

# BOOK SYMPOSIUM

---

## *What's the Point of Knowledge? A Function-First Epistemology*

BY MICHAEL HANNON

Oxford University Press, 2019. ix + 288 pp. £47.99

### Summary

MICHAEL HANNON

*What's the Point of Knowledge?* has an ambitious goal: to reorient epistemological enquiry by suggesting a new starting point.

Traditionally, epistemologists have sought to analyse knowledge by using intuitions about cases to determine the necessary and sufficient conditions for knowing. This way of doing philosophy dates back at least to Socrates and it reached its climax in the late 20th-century. It is also a method that Wittgenstein believed 'shackled philosophical investigations' (1958: 19) because it relies on mistaken assumptions about the nature of language.

As most philosophers know, this research programme has so far failed on its own terms: we have no satisfactory analysis of knowledge. This is not to say, as uncharitable critics do, that attempts to analyse knowledge (especially after Gettier) were a waste of time. These enquiries did bear fruit. Nevertheless, the traditional approach has been largely overshadowed by newer, more productive ways of doing epistemology, such as knowledge-first epistemology, formal epistemology, experimental epistemology and naturalistic epistemology.

*What's the Point of Knowledge?* is a further attempt to release us from the shackles of an outdated philosophical ambition. It aims to develop a rigorous way of doing epistemology outside the paradigm of the traditional view. Following in the footsteps of Craig (1990), the book's central thesis is that pivotal issues in epistemology can be clarified by reflecting on the role of epistemic evaluation in human life. I call this approach *function-first epistemology*.

When theorists recommend new starting points, what sometimes happens is a changing of all the questions. We come to see that the traditional problems aren't really problems but rather are the result of misguided ways of thinking. For example, some pragmatists and early 20th-century positivists came to regard philosophical problems as 'pseudo-problems' that cannot be solved but only *dissolved*. Attempts were then made to replace these misguided questions with new points of interest.

Although *What's the Point of Knowledge?* recommends a new method of enquiry for epistemologists, it does not involve a changing of traditional epistemological concerns – it is still a starting point *for epistemology*. The book aims to make headway on easily recognized problems such as scepticism, epistemic relativism, the evidential standard for knowledge and other topics that are found in almost every introductory textbook on epistemology.

The book has three main goals.

First, I outline the method of function-first epistemology and defend it from objections (Ch. 1). A function-first epistemologist claims that we will better understand our epistemic concepts, norms and practices by investigating what they are *for*. Proponents of this method will ask questions like: Why do we speak and think in terms of ‘knowing’, ‘understanding’ and ‘rationality’? What epistemological norms and standards would best facilitate human survival, cooperation and flourishing? What role does epistemic evaluation play in science, philosophy or daily life? I also compare this method with some alternative approaches, such as conceptual analysis, epistemological naturalism and knowledge-first epistemology.

Second, I articulate and defend a specific hypothesis about the purpose of knowledge (Ch. 2). Drawing on [Craig \(1990\)](#), I argue that we speak of knowing primarily to *identify reliable informants* to members of our epistemic community. To some, this sounds utterly obvious; to others, it seems obviously false. I try to show that this simple idea has profound implications for epistemology. One such implication is that epistemologists should be pragmatists and almost all contemporary accounts of knowledge are mistaken about the metaphysical status of knowing.

Third, I use this hypothesis to make progress on a variety of issues in epistemology (Chs. 3–9). For example, I argue that function-first epistemology casts new light on the nature and value of knowledge, the power and limits of scepticism, the relationship between knowledge, assertion and action, the semantics of knowledge claims, the Gettier problem, fallibilism, epistemic relativism, and the nature and value of human understanding. To briefly highlight one of these issues: I claim that radical scepticism is implausible because it undermines the practical requirements that explain why we have a concept of knowledge in the first place.

In short, the book proposes a *method* of epistemological enquiry, a *hypothesis* about the role of knowledge and then *applies* this method and hypothesis to familiar areas of the epistemological landscape.

These three features of the book are not inseparable. You might endorse the function-first method while rejecting my hypothesis about the function of the concept of knowledge; or you might reject the methodological approach and still think there is an important conceptual connection between knowing and being a reliable informant; or you might accept both the method and the hypothesis but think they’ve been misapplied to the epistemological issues considered in the book. It’s worth emphasizing, however, that the usefulness

of the method would not be compromised by a mistaken hypothesis about the point of knowledge. We don't doubt the methods of science just because particular hypotheses go astray.

I wrote this book in the hope of showing that some pivotal epistemological issues can be resolved by taking a function-first approach. Whether it succeeds in resolving these thorny philosophical problems or not, I hope it convinces you that reflecting on the point of epistemic evaluation is a useful starting point for epistemology. We gain new insights by thinking about the point of our epistemic concepts, norms and practices, which illustrates the unappreciated role this method can play in philosophy.

*The University of Nottingham*  
*Nottingham NG7 2RD*  
 UK  
*michael.hannon@nottingham.ac.uk*

### References

- Craig, E. 1990. *Knowledge and the State of Nature*. Oxford: Clarendon Press.  
 Wittgenstein, L. 1958. *The Blue and Brown Books*. Oxford: Blackwell.

## What's the Point of Knowledge?

BY DAVID HENDERSON

Michael Hannon advocates an epistemological methodology – tracing its roots, articulating refinements, distinguishing it from alternative methodologies and giving reasons for preferring it to the alternatives. He also advances an account of knowledge as a compelling application of this methodology. As reflected in his title, both projects are pivotal to the work and intimately related. In its general outlines, I judge that that case for the method should be taken to heart – although details could stand for further attention – and that its application does advance our epistemological understanding.

### 1. *Hannon's function-first epistemological method*

In Chapter 1, 'Methodologies in Epistemology', Hannon outlines the method of function-first epistemology:

This methodology involves three broad steps: we start with a *prima facie* plausible hypothesis about the role of some epistemic concept (norm,

- Kratzer, A. 2002. Facts: particulars or Information Units? *Linguistics and Philosophy* 25: 655–70.
- Lawlor, K. 2013. *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford: Oxford University Press.
- Travis, C. 2008. *Occasion-Sensitivity: Selected Essays*. Oxford: Oxford University Press.

## Replies to Henderson, Elgin and Lawlor

MICHAEL HANNON

I acquired many intellectual debts while writing *What's the Point of Knowledge?*, but I am especially indebted to my three symposiasts. David Henderson's work helped me to appreciate the value of thinking about the point of epistemic evaluation; Catherine Elgin's writings prompted me to investigate the purpose of the concept of understanding; and Krista Lawlor's 2013 book revealed important connections between three of my primary epistemological interests: the role of epistemic evaluation, the semantics of knowledge claims and the work of J.L. Austin. It is therefore an honour to have such personally influential (and highly esteemed) scholars engage with my work. Their thoughtful, generous and philosophically rich comments have provided yet another opportunity to clarify my thinking and develop some ideas further.

*What's the Point of Knowledge?* is guided by one overarching idea: we can answer many interesting and difficult questions in epistemology by reflecting on the role of epistemic evaluation in human life. To make good on this claim, I pursue three interrelated goals. First, I outline the method of function-first epistemology. Second, I advance an account of knowledge as an application of this method. Third, I use this account of knowledge to make progress on a number of issues in epistemology.

Interestingly, my three symposiasts have each chosen to focus their comments on a different one of these three goals. While Henderson expresses sympathy with my account of knowledge and its application to epistemological issues, his critical remarks are primarily directed at my characterization of the method of function-first epistemology. By contrast, Elgin doubts the plausibility of my account of knowledge, but she does not question the general method or its application. In further contrast, Lawlor grants both the method and my main hypothesis about the point of knowledge, but she raises doubts about how I apply it to two epistemological issues: the standard of evidence required for knowledge and the semantics of knowledge claims. As these concerns reflect the logical structure of the book, I will reply to their comments in this order.

### 1. Replies to Henderson

Henderson provides a lucid snapshot of the overall project in *What's the Point of Knowledge?*, in addition to raising two main issues. First, he argues that my own characterization of function-first epistemology is 'too linear'. As a result, he says that I misleadingly suggest 'a Popperian cartoon of function-first epistemology' which 'goes wrong in several ways'. I take this to be his main critical remark about the book. Second, he focuses on the connection between the standards for knowing and the stakes of enquirers in an epistemic community, and he asks how straightforward this connection is.

To the charge of making function-first epistemology appear too linear, I plead (partly) guilty. In the book, I characterize the method as involving three broad steps: we start with a prima facie plausible hypothesis about the role of the concept of knowledge; then we try to determine what a concept having this role must be like; finally, we examine the extent to which the concept we have described matches our intuitive judgements (4–5 and 13–14; see also [Craig 1990](#): 2–3). This makes the first step seem like a discrete starting place: we form a hypothesis, derive predictions and then test against judgements. As Henderson rightly points out, it is unclear how we could settle on *any* prima facie *plausible* hypothesis (i.e. the first step) without drawing on a lot of background information to narrow the possibility space. We want to avoid mere conjecture, so our hunches about the purpose of knowledge must be informed by 'background information about the communities and projects regulated'.

Henderson is exactly right on this point. I do briefly address this issue in the book, where I say:

In order for this hypothesis to be plausible, it must be compatible with certain facts about human life, such as facts about our physical environment, our social organization, our cognitive capacities, and the basic aims and interests humans typically have. These facts about humans and their circumstances will then give rise to a certain conceptual need that is supposed to be satisfied by the purpose described by our hypothesis. (13–14)

Admittedly, this is quite vague. I probably should have said more about precisely *how* this is done. In Chapter 2, I attempt to *show* that we can arrive at a plausible hypothesis by appealing to fairly uncontroversial assumptions about the need for true beliefs, our epistemic dependence on others and other factors. By acknowledging such antecedent constraints, we move away from the linear model and towards the more 'holistic' and 'extended' epistemological reflection recommended by Henderson.

Henderson suggests that I drop the label 'function-first epistemology' (which implies a linear model) and replace it with 'purpose *centred*' or 'function *weighted*' epistemology. For whatever it's worth, I do regret the label 'function-first'. I had originally wanted to subtitle the book 'a purpose driven epistemology', but was convinced not to because it sounded too much like

Rick Warren's international bestseller *The Purpose Driven Life*. In hindsight, this might not have been a good reason to switch labels.

But this isn't Henderson's main complaint. As he says, it is a 'small refinement' to suggest that my starting point is actually 'one moment in an extended course of reflection'. His main complaint is that my characterization of function-first epistemology leads me to draw artificial contrasts with allegedly 'alternative' epistemological methods.

In particular, I contrast my own approach with attempts to 'reverse engineer' our epistemic evaluations (see [Dogramaci 2012](#)). The latter approach initially brackets speculation about what our epistemically evaluative practices are *for* and instead looks at what these practices actually *are*. For example, instead of starting with some hypothesis about the point of the concept of knowledge, we start by looking at how people actually use the word 'knows' and then draw conclusions about the function of the concept of knowledge from such facts about usage.

In the book, I claim that function-first epistemology and reverse engineering 'have different starting points' (23). The former method starts with a hypothesis about the role of some epistemically evaluative aspect of our language and then examines the extent to which it matches our everyday judgements; the latter method reverses the direction of investigation: we first look at some aspect of our actual evaluative practice and then try to infer the purpose of this practice. According to Henderson, my emphasis on these as 'alternative approaches' with 'different starting points' obscures an important truth: both methods are better understood as complementary components of a purpose-centred epistemology.

I'm happy to take this suggestion on board. Still, I want to insist on an important difference between my epistemological project and attempts to reverse engineer epistemic evaluation. Proponents of reverse engineering will take facts about our practice of epistemic evaluation as their primary input; thus, the method has a purely descriptive aim: to reveal the purpose of our actual epistemic practice (or some aspect of it). In contrast, a function-first epistemologist may investigate what epistemic norms, concepts or practices would best serve our interests and goals. This creates space for a normative project of evaluating how well or poorly our epistemic practices actually satisfy our needs and goals. We can think of this as conceptual *re-engineering* (25). What is unsatisfying about attempts to reverse engineer epistemic evaluation is 'they simply assume the propriety of our ordinary practices of epistemic evaluation; they do nothing to say why we should think these practices are worthy of our endorsement' (25).

Henderson is undoubtedly correct that the method of reverse engineering can be part of a broader normative project. But when a reverse engineer shifts from the purely descriptive project to the normative one, they thereby become a function-first epistemologist (or what Henderson would prefer to call a purpose-centred epistemologist).

There is another reason for my emphasis on putting function ‘first’. It is not because we should regard the first step in the methodological process as a discrete starting place. It is rather because the methodology itself provides a new ‘starting point’ for epistemological enquiry. Traditionally, epistemological reflection begins with intuitions about cases and results in analyses of concepts (some scholars claim to be interested in phenomena rather than concepts, but they still use conceptual analysis to somehow understand the world). By contrast, I suggest that epistemologists start by reflecting on the purpose of epistemic evaluation. We should then use our reflections about the role of the concept of knowledge as an adequacy constraint on theorizing about the nature and value of knowledge. Whatever our theory tells us, it had better be able to underwrite the role(s) that our concept of knowledge plays in epistemic evaluation. Elsewhere I call this the ‘functionalist turn’ in epistemology (2019: 145).

Henderson concludes his reflections on philosophical method by suggesting that function-first epistemology ‘could be understood with equal justice as finding a home within an ongoing abductive conceptual analysis’. This might be true, but I want to quibble with his characterization of conceptual analysis. He says that ‘a hallmark of conceptual analysis’ is that ‘our initial judgments about cases provide some evidence – but, on their face, some judgments seem to square poorly with others, and some will likely need to be rejected on reflection’. This doesn’t sound like conceptual analysis to me. Indeed, a hallmark of traditional conceptual analysis is that our intuitive judgements about cases bear so much evidential weight that a single counterexample will suffice to falsify even the most plausible account of knowledge. (Think of the attempts to show that there are fairly isolated cases of knowledge without belief or truth.) Thus, while I completely agree with Henderson that we shouldn’t put too much emphasis on such ‘counterexamples’, this idea seems incompatible with traditional conceptual analysis. By contrast, it is precisely what a function-first epistemology suggests (18 and 242).

Henderson then shifts from reflecting on metaepistemology to consider the main function of the concept of knowledge: to identify reliable informants to members of our epistemic community. About this proposal, he asks: how would the concept of knowledge develop so as to respond to diversity in a community? In any community of enquirers, you will find people with diverse interests, projects and stakes. How, then, do we settle on a standard of reliability that is ‘good enough’ for the many?

The answer, in brief, is that we need a minimum threshold of quality of epistemic position that is high enough to respect the ‘typical’ or ‘normal’ stakes of people in the epistemic community. Precisely how high these stakes are, and who counts as a member of the epistemic community, are difficult theoretical questions that the book attempts to answer. I tend to think of the epistemic community rather expansively:

an ‘epistemic community’ is much broader than, say, a group of people living in roughly the same area (i.e., a tribe, village, town, etc.). It is also broader than a group of people who share a language, a religion, and so forth. Although there are varieties of social, geographical, linguistic, and religious communities (among others), I think there is a sense in which we are all part of the same epistemic community. (93)

Henderson and I used to disagree on this point. His work (2009, 2011) has been informed by the idea that people do not inhabit one general epistemic community but rather many specific (ad hoc, situational) practical communities. While this idea might sound plausible, it runs into worries about widely varying epistemic standards that make it difficult, if not impossible, for enquirers to engage in the coordinated transmission of reliable information (see 2015). Happily, it seems that Henderson has, in more recent work, ‘come to think of engaged practical communities in a less situational way’, and in a way that is closer to my own understanding of a general epistemic community. In her comments for this symposium, Krista Lawlor presses me to say more about this ‘hard to pin down’ feature of my account. So, I will return to this issue when replying to Lawlor.

Henderson raises this issue about the epistemic standards for a community, in part, because it bears directly on the debates between invariantism and contextualism in epistemology. On the one hand, we need informants to meet some minimum standard of epistemic position to qualify as sufficiently reliable sources of information. This suggests an invariant epistemic standard. On the other hand, it seems that in cases where an enquirer confronts abnormally high stakes (e.g. a life or death situation), we are reluctant to volunteer others as having knowledge. Presumably, this is because the informant no longer qualifies as sufficiently reliable in the context. This suggests contextualism.

In older work, I agreed with Henderson that functional considerations provide support for contextualism over invariantism (Hannon 2013), but now I’m not so sure. This is not to say, as Henderson does, that I am now ‘attracted to a form of invariantism – insensitive invariantism’. While I think such a view is perfectly coherent with the functionalist story outlined in the book, I now think that putative facts about the function of knowledge claims will *not* instruct us as to whether contextualism or some form of invariantism is true. Instead, I tentatively endorse a form of epistemic pragmatism. Henderson says he is strongly attracted to this view, while Lawlor argues that I provide no good argument for it. I will discuss Lawlor’s objections shortly.

## 2. Replies to Elgin

Elgin goes along with the idea that we should guide epistemological enquiry by reflecting on the point of epistemic evaluation, but she is unconvinced that I have identified the correct function of the concept of knowledge. In the book, I argue



that we think and speak of knowing for many reasons, but identifying reliable informants is the most fundamental of them. I call this the ‘informant-flagging function’ (54). Elgin calls it the ‘information transfer view’. Against this view, Elgin says it is more plausible that we speak of knowing to certify that our enquiry has reached a point where we can reasonably stop. Elgin calls this ‘the certification view’. I call it the ‘inquiry-stopping function’ (109).

In addition, Elgin raises a complaint about my handling of the Gettier problem, and she registers two doubts about my overall project: first, she is ‘not convinced that knowledge has a single fundamental function’; second, even if the concept of knowledge *does* have a primary function, she doubts that this function ‘stems from epistemic interdependence’. Although Elgin does not argue for these two concerns, they loom significantly large to merit a brief response.

I’d like to start by keeping the big picture in view. Elgin and I agree that the concept of knowledge has more than one function. We also agree that an important function of this concept is to signal when to reasonably end enquiry. I also argue that the ‘inquiry-stopping’ function of knowledge is compatible with our need to identify reliable informants. Indeed, I say these two functions are ‘just different sides of the same coin’ (109). Moreover, Elgin does not deny that identifying reliable informants is an important aspect of our social-epistemic practices. She simply doubts whether this is the *most* fundamental function. Thus, I would not consider it a devastating blow to my account if it turned out that the concept of knowledge was geared most fundamentally towards signalling the point of legitimate enquiry closure.

But the certification (enquiry-stopping) view doesn’t provide a better account of the function of knowledge. I’ll start by briefly reviewing Elgin’s reasons to prefer the certification view to the information transfer (informant-flagging) view, and then I’ll explain why my own view has all the benefits of the certification view and more.

Elgin says the certification view has an edge over the information transfer view because it ‘readily accommodates cases of knowledge’ that my own view ‘has a harder time handling’. She lists a series of examples where an agent is looking for a knower, not because they are seeking a *reliable informant* but rather because they want a *reliable practitioner*. Her examples are:

- *An auto mechanic who is able to fix the ominous rattle in my car’s engine.* I am not looking for mere information; I want the car fixed. They could give me the information I need to fix the car myself, but I don’t care about that – I want the job done.
- *A dentist who is able to fix my tooth.* I want more than information; I want the tooth fixed. The dentist could not even tell me how to drill my own tooth. I need to rely on her knowledge.
- *A surgical team who is able to repair an aneurism, remove a spleen etc.* I want to get healthy and need to rely on a team of experts to help me.

Not only am I dependent on them, but each member of the team is also epistemically dependent on each other.

Elgin uses these cases to illustrate that epistemic interdependence runs far deeper than our depending on others as sources of information. Moreover, these kinds of cases are ‘widespread and vital’. Thus, a viable functional account of knowledge should be able to accommodate them.

Thankfully, my own ‘informant-flagging’ view *can* easily accommodate these cases. Indeed, Elgin’s own remarks indicate how.

In the dentist case, Elgin writes that I need to ‘rely on [the dentist’s] knowledge because it enables her to perform an action that I cannot perform for myself’. Thus, the dentist must first *be* someone who possesses the relevant information in order carry out the actions I desire. To identify her as a knower, then, is to say she has the information needed to satisfy my own goals. Even though I am not seeking this information myself, I am still searching for an agent who is a sufficiently reliable source of that information, since this information enables her to perform the relevant action.

The same can be said for Elgin’s other examples. In fact, Elgin does say it. About the surgical team case, she writes, ‘The team members have well defined roles, backed by appropriate credentials which they earned by mastering different bodies of knowledge’. To use the language of my book, the team members have mastered different bodies of *information*. As such, they become reliable informants on the relevant issue. This information, in large part, is what makes them expert practitioners. If they lacked this mastery, they would not qualify as knowers. It does not matter that I am not looking for them to transfer this information to me. What matters is that I am looking to rely on the expertise of people who have such information.

While my account can accommodate the cases outlined by Elgin, her comments do point to a shortcoming of my book: I focus primarily on the situation of an *information-seeking enquirer* at the expense of other, equally important epistemic situations. We often relying on others to *exercise their knowledge*; or we rely on others *for instruction* (not just information). Elgin provides several nice examples involving parents who show their child how to tie shoelaces; lab instructors who teach novices how to use a pipette etc. In all these cases, ‘information transfer may be part of the method, but it is not the goal’. Still, these cases are compatible with my general account because, as Elgin observes, ‘the instruction will succeed only if the instructors know that which they teach’. Thus, these examples provide no reason to prefer the certification view over the informant-flagging view.

Are we left at an impasse? I think not. It is at this point where the informant-flagging view has a significant edge over the certification view.

According to the certification view, ‘the function of the concept of knowledge is to set the point where enquiry can properly stop’ (Elgin, above). But where is that point? All the certification view says is that the concept of

knowledge is used to signal the appropriate end of enquiry, but this leaves unanswered precisely how much justification is needed for knowledge. It is here that the informant-flagging view has a significant advantage. As I say in the book, ‘the way to reasonably terminate enquiry is by identifying a sufficiently reliable informant’ (109).

Precisely how much justification it takes to qualify as a reliable informant is a question that the book spends a significant amount of time attempting to answer (see Chs. 3 and 4). The basic idea is that we attribute knowledge to someone as a way to indicate that the informant’s epistemic position (with respect to a given proposition) is good enough for us to stop further enquiry. This is why I claim that (i) the informant-flagging function and the enquiry-stopping function are ‘just two sides of the same coin’, but also (ii) that the informant-flagging function is more fundamental, since it provides a principled answer to when enquiry has gone on long enough. All the certification view says, in contrast, is that ‘inquiry stops where it does’ (Elgin, above).

Putting aside the function of knowledge, Elgin also takes issue with my handling of the Gettier problem. According to Elgin, I dismiss Gettier cases on two grounds. First, I do not aim to provide necessary and sufficient conditions for knowledge, so these ‘counterexamples’ do not discredit my view. Second, I insist that everyday knowledge situations do not give rise to Gettier cases, so they can be ignored. She is willing to grant the first point, but she (rightly) objects to the second. Elgin writes, ‘I am not convinced that Gettier cases do not occur in ordinary life. But even if they do not, the fact that they can be contrived using the everyday concept of knowledge shows that the concept . . . is inadequate’.

Now, I don’t recall saying that Gettier cases do not occur in ordinary life. In fact, I grant that we may encounter a Gettier scenario ‘in the course of our daily experience’ (76). However, I do attempt to *downplay* the importance of the Gettier problem. In particular, I argue that our standard situation with respect to knowledge is as *enquirers* who do not yet know whether *p* but want to; not as *examiners* who already know whether *p* but are trying to decide whether to apply the word ‘know’ to some unlucky Gettier victim (see Williams 1973: 146).<sup>1</sup> I use this idea to argue, following Mark Kaplan (1985), that solving the Gettier problem will do nothing to advance or clarify the proper conduct of enquiry – which is precisely the practice that our concept of knowledge serves to regulate (76). Thus, I conclude that little hangs on solving the Gettier problem.

However, I am aware that some epistemologists may not find this diagnosis compelling. Even if our practice of knowledge ascription works fairly well in everyday life, it may be inadequate because ‘the concept of knowledge seems internally inconsistent . . . Our criteria for knowledge are satisfied in circumstances where, we are convinced, the agent does not know’ (Elgin, above). This

1 Elgin seems to endorse this idea in the first paragraph of her commentary.

is what Gettier cases allegedly show. Thus, Elgin says, it is not enough merely to point out that our concept of knowledge serves a variety of practical purposes: ‘it may be inadequate in ways that do not typically impede its practical usefulness’.

Richard Rorty (in)famously said: ‘Pragmatists think that if something makes no difference to practice, it should make no difference to philosophy’ (1995: 281). As a pragmatist sympathizer, I’m inclined to agree. But instead of letting the issue hang on pragmatism (something even fewer people are willing to endorse), the book *also* provides an answer to the Gettier problem – one that is compatible with the informant-flagging account of knowledge.

The answer, in short, is that it makes sense to deny knowledge to Gettier victims even though they might *seem* to meet the criteria to be a reliable informant. When we discover that the correlation between an agent’s justification (on one hand) and their being right about *p* (on the other hand) was accidental or lucky, this rightly affects our attitude towards their reliability as to whether *p*. As Craig says, relying on a Gettier victim would ‘produce the retrospective feeling of having run a risk, of having done something that one would not have done had one been just a little better informed at the time’ (Craig 1990: 76). A Gettier victim may use a generally reliable method to acquire true belief, but that method does not play an appropriate part in their success. And it is in the interest of truth-seeking enquirers to want true beliefs that are not accidental relative to the method used by the informant (e.g. reasoning by a false premiss, or whatever). This affects our attitude towards their status as good informants. We regard Gettier victims as unreliable because they easily could have gotten things wrong, had things been slightly different; for instance, Jones might not have had 10 coins in his pocket (Gettier 1963); Henry might have easily identified a fake barn (Goldman 1976); etc. This affects their status as reliable informants, and thus as knowers.

For the sake of argument, Elgin takes on board two assumptions that she finds dubious. First, she questions whether the need for, and contours of, a concept of knowledge is entirely grounded in ‘epistemic interdependence’. Second, she remains unconvinced that the concept of knowledge has a single fundamental function.

About the first point, Elgin says ‘there is something odd about not insisting that individualist knowledge can be accommodated’. While I do defend a deeply social picture of epistemology, nothing in my account rules out that 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. Moreover, we can be reliable sources of information even when we’re alone on an island and nobody cares what we have to say. The point is that a system of epistemic evaluation would be rather pointless for such an individual (that said, even this individual may serve as a reliable informant for themselves at some future time). Language is a social phenomenon

and we use epistemic vocabulary to solve epistemic coordination problems that are faced exclusively by individuals in groups (see 44–5).

About the second point, I agree with Elgin that our concept of knowledge serves a variety of purposes. In this sense, it is like a Swiss army knife. However, in Chapter 5 of the book (and also my comments above) I try to explain why the informant-flagging function has explanatory priority over other functions, such as certifying the end of enquiry, assuring and blaming others etc. I argue that the informant-flagging function is more fundamental because it can explain these other functions, while these other functions cannot explain it. Moreover, it is highly unlikely that the concept of knowledge came into existence with all these other functions simultaneously present. If it were crafted from a variety of existing materials, like a Swiss army knife, then we'd need a story about how we already had different epistemic concepts to serve this plurality of functions, and then knowledge came along to amalgamate these various functions into an interconnected single concept. In the absence of such a story, I think it's more reasonable to think that knowledge is more like a hammer than a Swiss army knife. It came into existence to serve a fundamental function, and that function now explains why it can serve a variety of other useful purposes.

### 3. Replies to Lawlor

Lawlor and I agree on the big picture. We both think that epistemologists should investigate knowledge by reflecting on the role that knowledge claims play in our social interactions. On Lawlor's (2013) view, the act of claiming to know is a way to *assure* others of the truth of our claim. On my view, we speak of knowing to identify agents as reliable sources of information. Although we disagree about what one tries to do in claiming to know things, we agree that asking about the function of knowledge claims is a valuable epistemological tool. Both of us try to demonstrate this value by using our respective hypotheses to shed light on current problems in epistemology, including scepticism, fallibilism and the semantics of knowledge claims.

Lawlor raises two concerns about how I execute this project. First, a large part of my book is devoted to providing an account of the level of justification required for knowledge. Lawlor agrees that we should answer this question by asking what level of justification it makes sense for us to adopt and coordinate around, given our epistemic needs. But she says my solution includes some 'hard-to-pin-down features'. Second, she questions a 'big leap' in my argument for *epistemic pragmatism*. In the book, I make two main claims about the semantics of knowledge claims: (i) appealing to function will not help us decide which semantic theory (e.g. contextualism or invariantism) is correct; (ii) none of these semantic theories are correct. Lawlor seems happy to grant the first claim, but she wonders what motivation there is for the second, more radical claim.

Lawlor is right to ask for clarity on these issues. In what follows, I will try to spell out my view in a little more detail. Let's start with the issue of the standard of evidence for knowledge.

To qualify as a knower, I argue that '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' (68). I call this the *reliable informant standard for knowledge*. Lawlor's first question is: why do I allude to 'reliability' when (i) the notion of reliability doesn't figure in the statement of the standard and (ii) reliability may not be required to be in a position that eliminates relevant not- $p$  possibilities?

In response to (i), I refer to reliability because the level of justification needed for knowledge will be that which puts the agent in a strong enough epistemic position for her to fittingly serve as a *reliable* source of actionable information. We expect our sources to be reliable in order to rule out cases where one happens to have true beliefs as a matter of luck. Of course, this raises the question of what level of reliability is required to qualify as knower. I attempt to articulate this idea by using the relevant alternatives framework, where I say the agent's epistemic position must be sufficiently strong with respect to  $p$  to eliminate all of the not- $p$  possibilities that are relevant alternatives to members of the epistemic community. Thus, while 'reliability' doesn't figure into the statement of the standard, the standard itself is taken to be a statement of what it means to be reliable.

This leads to (ii), Lawlor's claim that reliability may not be required to be in a position to eliminate relevant not- $p$  possibilities. In support of (ii), Lawlor provides the following example: 'Mary may not be a reliable judge of cat varieties, but still be in a position to rule out the few alternatives her audience needs eliminated, because their background knowledge already makes a wide swathe of alternatives irrelevant.' Lawlor asks: 'Does Hannon intend to count Mary's case as one of knowing or not?'

A lot turns on the details of this case. Although Mary may not be a reliable judge of cat varieties in general, she may be reliable enough to rule out the few alternatives her audience needs eliminated. Suppose Mary cannot tell the difference between a Maine Coon, a Norwegian Forest cat, a Ragdoll and a British Shorthair, but she *can* reliably tell the difference between a Sphynx and a Persian cat. If she knows that her audience has eliminated all other cat varieties on good evidential grounds, leaving her to tell the difference only between a Sphynx and a Persian cat, I would count Mary as having knowledge.

Another hard-to-pin-down feature of my view is: who counts as a member of the 'epistemic community'?

Here is what I wrote in the book: 'we should think of the epistemic community as roughly comprised of anyone who might actually draw on the relevant information, where 'might' tracks the notion of a possibility that

could *reasonably* be expected to occur' (68). This was intended to rule out alternatives that we normally do not take to be likely counterpossibilities to what the subject knows (e.g. brains in vats, stuffed goldfinches etc.). Without such a restriction, the standards for knowledge would skyrocket.

As Lawlor points out, this characterization of the epistemic community differs from her own 'reasonable person' standard. According to the reasonable person standard, 'being in a position to know requires evidence sufficient to eliminate the alternatives a reasonable person in the situation of the knowledge attributor would want eliminated before taking *p* to be true'. This may differ from the expectations of those *who could reasonably be expected to use* the agent's information. After all, it can be reasonable to expect unreasonable people to use one's information. To illustrate, Lawlor provides the example of an agent who is surrounded by people who believe in demon possession. A *reasonable person* would not require demon possession to be eliminated in order to know that a virus caused the illness, but the people *who are reasonably expected to draw on the agent's information* might want this possibility ruled out.

Upon closer inspection, however, our views are not so different. Elsewhere I say that a knower must rule out those alternatives that are *reasonable* to the members of the epistemic community (69). To me, this was equivalent to saying that a knower must have evidence sufficient to eliminate the alternatives a reasonable person would want eliminated (at least, that's how I intended it). After all, the reasonable person standard presumably reflects whatever epistemic standard is taken to be reasonable by the relevant community's judgement.

Moreover, as Lawlor acknowledges, I deny that the particular interests of some local community (such as a cult that believes in demon possession) will decide which alternatives count as relevant for knowledge. Our practice of epistemic evaluation takes into account a much broader set of interests, which enlarges our understanding of who counts as a member of the epistemic community (93).<sup>2</sup> This allows us to say two plausible things. First, we can understand why the believers in demon possession will not ascribe knowledge to a competent scientist who fails to rule out this possibility, since these believers take demon possession to be relevant. Second, we can explain why they would be *mistaken* to deny knowledge to the competent scientist, since this community's understanding of which alternatives count as relevant is not properly anchored in the broader interests that shape the application conditions of our concept of knowledge. (This is also why we can criticize communities of conspiracy theorists for having bad epistemic standards.) I am here assuming that demon possession is not something believed more widely. If it *were* widely

2 I also leave open the possibility that the concept of knowledge can be *culturally elaborated* in different ways, just as the notion of a reasonable person may vary culturally.

believed, then it *would* be a relevant alternative and thus affect what it takes to know.

Lawlor worries that by enlarging the epistemic community in this way, I make the proposed standard too demanding. I found this remark confusing in the context of her example. While expanding the epistemic community does typically push the standard of knowledge upwards, it also sets a reasonable limit on how high it can go, precisely to rule out the kind of example she envisions: where people in a local community hold implausible beliefs that would push the standards too high. It is for this reason that I say the interests of *some actual individuals* will not affect the relevant epistemic standard (69). If these beliefs do not reflect the ‘typical’ or ‘normal’ stakes of information-seeking enquirers, we needn’t worry about them shaping our epistemology.

So I think Lawlor and I agree even on some of the finer details, not just the big picture. Her commentary has helped me realize that any indications to the contrary are mostly the result of sloppy formulations on my part.

That said, there is one significant issue over which we disagree. I endorse the ‘impurist’ idea that an individual’s practical reasoning situation can sometimes influence the evidential standard for knowledge. When the stakes are especially high, the situation may require us to crank the epistemic standard up at least another notch.

At first blush, this idea seems to be in tension with Craig’s insight that we all must rally around a communal standard to coordinate and share reliable information. As Lawlor says, ‘Hannon would like to put two competing demands together into one epistemic standard’. The resulting view is as follows:

Someone who meets the communal threshold for knowledge will usually qualify as a knower, but she will not qualify as a knower if she isn’t reliable enough to meet the more demanding expectations in a high-stakes practical reasoning situation. (89)

Lawlor makes two objections against this more complex view. First, she says I provide no argument for the impurist element of the standard. Second, she says that I’ve lost sight of the main motivation for a communal standard: by allowing an individual’s practical stakes to defeat the communal standard, I make it too difficult to coordinate expectations around who counts as a knower. I’ll take these points in turn.

Is there any reason to think individual stakes can override the communal standard when the concerns of some individual are more pressing than usual? I think there are two reasons. First, there is evidence that we *do* withhold ascriptions of knowledge when an enquirer faces incredibly high stakes. As various examples in the epistemology literature show, we often allow the high stakes of an individual to trump the ordinary standard for knowledge. Thus, my account offers a plausible diagnosis of the commonly held intuitions about low-stakes and high-stakes cases. Second, there is a principled reason for this:



knowledge ascriptions are used to identify sufficiently reliable informants, but we should not recommend an informant if we are aware that the enquirer's purposes are particularly pressing and so the informant will fall short of the enquirer's heightened demands. To flag the informant as a knower would misleadingly suggest that she is sufficiently reliable for the enquirer's purposes.<sup>3</sup>

Does this make it difficult to promote the kind of epistemic coordination at the heart of Craig's proposal? I have elsewhere (Hannon 2015) attempted to deal with this objection, so I will only say a few words about it here.

First, it is unclear that we hinder widescale epistemic coordination by temporarily allowing an individual's unusually high stakes to influence our knowledge claims. It might be rare that we encounter individuals who demand more than usual, in which case we would rarely encounter the type of situation described by Lawlor. Moreover, Lawlor saddles me with the view that *anyone* in the community with elevated needs will lead to a raising of standards. While some (see Grimm 2015) defend this 'rising tides' view, I find it implausible. My claim is not that the relevant epistemic standard is set by *whoever* has the highest stakes in the community (that way lies scepticism); my claim is that we *temporarily* heighten the standard *in a particular context* when we are recommending an informant to an enquirer with recognizably high stakes.

Second, impurism is not the only view that would make it tricky to coordinate expectations around who counts as a knower. Suppose we adopt the traditional purist view that there is a fixed standard for knowledge that meets the demands of the community, even though this standard is occasionally too weak for individuals with particularly high stakes. In these high-stakes cases, purists often say that it is *appropriate*, even if strictly speaking false, to deny knowledge (see Rysiew 2001, Brown 2006). But this also leads to coordination problems. When retrieving information from knowledge attributions made in other contexts, we must be careful not to be misled by our false-but-appropriate knowledge denials. For example, a purist might maintain that in a high-stakes case when Hannah says 'I don't know that the bank will be open tomorrow' she is saying something contextually appropriate but strictly speaking false. Now, assume that Sarah wants to preserve this information, so she stores the sentence 'Hannah doesn't know that the bank will be open on Saturday' in her belief box. There is a worry that in some future low-standards context, this sentence may mislead people into thinking that Hannah lacks information that she in fact possesses.

3 Pragmatic insensitive invariantists will argue that the standard for knowledge firmly settles at a level high enough to meet the community's demands, while the alleged context-sensitivity of our knowledge ascriptions ought to be dealt with at the level of pragmatics, not semantics (see Rysiew 2001 and Brown 2006). But I am not making any claims about the semantics of 'knows' in this section of the book. Rather, I am making a claim about the epistemic standards *operative* when evaluating others epistemically. It seems clear that the epistemic standards to *merit* a knowledge ascription change in the way the impurist describes.

Thus, the purist who goes in for a pragmatic story will have to say that we store and recall complex information about the context in which the assertion was made, which is the same problem the impurist allegedly faces. This suggests that the epistemic scorekeeping required to retrieve epistemically useful information from knowledge attributions is not a consequence of accepting impurism but rather a consequence of certain empirical claims about our ordinary knowledge attributing practices.

Finally, Lawlor raises a concern about the most tentative proposal in the book. She seems willing to grant my conclusion that putative facts about the function of knowledge ascriptions will not instruct us as to whether contextualism, sensitive invariantism, or insensitive invariantism is true; but she doubts the more radical claim that this entire debate about the semantics of 'knows' is misguided because it wrongly assumes there is a determine answer to the question of whether knowledge claims have an invariantist or contextualist semantics.<sup>4</sup>

As Lawlor says, this is a 'big leap'. In the book, I confess 'a good deal of uncertainty' about this proposal and admit to not providing a detailed argument. My aim, rather, is to 'provide a sketch of a position in order to make it more visible' (173). Still, I want to insist that this leap only looks big if we stay within the grip of the traditional framework. Scholars who debate the semantics of knowledge ascriptions tend to think there is an objectively true answer to the question: do knowledge claims have a contextualist or invariantist semantics? While they disagree about whether particular uses of 'knows' are true vs. appropriately asserted yet false, they nevertheless think there is an objectively right answer to this question. Drawing on insights from J.L. Austin, I try to cast doubt on this presupposition.

I argue that our practice of epistemic evaluation achieves its purpose whether it is accounted for in terms of what is pragmatically conveyed or semantically expressed. Each semantic view proposes a more or less equivalent way to achieve our epistemic goals. Thus, the issue of whether 'knows' takes an invariantist or contextualist semantics has no bearing on our actual practice of epistemic evaluation. From this, it does not logically follow that neither contextualism nor invariantism is true. I happily grant that. However, it does give us a reason to wonder why we should *assume* that the semantics must settle one way or the other.

Lawlor interprets me as arguing as follows: when theorists disagree about semantic theories that seem to provide equally good accounts of all our usage, we should conclude that this is because there is no fact of the matter as to which semantic theory is true. But that is not my argument. Rather, my argument is that the function and use of epistemic vocabulary plausibly shape the semantics of these terms, yet there is no good reason to assume the semantics must determinately settle either in favour of contextualism or invariantism

4 I set aside relativism in the book.

*precisely because* our practice of epistemic evaluations works perfectly well in the absence of such an answer; thus, we should question the motivation for thinking there *is* such an answer.

The University of Nottingham  
Nottingham NG7 2RD, UK  
michael.hannon@nottingham.ac.uk

### References

- Brown, J. 2006. Contextualism and warranted assertability maneuvers. *Philosophical Studies* 130: 407–35.
- Craig, E. 1990. *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- Dogramaci, S. 2012. Reverse engineering epistemic evaluations. *Philosophy and Phenomenological Research* 84: 513–30.
- Gettier, E. 1963. Is justified true belief knowledge? *Analysis* 23: 121–23.
- Goldman, A. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy* 73: 771–91.
- Grimm, S. 2015. Knowledge, practical interests, and rising tides. In *Epistemic Evaluations*, eds. J. Greco and D. Henderson, 117–37. Oxford: Oxford University Press.
- Hannon, M. 2013. The practical origins of epistemic contextualism. *Erkenntnis* 78: 899–919.
- Hannon, M. 2015. Stabilizing knowledge. *Pacific Philosophical Quarterly* 96: 116–39.
- Hannon, M. 2019. Skepticism: impractical, therefore implausible. *Philosophical Issues* 29: 143–58.
- Henderson, D. 2009. Motivated contextualism. *Philosophical Studies* 142: 119–31.
- Henderson, D. 2011. Gate-keeping contextualism. *Episteme* 8: 83–98.
- Kaplan, M. 1985. It's Not What You Know that Counts. *The Journal of Philosophy* 82: 350–363.
- Lawlor, K. 2013. *Assurance*. Oxford: Oxford University Press.
- Rorty, R. 1995. Is truth a goal of enquiry? Davidson vs. Wright. *The Philosophical Quarterly* 45: 281–300.
- Rysiew, P. 2001. The context sensitivity of knowledge attributions. *Noûs* 35: 477–514.
- Williams, B. 1973. *Problems of the Self*. Cambridge: Cambridge University Press.