

The Value of Knowledge

Alan Millar

Mainstream theory of knowledge in the second half of the twentieth century was in large measure shaped by the *traditional analytical project*: to provide an analysis of the concept of knowledge that is brought into play by instances of the schema *A knows that p*. The analysis was to take the form of the provision of necessary and sufficient conditions for the truth of instances of this schema. Notoriously, this project ran up against Gettier examples. While it seemed plausible that knowledge was justified true belief, the Gettier examples suggested that there were cases in which a subject has a true belief that is in some sense justified but does not on that account know. The hunt was on for a tightening up of the notion of justified belief or for a requirement for knowledge that would replace or supplement the justification condition. A good deal came out of this project. For instance, it is widely accepted that any plausible account of knowledge should build in the idea that when you know that p you have a true belief acquired in a manner that yields true beliefs with a high degree of reliability. Other notions have come to the fore that seem relevant to understanding what knowledge is— notions like those of sensitivity and safety. However, the traditional project has come under close critical scrutiny.

Internal problems for the traditional project

A number of problems have arisen that are internal to the traditional project. I'll single out the following:

- (1) Attempts to provide a traditional analysis of knowledge appear not to have been successful. Each analysis seems to be vulnerable to counterexamples. So there are at least inductive grounds for wondering whether it is likely that any analysis will succeed. (Williamson 1995; 2000.)
- (2) There are grounds for doubting that any traditional analysis could succeed if it analyses knowledge as true belief plus something else the obtaining of which does not guarantee the truth of the belief. For *a priori* it looks as if any such account will be vulnerable to Gettier cases. There will be a logical space for cases in which the truth requirement is

met, and the further requirements, but it seems to be a matter of luck that the subject forms a true belief. (Zagzebski 1999.)

- (3) It is a mistake to suppose that all concepts of which we have mastery admit of analysis in terms of necessary and sufficient conditions. So there is a question as to why we should suppose that knowledge is a concept that admits of such an analysis. Perhaps we grasp the concept through grasping certain central cases of knowledge, for instance, perceptual knowledge, and apply the concept to such cases and those sufficiently like them.
- (4) Some traditional analyses are complex. This raises a question about how analyses are supposed to relate to our actual mastery of the concept. On the face of it we do not represent such complex analyses to ourselves. It is therefore implausible that we are guided in any explicit way by any such analysis. So if the traditional project is still feasible there must be some subtle relation between the analysis and what goes into our mastery of the concept. It is not clear how any such relation should be explicated.
- (5) Under the terms of the traditional project, knowledge can be had on the basis of evidence the obtaining of which is compatible with the falsity of the implicated belief. (Call this fallibilism.) For instance, inductive evidence for the truth of an empirical generalization does not entail that the generalization is true. Suppose that such generalizations can be known to be true and that part of what enables one to know them to be true is having sufficient inductive evidence. Then presumably there has to be some amount of such evidence that provides justification that is strong enough for knowledge. It is obscure just how much accumulating evidence is needed to tip a case of reasonable belief that falls short of knowledge over into a case of knowledge. This presents, if not actually an objection to traditional approaches that are fallibilist, then at least a puzzle about how to think about the role in epistemic evaluation of justification that does not entail truth.

Alternative views about knowledge that yield criticism of the traditional project

Some criticisms of the traditional project presuppose an alternative view about knowledge, to which traditional analyses do not seem to do justice. Here are some such views.

- (6) Knowing that *p* is a status that one has in virtue of the exercise of rational capacities. If the truth requirement is ‘external’ to what the subject has achieved through the exercise of rational capacities—if for all that the subject has achieved through the exercise of such capacities the implicated belief could be false—then knowing will not be a status that one has simply in virtue of the exercise one’s rational capacities. Indeed, it will not be an achievement at all since it depends too much on a favour from the world that is not guaranteed by one’s having rightly formed one’s belief. (McDowell 1995)
- (7) Reflection on our concepts of knowledge, justification and rationality suggests that knowledge is conceptually prior to the other two. (Unger 1975; Williamson 1997, 2000; Hyman 1999.) Arguably, to understand what it is for a belief to be justified, when ‘justified’ means well-grounded, and not just reasonable (blameless), we need to understand what it is to know. When we provide justifications for what we believe in the form of considerations that entail or provide empirical evidence for the truth of what we believe, we presuppose that these considerations are true and that we know them to be so. At the point at which any such consideration is shown to be false the case collapses. So even if in some sense the belief in question was reasonable, the subject having been in no position to appreciate that the consideration was false, it was nonetheless not well grounded. And if it can be shown that we do not *know* that the considerations in question are true there is a question whether the justification is one that we have, just because we are not in a position to regard it as a justification.
- (8) Knowledge is not analysable as a composite of an internal mental state and an external truth requirement. (Williamson 1995, 2000) Knowledge is not thus decomposable (it is prime) because we can exhibit triples of cases like the following (Williamson 2000: 72):
- (□) A knows that the election was rigged by testimony from Smith who is trustworthy and A trusts. Brown tells A that the election was rigged but he is not trustworthy and A does not trust him.

- (□) A knows that the election was rigged by testimony from Brown who is trustworthy and A trusts. Smith tells A that the election was rigged but he is not trustworthy and A does not trust him.
- (□) A trusts Smith and believes that the election was rigged. Smith is not trustworthy yet the election was rigged. So the situation is internally like (□) but it is externally like (□). A does not know.

The upshot is that A's knowledge that the election was rigged cannot be a composite of the condition that the election was rigged and the internal condition that is satisfied in case (□). Whenever you have an internal condition that is supposed to be such that when both it and the truth requirement are satisfied one has knowledge, then there are possible situations in which the internal condition and the truth requirement are satisfied, but the subject does not know.

Cases like (□) in Williamson's discussion are akin to Gettier cases even though it is not explicitly built into the scenario that the subject has a justified true belief and does not know. They are cases in which the truth requirement is satisfied and a condition the satisfaction of which is supposed to turn true belief into knowledge, and does not entail that the truth requirement is met, but in which the subject not have knowledge. In the light of this Williamson's point about the primeness of knowledge is rather like Zagzebski's to the effect that there will always counterexamples to accounts of knowledge on which knowledge is conceived as true belief plus something else the satisfaction of which does not entail that the truth requirement is met. But Williamson's argument is specifically designed to establish that knowing is a prime *mental* state. It is prime because it is mental (involves mentality) yet is not decomposable into the satisfaction of an internal mental condition and the satisfaction of the truth requirement.

There is also some kinship between Williamson's point and that made by McDowell (1995) in the context of his discussion of the interiorisation of the space of reasons. Indeed, McDowell suggests an explanation of why knowledge is prime. The explanation is that if knowledge were a composite of internal mental state arrived at through the exercise of rational capacities conceived internalistically, and the satisfaction of the truth requirement, then knowledge could not be the direct upshot of the exercise of rational capacities.

Williamson's argument is telling against a certain conception of knowledge, as being the satisfaction of an internal mental condition and the truth requirement. Recognising the problem, some may add a causal condition to the account of knowledge, at least for some cases of knowledge, such as perceptual knowledge. Suppose then that to know that p by looking, my belief that p must be causally explained in some suitable way by the fact that p . Then it seems that we can explain why case (□) is not a case of knowledge. In that case the fact that the election was rigged does not causally explain in a suitable way the subject's believing that it is, since the chain of explanation does not go through a suitable informant. McDowell's point (noted above) suggests that adding the causal-explanation condition will not without further ado explain why knowledge is acquired when it is met. The problem is that, taken by itself, the causal-explanation condition still puts the acquisition of knowledge at a distance from the exercise of the relevant rational capacities. At any rate, if the capacities can have been exercised when the causal-explanation is not satisfied then for all that those capacities have been exercised it might be that the subject does not know.

Does it matter how we demarcate knowledge from other statuses?

The terms of the debates and issues described so far are such that it is assumed or argued that knowledge is an important and philosophically interesting concept on which it is worth trying to shed light. Indeed, some of the criticisms of the traditional project can be seen as urging us to accept that knowledge is fundamental and central in ways that the traditional project cannot adequately accommodate. If justification, so to speak, comes out of knowledge then knowledge is surely pretty important. However, there are other directions in which one might go that, on the face of it, look very different from those taken by theorists like Williamson and McDowell. There is a pragmatist tradition in which the focus of interest is on the rational fixation of belief often conceived in a manner that does not make the concept of knowledge or perhaps even the concept of justification central. The interest of this tradition for us is not that it deals with a different subject matter from traditional epistemology; it lies rather in the assumption that it is at least an open question whether concepts like that of knowledge and justification are central to the evaluations we make of our beliefs and theories. What is crucial for the rational fixation of belief and the rational conduct of enquiry is being open to correct one's beliefs while being resistant to doubts that are not well-motivated.

Drawing on this tradition Chris Hookway has suggested that we should take seriously the idea that the task of epistemology should be

to describe and explain our practice of epistemic evaluation; to investigate how far our epistemic goals are appropriate and how far our evaluative practice enables us to achieve our epistemic ends. (Hookway 2003)

Hookway distinguishes between those evaluations that concern people's cognitive *states* and those that concern their *activities* of enquiring and deliberating. He suggests that we should consider whether evaluations of the latter sort are bound to presuppose that concepts like those of knowledge and justified belief are central. The plausibility of this depends in part on how we view enquiry and deliberation. If as Hookway suggests the goal of enquiry and deliberation is to find things out then, one might think, it is already settled that knowledge is fundamental, for what is it to find things out, if not to come to know that something is the case? That reaction is, I think, too swift. It focuses attention on an interesting area of philosophical enquiry that concerns the extent to which it is possible to describe the goals of enquiry and the values that should guide enquiry in ways that do not implicate the concepts of knowledge and justification.

I shall not pursue this matter here. What I'd like to underline is the fact that there are ways of reacting to the sense that the traditional project has run out of steam that are different from those adopted by Williamson or McDowell in that they question the centrality of the concept of knowledge for a theory of rational enquiry and the rational fixation of belief

We have then two interesting contrasts. There is

(a) the contrast between those pursuing the traditional project and those who think that this project is suspect, and

(b) the contrast between two camps among the latter—that comprising those who think that knowledge is conceptually prior to other interesting notions, like that of justified belief, and that comprising those who think or suspect that evaluative notions other than knowledge and justification are central to the evaluation of belief and enquiry.

How can we expect to make progress in this area? It is at this point, I suggest, that issues about the value of knowledge come to the fore. If it can be shown that it matters to people in various social interactions whether or not they or others know whether or not something is so then there will be reason to think both that knowledge has a certain value and that telling

whether or not people have it plays an important role in people's lives. We can then raise the further question whether the value it has sheds light on what knowledge must be if it is to fulfil the roles described. That should in turn illuminate the shape of a theory of knowledge. (There is an evident kinship with the approach taken by Edward Craig 1990.)

Why knowledge matters

We should before proceeding further distinguish sharply two issues: (a) whether or not it matters that some subject knows something; (b) whether or not the concept of knowledge plays an important role in our lives. As we shall see, the two issues are connected because certain ways in which knowledge is valuable to us depend on our being able to recognise that people have it or lack it. The recognition, of course, involves application of the concept.

I shall identify two ways in which knowledge, as opposed to something that falls short of knowledge matters to us.

- (1) That we ourselves or somebody else knows that something is so can matter to us because for practical purposes we need well-grounded assurance as to whether or not it is so.
- (2) Knowledge can matter to us because for practical purposes we need our true beliefs to be secure.

I shall illustrate both of those points.

It is not hard to think of situations in which it matters whether or not someone else knows whether or not something is so, and situations in which it matters whether or not we know whether or not something is so. Here is an example illustrative of both. I am an electrician working with a partner who is nearer to the mains box than I am. I want to know whether the power is off because I need to handle some bare wires and I don't want to be electrocuted. So I ask my partner to find out. He checks the mains box, finds out that the mains switch is off, and tells me. I handle the wires. In this situation it was important to me to know that the power was off because I might be killed if it were not. My partner being a decent man, it was important to him to come to know either that the power was off or that it was not, because he would not want to tell me that it was off when it wasn't. That is why he found out by looking.

In this situation we have two simple enquiries: my enquiry of my partner as to whether the switch is off and his enquiry into whether the mains switch is off by going to the box and looking. Notice how natural it is in this situation to describe the enquiries as aiming at finding out, this being conceived, quite naturally, as requiring nothing less than knowing. I want nothing less than to know that the switch is off. Being assured that it probably is would not be good enough. So I want my partner to have nothing less than knowledge that the switch is off if he tells me it is. My partner for his part wants nothing less than to know whether or not the switch is off so he will look and thereby find out. If it's on he'll switch it off before telling me that it is off. If he tells me that it is off, I shall thereby come to know that it is.

The situation just depicted is a good one to illustrate how knowing enables one to be properly assured and how knowing enables one properly to assure. In both cases the assurance depends on treating something as known in some manner by oneself. If I lacked confidence that my partner knew the mains was off and thereby doubted that I had acquired the knowledge that the mains was off through his testimony, I would not have touched the wire. If he had lacked confidence that he knew—if, for instance, he thought he'd seen that it was off ten minutes ago but was not sure—he would not have told me that it was off without checking. (The importance of self-ascriptions of knowledge is stressed by Ward Jones (1997).) I know that I know once I am in a certain position. The same applies to my partner. In his case the position is *looking at the box and seeing that the mains switch is off*. That he sees and thus knows that the switch is off is itself something that he knows. It is that knowledge that provides him with the assurance to tell me. In my case the position is my knowing that he, who knows whether the switch is off, has told me that it is off. I know that I have thereby found out, and thus come to know, that the switch is off.

What does the example suggest about how we should think of knowledge? Am I hinting that if knowledge is to play the role it usually does then we must always know that we know when we do know? The view I want to commend does not commit us to this implausibly strong thesis. The point is rather that if knowledge is to play the role in social interactions that I have illustrated then there must be forms of knowledge such that we can often readily know that we have knowledge of those forms. The point of ascribing knowledge to oneself or others would be lost if we could not often readily tell whether we know or others know. This is one reason why perceptual knowledge is a paradigm of knowledge. For, in the cases that are likely to matter to us, when we know by perceiving that something is so, we generally know that we know, and by a process that is not very different from that by which we know

the thing in question. We can know that we know that something is so in virtue of seeing that it is so, and exercising a recognitional ability to tell whether or not we see that it is so.

The other way in which knowledge can matter to us, which I would like to highlight, arises from our need for our beliefs to be *secure*. Suppose that I believe correctly that I have sent off a letter of recommendation for a student, but I do so, not because I remember sending it, but on the grounds that a colleague has told me that I told him that I had sent it just after doing so. But my colleague is confused. *I* did not tell him I had sent off the letter, but someone else had told him that *he* (this someone else) had sent off a letter. When my colleague spoke to me he thought it was I who had told him. As things stand my only basis for thinking that I have sent off the letter is the false report from my colleague. So my belief is liable to be given up if I learn that my colleague was confused, for then I realise that I don't know. But if I distinctly remember putting the letter in the mail tray it would take a lot more than this to dislodge my belief that I have. (Compare Williamson's burglar example (2000: 62). Williamson was making a point about the explanatory power of the fact that someone knows as opposed merely to having a true belief. The idea was that there can be a scenario such that the burglar's knowing that the diamonds are in the house does more to raise the chance that he will persist in trying to find them than would his merely having a true belief. That's because the circumstances can be such that if the burglar knew he would not readily abandon the belief that the diamonds are in the house, whereas if he only believed truly he might readily abandon the belief.)

The security of a belief—its potential to remain stable over time in a range of possible circumstances—can matter because when we have the truth on a matter, it can be practically important that we not lose our hold on it. My true belief that I sent the letter was easily abandoned in the face of the realization that my colleague's report was false. I then had to make time-consuming enquiries to settle whether I did or did not send it. But if I knew that I'd sent it because I distinctly remember putting it in the tray for outgoing mail, I would avoid such enquiries.

I have in effect been illustrating ways in which knowledge can be instrumentally valuable to us. Both link the instrumental value of knowledge to assurance but in both cases the value depends not just on the combination of subjective certainty and truth. It depends on having come by the certainty through having acquired knowledge and knowing that one has. The value of the knowledge is not simply a matter of the instrumental value of the implicated true belief. It depends on the security that is grounded in knowledge.

Recognising knowledge

Given that knowledge has the kind of instrumental value I have described what does this tell us about the nature of knowledge? I have already given some indication of this in my comments about the scenario in which I need to know that the bare wire is safe. I observed that the interaction in that case depends on there being forms of knowledge that are readily recognisable as such. Fortunately there are such forms. In the case in hand these are knowing by looking and knowing through being told by someone who knows. I am aware that much more needs to be said about these forms of knowledge. But I take it to be part of the data we need to consider that people often acquire knowledge of these forms and can recognise when they or others have knowledge of these forms.

The ease with which we recognise what others know can be further illustrated in situations in which we make sense of how others behave on the basis of facts about their immediate environment. Imagine you are watching a game of tennis. Federer is about to address a high ball. He sees his opponent just begin to anticipate his (Federer's) shot by moving left from Federer's point of view. Federer wrong-foots his opponent, smashing the ball to the right. We make sense of this from knowledge of (a) the practice in which the players are engaged and (b) the epistemically available circumstances at the relevant stage of play. Federer's opponent was starting to run leftwards. We see this and take it that Federer must have seen it too and thus known that his opponent was heading to the left. That he smashes to the right confirms this. No doubt his smash is explained in part by his belief that his opponent was starting to move leftwards. But we have no independent knowledge of Federer's beliefs. We appreciate his take on the situation from facts about the situation and from the epistemic accessibility of these facts to Federer.

The scenario just described is not meant to illustrate why knowledge can matter as opposed to something less than knowledge. Rather it illustrates how important the concept of knowledge is in making sense of others. Seeing what people know from their situation and from facts about what they can tell from looking, hearing, and the like, is crucial for making sense of their behaviour. A vast amount of the behaviour is intelligible to us through this kind of route. As I have stressed already in relation to cases in which it matters that somebody knows, it is just as well that paradigm cases of knowledge are readily ascribable to others and to ourselves. This places a constraint on an adequate theory of the nature of knowledge. It must give a central place to these paradigm cases of knowledge and account for their character in a manner that sheds light on the ease with which we are able to ascribe them.

Although I have emphasised that there are ways in which knowledge matters to us that require that we can recognise when people have it, the view does not present any special problems for knowledge on the part of non-human animals. It makes sense to ascribe knowledge to animals because it is through such ascriptions such that we make sense of their behaviour.

The upshot is

(a) that knowledge is instrumentally valuable because it enables us to have true beliefs of the truth of which we are assured, enables us to give assurance that can be relied upon, and enables us to have stable true beliefs;

(b) that knowledge can fill this role only if we can readily tell when we have it.

My view is that (b) places significant constraints on an adequate theory of the nature of knowledge. In gaining knowledge by looking, for instance, we deploy discriminative capacities, whereby objects become available to us as possible objects of demonstrative thought, and we deploy recognitional capacities, whereby we bring those objects under concepts. The recognitional capacities are nothing less than capacities to acquire knowledge. That is the natural way to think of them.

When we tell that we know something by looking we deploy further capacities beyond those deployed in bringing worldly objects under concepts. We deploy capacities to recognise that we see and thus know that something is so. Just as we learn how to tell that a bird is a swallow in response to appropriate visual experiences, so we learn to tell that we are seeing that a swallow is there in response to experiences in the very same range as those to which our judgements about swallows are responsive. It is not a matter of establishing inductively that we have reliable judgement-forming procedures, but simply of deploying appropriate recognitional capacities.

There is no space here to develop the account of recognitional capacities in any detail. I merely float the idea that making recognitional capacities central to a theory of perceptual knowledge would fit with the constraint of explaining how we can readily tell when we know by looking that something is so.

If I am right in thinking that paradigm cases of knowledge are instances of the various forms of perceptual knowledge then we should not be too surprised if we have less empirical knowledge than we usually take ourselves to have. If I left my car in the car park, and have not seen it since, then perhaps I don't know that it is there. But then for everyday practical purposes I don't need to know it is there. I take it for granted that it will be, and this will help steer me to the car when I leave, but I am unlikely to be in a situation in which discriminating between whether I know it is there or simply have a reasonable belief that it is there is called for. This raises a nice question about the relation between knowledge and assertion. I would certainly be prepared to tell someone without qualification where in the car park the car is, even if I have not seen it since I left it there. Should we then reject the idea that in telling I give it to be understood that I know? Or should we retain that idea, but accept that in very many situations in which we give it to be understood that we know we don't, not just because of bad luck in those cases but because in giving to be understood that we know, on the grounds in question, we never speak from knowledge?

BIBLIOGRAPHY

- Craig, Edward (1999) *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Clarendon Press).
- Hookway, Christopher (2003) 'How to be a Virtue Epistemologist', in Michael DePaul and Linda Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*
- Hyman, John (1999) 'How Knowledge Works', *Philosophical Quarterly* 49, 433-51.
- Jones, Ward E. (1997) 'Why Do We Value Knowledge?' *American Philosophical Quarterly* 34, 423-39.
- McDowell, John (1995). 'Knowledge and the Internal'. *Philosophy and Phenomenological Research* 55, 877-93. References are to the version in McDowell 1998.
- McDowell, John (1998). *Meaning, Knowledge, and Reality*. Cambridge, MA.: Harvard University Press.
- Unger, Peter (1975) *Ignorance: A Case for Scepticism* (Oxford: Clarendon Press)
- Williamson, Timothy (1995) 'Is Knowledge a State of Mind?', *Mind* 104, 533-65.
- Williamson, Timothy (1997) 'Knowledge as Evidence', *Mind* 106, 717-41.
- Williamson, Timothy (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press.