

- Grimm, S.R. 2015. Knowledge, practical interests, and rising tides. In *Epistemic Evaluation*, eds. D. Henderson and J. Greco. Oxford: Oxford University Press.
- Hannon, M. 2013. The practical origins of epistemic contextualism. *Erkenntnis* 78: 899–919.
- Henderson, D. 2009. Motivated contextualism. *Philosophical Studies* 142: 119–31.
- Henderson, D. 2011. Gate-keeping contextualism. *Episteme* 8: 83–98.
- Henderson, D. 2019. Epistemic norms as social norms. In *Routledge Handbook of Social Epistemology*, eds. Fricker, M., Graham, P., Henderson, D., and Pederson, N., 423–34. New York: Routledge.
- Henderson, D. 2020. Are epistemic norms fundamentally social norms? *Episteme* 17: 281–300.
- Henderson, D. ms. The place of non-epistemic matters in epistemology: norms and regulation in various communities. *Synthese*.
- Henderson, D. and T. Horgan. 2015. What's the point. In *Epistemic Evaluation*, eds. D. Henderson and J. Greco. Oxford: Oxford University Press.
- Kappel, K. 2010. On saying that someone knows: themes from Craig. In *Social Epistemology*, eds. A. Haddock, A. Millar and D. Pritchard. Oxford: Oxford University Press.
- Kelp, C. 2011. What's the point of 'knowledge' anyway? *Episteme* 8: 53–66.
- Kvanvig, J.L. 2009. The value of understanding. In *Epistemic Value*, eds. A. Haddock, A. Millar and D. Pritchard, 95–111. Oxford: Oxford University Press.
- Rysiew, P. 2012. Epistemic scorekeeping. In *Knowledge Ascriptions*, eds. J. Brown and M. Gerken, 270–94. Oxford: Oxford University Press.
- Tomasello, M. 2009. *Why We Cooperate (A Boston Review Book)*. Cambridge, MA: MIT Press.
- Tomasello, M. 2020. The ontogenetic foundations of epistemic norms. *Episteme* 17: 301–15.

The Function of Knowledge

BY CATHERINE Z. ELGIN

Human beings are epistemically interdependent. Much of what we know and much of what we need to know we glean from others. Being a gregarious bunch, we are prone to venturing opinions whether they are warranted or not. This makes information transfer a tricky business. What we want from others is not just information, but reliable information. When we seek information, we are in the position of enquirers not examiners (Williams 1973: 146). We ask someone whether *p* because we do not ourselves already know whether *p*. So we cannot check our informant's bona fides by seeing whether she is speaking the truth about *p*. We need grounds for trust. In *What's the Point of Knowledge?* Michael Hannon maintains that this is the predicament

that underlies the concept of knowledge. The *we* here is comprehensive. Evidently, every language has a word for knowledge. This, Hannon suggests, indicates that the concept of knowledge satisfies a universal human need. A goal, if not *the* goal, of epistemology is to identify that need and ascertain how it is satisfied. Following [Craig \(1990\)](#), Hannon argues that the point of knowledge is to enable us to identify good – that is, suitably reliable – informants. Let us call this the *information transfer view*.

I am not convinced that knowledge has a single fundamental function. Nor am I convinced that, if it does, that function stems from epistemic interdependence. But I will not argue for either of these concerns here. Even if we grant that the need for a concept of knowledge stems from our ineluctable epistemic interdependence, I will urge, we ought not conclude that the fundamental function of knowledge – or more strictly, of true ascriptions of knowledge – is to identify reliable informants. I begin by reviewing Hannon's argument. I then go on to suggest that an alternative, what [Rysiew \(2012\)](#) calls *the certification view* – provides a better answer. If I am right, function-first epistemology does not vindicate the idea that the main function of true ascriptions of knowledge is to identify good informants.

Hannon's method is to venture his hypothesis and argues that if it is true, the concept of knowledge with the familiar profile emerges. He does not, of course, maintain that only if it is true does the familiar concept of knowledge emerge. But he compares it with, and contends that it does better than, rival proposals. Still, his argument is vulnerable if an alternative does as well; it is undermined if an alternative does better.

Hannon does not hold that flagging reliable informants is the sole function of knowledge; but he contends that it is the fundamental function. Although this function does not exhaust the utility of knowledge, he maintains that it subsumes or underwrites a variety of other important functions. In particular, the need to identify good informants fixes the threshold on justification. An agent's justification for *p* is sufficient for knowledge when his having that level of justification makes it appropriate for him to serve as an informant vis à vis *p*. Enquiry stops where it does, Hannon contends, because the point where it stops is the point where its result can responsibly be conveyed to others.

A principled threshold is important. Once we cease to demand certainty, we run the risk that any particular threshold on justification looks arbitrary. Why here rather than a bit higher or a bit lower? Hannon's answer is that a non-arbitrary answer emerges if we consider what knowledge is for.

Some information transfer relies heavily on circumstantial factors. When Sophie asks Tim whether *p*, Tim may have cues he can draw on – cues that enable him to formulate an appropriate answer. He may be aware of Sophie's cognitive and practical situation – what background she brings to the exchange, why she wants to know whether *p*, what has already been established in the conversation, what can be assumed in the context.

In formulating his answer, he can call on his epistemic empathy – his sensitivity to her epistemic predicament (see [Elgin 2019](#)). Given his highly textured appreciation of the circumstances, it might be reasonable for Tim to speak loosely or metaphorically, or to set unusually high, unusually low, or somewhat skewed standards whose rationale does not generalize. He can do so because his decision derives from the peculiarities of the context in which the question was raised or from their shared, even if idiosyncratic, background assumptions. Where circumstantial factors dominate, pragmatic features bear a considerable part of the burden in assuring that the enquirer gets useful information.

Not all information transfers are a product of face-to-face communication with someone who appreciates the details of the enquirer's situation and knows first-hand whether *p*. Information gleaned from our exchanges often is stored and passed on without contextual indicators. Tim remembers whether *p* and conveys that information. He ordinarily does not remember, and rarely conveys, whether the conversational context in which he gleaned the information was one where it was appropriate for his informant to speak loosely or to invoke particularly demanding or particularly undemanding standards. Ordinarily, he does not even remember how he learned that *p*. Information suitable for widespread use should satisfy a general-purpose standard. Pretty much anyone who needs to know whether *p*, for pretty much any reason, should be able to rely on the public standard, confident that information that satisfies it is sufficiently justified. That standard, Hannon maintains, is the threshold for knowledge. If Tim's true belief satisfies the community standard for justification, Tim knows that *p*, and can responsibly convey the information that *p*.¹

The public standard reflects a trade-off. Set the bar too low, and the community is subject to taking undue epistemic and practical risks. In that case, for some expected and legitimate purposes, the level of justification will not be high enough to make acting on that information reasonable. Set it too high, and the community lacks access to information whose justification is good enough for the uses to which that information is likely to be put. Although Hannon recognizes the importance of knowledge for action, his criterion keys knowledge directly to information transfer. To know that *p* is to be in a position where it is epistemically proper to convey to others the information that *p*. Knowledge is indirectly tied to action in that it is appropriate to convey the information that *p* (hence the would-be informer knows that *p*), only if for any likely situation where *p* is relevant, others could reasonably take action on the basis of that information.

Knowledge, Hannon maintains, is something we impose on the world, not something we discover that the world contains. It is not a natural kind. We

1 A caveat is embedded in the 'pretty much'. If the stakes are very high, Tim ought not convey the information that *p* unless his justification is well above the normal threshold.

contrive the concept of knowledge to suit our purposes. That concept is not only artefactual, it is thick; it has both descriptive and normative dimensions. Assume, for the purposes of illustration, that knowledge is justified, un-Gettierized, true belief. Then the truth and belief elements are primarily descriptive. Knowledge involves a relation in which an epistemic agent stands to a fact. If there were no such fact, there would be no knowledge of it. Were the agent's attitude towards the fact something other than or less than full belief, it would not qualify as knowledge.² 'True belief' then describes his relation to the world. The evaluative dimension is largely a function of the justification condition. Having sufficient evidence, good enough reasons or being suitably reliable is what makes the belief epistemically proper. The anti-luck condition that blocks Gettier cases, whatever it is, plays a supporting evaluative role. It assures that nothing interferes with the justification's being properly connected to the truth.

Hannon says virtually nothing about what it takes to impose a concept. Still, we cannot impose whatever we like. Minimally, the concepts we impose have to satisfy conditions of adequacy. We cannot just stipulate that they do. If we devise a concept of a set, stipulating that sets can take anything as members, we end up with the Russell paradox. What assures us that the everyday concept of knowledge does not give rise to similar difficulties? That it works fairly well in everyday cases is not enough. So does the naïve concept of a set. It may be true that a concept *c* is widely used, and serves a variety of practical purposes. Like the naïve concept of a set, it may be inadequate in ways that do not typically impede its practical usefulness. But it still may be inadequate. The concept of knowledge seems internally consistent. Its inadequacy lies elsewhere: it gives rise to Gettier cases. Our criteria for knowledge are satisfied in circumstances where, we are convinced, the agent does not know. Hannon dismisses Gettier cases on two grounds. First, since he does not purport to offer necessary and sufficient conditions for knowledge, his position is not discredited by counterexamples. Let us grant that point. But second, he insists that everyday knowledge attributions do not give rise to Gettier cases, so they can be ignored. Here I disagree. 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 that concept, like the naïve concept of a set, is inadequate.

We might conclude that there's no such thing as knowledge. Alternatively, we might think that the contours set by our everyday concept need to be revised. That concept roughly indicates something worth marking out. To think this would not require treating knowledge as a natural kind whose borders we are trying to discover. Rather, it would require getting clearer

2 The full belief criterion can be weakened slightly, but not much. Unless the agent is quite confident that *p*, she does not know that *p*.

on what characteristics the thick, artefactual concept needs to have to serve our purposes while satisfying rigorous criteria of adequacy.

Another worry is this: does the concept we impose perform the function(s) we want it to perform? That, of course, depends on what those functions are. Hannon has identified his preferred function – marking out good informants. He acknowledges that some who know are not good informants. They are unwilling or unable to say what they know. This is not a problem. All he needs is that those who are unwilling or unable to speak qualify as knowers so long as they would be good informants about the matter if they were informants at all. The more serious question is whether marking out good informants largely exhausts the major functions of the concept of knowledge. I will argue that it does not.

Let us provisionally grant that the concept of knowledge is grounded in the fact of our ineliminable epistemic interdependence. No human being independently knows enough to survive for long. Still, I maintain that epistemic interdependence runs far deeper than our depending on one another for reliable information.

The certification view maintains that the function of the concept of knowledge is to set the point where enquiry properly can stop. In that case, it is because enquiry stops where it does that at that point it is appropriate to convey a conclusion to others. This might seem to leave us at an impasse. Hannon says that because the threshold for responsibly conveying information is reached, it is reasonable to stop enquiry; certificationists say that because it is reasonable to stop enquiry at that point, it is reasonable to convey information whose justification reaches that threshold. I will urge that certificationists have the edge. I will maintain that the enquiry stopping criterion readily accommodates mundane cases of knowledge that Hannon's criterion has a harder time handling.

I do not know what causes the ominous rattle in my car's engine. Being a reliable source of information about cars, my auto mechanic can tell me. But I want more than mere information. I want the car fixed. Conceivably, he could convey to me the information I need to fix my own car. Given my level of competence, that's not a promising strategy, but perhaps it could be done. Still, I'd rather rely on his expertise. He has the knowledge to do the job himself.

Besides car problems, I also have a worrisome twinge in my tooth. Here I rely on my dentist's expertise. She has the knowledge needed to examine the tooth and read the X-ray. She can convey the information that I have a cavity that needs to be filled. Again, I want more than information. I want the tooth fixed. Unlike the auto mechanic, however, my dentist cannot tell me how to drill my own tooth. This is an area where do-it-yourself-under-suitable-guidance simply will not work. I need to rely on her knowledge because it enables her to perform an action that I cannot perform for myself. Epistemic interdependence often consists in our relying on knowledgeable others to do things that we cannot do for ourselves.

Organized groups display and rely on epistemic interdependence. A surgical team typically consists of (at least) a surgeon, a surgical assistant, a nurse, an anaesthesiologist, a radiologist and one or more medical technicians. Depending on the surgery, the team may have additional members. Each has expertise the others lack. Collectively, they have the knowledge needed to, e.g., remove a spleen or repair an aneurism. None of them has the knowledge to do it alone. Moreover, the surgery is and must be a team effort. There are many surgical operations that no one person could do alone. To perform them, the team needs to work together. Each member has to know when, how, to what extent and in what respects to rely on which other members of the team. The requisite knowledge goes far beyond relying on others as sources of information. The team members have well-defined roles, backed by appropriate credentials which they earned by mastering different bodies of knowledge. Their training equips them with different ranges of information, different skills and differences in the selective attention they bring to the task. Each is apt to notice things that the others don't notice or can't discern. Each is well-trained to do her particular job. The sort of epistemic interdependence a surgical team displays is characteristic of other organized groups whose members work together while having distinctive roles. With or without credentials, members differ in the expertise they contribute to their joint venture. They rely on each other to have and exercise the knowledge they need to perform their several functions.

We are epistemically dependent on others for instruction as well as for information (see [Schmitt 1992](#): 558). A parent shows a child how to tie his shoes. A lab instructor teaches a novice how to use a pipette. A philosophy professor shows her students how to recognize a fallacy. What is conveyed in these cases is a generalized skill that goes beyond a particular instance. Information transfer may be part of the method, but it is not the goal. For our purposes, the main point is simply that humans are profoundly dependent on others for instruction. And, barring accident, the instruction will succeed only if the instructors know that which they teach.

The sorts of epistemic interdependence I've mentioned are widespread and vital. A viable functional account of the function of knowledge should accommodate them. A certification model that sets the threshold for justification for knowledge at the point where enquiry can responsibly stop can handle all of our cases. Information transfer and storage is epistemically responsible when the justification for the information being transferred or stored reaches the threshold for ending enquiry. If there is no need to enquire further, *S* can responsibly convey *p* to pretty much anyone who seeks the information. Action on behalf of others is epistemically responsible when the reasons that underwrite that action reach the threshold for ending enquiry. *Ceteris paribus*, if it is appropriate for the auto mechanic to end enquiry into my engine's rattle, or my dentist to end enquiry into the cause of my toothache, or my cardiologist to end enquiry into the source of my arrhythmia, it is appropriate for each of

them to take action on the basis of what their enquiry revealed. Teams work together in contexts of trust when each member can be confident that her own and her teammates' actions are justified by reasons that are good enough that no further enquiry is needed. Instruction is epistemically responsible when the instructors' words and actions are justified by reasons that reach the same threshold.

I said earlier that the threshold represents a trade-off between views that are too weakly justified to be acceptable for some of the community's legitimate needs and being too strong so that valuable information would be lost. This needs to be extended. The threshold remains a trade-off, but the trade-off is a trade-off in the actions as well as the beliefs it would vindicate. Set the standard too high and no one will purport to be a knowledgeable auto mechanic, dentist or cardiologist. Hence our prospects of getting our cars, teeth or hearts repaired will diminish. Set it too low, and many with dubious abilities will purport to be able and think they are able to fix our cars, our teeth, our hearts. Again our prospects are not good. But if we set a reasonable threshold, we are in a position to benefit from the expertise of others, both their words and their actions. This is what licensing requirements in the trades and the professions claim to do. They certify the competence of those they license. In other areas, the criteria are less regimented, but no less real.

So far, I have followed Hannon and assumed that the need for, and contours of, a concept of knowledge are grounded entirely in our epistemic interdependence. I've argued that the certification view does a better job of accommodating the various ways human beings depend on one another epistemically than the information transfer view. But there is something at least odd about not insisting that individualist knowledge be accommodated as well. We can readily grant that epistemology has long been too focused on what an individual knows or can know without relying on others. Still, each of us on her own knows a lot. We have, for example, vast amounts of perceptual knowledge, inferential knowledge, introspective knowledge and a priori knowledge. I do not see how the information transfer model can accommodate these. Nor is it plausible that they are peripheral or unimportant.

Here again, the certification view does better. It marks out a suitable end for private enquiry as well as for public enquiry. It certifies that, for example, an apparent perception as of a familiar cat in the centre of our visual field, 10 feet away, in broad daylight is an appropriate place to end enquiry into a question where the cat is. It certifies that checking our sums and arriving at the same, plausible answer is an appropriate end to enquiry into the question of how much the groceries cost. And so forth. I suggest then that if we want a function-first epistemology, where the concept of knowledge is fixed by its practical function, we would do better to endorse a certification view.

Still, I'm not convinced that the concept of knowledge has a single fundamental function. Hannon maintains that a function-first approach is pragmatic. The concept of knowledge gains its identity from its primary function. In

this, he says, it is like a hammer. Although a hammer can be and is used for many purposes, its primary purpose is to pound nails. That purpose determines its design. If the concept of knowledge is a tool like a hammer, then by identifying its purpose we can determine its contours. Although it is plausible that the concept of knowledge is a tool, it may not be a tool like a hammer. Given the variety of purposes it fills and fills well, I think it is more likely to be the epistemological analogue of a Swiss army knife.

Graduate School of Education
Harvard University
Cambridge, MA 02138, USA
catherine_elgin@harvard.edu

References

- Craig, E. 1990. *Knowledge and the State of Nature*. Oxford: Clarendon Press.
Elgin, C. 2019. The mark of a good informant. *Acta Analytica* 35: 319–31.
Rysiew, P. 2012. Epistemic scorekeeping. In *Knowledge Ascriptions*, eds. J. Brown and M. Gerken, 270–94. Oxford: Oxford University Press.
Schmitt, F. 1992. Review of *Knowledge and the State of Nature* by Edward Craig. *Mind* 101: 555–59.
Williams, B. 1973. *Problems of the Self*. Cambridge: Cambridge University Press.

Comments on *What Is the Point of Knowledge?*

KRISTA LAWLOR

The point of knowledge is to answer our need for information that will let us successfully navigate our world. So says Edward Craig in *Knowledge and the State of Nature* (Craig 1999). This claim may sound anodyne, but according to Craig, it is crucial we keep this fact uppermost in our minds as we theorize about knowledge. Craig argues that our concept of knowledge begins its life by answering our need to mark out those who have the information we seek. A priori analysis may not reveal the application conditions of our concept of knowledge, but reflection on its function can.

Michael Hannon's book seeks to build on Craig's insights about the function of the concept of knowledge, by showing how Craig's perspective can resolve some important debates in epistemology. Hannon scopes out an array of issues in epistemology where a functional approach might bear fruit, and his book is chock full of discussions of key issues in epistemology. My comments necessarily leave much out. I will focus on just two sets of issues.