

On Williamson's Arguments that Knowledge is a Mental State

Is knowledge a mental state? For philosophers working within the idealistic tradition, the answer is trivial: there is nothing else for knowledge to be. For most others, however, the claim has seemed *prima facie* implausible. Knowing that *p* requires or involves the fact that *p*, or *p*'s truth, and that – with certain specifiable exceptions – is quite independent of my (or anyone's) mind; so while knowledge may require or involve certain mental states, it itself is not a state of mind.

More generally, it is very natural or intuitive to think in the following terms. On the one hand, there is the world apart from my mind. On the other hand, there is my mind. In many cases in which I have knowledge, I have it because of something about how the world is apart from my mind and because of something about me (my mind) which could be as it is even if the world were not that way. For instance, consider the well-known example of Henry who is driving down the road and observes a barn in a field. In the ordinary case, he thereby comes to know that there is a barn in the field. But he does not come to know this in an unusual case in which, unbeknownst to him, there are barn facades in the vicinity which are not visually discriminable from real barns when viewed from the road.¹ Intuitively, the difference between the two cases isn't a difference

regarding Henry's mental states. There is something about him or his mind which remains the same in the two cases; what changes is the situation in the world apart from his mind. As the traditional conception of knowledge has it, then, one's knowing a particular proposition about the world is often a complex state or condition comprising both purely mental factors (such as belief) and non-mental, environmental factors. The term "mental state" can, of course, function as a term of art and may reasonably be extended in certain ways for theoretical purposes. But we have a pretheoretical conception of the mental which prompts ready agreement with the traditional view. The claim that knowledge is nothing but a mental state comes as a surprising philosophical thesis.

Timothy Williamson has recently argued forcefully in favor of this thesis.² As he sees it, its merits are many (in addition to its supposed truth). It prompts a certain reorientation of epistemology in relation to the philosophy of mind: epistemology becomes a branch of philosophy of mind, and philosophy of mind and action are freed to make explanatory use of the notion of knowledge. Moreover, this thesis enables the solution of several epistemological problems. For instance, its truth would undercut one familiar form of skeptical argumentation. Skeptical arguments sometimes involve the following premise: Even in the best possible case for knowing something about the world, your total mental state is *neutral* between knowing and not knowing; it could be just as it is and yet because of conditions external to your mind you might not have

knowledge of the world (p. 6). If knowledge itself is a mental state, however, then this premise is obviously false.

Williamson's primary strategy against the traditional conception is (1) to demonstrate the theoretical resources gained by rejecting it and (2) to defang a number of *prima facie* objections to the claim that knowledge is a mental state (such as objections arising from the supposed "transparency" of the mental). However, three considerations against the traditional conception also figure in his discussion: the non-analyzability of the concept of knowledge, the "primeness" of knowledge, and the (alleged) inability to satisfactorily specify the "internal" element involved in knowledge. It is consequently worthwhile to consider whether a convincing argument for Williamson's view can be derived from these considerations. As I will argue, it cannot. To that extent, then, Williamson has not yet made his case.

1. The non-analyzability of the concept of knowledge

It is a relatively uncontroversial claim that no satisfactory non-circular statement of necessary and sufficient conditions for the application of the concept *knowledge* has ever been provided. Moreover, as Williamson rightly stresses, there is no particular reason to expect that one will be (pp. 30ff). Few concepts that are not introduced through explicit definition yield such analyses.

What should we conclude from such considerations? Williamson's suggestion seems to be this: our concept of knowledge "cannot be analysed into concepts of narrow and environmental conditions" (p. 91), so our concept of knowledge is not a concept of a complex state which comprises factors in the person's mind as well as factors in the world outside of the person's mind. Williamson is not explicit about how the argument from premise to conclusion might go, but the following line of thought is suggested by his discussion (pp. 2 – 5, 27 – 33). Suppose that we conceived of knowledge as a complex state involving both purely mental factors and non-mental, environmental factors. Then it should be possible to decompose our concept of knowledge into more basic concepts in such a way that anything that satisfies certain conditions stated only in terms of the more basic concepts would be an instance of knowledge and nothing that fails to satisfy them would be an instance of knowledge. But we can't (or at least it seems quite likely that we can't) decompose our concept of knowledge in this way, so (it seems quite likely that) our concept of knowledge is not a concept of a complex state understood in the traditional way, as involving both purely mental factors and non-mental, environmental factors. So, (it seems quite likely that) knowledge is not a complex state of this sort.

The conclusion does not follow, however. Even if no satisfactory analysis is forthcoming, it doesn't follow that knowledge is not (often) a complex state involving both mental and environmental factors. The obstacle is not the potential difficulty, noted

by Williamson (pp. 28-9), in drawing conclusions about the nature of knowledge from considerations about the concept of knowledge. Rather, it is that even if non-circular necessary and sufficient conditions for knowledge cannot be provided, this would not show anything about whether we conceive of knowledge as (often) being a complex state involving both purely mental factors and non-mental, environmental factors.

The crucial issue here concerns the source of the failure of the project of conceptual analysis. For this failure can be explained in ways which are compatible with the traditional conception of knowledge as often being a complex state involving both mental and non-mental factors. For instance, suppose (as seems plausible) that the idea that concepts are individuated by necessary and sufficient conditions for their application is, for the most part, incorrect: in many cases our concepts are rather “family resemblance” or “open texture” concepts for which not even disjunctive necessary and sufficient conditions can be given. If that were so, then it might still be the case that in a good many paradigmatic cases knowledge is, and we conceive of it as, a complex state involving features of a person’s mind as well as circumstances in the non-mental world. Likewise, suppose that the concept of knowledge involves features of a person’s mind and circumstances in the non-mental world coordinating in just the right sort of way, but we lack the resources to provide a non-circular, general specification of that way which satisfactorily covers all the cases: the best we can do is to say, “It’s the way that yields knowledge.” That possibility, too, would explain our inability to analyze the concept of

knowledge while allowing for the correctness of the traditional conception. These are just two examples. The point here is a general one. In general, it is implausible to claim that if a concept of a certain state or condition cannot be provided with informative necessary and sufficient conditions for its application, then we do not conceive of that state or condition as a complex involving different kinds of factors. There is no reason to think that knowledge would be exceptional in this regard.

It might be suggested in response that the claim that knowledge is a purely mental state provides the *best explanation* for the sorry history of attempts to analyze knowledge. But this idea is dubious. There is no particular reason to think that if knowledge were a mental state, then it would be less likely that the concept of knowledge would succumb to philosophical analysis. If one did not have some other reason to expect the project of analysis to fail, then if knowledge were a mental state, one would (*prima facie*) expect the concept of knowledge to be analyzable in purely mental terms.³

2. *Primeness*

Primeness is the thesis that having knowledge is not a composite state or condition (p. 66). A state or condition is composite just if it is a conjunction of a narrow (internal) condition and an environmental condition (p. 66). Williamson characterizes the internal as what is part of the “total internal physical state of the agent” (p. 66), but this is inessential to the argument. The thesis of *primeness* is just the thesis that possessing

knowledge cannot be a complex state involving purely mental factors and purely environmental factors, on any plausible understanding of the terms “mental” and “environmental”. Thus, the argument for primeness is aimed at showing that knowledge is a mental state by showing “on structural grounds that the envisaged separation of internal and external factors is impossible” (p. 64). In fact, however, the argument does not show this.

The argument for primeness goes like this. Clearly define what is supposed to be involved on the internal side and on the external side. Then, “we can show that [a condition C] is prime simply by exhibiting three cases A, B, and G, where G is internally like A and externally like B, and C obtains in A and B but not in G” (p. 68). Williamson accordingly describes three cases which meet this description: in two of them one has knowledge, but in the third, which combines the internal elements of one case with the external elements of the other, one does not have knowledge (p. 72). Let’s grant that this example succeeds. What does it show? Just this: you can take the internal components from one instance of knowledge, add them to the external elements of another, and get something which isn’t a case of knowledge. Or, to put the lesson fully generally: Take any informative general characterization of the purely mental conditions in cases A and B which specifies those conditions independently of the environmental conditions, and take any informative general characterization of the environmental conditions in cases A and B which specifies them independently of the person’s mental state(s); there will always

be some case in which the mental and environmental conditions meet these characterizations but which isn't a case of knowledge. But that doesn't tell us anything more than non-analyzability did about whether any particular case of knowing is a complex state composed of purely mental and environmental factors.

To see the problem here, consider another example: the state or condition of *driving lawfully*. In any particular case, this state or condition comprises two factors: how one is driving, and how the law states one should drive. Now here are three cases of just the sort Williamson describes.

Case	Driving	Law
A (America)	on the right	drive on the right
B (Britain)	on the left	drive on the left
G	on the right	drive on the left

Cases A and B are both cases of driving lawfully. Case G is not. So the general state or condition picked out by the term 'driving lawfully' is prime in the following sense: it cannot be factored into two components which are (a) independently specified and (b) also characterized in such a way that any third case meeting those characterizations will also be a case of driving lawfully. But it would be a fallacy to conclude from this that driving lawfully is not a complex state or condition involving two sorts of factors. In any particular case, driving lawfully consists of an appropriate combination of factors having

to do with how one is driving and factors having to do with what the law states about where one should drive.⁴

We can grant, then, that Williamson has shown that knowing can't be factored into informative general characterizations of two independently-specified components such that any third case meeting those specifications will also be a case of knowledge. It still wouldn't follow that when one knows that *p*, this is not a complex state or condition which comprises mental factors and non-mental, environmental conditions. So it would not follow that knowing is a purely mental state. It wouldn't even follow that *in any particular case* one could not cleanly separate and describe the "internal" and "external" factors.

In the end, Williamson in effect grants as much by noting that a prime condition "may be a truth-function or some subtler function of narrow and environmental conditions [such as an infinite disjunction of conjunctions of narrow and environmental conditions]" (p. 91). He claims that though "the arguments for primeness ... do not show that our [concept of knowledge] cannot be analyzed into concepts of narrow and environmental conditions," this was to be shown by the earlier argument from non-analyzability (p. 91). But showing that the concept of knowledge can't be analyzed in this way isn't enough to yield the conclusion that knowledge is a mental state. As we've seen, the non-analyzability of the concept of knowledge would not show, any more than does the argument for primeness, that when a person knows a particular proposition, the

person is not in a complex state comprising both mental and non-mental, environmental factors.

3. *Specifying the Internal*

Williamson offers a challenge to the traditionalist. “Factive constructions are held [by the traditionalist] to characterize the subject by reference to a mixture of genuinely mental states and conditions on the external environment. ... the challenge to the internalist is to make good this claim by isolating a level of description that is both narrow and genuinely mental” (p. 54), that is, by saying “what mental state knowing adds nothing mental to” (p. 57). Williamson argues that this challenge cannot be met. His argument is couched in terms of a particular conception of the mental as what is “internal” to one’s body, but again this is quite inessential to the argument. If successful, the argument would equally apply if one held a thorough-going dualist view.

It is worth pausing to ask whether Williamson’s challenge is really fair. The traditional claim that knowledge is not a mental state is just the claim that it is (often) a complex state involving environmental as well as mental conditions. In order to maintain that claim, why must one be able to specify a mental state, present in every case of knowing, to which knowing adds nothing mental? Granted, it would be nice to be able to do so. But given that knowledge manifestly involves mental states, it would seem that in order to “make good” on the traditional claim it is sufficient to state the *non-mental*

component of some particular case of knowing, for instance, the environmental conditions that p is so, and that (as it may be in the particular case) there aren't any barn facades in the vicinity.⁵

It might be objected that specifying such conditions would do nothing to vindicate the traditionalist's position, since Williamson's view is perfectly compatible with the claim that knowledge has non-mental necessary conditions. Externalist views of content have familiarized us with the idea that mental states, such as believing that there are tigers nearby, can have non-mental necessary conditions. Williamson proposes to extend this idea from the *content* of the state to the *attitude* itself. Merely stating some non-mental conditions which are somehow "involved" in a particular case of knowing does not tell against this proposal in any way.⁶

However, this objection — though true enough — misinterprets the dialectic at this point. Williamson's charge isn't that the traditionalist can't provide a good argument against his proposal that knowledge is a mental state; his charge is rather that the traditionalist cannot fulfill a demand imposed *by the traditional view itself*. The question of fairness, then, amounts to the question of whether the traditional view actually imposes that demand — whether it stands or falls with our ability to provide an illuminating characterization of "what mental state knowing adds nothing mental to". It is doubtful that it does.

Still, Williamson does not show that the challenge cannot be met. The challenge is to capture the idea that for each case of knowing, some purely mental state which is not yet knowledge is involved. It is natural to try to characterize this state in terms of belief. After exploring several abortive attempts, Williamson suggests the following strategy: try to capture the idea that the difference between belief and knowing is not (always) mental by appealing to the thought that *not* knowing *p* (sometimes) adds nothing mental to believing *p* (p. 59). Williamson suggests that this idea be characterized as follows:

“(8) For all propositions *p* and cases *A*, if in *A* one believes *p* then in some case *B* one is in exactly the same mental state as in *A* and one does not know *p*” (p. 59).

As Williamson points out, on the assumption that knowing *p* is compatible with believing *p*, (8) implies that knowing *p* is not a mental state (p. 59).

(8) has some initial plausibility. Right now I believe that George W. Bush is the U.S. President. I presumably also know it. But suppose that unbeknownst to me, George W. Bush died of a heart attack just minutes ago. Then I don't know that George W. Bush is the U.S. President. But it seems as though my mental state would be just the same either way. The difference between the two cases is not mental, but rather environmental – it concerns a certain man thousands of miles away from me.

Williamson argues that despite its surface plausibility, (8) is false. It fails for cases in which *p* concerns the subject's own mental state or a necessary (mathematical) truth (p. 59-60). Fortunately for the traditionalist, however, this observation does not

save Williamson's view. The traditionalist does not need to accept (8), which generalizes over all propositions; it is sufficient for the traditionalist's purpose that a version of (8) be true for *some* proposition or propositions. Consider, then, a version of (8) in which the relevant class of propositions is the class of propositions concerning contingent matters in the world outside the person's mind. Here, the principle is quite plausible, as is shown by the above example regarding my belief that George W. Bush is President. In fact, the single example is all the traditionalist needs. It is a case where the difference between knowing and believing (in the way that one does) is not mental. It is therefore a case in which knowing is not purely a mental state. So it is not true that knowledge is a purely mental state.

This argument can be made more formally. Let's follow Williamson in using the following thesis (1) to characterize the claim that knowledge is a mental state.

(1) For every proposition p , there is a mental state S such that in every case A , one is in S iff one knows p (p. 54).

As Williamson points out, (1) entails the following:

(2) For all propositions p and cases A and B , if one is in exactly the same mental state in A as in B , then in A one knows p iff in B one knows p .

Now consider the following variant of (8):

(8*) For some proposition p and all cases A , if in A one believes p then in some case B one is in exactly the same mental state as in A and one does not know P .

On the modest assumption that knowing is nomically compatible with believing, which entails that there is some case A in which one knows p ⁷, (8*) is inconsistent with (2). For on this assumption, (8*) entails that for some proposition p , there are two cases A and B in which one is in exactly the same mental state though one has knowledge in only one of them. And that entails the denial of (2), which entails denial of (1), the claim that knowledge is a purely mental state. So if (8*) is true, as it surely is, then (given the above unproblematic assumption) knowledge is not a purely mental state.

4. Conclusion

To establish that knowledge is purely a mental state, we need an argument which shows that in any particular case, one's possession of knowledge is not a complex state or condition comprising mental factors and purely environmental, non-mental factors.

Williamson's arguments fail to show this. So far as I can see, then, we should maintain our pretheoretical conception of knowledge as something other than a purely mental state.⁸

¹ Goldman, Alvin. "Discrimination and Perceptual Knowledge," in *Essays on Knowledge and Justification*, Pappas and Swain, eds., Cornell University Press, 1978, pp. 120 – 145. The example is reported to be from Carl Ginet.

² Williamson, *Knowledge and Its Limits*, Oxford, 2000, esp. chapters 1 – 3.

³ It is plausible to see Williamson's overall argument as running like this: Given the failure of the project of providing non-circular necessary and sufficient conditions for knowledge, the best alternative for philosophical theory-building purposes is to regard knowledge as an unanalyzable mental state. So understood, Williamson's argument isn't any shorter than his book itself. I do not mean to be challenging this line of thought here. My point here is simply that there is no more direct route from the failure of the project of analysis to the conclusion that knowledge is an unanalyzable *purely mental* state. (I am grateful to an anonymous referee for helping me clarify my thinking on this point.)

⁴ I am grateful to [] for help with this example.

⁵ This response is further motivated by reflection upon the failure of attempts to analyze the concept *knowledge*. Perhaps there is no illuminating general characterization of a purely mental factor which is involved in every case of knowing and which "knowing adds nothing mental to" – that is, a state which would suffice, with the addition of only appropriate external factors, for knowledge in every case. Still, if, for some given case, we could specify some particular purely non-mental condition which is involved, wouldn't that be enough to "make good" on the traditionalist's claim? (Of course, such a response won't be satisfying if one is after an analysis. But the claim that knowledge is not a purely mental state does not commit one to the possibility of any familiar sort of analysis, not even a disjunctive one.)

⁶ I am grateful to [] for raising this concern.

⁷ A and B are supposed to range over all possible cases, or at least all nomically possible cases (p. 52). (The argument also depends upon the assumption, presumably unproblematic in this context, that knowledge is possible.)

⁸ An additional argument could perhaps be extracted from Williamson's discussions of the role of knowledge in the explanation of action. Here's how it would go. Knowledge ascriptions can play an important role in action explanation (pp. 60 – 4, 80 ff.): the description "knows p" is sometimes the most relevant one for explanatory purposes (p. 63); moreover, the explanatory appeal to the person's knowledge is sometimes ineliminable, since appealing to anything less or more specific would bring about causal-explanatory loss (pp. 63 – 4); and explanations in terms of knowledge enable generalizations which we would otherwise miss (p. 81ff.). Conclusion: because knowledge ascription plays this sort of role in action explanation, knowledge must be a purely mental state. However, this conclusion does not follow. The explanatory role of knowledge ascriptions forces the conclusion that knowledge is a purely mental state only if whatever plays a role in person-level action explanations must be a mental state, and that thesis is extremely dubious. Uncontroversially non-mental circumstances, too, can sometimes provide the best commonsense explanation of a person's action, and they are sometimes ineliminable. "Why did he move his knight?" "Because it was threatened by

the queen's pawn." "Why did he go to the opera last night? Doesn't he usually hate opera?" "Because it was *Tosca*." Given the questioner's purposes and the background considerations in the particular explanatory context, these explanatory appeals to the non-mental circumstances may get things exactly right. It might be objected that these are not the *proximal* explanations of the person's action, and that the immediate cause of the action must be something "internal" to the person's mind. But even if that is so, merely adverting to some such state would inappropriately weaken the proffered explanation, if it would leave open the possibility that the person was acting on a mistaken impression – which he wasn't – and so would generate incorrect counterfactual implications. And it will not do to suggest that in all such cases what is really doing the explanatory work is the person's *knowledge* of the relevant fact, since there can be cases in which the difference between knowledge and true belief is explanatorily irrelevant or in which the truth of a person's belief plays an explanatory role even though the person lacks knowledge. Consequently, the role of knowledge in person-level action explanation won't show, by itself, that knowledge is a purely mental state rather than a complex state involving environmental conditions. To get that result out of such considerations, you would need to impose dubious constraints on the sorts of things which can play a role in person-level action explanation.