evaluation as primary in our theorizing, we may realize there is something confused about the standard attempts to determine the semantics of knowledge ascriptions by identifying their truth conditions.

## 1.2. Skepticism

Gardiner says I am too concessive to the skeptic. In my book, I do not attempt to resolve the clash between skepticism and fallibilism by endorsing one perspective over the other. Instead, I appeal to considerations

about the function of epistemic evaluation to explain (a) the force of the skeptic's argument and (b) the legitimacy of our ordinary knowledge claims. I try to explain why knowledge is possible without assuming there is something drastically mistaken about the skeptical challenge. In contrast, Gardiner says I should endorse anti-skepticism over skepticism. Her argument, roughly, is that skepticism goes against the point of our concept of knowledge, which is to identify reliable informants in ordinary contexts.

I am sympathetic with this idea. In fact, I published an article titled 'Skepticism: Impractical, Therefore Implausible' in the same year as my book. So, I really am of two minds about this issue. What I can say with confidence is this: if I were forced to choose between skepticism and fallibilism, I would choose fallibilism. As David Lewis says, 'Better fallibilism than skepticism'. But Lewis also adds, 'it would be better still to dodge the choice' (1996, 550). So, here's a question: if we can rescue knowledge from the skeptic's clutches while *also* explaining the rational intelligibility of skeptical doubt, why not dodge the choice?

To cast our lot with anti-skepticism, as Gardiner recommends, would leave us with a puzzle: what explains our temptation towards skepticism in the first place? According to Gardiner, the skeptic is illicitly importing unduly high standards for knowledge. The skeptic takes knowledge to require something that it does not in fact require. But if this were true, then why have we not simply dismissed the skeptic long ago? Why do we feel threatened by the skeptic's denial of knowledge? Upon reading the *Meditations*, many are tempted by the thought that Descartes's meditator must know that he is not dreaming if he is to know that he is sitting by the fire. We find this reasoning gripping and we recognize that we share the meditator's epistemic shortcomings. An adequate theory of knowledge must shed light on the source of this temptation. So even if Gardiner is right to insist on the practical value of fallibilism over skepticism, this would provide no insight into how we (and the skeptic) are led astray.

To resolve this issue, I distinguish two perspectives: our practical, everyday outlook and the attitude to which we are led by philosophical reflection. We find this distinction in the work of Descartes, Hume, J. L. Austin, Bernard Williams, Marie McGinn, and others (as I argue in the book). Borrowing from Bernard Williams, I label these 'the project of practical inquiry' and 'the project of pure inquiry'. I argue that skeptical standards are inappropriate in contexts of 'practical inquiry', for they would frustrate our communal epistemic practices, but skeptical

standards may be legitimate in contexts of 'pure inquiry', when we bracket the practical considerations and goals of daily life. Everybody wins (sort of).

In her comments, Gardiner gives an interesting account of where the skeptic goes wrong. She accepts the distinction between 'practical inquiry' and 'pure inquiry', but she argues that contexts of pure inquiry aim at 'rational certainty', not knowledge. According to Gardiner, we have the concept of knowledge to govern the everyday circumstances of practical inquiry, whereas the concept of certainty is used for pure inquiry. The skeptic's error (and presumably my own) is to take pure inquiry to aim at knowledge, not certainty. If the concept of knowledge is designed for everyday situations, we should not expect the circumstances of pure inquiry to leave their mark on this concept.

I think our practice of epistemic evaluation *could* have turned out this way. Our conceptual repertoire could have been such that our knowledge concept did not generate pressure toward skepticism. Also, I agree with Gardiner that we rarely have rational certainty; thus, the skeptic is right to doubt our ability to achieve this cognitive goal. But here's the rub: almost nobody worries about skepticism as a threat to rational certainty. If that's all there were to skeptical arguments, our reply would be: 'Who cares?'. Few of us think we have much rational certainty, so we wouldn't feel challenged if we were denied it. What makes skepticism worrying and interesting is that it purports to show that we have *no knowledge*. In its strongest form, the skeptic argues that even our most foundational beliefs are completely *unjustified*. If Gardiner were right, why haven't we just dismissed the skeptic's challenge? It seems there is some feature of *our* concept of knowledge that generates pressure toward skepticism. Yet, the source of this temptation is a mystery on Gardiner's view. In contrast, my own account explains why the skeptic's argument is plausible without it being a threat to ordinary knowledge.

### **1.3. Pluralism and rival functions**

A common objection to my view is: why think the concept of knowledge is geared primarily toward identifying reliable informants? (DiPaolo also raises this objection in his commentary.) Plausibly, there are *many* uses of this concept, for example, to signal the proper end of inquiry, to distinguish between blameworthy and blameless behavior, to provide assurance to others, to encourage good testimony, and so forth. Is there any reason to favor one function over others?

I am a pluralist about knowledge functions; I do not claim our concept of knowledge has just *one* function. We speak of knowing for a variety of purposes. However, I also maintain that the *primary* function of the concept of knowledge is to identify reliable informants. The informant-flagging function is 'primary' in two senses. First, it can best account for a range of intuitive judgments about cases, linguistic data, and plausible epistemological theories. As Craig (1990) argues, the informant-flagging hypothesis explains several features of our concept of knowledge that have been identified by Nozick's tracking account, Goldman's causal theory, reliabilism, and internalism about justification (see also Kusch and McKenna 2020). Second, the informant-flagging function is *explanatorily prior* to other functions because we can use this hypothesis to explain why the concept of knowledge serves a variety of other functions, whereas those alternative functions do not better explain why the concept of knowledge is used to identify good informants.

Gardiner is not persuaded by this argument. She raises three objections to the idea that the informant-flagging function is explanatorily prior to other putative functions. First, she takes issue with my claim that the informant-flagging function can explain the 'inquiry stopping function' but not vice versa. Here's a relevant passage from my book:

the way to reasonably terminate inquiry is by identifying a sufficiently reliable informant. . . . If our informant knows, then there is no need to investigate further. Attributing knowledge to someone is a way of expressing the attitude that someone's epistemic position (with respect to a given proposition) is good enough to stop further inquiry. That's precisely what makes such an informant reliable *enough*. (Hannon 2019, 109)

In response, Gardiner argues that 'testimony is not *the* way to reasonably terminate inquiry; it is one of many'. I completely agree. But I nowhere claim that *testimony* is the only way to reasonably terminate inquiry. Instead, I claim that *finding a sufficiently reliable informant* suffices to terminate inquiry. We can do this in the absence of testimony; for example, I might become a reliable informant myself by coming to learn something via my perceptual faculties, powers of reasoning, or memory. While I think testimony is perhaps our *most common* source of information, I do not claim it is our *only* source.

Second, Gardiner maintains that the 'inquiry-stopping' function is more fundamental than the 'informant-flagging' function on the following grounds: 'good informants are typically those who have already reasonably ended inquiry'. The problem with this explanation, however, is that

it leaves unanswered how much justification is needed to reasonably end inquiry. It is here that the informant-flagging view has an advantage. In my book, I spend several chapters explaining how much justification it takes to qualify as a reliable informant (see chapters 2–4). This provides a principled answer to when inquiry has gone on long enough. To reverse the explanatory direction would leave unanswered when it is reasonable to stop inquiry.

Third, Gardiner argues that the informant-flagging function is not more explanatorily fundamental than the ‘threshold-marking function’. In my book, I argue that the concept of knowledge is used to mark the threshold that normally corresponds to appropriate assertion and practical reasoning. Gardiner says we should reverse the explanatory direction: ‘Appropriate action, including actions underwriting practices of testimony and ‘certifying’ one another, seem more foundational than certifying good informants’. Gardiner goes on to suggest that the ‘assurance-giving’ function is *also* more explanatorily fundamental than the ‘informant-flagging’ function.

While Gardiner and I agree that our knowledge concept plays numerous important functions, we disagree about which role is primary. I argue that the informant-flagging function can be used to explain various other functions, but Gardiner seems to think that *several* other functions are more foundational than the informant-flagging function. I find this puzzling. Are these other purported functions *all equally foundational*? Is the ‘inquiry-stopping’ function just as primary as the ‘assurance’ function, which is also just as primary as the ‘threshold-marking’ function? Gardiner does not say. Instead, she argues that my view exhibits ‘over-compatibilism’ because I am too eager to show how these various functions can all be explained by the more primary need to identify good informants. Yet Gardiner also seems to ‘absorb rival claims’ about the function of knowledge ascriptions into her own view. Instead of arguing that the informant-flagging function is explanatorily prior to these other functions, however, she claims those other functions are more fundamental than the need to identify reliable informants. While this reverses the explanatory direction, it does not make her proposal any less guilty of ‘over-compatibilism’.

## 2. Reply to DiPaolo

Along with Gardiner, DiPaolo claims that I have misidentified that primary function of the concept of knowledge. This concept is not primarily for identifying reliable informants, he argues, but rather for *identifying those who have eliminated all relevant alternatives*. DiPaolo’s argument has three



steps. First, he provides a series of hypothetical cases in which people ‘track knowledge for purposes unrelated to [the] need to identify reliable informants’. These examples are designed to discredit my ‘informant-flagging’ hypothesis. Second, DiPaolo argues in favor of an alternative functional hypothesis that allegedly has several advantages over my proposal. Finally, DiPaolo diagnoses why function-first epistemologists like myself were mistakenly led to endorse the informant-flagging hypothesis.

I will argue that DiPaolo’s argument goes wrong in two ways. First, his examples describe situations in which individuals care about *tracking knowledge*, which begs crucial questions against my account; in particular, these examples make use of the concept at issue and thereby presuppose what we are trying to explain. Second, DiPaolo presents the ‘relevant alternatives’ hypothesis as a competitor to my view, but this proposal is an important element of my view.

### **2.1. DiPaolo’s hypothetical cases**

DiPaolo provides a number of ordinary examples in which individuals use the concept of knowledge for purposes that seem to have nothing to do with identifying reliable informants. Here is one example:

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.’

In this case, Trina is not primarily concerned with identifying a reliable informant. Instead she is tracking George’s knowledge to keep him safe. She wants to know ‘who needs to be informed and who doesn’t’.

DiPaolo is surely right about this, but he misunderstands my account in two ways. First, it is perfectly compatible with my view that we can use the concept of knowledge for diverse purposes once we have it. As I say in the book,

Suppose the concept of knowledge is for flagging good informants. Even if this were correct, we might also ascribe knowledge for a variety of other purposes. For instance, I might try to comfort a friend who is experiencing hardship by saying, ‘I know you’ll get through this.’ I might say this even though I realize I do *not* know this. My use of ‘knows’ in this case is intended to provide reassurance, not to identify a reliable informant. (Hannon 2019, 23)

The aim of my book is not to explain the many ways in which we might use the concept of knowledge once we have it. (Analogy: someone

theorizing about why hammers were invented needn't be interested in the myriad functions this tool can serve once it exists, such as using it as a paperweight.) Instead, my book aims to explain why we have the concept of knowledge *at all*, by reflecting on the practical necessity of this concept, tracing it to its functional origins, and revealing what it does for us. Following Craig (1990), I model the emergence of our idea of knowledge by imagining how it serves the interests of those in a hypothetical situation who lack this concept. The model starts with some basic assumptions about the human need for true information, our mutual dependence on others, and so forth, and then attempts to derive needs that are less obvious (e.g. the need for a concept of knowledge) from more basic and uncontroversial needs. The aim, broadly speaking, is to take an abstract idea whose point eludes us, such as knowledge, and try to explain why we come to have this term by reconstructing the practical problems to which this idea provides a solution (see Queloz 2021). In contrast, DiPaolo's examples illustrate what we might *do* with our epistemic evaluations once this conceptual practice is already up and running.<sup>2</sup>

This takes me to DiPaolo's second error. He provides a number of hypothetical cases that are supposed to cast doubt on the following hypothesis: the *primary* purpose of our knowledge concept is identifying reliable informants. To undermine this hypothesis, DiPaolo argues that we regularly 'track knowledge for purposes unrelated to our need to identify reliable informants'. Yet his examples cannot achieve their purported goal. In each case, DiPaolo describes a situation in which some individual cares about *tracking knowledge* for various purposes, but I nowhere argue that we track knowledge to identify reliable informants. That way of putting things gets my account backwards. I argue that we identify others *as knowers* in order to *track their reliability* (and my account of reliability does not invoke knowledge). To say we are 'tracking their knowledge', as DiPaolo does, is to make use of the very concept at issue and thereby presupposes what we are trying to explain, namely the *communal acceptance* of this concept.

---

<sup>2</sup>This relates to another misunderstanding of my view. DiPaolo says, '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' (emphasis mine). But that's not exactly my view. Here's what I say in the book: 'We may contrast the point of *the concept of knowledge* with the point of *specific attributions of knowledge* (or the speech act of ascribing knowledge). The set of considerations for one might be different from the other, since the purpose of attributing knowledge in a certain situation might differ from the point of the concept in question. It is not a necessary condition that each proper attribution of knowledge serves its primary (or characteristic) purpose' (Hannon 2019, 106).

In TIGER, for example, Trina thinks to herself, ‘George doesn’t *know* there’s a tiger out there. . . I need to let him *know*.’ She is able to track George’s knowledge only because she *already has* this concept in her repertoire. But why did she come to think in these terms? And why is this concept so fundamental to human life and thought? DiPaolo’s examples do not shed light on these questions because they involve protagonists who are *engaged users* of the relevant concept. As such, DiPaolo cannot explain why our highly abstract idea of knowledge gained a hold in the first place. His examples focus on why specific individuals *apply* the concept the way they did. In contrast, my book focuses on the more fundamental question of why individuals—like Trina—participate in a practice of letting their thoughts, attitudes, and actions be shaped and guided by the relevant concept at all.

With this misunderstanding now clarified, we can see that DiPaolo’s examples *support* my preferred ‘informant-flagging’ hypothesis. Consider the TIGER case. 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 . . . keeping him safe.’ I agree with DiPaolo about this point. But *why* does George fail to qualify as a *knower*? Why does Trina think to herself, ‘George doesn’t know there’s a tiger out there?’ Here’s an explanation: George is not *a reliable source of information*—or, if you prefer, not a reliable believer—on the issue.<sup>3</sup> I won’t rehash all the criteria for being a reliable informant outlined in my book (see pp. 37–46), but I will say that George fails to satisfy even the most basic criterion: someone who has a true belief on the relevant matter. If I am inquiring as to whether *p*, a reliable informant is someone who, at minimum, can tell me whether *p*. This will typically involve the informant holding at least a true belief as to whether *p*. But if George really believed there was a tiger nearby, he wouldn’t be walking into the field! George is not even *likely to be right* about the relevant issue, which is another marker for being a reliable informant (see Hannon 2019, 38 and 44; Craig 1990, 91). Thus, the concept of knowledge

---

<sup>3</sup>DiPaolo grants this point. He writes, ‘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 he says this only demonstrates that knowledge is *positively correlated* with being a reliable informant (and not that identifying reliable informants is the primary purpose of our knowledge concept). There are two problems with this claim. First, DiPaolo’s argument relies on problematic examples, as I suggested above. Second, the whole point of my book is to explain why the connection between knowledge and reliable informants is not a mere correlation.

is playing exactly the role I claim. George does not qualify as a knower because he fails to meet the criteria to count as a reliable informant.

The same can be said for DiPaolo's other examples. In *DIRECTIONS*, Jeff tracks Tara's knowledge in order to decide whether Tara needs direction to the party, not to determine whether she's a reliable informant. But one cannot tell *whether* Tara needs directions without first determining whether she *is* reliably informed about the issue. If we conclude that she's not a sufficiently reliable informant, we will not say that she knows. This is exactly what my account predicts. To say 'Tara knows where to go' is to identify her as a sufficiently reliable source of information. Moreover, we may identify her as a reliable informant *even if nobody actually relies on her for the relevant information*. Individuals can qualify as reliable informants without informing anyone.<sup>4</sup> For example, a mechanic is a reliable informant about how to fix your car, even if you never speak to him. DiPaolo doesn't consider this point, which explains why he provides numerous examples of individuals who track knowledge for purposes unrelated to the need to identify reliable informants. As I've argued, the fact that we track knowledge for various purposes is no threat to my account. We still have a fundamental need to distinguish between individuals who are sufficiently reliable sources of information and those who are not.

## 2.2. *The relevant alternatives hypothesis*

I have defended the 'informant-flagging' hypothesis from DiPaolo's objections. If this defense succeeds, it undercuts DiPaolo's motivation to look for an alternative hypothesis about the primary function of knowledge. But I would like to consider his alternative proposal anyway, since there is far more agreement between us on this issue than he realizes. In fact, his preferred hypothesis about the primary function of knowledge is roughly the same proposal made in my book. In this section, I will briefly outline DiPaolo's argument and then explain why there's no real disagreement between us here.

What is the concept of knowledge for, if not identifying reliable informants? According to DiPaolo, 'The primary purpose of the concept of

---

<sup>4</sup>Here's what I wrote in my book: 'A knower must therefore meet a sufficiently high-quality epistemic position such that a wide range of potential inquirers could in principle rely on this person's information, even if nobody actually *does* ever seek such an informant. These facts explain why we might want to say that someone knows whether *p* even though, as it happens, that person does not actually function as an informant for anybody.' (Hannon 2019, 43)

knowledge is to identify those who have eliminated all relevant alternatives'. This hypothesis is rooted in some plausible theories of knowledge in epistemology (e.g. Dretske 1970; Stine 1976; Heller 1999). The rationale for a relevant alternatives approach is fairly uncontroversial: there seems to be an important connection between our discriminative capacities and knowledge, coupled with a desire to avoid skepticism (see Goldman 1976). Also, the relevant alternatives framework can throw light on the relationship between knowledge and action: it is plausible that one is in a good enough epistemic position to act on  $p$  if one knows that  $p$ , yet actionable information does not generally require the elimination of *all* possible alternatives.

My book defends a relevant alternatives theory of knowledge. (DiPaolo says this view can be 'derived' from my argument, but I will try to explain why it *just is* my argument.) The central hypothesis of my book is that the primary function of the concept of knowledge is 'to identify informants who are reliable enough to appropriately serve as sources of actionable information for members of our community' (Hannon 2019, 13). Who counts as reliable enough? I use the relevant alternatives framework to answer this question:

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. (Hannon 2019, 68)

I call this *the reliable informant standard for knowledge*. According to this view, who counts as a reliable informant depends on the ability to eliminate relevant alternatives. If one eliminates all the not- $p$  possibilities that count as relevant alternatives to members of one's epistemic community, then one knows that  $p$ .<sup>5</sup> So if, as DiPaolo argues, the primary purpose of the concept of knowledge is to identify those who have eliminated all relevant alternatives, and if, as I argue, those who have eliminated all relevant alternatives are precisely those who qualify as reliable informants, then it follows that the primary purpose of the concept of knowledge is to identify reliable informants. Why, then, does DiPaolo take the relevant alternatives hypothesis to differ from the informant-flagging hypothesis?

---

<sup>5</sup>As a function-first epistemologist, I also try to explain *why* the concept of knowledge would function this way. One of the major obstacles to the theory of relevant alternatives has been that most of these theories do not provide an explanation for what makes a particular alternative relevant. My view tries to give a plausible practical story about why this would be so.

As discussed in the previous section, DiPaolo relies on the idea that we often track knowledge for reasons unrelated to identifying reliable informants. Here's another example from his paper:

You want your daughter to gather berries, but you need to know whether she can distinguish poisonous from safe berries . . . So, you show her a poisonous berry and ask: is this poisonous or safe? . . . 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.

As I've argued, however, these cases do not threaten my account. I do not claim that we *keep track of knowledge* in order to figure out who is a reliable informant (as DiPaolo suggests), but rather that we *have a concept of knowledge* to keep track of who is a sufficiently reliable informant. So, we cannot rely on these cases to mark an important difference between the informant-flagging hypothesis and the relevant alternatives hypothesis. In each of DiPaolo's examples, the agent qualifies as a knower *because* they are a reliable informant, which is partly understood in terms of ruling out the relevant alternatives.

While my account can avoid DiPaolo's objections, his comments do point to a shortcoming of my book: I focus primarily on the situation of an *information-seeking inquirer* at the expense of other important epistemic situations. For example, we often rely on others to *exercise their knowledge* or to *provide us with instruction*, not just information. Catherine Elgin (2021) provides some illuminating examples in this regard. When I take my car to the auto mechanic, I want more than information: I want him to fix my car. Likewise, I may go to the dentist to get my tooth fixed without looking for information about how to fix my tooth. Our epistemic interdependence runs far deeper than our relying on others as sources of information. Thus, I agree with DiPaolo that our epistemic life is often more complicated than the relatively simple situation of an inquirer seeking an informant. I should have made this point clearer in my book.<sup>6</sup>

In addition to these examples of epistemic interdependence, DiPaolo says I ignore the *individualistic* aspects of our epistemic situation. On this point, I am less convinced. He writes,

---

<sup>6</sup>That said, I think my own 'informant-flagging' view can accommodate this insight. In the dentist example, I need to rely on the dentist's knowledge because it enables her to perform the action that I cannot perform for myself. Thus, the dentist must first *be* someone who possesses the relevant information in order to carry out the task. To identify her as a knower, then, is to say she has the relevant information to accomplish this task. Even though I am not seeking her information, I am still searching for an agent who is a sufficiently reliable source of that information, since this information is partly what enables her to perform the relevant action (see Hannon 2021, 120).

Hannon's implementation of the function-first approach dismisses important individualistic dimensions of knowledge. So much of what we count as knowledge simply *hits us*, without 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.

I agree that we acquire much of our knowledge passively rather than through active inquiry. In the book, I suggest that merely looking in the right direction at the right time may suffice to acquire knowledge (see Hannon 2019, 27). Why does DiPaolo think I cannot accommodate this point? It is because he interprets me as arguing that the only (or primary) way to achieve knowledge is through active inquiry. But this is not my view. While I maintain that the concept of knowledge is primarily geared toward the activity of inquiry, I do not claim we can only acquire knowledge via active inquiry. Each of us knows a lot via perception, inference, memory and other individualistic sources of knowledge. These faculties may put us in a position to reliably inform others. (When I lazily watch the news from my sofa, I may become a reliable source of information about what they are reporting.) Although my book foregrounds the social dimensions of knowledge over aspects of individualistic epistemology, it does not entail implausible conclusions about the individualistic dimensions of knowledge.

### **2.3. Traditional epistemology and function-first epistemology**

Toward the end of his piece, DiPaolo says that my view 'unnecessarily abandons key insights uncovered within [the] research program [of traditional epistemology]'. He also finds it 'surprising that Hannon and Craig argue that the *primary* purpose of the concept of knowledge is to identify reliable informants' given that 'no one working in the traditional paradigm has ever argued that knowledge should be analyzed and explained in terms of being a reliable informant'. I'll conclude by addressing these two claims.

First, I do not see myself as abandoning key insights from traditional epistemology. In fact, both Craig and I attempt to vindicate several lessons from the conceptual analysis of knowledge. In *Knowledge and the State of Nature*, Craig goes to great lengths to show that his account explains several features of knowledge that have been identified by those working within the traditional paradigm, namely:

- i. uses of 'knowledge' without belief (Radford) (Craig 1990, 15–6)

- ii. the role of counterfactuals (Nozick, Dretske) (Craig 1990, Ch. III)
- iii. the role of causal relations (Goldman) (Craig 1990, Ch. IV)
- iv. the role of methods (reliabilism) (Craig 1990, Ch. IV)
- v. the role of justifying reasons (internalism about justification) (Craig 1990, Ch. VIII)
- vi. that all analyses have counterexamples (Gettier) (Craig 1990, Ch. VI)
- vii. the contextual variation in standards (Unger) (Craig 1990, Ch. XII).<sup>7</sup>

As Kusch and McKenna point out, ‘These theories are often seen as excluding one another, but Craig thinks that his model can partially vindicate all of them’ (2020, 1060). In *What’s the Point of Knowledge?*, I also draw on insights from the Gettier problem (Gettier 1963), the justification condition (Chisholm 1977), relevant alternatives theory (Stine 1976), the lottery puzzle (Hawthorne 2004), the epistemology of understanding (Grimm 2006), the epistemology of testimony (Lackey 2007), and philosophical skepticism (Stroud 1984). Moreover, I argue that function-first epistemology is perfectly compatible with the goal of traditional conceptual analysis (Hannon 2019, 19). Although I abandon some aspects of the traditional method, it is misleading to say that ‘traditional insights are strikingly absent from [my] account of knowledge’.

Second, I can explain why those working within the traditional paradigm have not explicated knowledge in terms of being a reliable informant, namely: they were not investigating the *point* or *purposes* of the concept of knowledge. Traditional conceptual analysis aims to enumerate the necessary and sufficient conditions of the relevant concept. The guiding question is ‘What are the conditions for something to fall under the concept of knowledge?’. Those working within this tradition do not ask pragmatic questions about the *function* of this concept; they do not think of our concepts in terms of their point. So, it is unsurprising why these theorists failed to consider the role of knowledge in identifying reliable informants, given the paradigm within which they were operating. What *is* surprising, however, is why theorists failed to ask such pragmatic questions in the first place. Panayot Butchvarov considers this the most characteristic symptom of the inadequacy of traditional accounts of knowledge. He says that ‘an adequate account of the concept of knowledge must display its essential place in the conceptual framework

---

<sup>7</sup>I borrow this list from Kusch and McKenna (2020, 1060).



through which we would most perspicuously understand ourselves, our life, and the world in which we live' (1970, 25–26). Yet traditional accounts of knowledge, even if correct, fail to explain why knowledge should have the features it has (e.g. truth, justification, and belief), how these are related to one another, and why these must be present for knowledge to be present. There is something deeply unsatisfying with an account that leaves these questions unanswered. *What's the Point of Knowledge?* attempts to answer these questions.

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## Works Cited

- Butchvarov, P. 1970. *The Concept of Knowledge*. Evanston, IL: Northwestern University Press.
- Chisholm, R. 1977. *Theory of Knowledge*. Hoboken, NJ: Prentice Hall.
- Craig, E. 1990. *Knowledge and the State of Nature*. Oxford: Oxford University Press.
- Dretske, F. 1970. "Epistemic Operators." *Journal of Philosophy* 67 (24): 1007–1023.
- Elgin, C. 2021. "The Function of Knowledge." *Analysis* 81 (1): 100–107.
- Gettier, E. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–123.
- Goldman, A. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73: 771–791.
- Grimm, S. 2006. "Is Understanding a Species of Knowledge?" *British Journal for the Philosophy of Science* 57 (3): 515–535.
- Hannon, M. 2019. "Skepticism: Impractical, Therefore Implausible." *Philosophical Issues: A Supplement to Noûs* 29 (1): 143–158.
- Hannon, M. 2021. "Replies to Henderson, Elgin and Lawlor." *Analysis* 81 (1): 114–129.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Heller, M. 1999. "Relevant Alternatives and Closure." *Australasian Journal of Philosophy* 77 (2): 196–208.
- Kusch, M., and R. McKenna. 2020. "The Genealogical Method in Epistemology." *Synthese* 197: 1057–1076.
- Lackey, J. 2007. "Why We Don't Deserve Credit for Everything That We Know?" *Synthese* 158 (3): 345–361.
- Lewis, D. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549–567.
- Queloz, M. 2021. *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering*. Oxford: Oxford University Press.
- Stine, G. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29 (4): 249–261.
- Stroud, B. 1984. *The Significance of Philosophical Skepticism*. Oxford: Clarendon Press.