

CRITICAL STUDY

WHO CARES WHAT YOU KNOW?

BY MARK KAPLAN

Knowledge and its Limits. By TIMOTHY WILLIAMSON. (Oxford UP, 2000. Pp. xi + 340. Price £25.00.)

How important is knowledge? One could be forgiven if from a study of the last four decades of Anglo-American epistemology one concluded that it is not so very important. After all, most epistemologists writing since 1963 have written of knowledge as a derivative notion, definable in terms of belief, truth (perhaps also justification), and some condition tailored to handle Gettier-style cases. One could be forgiven if one concluded that so long as one pays proper attention to the more fundamental notions, knowledge will take care of itself. (If it even needs taking care of at all. There are prominent philosophers who have concluded from their study of the literature that knowledge is of *no* real importance. Why worry about scepticism about knowledge, they have asked, so long as the integrity of justified belief can be sustained?)

In *Knowledge and its Limits*, Timothy Williamson places on its head this way of thinking about knowledge. He argues that knowledge is a distinct mental state that, while involving certain other states (mental states included), is unanalysable in terms of other states. And, he argues, knowledge is of much more fundamental importance than we are wont to think it: knowledge plays an ineliminable role in the causal explanation of human behaviour; only by appeal to knowledge can one say when it is that a belief is justified; a person's evidence is to be identified with the totality of what he knows; knowledge, and only knowledge, warrants assertion.

I find Williamson's case against regarding knowledge as a derivative notion compelling. It seems to me that he is right in arguing (p. 32) that from the discovery that justified true belief is not sufficient for knowledge it does not follow that there must exist some condition which, conjoined to justified true belief, *is* sufficient for knowledge. It seems to me that he is entirely right in arguing (p. 31) that there is

nothing outside pure prejudice that says that knowledge is properly to be regarded as somehow less primitive than the terms in which philosophers have (ever since Gettier) been seeking to define it. It seems to me he is also entirely right to think that there may well be more profit in an approach to knowledge which instead of seeking to analyse it, seeks to say useful things about its import, about how it interacts with other things of interest to us, about how it functions in enquiry.

I am, however, less compelled by his account of how exactly knowledge functions in our intellectual ecology, and in particular by his claim to have found it occupying a fundamental role in the workings of that ecology. My hope is that I can give some sense of the sweep of this remarkable book by undertaking to say why.

Knowledge and the explanation of behaviour

The prevailing wisdom is that appeal to knowledge is not necessary for explaining rational behaviour; behaviour can be fully explained by appeal to full or graded belief (together with preferences and/or desires). Williamson wants to argue that this prevailing wisdom is mistaken: there are contexts in which knowledge figures ineliminably in the explanation of behaviour. At stake here is not just whether knowledge is important in the explanation of behaviour. Also in the balance, Williamson thinks, is the propriety of construing knowledge as a mental state. His worry is that it is at least plausible to suppose that something is a genuine state only if it is efficacious. But if the prevailing wisdom is correct, appeal to knowledge in the explanation of behaviour is completely unnecessary: it is belief that is doing all the work. And if that is right, then (p. 61; here Williamson is quoting Harold Noonan) knowledge 'is best regarded not as a psychological state, but as a complex consisting of a psychological state (belief) plus certain external factors – not because its status as knowledge is causally irrelevant in action explanation, but because it does not have to be cited, as such, in the psychological explanation at all'.

Williamson means to allay the worry by placing before his readers a case in which knowledge figures ineliminably in a causal explanation of behaviour (p. 62).

A burglar spends all night ransacking a house, risking discovery by staying so long. We ask what features of the situation when he entered the house led to that result. A reasonable answer is that he knew that there was a diamond in the house.

Why is his believing (or even believing truly) as he enters the house not an equally good explanation of the burglar's behaviour? Because, Williamson maintains, the probability of his behaving in that way is greater, given that he knows the diamond is there, than it is given that he merely believes (or even believes truly) that it is there. One of the very things that keeps true belief from being sufficient for knowledge (for example, the possibility that the belief is based on a justificatory argument with a false lemma) might have kept him from ransacking the house for as long as he did (suppose he discovered the falsity of the lemma when he looked for the diamond in the place where he thought it was, and gave up his claim to know, and his search). Indeed, for any condition that fails to be necessary and sufficient for knowledge and that one might be tempted to substitute for 'knew' in the explanation above, Williamson thinks one can construct a case 'in which the failure of necessity

or sufficiency makes a causal difference, making the proposed substitute not even causally equivalent to knowing' (p. 63). If knowledge is unanalysable, then the case is hopeless.

What of conditions like stubbornness in belief, or having the very best of evidence, or being certain that the diamond is there? All of them seem to do better. But they are ruled out on other grounds (pp. 63–4): they

cannot replace knowing in all causal-explanatory contexts, for the simple reason that those who know that p often lack a stubborn belief that p . The burglar's beliefs need not be stubborn. Similarly he need not feel certain of them; subjective certainty cannot always replace knowing. The same applies to believing truly on the best evidence, for the example can be so constructed that the burglar's evidence, although good, is not the best possible.... When one works through examples of this kind, it becomes increasingly plausible that knowledge can figure ineliminably in causal explanations.

I anticipate that some readers will share an immediate reaction to the argument, as I (following Williamson) have portrayed it: the burglar's knowing that there was a diamond in the house, they will think, cannot possibly explain the burglar's subsequent behaviour. At the very least he needs, in addition, to have wanted the diamond quite badly, to have been possessed of considerable *sans froid*, to be relatively free of moral scruples about engaging in theft, to be very confident indeed of his ability to find the diamond if it was there. Anyone who did not have those attributes would not have done what he did, even knowing as well as he did that there was a diamond in the house.

But there is nothing here that Williamson would be inclined to deny. His view is that there is one particular aspect of his behaviour that the burglar's knowing that the diamond was there as he entered the house explains better than any other condition one might substitute: the fact that he kept on looking all night. He presumably would not have done this unless his conviction that the diamond was there persisted throughout; presumably he would have given up. Williamson's brief is that his knowing when he entered the house that the diamond was there explains, better than any condition that is not necessary and sufficient for knowing, why his belief persisted. This is because, Williamson thinks, the probability that the burglar's belief would persist was greater, given that he knew when he went in that the diamond was there, than it would have been, given any other condition short of knowledge that one might want to substitute.

Why? Because 'Although knowing is not invulnerable to destruction by later evidence, its nature is to be robust in that respect' (p. 63). It is a view Williamson has occasion to repeat elsewhere. For example: 'Present knowledge is less vulnerable than mere present belief to *rational* undermining by future evidence.... If your cognitive faculties are in good order, the probability of your believing p tomorrow is greater conditional on your knowing p today than on your merely believing p today (that is, believing p truly without knowing p)' (p. 79). And: 'Knowledge is superior to true belief because, being more robust in the face of new evidence, it better facilitates action at a temporal distance' (p. 101). The view is that knowing entails

robustness of belief in the face of future evidence, and that this robustness, which follows from the burglar's knowing the diamond was in the house as he entered it, explains the persistence of his belief. This robustness is something that no condition short of knowledge entails – not even true belief supported by reasoning with no false lemma.

I am unconvinced. It seems to me that there is a condition, neither necessary nor sufficient for knowledge, that will explain equally well the persistence of the burglar's belief. It is, to be fair, not a condition that (so far as I know) one will find discussed in the traditional literature in epistemology. (No surprise here: few, if any, writers apart from Williamson have thought to pay attention to how belief persistence and knowledge might be importantly related.) But the requisite condition has long enjoyed a secure place (albeit in the service of rather different ends) in the Bayesian probabilist literature.

The burglar entered the house in a certain state of confidence that included being confident – confident enough to count as believing – that a diamond was there. Subsequently his state of confidence changed. He became confident, with respect to the first place where he looked for the diamond, that the diamond was not there; likewise for the second, third, fourth, and so on. He became confident that he was not going to find the diamond within the first hour; likewise for the second, the third, the fourth, and so on. Yet, even as his state of confidence changed substantially in these ways (among others), his confidence that there was a diamond in the house did not – or, at least, did not change enough to disqualify him as believing that the diamond was in the house. What feature of the situation as he entered the house explains this? A natural answer is that his state of confidence upon entering the house was marked by the following characteristic: his confidence that a diamond was in the house (p) was *resilient* with respect to the set s of propositions in whose truth he subsequently underwent a change of confidence. That is, upon entering the house, his state of confidence was such that for each subset of s , his confidence in p , given that subset, was pretty close to his unconditional confidence that p – or, in any event, great enough to qualify him as believing p given that subset of s . That is, he entered the house disposed not to have his belief in p extinguished by changes in the confidence he invested in the members of s .

That the burglar's state of confidence was thus marked by resilience is clearly insufficient for his knowing that p . It is entirely compatible with his state of confidence's being thus marked that p is false.

That his state of confidence was thus marked is also arguably unnecessary for his knowing that p . True, there is something to be said for the view that (at the very least) we do not credit a person's claim to knowledge if we think his state of confidence is completely unmarked by resilience – if we think his confidence in the proposition he has claimed to know is set to collapse at the slightest perturbation of his state of confidence. Still, suppose the burglar entered possessed of a state of confidence in which his confidence that p , given that he searched for four hours without finding the diamond, was significantly lower than his initial confidence that p . Suppose he had, after four hours' search, lost his confidence that p . Would we really want to say that this showed he had never known it? I suspect not.

But unlike his being stubborn, certain, or possessed of the very best possible evidence, his entering the house in this state of confidence *is* necessary if his knowing p is to explain the subsequent persistence of his belief. For if he knew p but was *not* disposed to have his belief in p untouched by changes in the confidence he invested in the members of s , we would have to conclude that what explained the persistence of his belief is (at least in part) something (say, an onset of stubbornness) that came over him *after* he entered the house.

Of course, there can be little question but that the burglar's state of confidence's being resilient in this particular way cannot replace his knowing that p in every causal-explanatory context in which his knowing that p might figure. We could simply modify the case so that the set of propositions with respect to which his confidence changes is disjoint from s . The explanatory power of his confidence's being resilient with respect to s is lost.

But I do not see why Williamson thinks it is a burden on his opponents that they must provide a *single* condition that will substitute equally well for knowing p in every causal-explanatory context in which the latter may figure. Perhaps it is because Williamson thinks that if it takes a motley of different conditions to do the explanatory work that knowing p does, then considerations of consilience will make knowing p in every instance a better explanation than the condition offered by way of substitute. But even if this thought about motley conditions is granted (and, I confess, I am not sure it should be), it falls far short of supporting Williamson's way of construing his opponents' burden.

Granted, knowing p will explain what resilient confidence in p with respect to s cannot, as the modified case of the burglar (in which the set of propositions with respect to which his confidence changes is disjoint from s) shows. On the other hand, there will be in this new case some other set s' such that the fact that his confidence in p was resilient with respect to s' will explain the persistence of his belief. And each of the two conditions of resilience will be capable of explaining cases of belief persistence that knowing p cannot – for example, cases in which p is not true.

So the situation is symmetrical. Resilient confidence in p with respect to s cannot substitute equally well for knowing that p in every causal-explanatory context in which the latter might figure. And knowing p cannot substitute equally well for having resilient confidence in p with respect to s in every causal-explanatory context in which the latter might figure. Under these circumstances, I see no grounds for thinking the knowledge explanation better. Indeed, I see a reason for thinking it worse. Appeals to resilient confidence clearly can explain belief persistence in cases in which appeals to knowledge cannot (e.g., cases of false belief). But (as noted earlier) it is hard to see how any appeal to knowledge could explain belief persistence where an appeal to resilience could not. Unless the burglar entered the house in a state of confidence suitably marked by resilience, his knowing p cannot explain the subsequent persistence of his belief.

The moral seems to be that Williamson's case for the importance of knowledge in the explanation of rational behaviour fails. (And, in so far as knowledge's being a mental state requires that it must ineliminably explain at least some behaviour, so too his case for knowledge's being a genuine mental state fails.) He is right to this

extent: the received view as stated by Noonan will not do. Belief alone (even rational belief, true belief, justified belief with no false lemmas) arguably cannot explain everything that knowledge can. But resilient confidence is another story.

Knowledge and justification

Fortunately for Williamson, the rest of his case for the importance of knowledge is not in any way dependent on his contention that knowledge is a mental state that figures ineliminably in the explanation of behaviour. Indeed, the most significant part of his case depends only on the rejection of the doctrine that knowledge is a derivative notion, defined, in part, in terms of justification. Williamson's brief is that 'knowledge, and only knowledge, justifies belief', that 'what justifies belief is *evidence*', and thus that 'knowledge, and only knowledge, constitutes evidence'. He means to defend a principle (he calls it ' $E = K$ ') that 'equates S 's evidence with S 's knowledge, for every individual or community S in any possible situation' (p. 185).

$E = K$ constitutes, Williamson thinks, a significant departure in epistemology, in so far as (p. 185) it

uses the concept of knowledge in partial elucidation of the concepts of evidence and justification. To some people it will therefore seem to get things back to front. For although knowledge is more than justified true belief, many philosophers still expect to use concepts such as *evidence* and *justification* in a more complex explanation of the concept *know*; it would then be circular to use the latter to explain the former. Others prefer to use concepts of a different kind, such as *causation* or *reliability*, to explain the concept *knows*; but even they are likely to regard the concept *knows* as so much in need of explanation that its pre-theoretic use would lack explanatory value.

It is here, Williamson thinks, that $E = K$ makes its departure. 'The order of explanation has been reversed in this book. The concept *knows* is fundamental, the primary implement of epistemological enquiry' (p. 185). Having already 'rejected the programme of understanding knowledge in terms of the justification of belief', he has bought the freedom 'to try the experiment of understanding the justification of belief in terms of knowledge' (pp. 185–6).

Indeed, Williamson suggests (p. 186) that 'standard accounts of justification have failed to deal convincingly with the traditional problem of the regress of justification – what justifies the justifiers – *because* they have forbidden themselves to use the concept of knowledge. $E = K$ suggests a very modest kind of foundationalism, on which all one's knowledge serves as the justification for all one's justified beliefs.'

It seems to me that Williamson is simply wrong in his characterization of the standard accounts of justification. The idea that the regress of justification stops with items of knowledge is a commonplace of the foundationalist tradition in the theory of justification. Chisholm, for example, holds that the regress of justification terminates with knowledge.

But all the same I think Williamson is right to regard his modest foundationalism as a significant departure from standard accounts. While Chisholm holds that the regress of justification ends with items of knowledge, he holds that only certain items will do. Only knowledge of propositions that do not imply the existence of a physical

object (e.g., knowledge of how one is being appeared to, or of what one believes) can constitute a stopping point in the regress of justification. That is, Chisholm holds that something's being an item of knowledge is necessary but insufficient for its playing the role of ultimate justifier.

Williamson's foundationalism is much more permissive (p. 9):

A belief is justified relative to some other beliefs from which it has been derived in some appropriate way (perhaps by deduction), but it is not justified absolutely unless those other beliefs are justified absolutely. Where does the regress end? On the assumption that it ends at evidence, the equation of evidence with knowledge implies that one's belief is justified absolutely if and only if it is justified relative to one's knowledge. The regress of justification ends at knowledge.

Williamson's suggestion appears to be that as soon as it is the case that each of the strands in a regress of justification for a belief has reached a proposition one knows, the justification of that belief is complete: the belief is justified absolutely. Something's being an item of knowledge is not merely necessary for it to play the role of ultimate justifier; it is also sufficient.

But if this *is* his suggestion, then it simply will not do as a solution to the traditional problem of the regress of justification. Or, at least, it will not do so long as we mean by 'the traditional problem of the regress of justification' the problem to which, say, Chisholm, Austin and BonJour offer competing solutions.

For, different as they are in their points of view, Chisholm, Austin and BonJour construe in exactly the same way the question as to whether epistemological foundationalism is correct. A belief is justified if and only if the believer is able to meet (in some, perhaps permissive, sense of 'able') all legitimate epistemic challenges to its propriety. Foundationalism is then the view that meeting an epistemic challenge requires ultimate appeal to beliefs to which there are no legitimate challenges. These foundational beliefs are variously described as beliefs to which an epistemic challenge 'makes no sense' or 'does not arise' or would be 'bizarre'. It is in this sense (the expressions quoted in the last sentence come from Chisholm) that Chisholm is a foundationalist. It is in the same sense that (at least in the early 1980s) BonJour is not: BonJour holds that there is no belief not open to legitimate epistemic challenge. It is in the same sense that Austin is a foundationalist, but of a different sort from Chisholm. Like Chisholm, he holds that, in any context, a regress of justification must terminate in claims that are not open to legitimate challenge. But, unlike Chisholm, he holds that one cannot determine by their content alone which claims will be foundational; context figures inexorably in the decision.

In the context of the debate just sketched, Williamson's modest foundationalism seems quite misguided, as the following dialogue reveals.

'I believe that Harry will be envious of me when I tell him what I am looking at right now in my back yard.'

'Why do you think he'll be envious of you?'

'Because I am right now looking at a goldfinch.'

'Really? How do you know it's a goldfinch and not some other sort of bird?'

'By its peculiar markings. No other bird has those markings.'

To all appearances, this dialogue tells a story in which I justify one of my beliefs (that Harry will be envious of me) by appeal to something I claim to know (that I am looking at a goldfinch), and thereupon justify that claim to knowledge by appeal to something else I claim to know (that the bird has markings peculiar to goldfinches). In so far as I am not a person who just knows a goldfinch when I see one – in so far as I am a person who tells that something is a goldfinch by deliberately attending to the distinguishing features I have learned of from a reputable book – the dialogue also seems to be recapitulating quite properly the regress of justification for my belief that Harry will envy me.

But if this right, then Williamson's modest foundationalism is mistaken. Far from stopping the regress with the first item of knowledge I reached (that it is a goldfinch), I justified that item by appeal to another item of knowledge. By the light of Williamson's doctrine, I should have done no such thing. I should have regarded the challenge to say how I know the bird is a goldfinch as bizarre or nonsensical, in the same way as I would the challenge to say how I know that I feel tired (to use an example that would suit Chisholm), or the challenge to say how I know I am typing right now (to use an example that, while not suiting Chisholm, would suit Austin). In so far as I did not so regard it – in so far as, since I know the bird is a goldfinch only by virtue of having noted its distinguishing features, I am *right* not so to regard it – Williamson's doctrine seems to be refuted.

Oddly enough, Williamson appears to concede the very observation that does the refuting. He considers quite explicitly (p. 9) the complaint that 'non-trivial questions appear to arise about the justification and evidence for much of our knowledge, especially that which is mediated by theory'. He thinks that it is a complaint that his proposal can accommodate (p. 10):

For even if one knows p , one can call that knowledge into question, provisionally treat p as though it did not belong to the body of one's knowledge, and then assess p relative to the rest of one's evidence – one's independent evidence. Non-trivial issues of evidence and justification will then arise for p . This procedure is a good test for some kinds of supposed knowledge, especially those mediated by theory. In such cases, given the manner of knowing p , one knows p only if the rest of one's knowledge justifies p .

I do not see how Williamson can concede all this and yet maintain that his modest foundationalism is correct. (For that matter, I do not see how he can concede that it is sometimes a necessary condition for knowing p that the rest of one's knowledge must justify p , yet say that his is a view that has abandoned the project of understanding knowledge in terms of the justification of belief. It seems that he is telling us here that at least *some* knowledge is to be understood in precisely this way.)

Perhaps it is because Williamson has some different view of what is at issue in the debate over what, if anything, stops the regress of justification. If so, he owes us an account of how else to take it. Otherwise the moral seems to be this. His proposal would, if correct, provide just what he promises. It would provide a clear way to spot the ultimate justifiers of belief, and thereby settle the traditional debate over how to characterize foundational beliefs properly: the foundation consists of everything that

the believer knows. But Williamson's proposal is not correct: the mere fact that I know that p does not suffice to stop a regress of justification. Thus to the question how to decide when one can be said to have arrived at an appropriate stopping point to a regress of justification for a particular belief (a question that, in a disciplined enquiry, amounts to the methodologically crucial question what we are to take as our data), $E = K$ suggests no helpful answer.

Knowledge and evidence

Fortunately, Williamson's case for the importance of $E = K$ does not rest completely on its suggesting that the regress of justification stops with knowledge. It rests in large part on a quite independent matter: how we are to interpret the requirement of total evidence. The requirement of total evidence enjoins us to proportion our beliefs to the support they receive from our total evidence. As Williamson (p. 189) points out, how we apply this injunction depends crucially on what we take to be our total evidence. $E = K$ gives us straightforward and unambiguous marching orders: we are to take as our total evidence the totality of what we know. As Williamson puts it (p. 189), 'one should proportion one's belief in a proposition to the support which it receives from one's knowledge'.

And how are we to gauge how much support a proposition receives from our knowledge? Williamson's answer (p. 209) is: by assessing its probability on the knowledge. 'Given a scientific hypothesis h , we can intelligibly ask: how probable is h on present evidence? We are asking how much the evidence tells for or against the hypothesis.' The upshot seems to be, then, that one should proportion one's belief in a proposition to its probability on one's body of knowledge.

But should one really do so? For any proposition p that you regard yourself as knowing, since p is part of your body of knowledge, your body of knowledge logically entails p . It is a consequence of the axioms of probability that the probability of a proposition, on a set of propositions by which it is logically entailed, is equal to 1. Thus the probability of p on your body of knowledge is equal to 1. By the lights of the foregoing, this is to say that p receives the maximum possible support from your body of knowledge: no evidence could ever tell more in favour of a proposition than your present evidence tells now in favour of p . The requirement of total evidence, as Williamson expresses it, therefore seems to enjoin you to invest the maximum degree of belief in p . You should invest as much belief in p as you do in an obvious tautology.

But do we really want to say that it is a condition of your knowing p that you should invest as much belief in p as you do in an obvious tautology? I think not. Nor does Williamson think so. He means $E = K$ to be entirely compatible with the fact that (p. 209) 'When we give evidence for our theories, the propositions which we cite as evidence are themselves uncertain'. Indeed, the doctrine that knowing p entails being certain that p is one that Williamson consistently repudiates throughout.

Perhaps the appearance that Williamson is none the less committed to the doctrine is the result of misreading 'belief', as it occurs in his formulation of the requirement of total evidence, as something intimately connected to subjective probability. It is a connection Williamson wants to deny. In a discussion earlier in

his book (p. 99, a discussion to which he draws the reader's attention in his chapter on evidential probability), he distinguishes degrees of outright belief from subjective probabilities: 'one's degree of outright belief in p is not in general to be equated with one's subjective probability that p ; one's subjective probability can vary while one's degree of outright belief remains zero'.

But to avoid the doctrine that knowing p entails being certain that p , one needs to effect a more radical distinction between outright belief and subjective probability than Williamson is prepared to make. For, on his view (p. 99), 'one believes p outright when one is willing to use p as a premise in practical reasoning'. How, then, can outright belief come in degrees? By virtue of the fact that 'one may be willing to use p as a premise in practical reasoning only when the stakes are sufficiently low'. His view is this (p. 99): 'Since using p as a premise in practical reasoning is relying on p , we can think of one's degree of outright belief in p as the degree to which one relies on p '.

The trouble here is that, however independently they may operate in other regions of their respective scales, the two measures of belief under discussion seem bound to agree on how they assign their respective maxima. To assign maximum subjective probability to p is to be willing to bet anything on p for even the most modest prize. What Williamson appears to have told us is that to have the maximum degree of outright belief in p is to be willing to rely on p – to use it as a premise in practical reasoning – no matter what the stakes. And these come to exactly the same thing: the very certainty that Williamson (rightly) wants to deny is a requirement for knowing. As he puts it (p. 86), 'I know many things without being prepared to bet my house on them'.

Can a more radical distinction between degree of outright belief and subjective probability – radical enough to avoid making certainty that p a necessary condition of knowing p – be plausibly established? Maybe it can, maybe it cannot. For my purposes here, it suffices to note that Williamson's case for knowledge's being important because of its role in the requirement of total evidence remains unmade until we see whether such a distinction *can* plausibly be made – and, if so, what turns on how we proportion our degrees of outright belief.

Likewise for Williamson's case for knowledge's being important to assessment of evidential probability. For suppose that so long as we interpret the requirement of total evidence as $E = K$ requires, there is no way, short of making certainty a necessary condition for knowledge, to construe outright belief so that it is even the least bit important that we proportion our outright beliefs to their probability given the evidence. Then, in so far as we continue to think (as Williamson does) that the requirement of total evidence is an important requirement, the conclusion seems inescapable: we have to think of the probability a hypothesis has, given the total evidence, in some way other than the one imposed upon us by $E = K$.

It is not as if there are no alternatives. We could, for example, construe your current evidence as consisting of just those propositions that are currently foundational by Austin's lights – just those propositions that (in the current context) are available to for you to cite in response to a legitimate challenge to some other claim, but are not themselves open to legitimate challenge. We could then construe

the probability of h , given your evidence, as the subjective probability it is appropriate to assign to h , given the subjective probabilities it is appropriate to assign to the propositions that constitute your current evidence. The principle of total evidence would then simply enjoin you to tailor your subjective probability assignment to that proper part of it that assigns probabilities to the propositions that make up your current evidence.

Here is a rendering of the requirement of total evidence (and the probability of a hypothesis given the current evidence) that makes no reference to knowledge, and that does not require that to constitute part of your evidence, a hypothesis must be in any sense maximally probable. It is not the only rendering of this sort available. But it should do to make vivid the possibility that it is no accident that Williamson encounters difficulty in finding a central role for knowledge in the story of how our opinions should be responsive to our evidence: knowledge *has* no such role to play.

The importance of knowledge

Suppose, then, that we do not need recourse to talk of knowledge to explain behaviour; that appealing to knowledge will not help us determine when a regress of justification ought to terminate; that knowledge has no role to play in spelling out the requirement of total evidence. Should we conclude that knowledge is unimportant?

In my view, no. It is entirely compatible with all this that it is a matter of genuine importance to us to determine what we know, that what we know determines, in an important sense, what we are to rely on in enquiry and decision. All we should conclude is that the entitlement (and obligation) to rely on p that is involved in knowing that p has to be understood in a way more nuanced than that in which Williamson understands it. It cannot, for example, involve an entitlement (or obligation) to rely on p 's being true for the purposes of determining what subjective probabilities are appropriate to adopt. (As I have shown, that would make certainty that p a necessary condition of knowing that p .) But that does not mean that we cannot make sense of why, even after we have decided what subjective probabilities to adopt, we would take an interest in what we know.

Suppose, as seems plausible, that knowing p in a given circumstance involves entitlement and obligation to act in that circumstance as if p is true, and, in that circumstance, to stop seeking further evidence as to p 's truth. It is perfectly understandable that we should take an interest in finding out (by considering how good a case we can make for p , what decision problems are on our plate, and how high the stakes are that they involve) whether we are thus entitled and obliged. True enough, the story about how we indulge that interest will be one on which what we know is, in a significant way, epistemologically less fundamental than what subjective probabilities we adopt. It will be a story on which it is (in part) by consulting our subjective probabilities that we determine what we know, but on which knowledge plays no role in determining what our subjective probabilities should be. But then one cannot expect everything important to be of *fundamental* importance.

In any event, whatever we choose to conclude about the importance of knowledge, there is no denying the importance of the book under review. *Knowledge and its Limits* brings together concerns in the philosophy of mind, epistemology and the philosophy of science in a way no book has ever done before. There are intriguing arguments, observations and insights on almost every page. It is a book that will challenge, and change, almost any philosopher who takes it on. And no one seriously interested in epistemology can afford *not* to take it on. There can be no question: *Knowledge and its Limits* is one of the most significant works of epistemology we have seen in many years.

Indiana University