

Can the Concept of Knowledge be Analysed?

Quassim Cassam

1. Overview

In Knowledge and its Limits, Timothy Williamson argues for what I am going to call the Unanalysability Hypothesis (UH), the hypothesis that ‘the concept knows cannot be analysed into more basic concepts’ (p. 33).¹ Williamson’s defence of UH constitutes the negative phase of his discussion. The positive phase consists in an attempt to give a modest positive account of the concept of knowledge that is not an analysis in the traditional sense. There is room for such an account because, as Williamson emphasizes, it does not follow from the fact that a concept cannot be analysed into more basic concepts that no reflective understanding of it is possible.

Williamson puts forward a range of arguments in support of UH. The first is what I am going to call the Distinct Concepts Argument (DCA). This argument assumes that every standard analysis of the concept knows equates it with some conjunctive concept like justified true belief. With this assumption in place, the aim of DCA is to show that every standard analysis of knows is ‘incorrect as a claim of concept identity, for the analysing concept is distinct from the concept to be analysed’ (p. 30). Another argument in support of UH is the Inductive Argument, according to which ‘experience confirms inductively.... that no analysis of the concept knows of the standard kind is correct’ (p. 30). A third argument is the False Expectations Argument, according to which one should not expect the concept knows to have a non-trivial analysis in more basic terms. Few concepts have such analyses, and there is no special reason to expect knows to be one of them.

In the positive phase of his discussion, Williamson argues that ‘knowing is the most general factive stative attitude, that which one has to a proposition if one has any factive stative attitude to it at all’ (p. 34). This leads to the suggestion that ‘if one knows that A, then

there is a specific way in which one knows; one can see or remember or ... that A' (p. 34). 'See' and 'remember' are examples of factive mental states operators (FMSOs). An FMSO is the realization in natural language of a factive stative attitude. In these terms, one principle to which Williamson subscribes is that 'know' is an FMSO. Another is that if Φ is an FMSO, then from 'S Φ s that p' one may infer 'S knows that p'. According to Williamson, these two principles 'characterize the concept of knowing uniquely, up to logical equivalence, in terms of the concept of an FMSO' (p. 39).

My discussion will be divided into three parts. In the first part I will concentrate on DCA. I will also comment briefly on the other arguments for UH but my main aim will be to show that DCA doesn't work. In the second part I will argue that Williamson's positive account of the concept of knowledge amounts to a kind of analysis. Traditional conceptual analysis is 'reductive' whereas Williamson's account is largely non-reductive. Nevertheless, it has some reductive elements, and I will attempt to bring these out. Finally, I will discuss the principle that if one knows that A then there is a specific way in which one knows. I will distinguish between different 'ways of knowing' and suggest that the sense in which seeing that A is a way of knowing that A is very different from the sense in which remembering that A is a way of knowing that A. What distinguishes seeing that A from some other 'ways of knowing' is that it is potentially a way of coming to know that A. The notion of a way of coming to know that A is significant because, as I will argue in conclusion, it holds the key to understanding how one knows that A and what it is for one to know that A.

2. The Distinct Concepts Argument

The Distinct Concepts Argument relies on the notion of a mental concept, so let's start by briefly considering this notion. Although Williamson himself doesn't attempt a formal definition of it, he does argue at one point that the concept true is not mental because 'it makes no reference to a subject' (p. 30). So a concept won't count as mental unless it refers to

a subject. This is obviously a long way from constituting a definition of the notion of a mental concept, but Williamson's idea is presumably that we have an intuitive grasp of what mental concepts are, and that this is enough for the purposes of DCA. Now consider the case of a concept C which is the conjunction of the concepts C_1, \dots, C_n . Williamson's proposal is that ' C is mental if and only if each C_i is mental' (p. 29). On this account, believes truly is not a mental concept of a state since true isn't a mental concept. By the same token, has a justified true belief is not a mental concept. These concepts are not mental because they have 'irredundant non-mental constituents, in particular the concept true' (p. 30).

Having accepted that believes truly and has a justified true belief aren't mental concepts, let's also accept, at least for the same of argument, that knows is a mental concept. What follows from this? What follows straightforwardly is that the concept knows can't be the same concept as the concept believes truly or the concept has a justified true belief. The point is that if C is a mental concept and D is not a mental concept, then they can't be the same concept. But, as Williamson sees things, every standard analysis of the concept of knowledge takes it that this concept is the very same concept as some conjunctive concept like has a justified true belief. So every standard analysis of the concept knows is incorrect.

Crucially, it doesn't matter for the purposes of this argument which particular conjunctive concept the concept of knowledge is equated with, as long as it has the concept true as a constituent. For example, suppose that instead of equating the concept of knowledge with the concept has a justified true belief one equates it with the concept has a reliably caused true belief. Williamson's argument would still go through since it 'it applies to any of the concepts with which the concept knows is equated by conjunctive analyses of the standard kind' (p. 30). As long as the analysing concept is not mental, it can't be the same as the concept being analysed, and this is the crux of DCA.

Here, then, is a breakdown of the main components of the Distinct Concepts Argument:

- (a) Every standard analysis of the concept knows equates it with some conjunctive concept which has the concept true as a non-redundant constituent.
- (b) The concept true is not a mental concept.
- (c) Any concept with a non-redundant non-mental constituent is not a mental concept.
- (d) So the conjunctive concepts with which the concept knows is equated by analyses of the standard kind are not mental concepts.
- (e) The concept knows is a mental concept.
- (f) A mental concept can't be the very same concept as a non-mental concept.
- (g) So the mental concept knows can't be the same concept as any of conjunctive concepts with which it is equated by standard analyses.
- (h) So every standard analysis of the concept knows is incorrect.

The question is whether this argument works, and I have already indicated that I don't think that it does. It's now time to explain why not.

To begin to get a sense of what might be wrong with DCA consider the following parallel line of reasoning: let us say that a marital status concept is one which, when applied to an individual, says something about that individual's marital status. So, for example, married, single, bachelor, separated and divorced all count as marital status concepts. Furthermore, where C is the conjunction of the concepts C_1, \dots, C_n , let us stipulate that C is a marital status concept if and only if each C_i is a marital status concept. On this account, unmarried man isn't a marital status concept, since man isn't a marital status concept. But bachelor is a marital status concept. So bachelor and unmarried man can't be the same concept.

Something has clearly gone wrong here, because bachelor and unmarried man are identical if any concepts are. But if we aren't prepared to conclude on the basis of the argument which I have just given that bachelor and unmarried man aren't the same concept, why should we conclude, on the basis of a parallel argument, that knows and has a justified true belief aren't the same concept? Perhaps the reason is this: the sense in which unmarried man isn't a marital status concept is that it isn't what might be called a pure marital status concept. It isn't a pure marital status concept because one of its constituents, the concept man, isn't a marital status concept. To put it another way, to describe someone as an unmarried man is to say something about his sex as well as his marital status. But if this is why unmarried man isn't a marital status concept, then bachelor isn't going to count as a marital status concept either; to describe someone as a bachelor is, after all, also to say something about his sex as well as his marital status. So there is no longer any basis for the claim that bachelor and unmarried man can't be the same concept.

On the face of it this is where the parallel with DCA breaks down. Williamson's point is that knows and has a justified true belief can't be the same concept because knows is a purely mental concept whereas concepts like has a justified true belief aren't 'purely mental' (p. 30). On this reading of DCA both (d) and (e) need to be slightly modified. Premise (d) should be read as claiming that the conjunctive concepts with which knows is equated by standard analyses aren't purely mental because they have at least one non-mental constituent. In contrast, (e) now needs to be read as the claim that the concept knows is purely mental. The argument still goes through but is only as compelling as the case for accepting this version of (e). So what is now needed is an argument for the view that the concept knows is purely mental which doesn't presuppose the non-identity of the concept knows and conjunctive concepts like has a justified true belief.

What is Williamson's argument for (e)? It is important to note at this point that the primary concern of chapter 1 of Knowledge and its Limits isn't to defend the thesis that the concept of knowledge is mental or purely mental. The primary concern of this chapter is to defend the thesis that knowing is a state of mind. This is a metaphysical rather than a conceptual thesis. Furthermore, Williamson doesn't argue for the metaphysical thesis from first principles. He maintains that 'our initial presumption should be that knowing is a mental state' (p. 22), and then tries to disarm a range of arguments against this presumption. He also points out that it doesn't follow directly from the fact that knowing is a mental state that the concept knows is mental in his sense, but he nevertheless argues that someone who concedes that knowing is a mental state should also concede that the concept knows is mental, that is, purely mental.

It might be helpful to divide this argument for (e) into two parts. First there is the presumption that knowing is mental state. Let's call this Williamson's Presumption (WP). Then there is the move from WP to (e). Assuming that the move from WP to (e) is defensible the interesting question is whether WP is correct. The problem is that WP is not just the presumption that knowing is a state of mind. It is the presumption that it is 'merely a state of mind' (p. 21). That is to say, it is the presumption that 'there is a mental state being in which is necessary and sufficient for knowing p'. Presumably, it is only because knowing is 'merely' a state of mind that the concept of knowing can plausibly be regarded as 'purely' mental. So everything depends on whether we should accept the existence of an initial presumption to the effect that knowing is merely mental.

In defence of his presumption, Williamson claims that 'prior to philosophical theory-building, we learn the concept of the mental by examples' (p. 22). Our paradigms include not just mental states such as pleasure and pain but also non-factive propositional attitudes such as believing and desiring, that is, attitudes that one can have to falsehoods. In contrast,

knowing is factive since one can only know that p if p is true. So how is it that factive attitudes also come out as mental given that they are different from non-factive attitudes and also from mental states which aren't attitudes at all? Williamson's answer to this question is that 'factive attitudes have so many similarities to non-factive attitudes that we should expect them to constitute mental states too' (p. 22). Indeed, he suggests that there are no pre-theoretical grounds for omitting factive propositional attitudes from the list of paradigmatic mental states. The idea that the mental includes knowing and other factive attitudes 'is built into the natural understanding of the procedure by which the concept of the mental is acquired'(p. 22).

What are the similarities between factive attitudes and non-factive attitudes? Trivially, factive and non-factive attitudes are attitudes, and attitudes are states of mind. So this is one respect in which factive and non-factive attitudes are similar to one another. But this is not enough for Williamson's purposes. What he needs to establish is that knowing is sufficiently similar to believing to sustain the presumption that knowing is, like believing, merely a state of mind. This is the point at which the idea that knowing is factive might appear to be in conflict with the idea that it is merely a state of mind. As Williamson's own discussion illustrates, it takes a good deal of sophisticated argument to weaken the prejudice that a factive attitude can't be merely a state of mind, and this is difficult to reconcile with the suggestion that we have a pre-theoretical commitment to the idea that knowing is merely mental.

There is also a question about the suggestion that WP is built into the procedure by which the concept of the mental is acquired. The procedure that Williamson has in mind is that of learning the concept of the mental by examples, but one might wonder whether this procedure is sufficiently well-defined to sustain the suggestion that WP is built into it. Prior to philosophical theory-building, what we acquire by example are concepts of particular types of

mental state rather than the concept of the mental as such. It's arguable that the procedures by means of which we acquire the concept of the mental leave it open whether knowing is mental in the bland sense that there is a mental state being in which is merely necessary for knowing or in the 'unexpected' (p.21) sense that there is a mental state being in which is necessary and sufficient for knowing. To acquire the concept of the mental as such is to abstract from the differences between different types of mental state, and this already involves taking on theoretical commitments which might be described as 'philosophical'. If this is right, then it is perhaps doubtful whether we have any conception of the mental as such, prior to some philosophical theory-building.

This is not an argument for the view that knowing isn't merely a state of mind or, for that matter, for the view that the concept of knowing isn't purely mental. What is at issue here is Williamson's overall strategy for defending these claims rather than their truth or falsity. He thinks that the default position in the metaphysics of knowing is that knowing is merely a state of mind, and that the burden of proof is therefore on those who reject this position to show what is wrong with it. On this account, defending WP is largely a matter of eliminating putative differences between knowing and non-factive attitudes that have mistakenly been thought to count against the idea that knowing is merely mental. What I have been suggesting, in contrast, is that the idea that the merely mental includes knowing isn't as clearly and deeply entrenched in our ordinary thinking about the mental as Williamson implies, and that this undermines any attempt to represent WP as the default position in this area. While this isn't a knockdown argument against (e), or against DCA, it does raise a question about Williamson's argumentative strategy.

Another potentially problematic aspect of DCA is premise (a). On the face of it, this premise misrepresents what analytic epistemologists are after when they analyse the concept of knowledge. Their aim is to uncover non-circular 'necessary and sufficient conditions for

someone's knowing a given proposition' (Gettier 1963: 121). Such analyses are given in the form 'S knows that p if and only if...', and there is no suggestion that what is on the right-hand side of such biconditionals is equivalent in meaning to what is on the left-hand side.² But if their two sides aren't equivalent in meaning then it is hard to see how they can reasonably be regarded as expressing one and the same concept. At any rate, it is hard to see how they can reasonably be regarded as expressing one and the same concept if we take it that concepts are meanings.³ So it seems that the first step of DCA is mistaken. Standard analyses of the concept of knowledge do not equate this concept with the conjunctive concepts which appear on the right-hand side of their biconditionals.

In fact, matters are a bit more complicated than this suggests. It is one thing to point out that standard analyses of knows do not explicitly endorse the first step of DCA, but it might nevertheless be the case that they are committed to (a), whether they realize it or not. To see why this might be so compare the following biconditionals:

(B) S is a bachelor iff S is an unmarried man.

(G) O is made of gold iff O is made of the element with the atomic number 79.

Both (B) and (G) are correct but there are two important differences between them. First, the concept expressed by the right-hand side of (B) is, on most views of concept-identity, the same as the concept expressed by its left-hand side. This is not true of (G), since made of gold and made of the element with the atomic number 79 are clearly distinct concepts. Second, the truth of (B) but not of (G) can be established solely by means of conceptual analysis. It looks like all one needs to do in order to establish the truth of (B) is to analyse the concept bachelor whereas one can't establish the truth of (G) just by analysing the concept gold.

These differences might lead one to suppose that it is only because the concepts expressed by the two sides of (B) are the identical that its truth can be established solely by means of conceptual analysis. If this supposition is plausible then the first step of DCA begins

to look a bit more plausible. For analytic epistemologists aren't just interested in identifying necessary and sufficient conditions for S to know that p. They are interested in identifying necessary and sufficient conditions that can be uncovered solely by analysing the concept of knowledge. Perhaps this means that its biconditionals will have to be like (B) rather than like (G). The suggestion, in other words, is that if we want to establish the correctness of biconditionals of the form 'S knows that p if and only if...' just by analysing the concept of knowledge it had better be the case that the concepts expressed by their two sides are identical. This is exactly what one would expect if one is serious about the analogy between analysing the concept knows and analysing the concept bachelor. Just as bachelor and unmarried man are the same concept so one would expect knows and any suitable conjunctive analysing concept to be the same concept.

But is it true that it would only be possible for conceptual analysis to establish the relevant biconditionals if the concepts expressed by their two sides are identical? Consider Grice's account of conceptual analysis in his paper 'Postwar Oxford Philosophy'. The basic claim of this paper is that 'to be looking for a conceptual analysis of a given expression E is to be in a position to apply and withhold E in particular cases, but to be looking for a general characterization of the types of cases in which one would apply E rather than withhold it' (1989: 174). Grice's assumption here is that analysing expressions is the best way of analysing the concepts which those expressions express. At the same time, he emphasizes that conceptual analysis isn't the same as lexicography. To see why not consider this pair:

(1) father-----male parent

(2) awe-----mixture of fear and admiration

Although both (1) and (2) give necessary and sufficient conditions for the correct application of an expression, or of the concept expressed by it, Grice thinks that only (2) is a genuine piece of conceptual analysis. Someone who knows the meaning of 'male parent' but who does

not assent at once to (1) shows that he doesn't know the meaning of 'father'. In contrast, someone who does not assent at once to (2) does not thereby show that he doesn't know what 'awe' means. Such a person might be able correctly to apply and withhold this expression in particular cases and yet not be convinced without further argument and examples that (2) is an accurate general characterization of the correct use of the expression.

The next issue is whether awe and mixture of fear and admiration are the same concept. This is a tricky issue since it raises questions about the nature of concept-identity that neither Grice nor Williamson address. Nevertheless, it's easy to see why one might think that awe and mixture of fear and admiration can't be the same concept. For the fact that someone can know the meaning of 'awe' without immediately assenting to (2) suggests that (2) is potentially informative, in a way that (1) is not. But (2) would not be as informative as it is if awe and mixture of fear and admiration were the same concept so awe and mixture of fear and admiration aren't the same concept. Yet the correctness of (2) can be established by conceptual analysis. So it is false that conceptual analysis, at least as Grice conceives of it, can only establish necessary and sufficient conditions for the application of a particular concept by finding a conjunctive concept which is identical to the concept being analysed.

This argument isn't decisive since one might have doubts about the suggestion that concepts are individuated in the way that it implies.⁴ Nevertheless, it remains difficult to avoid the suspicion that the issue of concept-identity which figures so prominently in Williamson's discussion is something of a red herring. We have so far found no clear evidence either that analytic epistemologists who rely on conceptual analysis to establish necessary and sufficient conditions for propositional knowledge actually think that concepts like knows and has a justified true belief are identical or that they are committed to thinking this. Why, then, does Williamson make so much of the issue of concept-identity in DCA? The reason is that he regards the programme of analysis as one that has its origins 'in great

philosophical visions' (p. 31). Russell's logical atomist conception of thought and language is one such vision, and Williamson implies that providing an analysis in Russell's sense of the concept of knowledge really would be a matter of finding a conjunctive concept with which this concept is identical. So if something like Russell's vision is needed to motivate the programme of analysis then it can't be wrong to represent conceptual analysis in the way that DCA represents it. At the same time, this reading makes the pursuit of analyses look even more like a 'degenerating research programme' (p. 31) given that 'the philosophical visions which gave it a point are no longer serious options' (p. 32).

Although this line of thinking goes some way towards explaining and justifying the opening premise of DCA it still falls short of vindicating it. What seems plausible is that a commitment to Russell's vision, or something like it, would be one way of motivating the programme of analysis. What is much less plausible is that some such vision is necessary to motivate the project of analysing the concept of knowledge. As far as most standard analyses of the concept of knowledge are concerned, the central question is: what is knowledge? They think that the best way of answering this question – call it (W) - is to analyse the concept of knowledge and that is why they proceed as they do. There is no suggestion here of some further overarching philosophical vision in which the project of analysis is grounded. In particular, the assumption that analysing the concept of knowledge is the best way of answering (W) seems independent of any commitment to Russell's programme. On a Gricean conception of analysis, it is also independent of any commitment to the idea that the analysing concept must be the same as the concept being analysed. Yet the project of analysing the concept of knowledge still has a point: the point is to help us understand what knowledge is.

But is it plausible that we can explain what knowledge is by analysing the concept of knowledge? We don't think that we can explain what water is just by analysing the concept of water so why should we treat knowledge any differently?⁵ The obvious answer to this

question is that knowledge, unlike water, is not a natural kind.⁶ So what kind of kind is it? Williamson claims that knowing is a state of mind and compares it with the property of being coloured. Anything that is coloured has a more specific colour property, so ‘being coloured is being red or green or..., if the list is understood as open-ended, and the concept is coloured isn’t identified with the disjunctive concept’ (p. 34). In the same way, ‘knowing that A is seeing or remembering or... that A if the list is understood as open-ended, and the concept knows isn’t identified with the disjunctive concept’ (ibid.). On this account, the fundamental problem with traditional attempts to analyse the concept of knowledge is not that they try to answer (W) by focusing on the concept of knowledge. The real problem is that what they try to do with this concept is to analyse it in terms of necessary and sufficient conditions. Just as it doesn’t make sense to explain what it is for something to be coloured by giving necessary and sufficient conditions for something to be coloured so it doesn’t make sense to explain what knowledge is by giving necessary and sufficient conditions.

At this point, however, concept-identity is no longer the main issue. Since it is no longer being suggested that the pursuit of analyses is pointless without Russell’s vision in the background it is irrelevant whether conceptual analysis in Russell’s sense is committed to identifying the analysing concept with the concept being analysed. Most current analyses of the concept of knowledge are Gricean rather than Russellian, and the present objection isn’t that Gricean analyses are unmotivated or that they can’t achieve what they set out to achieve without equating knows with conjunctive concepts like justified true belief. Instead, the objection is that finding necessary and sufficient conditions for someone’s knowing a given proposition is, in principle, a bad way of trying to achieve what many analytic epistemologists are trying to achieve, that is, a bad way of trying to explaining what knowledge is. Given that this worry has nothing to do with the somewhat intractable questions about concept-identity that are at the heart of DCA it probably makes sense to abandon this argument. Whereas DCA

looks like an attempt to knock out a version of analytic epistemology with a single blow the point that Williamson should be emphasizing is that simply that he has a better way of doing things.

Before moving on to the positive phase of Williamson's discussion we should have a brief look at the other arguments for UH. The Inductive Argument assumes that 'no currently available analysis in terms of belief is adequate' (p. 4). Although this assumption is not uncontroversial we can let it pass for present purposes. The question is whether it is reasonable to conclude on the basis of four decades of failure that no analysis of the concept knows of the standard kind is correct. On the one hand, there is the worry that four decades isn't a long enough period of time in the history of philosophy to warrant such pessimism. On the other hand, time isn't the only relevant factor here. Given that attempts to analyse the concept of knowledge have succumbed to a 'pattern of counterexamples and epicycles' (p. 31) it's not unreasonable to conclude that all such attempts suffer from a common underlying defect. So while a successful analysis might be around the corner there is no reason to believe that this is the case. It seems, therefore, that the Inductive Argument carries considerable weight unless it is wrong to assume that no currently available analysis is correct.⁷

The False Expectations Argument says that there is no special reason to expect an analysis of knows. Assuming that truth and belief are necessary for knowledge, 'we might expect to reach a necessary and sufficient condition by adding whatever knowing has which believing truly may lack' (p. 32). But that expectation is based on a fallacy, Williamson claims. For example, 'although being coloured is a necessary and sufficient condition for being red, we cannot state a necessary and sufficient condition for being red by conjoining being coloured with other properties specified without reference to red. Neither the equation 'Red = coloured + X' nor the equation 'Knowledge = true belief + X' need have a non-circular solution' (p. 3).

One question about this argument is whether the analogy with red is appropriate. Since Locke introduced the distinction between simple and complex ideas and insisted that simple ideas can't be broken down those who have gone in for conceptual analysis have been careful to argue that only complex concepts are analysable. From this perspective red is the paradigm of a simple concept. Its unanalysability should therefore come as no surprise but it doesn't follow that the concept of knowledge can't be analysed. More cautiously, it does not follow that the concept of knowledge can't be analysed if this concept is complex rather than simple. If knows is simple, or if there isn't a viable simple/ complex distinction, then the False Expectations Argument goes through. The fact is, however, that Williamson doesn't establish the simplicity of knows or the unsustainability of the distinction between simple and complex concepts. As things stand, therefore, the False Expectations Argument is inconclusive.

3. Williamson's Proposal

Let us now consider the positive phase of Williamson's discussion. His main positive proposal is that knowing is merely a state of mind. Specifically, knowing is 'the most general factive stative attitude, that which one has to a proposition if one has any stative attitude to it at all' (p. 34). A factive attitude is one which one can only have to truths, and we have already seen that the linguistic expression of a factive stative attitude is a factive stative mental state operator (FMSO). If Φ is an FMSO, it 'typically takes as subject a term for something animate and as object a term consisting of "that" followed by a sentence' (p. 34), the inference from 'S Φ s that p' to 'p' is deductively valid, and 'S Φ s that p' entails that 'S grasps the proposition that p'. In addition, if Φ is an FMSO, it is semantically unanalysable.

It is against this background that Williamson identifies the two principles which he represents as characterizing the concept of knowing uniquely, 'up to logical equivalence', in terms of the concept of an FMSO. There is the principle that 'know' is an FMSO and the principle that if Φ is an FMSO, then from 'S Φ s that p' one may infer 'S knows that p'. In

these terms, both 'sees that' and 'remembers that' are examples of FMSOs. As Williamson puts it 'if you see that it is raining, then you know that it is raining. If you remember that it was raining, then you know that it was raining' (p. 37). 'Know' is therefore 'the most general FMSO, one which applies if any FMSO applies' (p. 39).

How, then, does this account of the concept of knowledge differ from analysis in the traditional sense? Analysis in this sense is both decompositional and reductive; it proceeds by dismantling the concept of knowledge, and represents the concepts into which it analyses the concept of knowledge as more basic than this concept. Williamson's account is 'explicitly not a decomposition of the concept knows'. (p. 40). He is not claiming that this concept 'contains' or 'embeds' the concept of an FMSO in anything like the sense in which, on the traditional conception, the concept knows contains the concepts of truth or belief or justification. But does this mean that his account isn't in any sense a reductive account? Since he characterizes the concept of knowledge in terms of the concept of an FMSO the question is whether the concept of an FMSO is more basic than the concept of knowledge. If so, then we would have to conclude that Williamson's account is reductive even if it isn't decompositional.

What would it be for the concept of an FMSO to be more basic than the concept knows? One form of basicness is what might be called explanatory basicness. A concept C is more explanatorily basic than another concept D iff C can be explained without using D but D can't be explained without using C. So if it turns out that we can explain the concept of an FMSO without employing the concept of knowledge, but that we can't explain the concept of knowledge without employing the concept of an FMSO, then we will be forced to conclude that the concept of an FMSO is more basic than the concept of knowledge. Although explanatory basicness isn't the only form of basicness the distinguishing feature of reductive approaches to the concept of knowledge is precisely that they attempt to analyse this concept in terms of concepts which are more basic in the explanatory sense.

Is it true that the concept of an FMSO is explanatorily more basic than the concept of knowledge? We have already seen that Williamson explains the concept of knowledge in terms of the concept of an FMSO, but he explicitly denies that the concept of an FMSO has to be explained in terms of the concept of knowledge. In an important passage, he claims that the notion of an FMSO ‘has been explained without essential reference to the notion of knowing, although “know” is an example of an FMSO’ (p. 37). This suggests that the concept of an FMSO is explanatorily more basic than the concept of knowledge. And if this is right, then what Williamson is proposing is a non-decompositional but nevertheless reductive account of the concept of knowledge.

It is helpful to compare this approach to giving a positive account of the concept of knowledge with Strawson’s approach in Analysis and Metaphysics. In his discussion, Strawson is concerned to distinguish between two forms of conceptual analysis, reductive and non-reductive.⁸ The former tries to explain concepts like knowledge by breaking them down into explanatorily more basic concepts. In contrast, non-reductive conceptual analysis tries to explain such concepts by relating them to concepts at the same level, that is, concepts which are no more basic. From the perspective of reductive conceptual analysis, circularity is a major issue. If one analyses one concept C in terms of another concept D, it mustn’t turn out that D can’t be explained without reference to C. In contrast, non-reductive conceptual analysis is much more relaxed about circularity. Its aim isn’t so much to analyse concepts as to elucidate them, and the concepts which one uses to elucidate a given concept needn’t be any more basic. As far as non-reductive analysis is concerned, circularity needn’t be a vice.

In these terms, Strawson is a practitioner of non-reductive conceptual analysis. The same doesn’t appear to be true of Williamson. For a start, he isn’t relaxed about circularity. Thus, in the first few pages of Knowledge and its Limits, he considers the proposal that Gettier-type counter-examples to the traditional tripartite definition of knowledge can be

circumvented by strengthening the justification condition, but he rejects this proposal on the grounds that the required standard of justification is unlikely to be ‘independent of knowledge itself’ (p. 4). In contrast, Williamson’s own account of the concept of knowledge in terms of the concept of an FMSO won’t be circular in this way if, as he claims, the concept of an FMSO can be explained without essential reference to the concept of knowing.

How seriously should we take the suggestion that the concept of an FMSO is explanatorily more basic than the concept of knowledge? Perhaps Williamson’s thought is that the concept of an FMSO can be characterized by reference to the various marks listed above and that none of these marks makes any reference to knowledge. So this is the sense in which the concept of an FMSO can be explained without essential reference to the concept of knowing. On the other hand, the two principles which Williamson uses to characterize the concept of knowing ‘up to logical equivalence’ tell a different story. The second of these principles states that if Φ is an FMSO, then from ‘S Φ s that p’ one may infer ‘S knows that p’. This does not look like a principle which does not involve the concept of knowledge or which is independent of knowledge itself. What it suggests is rather that we will not have understood the notion of an FMSO unless we regard factive mental states as states of knowledge. But it is difficult to see how this can be reconciled with the suggestion that the concept of an FMSO can be explained without essential reference to the concept of knowing.

It seems, then, that Williamson’s modest positive account of the concept of knowledge contains both reductive and non-reductive elements. The reductive line of thinking is the one which uses the concept of an FMSO to characterize the concept of knowledge while insisting that the notion of an FMSO can itself be explained without essential reference to the concept of knowledge. The non-reductive line of thinking is the one which uses the concept of an FMSO to characterize the concept of knowledge while insisting that the notion of an FMSO cannot be explained without essential reference to the concept of knowledge. If we emphasize

the non-reductive dimension of Williamson's thinking, then we will see him as elucidating the concept of knowledge in something like Strawson's sense. If, on the other hand, we emphasize the reductive dimension of Williamson's thinking, then we won't be able to see him as doing non-reductive conceptual analysis. Either way, it appears that what Williamson is doing can legitimately be described as a form of conceptual analysis.⁹

4. What is Knowledge?

The final issue concerns the suggestion that if one knows that A then there is a specific way in which one knows. What is right about this is that if a subject S is said to know that A then one question which can always legitimately be asked is: how does S know?¹⁰ It might seem that it is in connection with this question – call it (H) – that Williamson's 'ways of knowing' come into our picture of knowledge. For once we have identified a specific way in which S knows that A it would appear we have answered (H). What we are reluctant to accept is that it can be a brute fact that S knows that A without there being a specific way in which he knows it; it can't be that A 'just knows'. On the present reading, this is the sense in which if one knows that A there is a specific way in which one knows it.

At this point, however, we run into the following difficulty: a satisfactory answer to (H), as this question is usually understood, will need to explain how S came to know that A, that is, how S acquired this piece of knowledge. For example, we can say that S came to know that A by seeing that A. Seeing that A is a way of coming to know that A, and we can explain how S knows that A by reference to a specific way of coming to know. But are ways of coming to know that A the same as Williamson's 'ways of knowing'? The problem is that, for Williamson, remembering that A is a way of knowing that A yet it is hard to see how it can be a way of coming to know that A. Remembering that A is a way of retaining the knowledge that A, not a way of acquiring it. This suggests that Williamson's ways of knowing needn't be ways of coming to know.

So what does it mean to say that seeing that A and remembering that A are both ways of knowing that A even though they are not both ways of coming to know that A? All it can mean is that certain entailments hold, specifically, the entailment from ‘S sees that A’ to ‘S knows that A’ and from ‘S remembers that A’ to ‘S knows that A’. The problem, however, is that there are many such entailments that have no bearing on (H). For example, ‘regret’ and ‘realize’ are both FMSOs, which means that one can infer ‘S knows that A’ from ‘S regrets that A’ as well as from ‘S realizes that A’. Yet regretting or realizing that A aren’t ways of coming to know that A. Indeed, it’s not even clear that they are ways of knowing that A; it sounds distinctly odd to say that realizing that it is raining is a way of knowing that it is raining or that regretting that one never learned to play the piano is a way of knowing that one never learned to play the piano. This suggests that a problem with the notion of an FMSO is that it is far too indiscriminating. There are important differences between ‘see’, ‘remember’ and ‘realize’ that the description of all of them as FMSOs overlooks, and this is one of the limitations of Williamson’s framework.

Suppose, then, that we insist on distinguishing between different kinds of FMSO and focus on those which can be used to answer (H). I have suggested that seeing that A is one such FMSO. Because seeing that it is raining is potentially a way of coming to know that it is raining I can explain how I know that it is raining by reference to the fact that I can see that it is raining. But seeing that it is raining is obviously not the only way of coming to know that it is raining. I can also come to know this by hearing the rain lashing against the windows of my apartment, by listening to a weather report on the radio, and so on. With this in mind, we also find ourselves in a position to explain what it is to know that it is raining. Instead of trying to explain this by specifying non-circular necessary and sufficient conditions for knowing that it is raining we now have another possibility. The alternative is to identify some of the many different ways of coming to know that it is raining and thereby to explain what it is to know

that it is raining. More generally, the suggestion is that we can explain what this kind of knowledge is, or what it consists in, by identifying acceptable answers to the question ‘how do you know?’.

In fact, this can only be part of the story. As well as being acquired, knowledge can also be retained and transmitted. A more complete account of what knowledge is will therefore also say something about the different ways of retaining and transmitting it. For example, it will point out that keeping a written record of what one knows about a particular subject can be a way of retaining one’s knowledge as well as a way of transmitting it to others. Nevertheless, acquisition remains at the centre of this conception of knowledge. If one has no conception of the different ways in which knowledge can in principle be acquired one doesn’t fully understand what knowledge is. Yet there is no suggestion that particular ways of coming to know that A are necessary conditions for knowing that A. Although I can acquire the knowledge that it is raining by seeing that it is raining seeing that it is raining isn’t a necessary condition for knowing that it is raining. It is a sufficient condition but this still doesn’t mean that knowledge is being explained by reference to conditions for knowing as distinct from ways of coming to know.

Suppose that ways of coming to know that A are described as means of knowing that A. Then what we now have is what might be called a Means Response to (H) and to (W). This response doesn’t treat knowledge as a natural kind or as something that is to be studied by analysing the concept of knowledge into more basic concepts. Instead, the Means Response explains how one knows that A by identifying the means by which one actually came to know it, and it explains what it is to know that A by identifying different possible means of knowing it, including the means by which one actually came to know it. Since there may be countless different means of coming to know that A the Means Response doesn’t try to come up with a complete list. To understand what it is to know that A all one needs is an open-ended list of

means of knowing that A, perhaps together with some indication of whether some means of knowing that A are more basic than others. In drawing up its list of means the Means Response relies on armchair reflection rather than empirical science. It can still be described as ‘analysing’ the concept of knowledge but it isn’t engaged in the kind of reductive conceptual analysis which Williamson criticizes.

One worry about the Means Response is that it puts the cart before the horse. The worry is that the question ‘what is knowledge?’ is prior to the question ‘what are the different means by which it is possible to come to know?’, and that (W) must therefore be answered before (H) can be answered. Intuitively, one can’t figure out how something is acquired unless one already knows what it is, and this implies that the Means Response is wrong to try to answer (W) on the basis of its response to (H), that is, by helping itself to the notion of a means of knowing. On reflection, however, it is clear that this objection fails. Consider the question ‘what is cricket?’. A good way of answering this question would be to explain how cricket is played, and it would be bizarre to object that this kind of explanation fails because it isn’t possible to explain how cricket is played unless one already knows what cricket is. The position is rather than one explains what cricket is by describing how it is played. The ‘what’ question isn’t prior to the ‘how’ question, and the same goes for (W) and (H).¹¹ Describing how knowledge is acquired, retained and transmitted is as good a way of explaining what knowledge is as describing how cricket is played is of explaining what cricket is. In both cases, the focus is on means, means of playing cricket in the one case and means of acquiring, retaining and transmitting knowledge in the other.

The Means Response to (W) isn’t ultimately very different from Williamson’s response to this question. There is barely any difference between Williamson’s proposal that ‘knowing that A is seeing or remembering or... that A if the list is understood as open-ended, and the concept knows isn’t identified with the disjunctive concept’ (p. 34) and the suggestion

that (W) can be answered by producing an open-ended list of means of knowing. The one difference is that means of knowing are, in the first, instance, means of coming to know and so will exclude many of Williamson's 'ways of knowing'. Nevertheless, the Means Response has the same basic structure as Williamson's account. The interesting final question that needs to be considered, therefore, is whether the Means Response is the end of the story or whether a philosophically satisfying response to (H) and to (W) needs to go beyond the identification of means.

The most straightforward reason for thinking that a Means Response to (H) and (W) is insufficient is this: suppose that a person S is said to know that it is raining by seeing that it is raining, and that this is how S's knowledge that it is raining is explained. This explanation is no good if it isn't possible for S to see that it is raining, that is, if there is some obstacle that prevents S from seeing that it is raining. Some obstacles are comparatively superficial. For example, it might be too dark for S to see or his eyesight might not be good enough. On the other hand, it is sometimes suggested that there are deep epistemological obstacles that stand in the way of anyone being able to see that it is raining, and that that is why the proposed explanation of S's knowledge doesn't work. For example, sceptics argue that S can't see and thereby know that it is raining unless he can eliminate the possibility that he is dreaming that it is raining, and that this is not a possibility which S can possibly eliminate.¹²

If there are obstacles to seeing that it is raining that take the form of epistemological requirements that can't be met then it won't do just to say that S knows that it is raining because he can see that it is raining. It also needs to be shown that the obstacle-generating requirements aren't genuine and that there is therefore nothing that prevents S from seeing that it is raining.¹³ This would be what might be called an obstacle-dissipating response to (H) and to (W). For example, one way of dissipating the alleged obstacle would be to insist on a distinction between what is necessary for seeing that it is raining and what is necessary for

knowing that one sees that it is raining. With this distinction in place, we can argue that being able to eliminate the possibility that one is dreaming is only necessary for knowing that one sees that it is raining rather than for seeing that it is raining. This is an obstacle-dissipating supplement to the Means Response, and the present suggestion is that this response won't be philosophically satisfying unless it is supplemented in this way.

This is still not the end of the line. In addition to showing that there is nothing that prevents *S* from seeing that it is raining, a philosophically satisfying account of *S*'s knowledge might also need to include an account of what makes it possible for *S* to see that it is raining and thereby to know that it is raining. For example, it is arguable that *S* couldn't see that it is raining without a capacity for spatial perception or grasp of the concept rain and other related concepts.¹⁴ These are what might be described as enabling conditions for seeing that it is raining.¹⁵ They enable or make it possible for *S* to see that it is raining and thereby also enable or make it possible for him to know that it is raining by seeing that it is raining. So the proposed explanation of *S*'s knowledge that it is raining is a multi-levels explanation, one that operates at the level of means, the level of obstacle-dissipation and the level of enabling conditions.

In these terms, Williamson's modest positive account of the concept knows operates largely at the level of means. This is a reflection of his somewhat minimalist conception of what it would be to answer (W) and (H). By Williamson's lights, it is not that it isn't possible to go beyond the identification of means or 'ways of knowing' but that it is not necessary to do so. What is necessary is that we are prepared to tackle obstacles to knowing as and when we encounter them. This does not mean that we are under any obligation in epistemology to anticipate and rebut every potential obstacle or to provide a substantive account of what makes it possible to see that *A* or to know that *A* by seeing that *A*. This is the sense in which Williamson's positive account amounts to a form of explanatory minimalism. It is content to

represent knowing as the determinable of which seeing, remembering, realizing and so on are determinates without attempting to explain what makes these attitudes ways of knowing or of coming to know.

For all its attractions, explanatory minimalism is too austere to be the source of much philosophical satisfaction. What we seek in epistemology is, as Stroud puts it, ‘some kind of explanation of our knowledge - some account of how it is possible’ (2000: 4). Different epistemologists have different conceptions of what a good explanation should look like but it is hard not to think that explanatory minimalism takes too much for granted. Once we agree that ‘S sees that A’ entails ‘S knows that A’ it is all too easy to ‘explain’ S’s knowledge that A by pointing out that he sees that A. The problem is that S’s ability to see that A is itself something which we should be trying to understand and explain.¹⁶ Perhaps it is only in the context of philosophy that this needs explaining but philosophy is, after all, what we are doing. What needs explaining is how there can be such a thing as seeing that A, and this means that the philosophy of knowledge must be prepared to think in general terms about potential obstacles to seeing that A as well as the enabling conditions for this kind of seeing. Viewed in this light, Williamson’s ingenious account of knowing is the beginning rather than the end of the story in epistemology.¹⁷

References

- Austin, J. L. (1979), 'Other Minds', in Philosophical Papers (Oxford: Oxford University Press).
- Dretske, F. (1969), Seeing and Knowing (London: RKP).
- Fodor, J. (1998), 'Review of Peacocke's A Study of Concepts', in In Critical Condition (Cambridge, Mass.: The MIT Press).
- Gettier, E. (1963), 'Is Justified True Belief Knowledge?', Analysis 23.6.
- Goldman, A. (1992), 'A Causal Theory of Knowing', in Liaisons: Philosophy Meets the Cognitive and Social Sciences (Cambridge, Mass.: The MIT Press).
- Grice, H. P. (1989), 'Postwar Oxford Philosophy', in Studies in the Ways of Words (Cambridge, Mass.: Harvard University Press).
- Kant, I. (1932), Critique of Pure Reason, trans. Norman Kemp Smith (London: Macmillan).
- Kornblith, H. (2002), Knowledge and its Place in Nature (Oxford: Oxford University Press).
- McDowell, J. (1998), 'Singular Thought and the Extent of Inner Space', in Meaning, Knowledge and Reality (Cambridge, Mass.: Harvard University Press).
- Strawson, P. F. (1992), Analysis and Metaphysics: An Introduction to Philosophy (Oxford: Oxford University Press).
- Stroud, B. (2000), 'Scepticism and the Possibility of Knowledge', in Understanding Human Knowledge (Oxford: Oxford University Press).
- Williamson, T. (2000), Knowledge and its Limits (Oxford: Oxford University Press).

¹ All references in this form are to Williamson 2000.

² For example, Alvin Goldman says at the end of ‘A Causal Theory of Knowing’ that he is ‘not interested in giving the meaning of “S knows p”; only its truth conditions’ (1992: 82).

³ See Fodor 1998 for the idea that concepts are word meanings.

⁴ There are some trenchant criticisms of the informativeness test for concept identity in Fodor 1998.

⁵ Some epistemologists argue that we shouldn’t treat knowledge differently and that ‘the subject matter of epistemology is knowledge itself, not our concept of knowledge’ (Kornblith 2002: 1).

⁶ For a contrary view see Kornblith 2002.

⁷ As I read it the Inductive Argument is Williamson’s best argument for UH. This isn’t made explicit in the text and the fact that so little time is spent on this argument might easily mislead one into thinking that it is less important in Williamson’s eyes than, say, DCA.

⁸ See Strawson 1992, chapter 2 for an account of this distinction.

⁹ Not if conceptual analysis must be decompositional and reductive but Strawson’s discussion suggests that this characterization of conceptual analysis is too narrow.

¹⁰ As Austin remarks, ‘when we make an assertion such as “There is a goldfinch in the garden” or “He is angry”, there is a sense in which we imply that we are sure of it or know it.... On making such an assertion, therefore, we are directly exposed to the questions (1) “Do you know there is?” “Do you know he is?” and (2) “How do you know?” If in answer to the first question we reply “Yes”, we may then be asked the second question, and even the first question is commonly taken as an invitation to state not merely whether but also how we know’ (1979: 77). Austin goes on to make several very useful points about (2).

¹¹ The analogy isn’t perfect. There are lots of different ways of acquiring knowledge but it isn’t true in the same sense that there are lots of different ways of playing cricket. In addition, describing how cricket is played won’t tell someone what cricket is if it isn’t understood that cricket is a game.

¹² Barry Stroud is someone who emphasizes the importance of dealing with obstacles to knowing . See, for example, Stroud 2000.

¹³ Another possibility would be to show that the epistemological requirements can be met. For example, McDowell argues that it is possible to know that one is not dreaming: ‘one’s knowledge that one is not dreaming, in the relevant sort of situation, owes its credentials as knowledge to the fact that one’s senses are yielding one knowledge of the environment – something that does not happen when one is dreaming’ (1998: 238). As McDowell goes on to admit, however, this line of thinking does not meet the obstacle-generating

requirement on its own terms; it does not show that one is not dreaming 'independently of the epistemic status of whatever putative perceptual knowledge is in question' (1998: 238-9).

¹⁴ A capacity for spatial perception is the capacity to perceive spatial properties such as location. It isn't possible to see that it is raining without seeing that it is raining somewhere, and the perception of location is a form of spatial perception. See Kant 1932 for a defence of the idea that the perception of space is required for the perception of any objective state of affairs. What Kant says is much less plausible for auditory than for visual perception.

¹⁵ For more on the notion of an enabling condition see Dretske 1969: 82-3. The enabling conditions that Dretske is interested in can only be established empirically but not all enabling conditions are like that. There are also a priori enabling conditions for seeing that it is raining, ones that can be established non-empirically. For example, we can know just by armchair reflection that it wouldn't be possible to see that it is raining if one lacked the concept rain. It doesn't follow, of course, that one can't see rain without the concept rain. See Williamson 2000: 38 for an explanation and defence of this distinction.

¹⁶ Perhaps the entailment from 'S sees that A' to 'S knows that A' also calls for some kind of explanation. There is a very interesting attempt to explain it in chapter 3 of Dretske 1969, though Dretske's explanation relies on a reductive analysis of the concept of knowledge in terms of the concept of a conclusive reason. See, for example, Dretske 1969: 124.

¹⁷ I thank Tim Williamson and other members of David Charles' discussion group for useful comments. Thanks also to Ciara Fairley for helpful discussions and comments.