

CRITICAL NOTICE OF *KNOWLEDGE AND ITS LIMITS*  
BY TIMOTHY WILLIAMSON<sup>1</sup>

Frank Jackson

What did we learn from the Gettier incident and its fall-out?<sup>2</sup> A few say that we learnt nothing. Knowledge is true justified belief despite Gettier's and like examples. Some say that we learnt that knowledge is not true justified belief and we must seek a fourth condition to add to truth, belief and justification. Perhaps they have a candidate ready to hand. Some say that we learnt that knowledge is not true justified belief but that Gettier's examples mislead us. We should not look for a fourth condition but should rather look for a suitable replacement for the justification condition. Some say that we learnt as much from the fall-out as we did from Gettier's examples. His examples prompted a search for an analysis of knowledge that can only be described as a failure. This tells us that there is no unitary concept of knowledge but rather a series of interlocking and to some extent vague notions, each of interest and value in its own way. Finally, some agree that we learnt as much from the fall-out as from the examples but urge that what we learnt is that our single concept of knowledge is unanalysable.

Williamson belongs to the last group. He combines the thesis that knowledge is unanalysable with four other theses. The first is that knowledge is a key concept in epistemology. Many have been attracted to the position that true justified belief is good enough for many purposes. If I have true justified belief about where the beer is, how to cure AIDS, or my not being a brain in a vat, who cares whether or not I have knowledge? This is why some view the Gettier industry with bemusement. Williamson argues, however, that knowledge does important theoretical and explanatory jobs that cannot be done by true justified belief or variants thereon. He argues, that is, for an indispensability thesis. Secondly, he argues that we should rethink knowledge's position in epistemological hierarchies. There is a long tradition of treating knowledge as something we end up with: the theory of knowledge is the theory of how to put more basic ingredients together in the right way to reach knowledge. Broadly speaking, Williamson thinks we should work in the opposite direction. Evidence does not produce knowledge; our evidence *is* what we know; as he puts it,  $E = K$ . The symbolism is a little misleading. Not all knowledge is evidence obviously; but on his view every item of evidence is something we know, sometimes but not always knowledge that things seem thus and so. This is how justified belief can come from knowledge. Moreover, we should understand belief in terms of knowledge. Instead of viewing belief as a state that essentially aims at truth, we should view it as essentially aiming at knowledge (and, incidentally, desire at action rather than at

<sup>1</sup> Oxford: Oxford University Press, 2000, xi + 340, £25 hb.

<sup>2</sup> E. Gettier, 'Is Justified True Belief Knowledge?', *Analysis* 23 (1963), 121–3. The fall-out defies cataloguing but note the subtitle of Robert K. Shope, *The Analysis of Knowing: A decade of research* (Princeton, NJ: Princeton University Press, 1983), and it didn't stop then.

satisfaction). Thirdly, we should rethink the place of knowledge in the philosophy of mind. Knowledge is just as much a mental state as is belief (or pain). It is of course a broad state but, post Tyler Burge and Hilary Putnam, that should not be seen as a point against its status as mental. In particular, it is wrong to think of knowledge as a conjunctive state with a belief part as the mental part strictly speaking. Williamson takes a similar position on a number of factive mental states. For example, he rejects the common view that remembering that *P* and seeing that *P* should be thought of as conjunctions of that *P* with mental states that do not imply that *P*.

Finally, Williamson's view that knowledge is unanalysable is part of a general scepticism about conceptual analysis. I take it that everyone thinks (should think) that not every concept is analysable. There are fundamental concepts. I take it that everyone thinks (should think) that some concepts are analysable, though they may prefer to use the term 'paraphrase'. Mathematics and formal logic contain non-controversial examples of illuminating analyses (paraphrases). The live issue is whether we should, as a rule, seek analyses or, as a rule, hold that most of the concepts we employ are unanalysable. Williamson makes it clear that his view is that we should do the second ('Most words express indefinable concepts, . . . ' p. 100; ' . . . attempts to provide non-circular necessary and sufficient conditions for ordinary concepts have a miserable record of failure' p. 77.)

In sum, this book is radical and challenging. Although I will express some misgivings, the book is without question an important exercise of the 'let me show a new way of looking at things' kind; something we sorely need in epistemology. Quite a bit has appeared in journal articles but the book brings it all together. Williamson says in the Preface that at one stage he envisaged a collection of previously published papers. It is good that he changed his mind.

I will say something about his claim that knowledge is a mental state, and something about difficulties that arise from holding that knowledge is unanalysable in the strong sense in which he holds the doctrine. There is a cost in going down the *sui generis* route. There is a lot I won't discuss at all.

## I.

Williamson argues at length that knowledge is a mental state and regards it as a pivotal, highly distinctive and potentially controversial part of his overall position. But there are two ways of reading his claim. On one, his view is more standard than it might seem. Many hold that knowledge is a special kind of belief. Persons' states of knowledge are a proper subset of their states of belief; namely, those which satisfy the condition of being true and *C*, where the nature of *C* is highly controversial. For these philosophers, knowledge is a mental state for the reason that belief is. Of course, these philosophers do not think that knowledge is a mental state in the sense that being in it implies nothing about the non-mental world. Knowing that diamonds exist implies that they exist, and diamonds are non-mental. But no-one, including Williamson, thinks that knowledge is a mental state in that sense. That's not the sense at issue.

Now Williamson is explicit that knowledge entails truth, and accepts (in a slightly fence sitting way, see p. 42) the view that knowledge entails belief. He is explicit that he opposes conjunctive analyses of knowledge with belief as one of the conjuncts. But such a position is compatible with holding that knowledge is a special kind of true belief. Take

the example of colour, an example he mentions a number of times to support his denial of conjunctive treatments of knowledge (and other factive mental states). He points out that being red implies being coloured but that it is impossible to analyse being red as a conjunction of being coloured with something else. Impossible to do in an interesting way that is; it is not interesting to say that  $X$  is red iff  $X$  is red and coloured. All the same, being red is a special case of being coloured. In the same way, Williamson could hold that knowledge is true belief that meets condition  $C$  consistently with his denial of analysability, provided he insisted that condition  $C$  has no illuminating non-circular analysis. Holding that knowing is a kind of believing does not commit us to holding that knowing is analysable in terms of believing.

The reading of his claim that knowledge is a mental state that does make his view highly distinctive interprets him as holding that although, necessarily, if someone knows that  $P$ , they truly believe that  $P$ , their state of knowledge is a distinct state from their state of belief. Of course ambiguity abounds hereabouts. There is a sense in which true belief is a different state from knowledge because the conditions for being in the former differ from those for being in the latter. But it is not a different state in the sense that a subject who believes  $P$  and truly believes  $P$  is in two states that might differ in causal powers, say.

I think it would have been good if Williamson had distinguished explicitly these two readings. On the first, we should agree with him that knowledge is a mental state, whereas the second raises a series of difficult issues. If there are two states, which one causally explains a subject's actions when they are in both? Suppose I know that there is a tiger outside my door. I will also believe that there is (necessarily according to Williamson and I agree, but those who dissent should take the case to be one where in fact I both know and believe). Which state combines with my desire to go on living to stop me from opening the door? Or is there a degree of over-determination? Or is it a joint effort? Or does one pre-empt the other, and if so which does the pre-empting, or does it vary case by case? And if there are two states, how are they related to inputs from the environment? If the growl of the tiger is responsible for my believing that there is a tiger outside the door, is it also responsible for my knowing this? Or does one state cause the other?

Some short remarks on page 80 suggest the 'two state' interpretation. There Williamson says, ' . . . the mental state of knowing makes a distinctive contribution to the causal explanation of action' [distinctive by contrast with that made by belief]. But how the distinct contributions from knowledge and from belief on the two state interpretation might fit together is left obscure, and it may be that he meant to say not so much that the states as such make distinctive contributions but that the *information* that someone knows as opposed to the *information* that they believe plays a distinctive role in causal explanations of action. The latter would be compatible with the less controversial reading of his view that knowledge is a mental state.

## II.

There is a weak and a strong sense in which you might hold that knowledge is unanalysable. On the weak sense, the view is, as we put it above, that there is no unitary concept of knowledge but rather a series of interlocking and to some extent vague notions. Each individual notion has its analysis but there is no illuminating formula that covers the lot. The Gettier literature has, on this view, been an exercise in displaying the variety of

notions and its failure to induce consensus gives us defeasible reason to hold that there is no unitary notion underlying our folk use of 'knowledge'. This is a weak sense because it does not require us to enlarge our stock of unanalysable notions. Knowledge on this view is not an unanalysable extra over and above the discerned variety. Williamson's view, however, is that knowledge is unanalysable in the strong sense. There is a single notion, it is unanalysable, and it is an addition to our stock of unanalysable notions.

The cost associated with going down the *sui generis* route is that it is hard to respond to a certain kind of request for why we should hold that we ever know anything. I do not mean here the kind of unduly demanding requests we are only too familiar with from traditional sceptics. I mean the kind of request we might get from someone who is agreeably anti-sceptical about a great deal. They agree that we have true justified belief; they agree that we have true beliefs reached by reliable processes; they agree that we have true justified beliefs whose truth and justification is not an accident; they agree that we have true justified beliefs with desirable anti-defeasibility properties; they agree that we have true beliefs in situations where the possibility that they might be false has been excluded; and so on. They agree, in short, that we have on occasion true beliefs that satisfy all the kinds of additional constraints that have been suggested, in one form or another, as what makes true belief into knowledge by the many who have sought an analysis of knowledge. Their worry is that they do not see why they should, in addition to believing all of the above, believe that there is knowledge. The analysers of knowledge have an answer. They are not asking our enquirer to accept anything extra: the list, or some similar list, automatically includes knowledge somewhere or other, though not under the name of knowledge.

Williamson does not confront this challenge as such but we can speculate that he would insist that knowledge has an explanatory and predictive value that belief lacks. At many points in the book he argues for the explanatory and predictive virtues of knowledge over belief in respect of action for one or another example. If he is right, he might have a theoretical justification for positing knowledge: it earns its keep by playing important explanatory and predictive roles that cannot be handed across to the varieties of belief. However, the examples he gives are not especially compelling. Here is the one he gives on pages 86–7.

How long would we expect a fox to be willing to search for a rabbit in the wood before giving up, assuming initially (a) that the fox knows that there is a rabbit in the wood, or (b) that the fox believes truly that there is a rabbit in the wood? In (b) but not (a), the fox's initial true belief may fail to constitute knowledge because the true belief is essentially based on a false one, for instance, a false belief that there is a rabbit in a certain hole in the wood. When the fox discovers the falsity of that belief, the reason for the search disappears. That will not happen in (a), because a true belief essentially based on a false one does not constitute knowledge. Thus, given plausible background conditions, more persistence is to be expected in (a) than in (b). In many such cases, lengthy persistence is better explained by initial knowledge than by initial true belief.

The trouble with this example is that Williamson tells us the problem with the explanation in terms of (a) *without mentioning knowledge*. The problem is that 'the true belief is essentially based on a false one'. In consequence, the case gives no reason for

favouring a knowledge story over a belief one with the defect remedied in terms that make no mention of knowledge and which Williamson himself provides. In terms borrowed from his final sentence: in many such cases, lengthy persistence is better explained by initial true belief that is not essentially based on a false one than by initial true belief *simpliciter*. The same is true for the example on page 86.

You see someone coming to your door; he is about to knock loudly. You are tempted not to reply. How would he react? You ask yourself, ‘Does he know that I am in?’ not, ‘Does he believe that I am in?’ If before knocking he does know that you are in, then he is unlikely to abandon his belief if you fail to reply; he will probably take offence. If before knocking he believes truly without knowing that you are in, then he is much more likely to abandon his belief if you fail to reply; he will probably not take offence. . . . Whether he would take offence is better predicted by whether he knows than by whether he believes.

As before, Williamson tells us the key point in a way that does not require the mention of knowledge: it is how likely or unlikely ‘he is to abandon his belief if you fail to reply’. To borrow from his last sentence, whether he would take offence is better predicted by whether he believes resiliently than by whether he believes *simpliciter*.

Here is a final example from pages 101–2.

Some hunters see a deer disappear behind a rock. They believe truly that it is behind the rock. To complete their kill, they must maintain a true belief about the location of the deer for several minutes. . . . If the hunters know that the deer is behind the rock, they have the kind of sensitivity to its location that makes them more likely to have future true beliefs about its location than they are if they merely believe truly that it is behind the rock. If we are to explain why they later succeeded in killing the deer, given the foregoing situation, then it is more relevant that they know that the deer is behind the rock than that they believe truly that it is behind the rock.

But, as before, Williamson tells us how to identify the key point without mentioning knowledge. The example illustrates the value to the hunters of their being sensitive to the location of the deer, which they will be if their *beliefs* are sensitive to the location of the deer.

At a number of points in the book, Williamson points to the explanatory value of finding the right generalization to cite. ‘Good explanations have an appropriate generality’, as he says on page 81. He points out that though it might be true that someone died because they were run over by a red bus, if the colour was irrelevant to the death, it would be better simply to say that they died because they were run over by a bus. This suggests that he might respond to our discussion of the above examples by granting that in each example taken individually there is no need to mention knowledge to account for the difference in explanatory or predictive value, but we need to mention knowledge to capture the generality across all the cases. There are, however, two possibilities to distinguish. One is that there is a projectible pattern discernible in the cases when described in the terms that make no mention of knowledge. But then we can use this pattern—the one we can capture without mention of knowledge as such—to generalize

appropriately. (We may want to go on and suggest that it gives us the long sought after analysis of knowledge but that's another question.) The other possibility is where there is no relevant pattern. In this case it would be wrong to generalise. Using the word 'knowledge' does not create patterns where none exist. In the first case, we do not need knowledge to give an explanation with suitable generality. In the second case, there is no relevant pattern to underpin the generalisation.

I conclude that those who hold that knowledge is strongly unanalysable have a serious problem explaining how they can avoid scepticism about knowledge. Let  $K_1, \dots, K_n$  be the sum total of all the sensible suggestions that have been or might be put forward as analyses of knowledge. Let  $\text{knowledge}_i$  be knowledge analysed according to  $K_i$ . We can be confident that we sometimes have  $\text{knowledge}_i$  for a number of the  $i$ 's, but what reason can the denier of the analysability of knowledge give for saying that, moreover, we sometimes have *knowledge*?