KNOWING WHAT THOUGHT REQUIRES

Part 3: The Problem of Armchair Knowledge

MARTIN DAVIES

Introduction: The problem and a plan for its solution

In the first two lectures, I developed two examples of ways in which being a thinking person depends on being embodied and embedded in the right way. In the first lecture, I defended a claim about embodiment, that thought requires a particular kind of internal cognitive machinery. In the second lecture, I defended a claim about being embedded, that there are requirements that our environment must meet if we are to have thoughts with certain contents. Both the architecturalist argument of the first lecture and the externalist argument of the second lecture give rise to instances of the problem of armchair knowledge when they are combined with an independently plausible claim about self-knowledge.

In Lecture 2, I said that the general form of the problem of armchair knowledge is this. Armchair philosophical theorising supports an inference from A to B. A can be known from the armchair; B cannot be known without detailed empirical investigation of the world. So a thinker can know that A without detailed empirical investigation of the world, and can likewise know that if A then B. But it is obvious that B follows; no empirical investigation is needed to see that. So the thinker's knowledge that A and knowledge that if A then B, and the thinker's ability knowledgeably to draw the obvious inference, seem together to provide a route to knowledge that B – and still without any need to rise from the armchair. Yet B was supposed to be something that could only be known by way of detailed empirical investigation of the world. Knowledge that B was supposed to require an investigative, rather than an armchair, methodology.

My response to instances of the problem of armchair knowledge involves placing limitations on the transmission of epistemic warrant, even from the premises to the conclusion of a palpably valid argument. But, in Lecture 1, I acknowledged that limiting the transmission of epistemic warrant might

strike you as an extreme measure. It may seem much more promising to reject the architecturalist and externalist arguments that generate instances of the problem of armchair knowledge or, if one of the arguments withstands critical assessment, simply to bite the bullet. In this final lecture, I first show that instances of the problem of armchair knowledge, or closely related problems about transmission of epistemic warrant, are relatively widespread. It would not be right to suppose that they are only generated by a couple of idiosyncratic philosophical arguments. Then, second, I show that the limitations on transmission of warrant that are needed are far from being *ad hoc*. Failure of transmission of epistemic warrant is the analogue, within the thought of a single subject, of the dialectical phenomenon of begging the question.

1. Failures of transmission of epistemic warrant

In this section, I first consider instances of the problem of armchair knowledge and then turn to closely related cases which, while not strictly instances of the problem of armchair knowledge, nevertheless illustrate problems about transmission of epistemic warrant.

1.1 Instances of the problem of armchair knowledge

In this sub-section, I very briefly review the two instances of the problem of armchair knowledge that arise from the arguments in the first two lectures and also add two new instances.

Example 1: Subpersonal requirements for thought

The architecturalist argument in Lecture 1, Aunty's own argument for the language of thought, supports the second premise of the argument:

LOT(1)   I am a thinking being.

LOT(2)   If I am a thinking being then I am an LOT being.

Therefore:

LOT(3)   I am an LOT being.

The first premise can, if it is true, be known from the armchair. But the conclusion cannot, even if it is true, be known to be true without empirical investigation.

I did not offer a solution to the problem of armchair knowledge in the first lecture. But in addressing a worry about eliminativism I proposed that the concept of a thinking being has at least two components. There is an exemplar component that specifies a sufficient condition: *we*, at least, are thinking beings. And there is a more theoretical component which, according to Aunty's argument, imposes a necessary condition: thinking beings are LOT beings. There is no logical guarantee that the items that meet the sufficient condition also meet the necessary condition. In a disobliging world, the two components would lead to contradictory verdicts on cases and our current conception of a thinking being would be of no use to us. If we turn out not to be LOT beings then we must negotiate our way to a revised conception of what it is to be a thinking being.

Example 2: Meaning and tacit knowledge

The first new example to be introduced in this lecture concerns the meaning of sentences that are never used. Ordinary speakers of English are credited with speaking a language in which sentences that no one ever gets around to using nevertheless have determinate meanings. But what facts about ordinary speakers and their language use could make it correct for us to describe them in this way? This is the problem of *meaning without use*.

A number of philosophers of language, including Brian Loar (1981) and Stephen Schiffer (1993),[32] have argued persuasively that this problem cannot be solved without appeal to the structure of the mechanisms of language processing in speakers' heads. I myself would argue specifically that our assignments of meaning without use are correct only if speakers have subpersonal-level tacit knowledge of a compositional semantic theory for their language.[33]

Suppose, for a moment, that Loar, Schiffer and I are right about this. Then the *modus ponens* inference to be considered is as follows:

MEANING(1)   Sentence *s* means that *p* in my language and would do so whether or not I ever used it.

MEANING(2)   If sentence *s* means that *p* in my language and would do so whether or not I ever used it then I have tacit knowledge of a compositional semantic theory for my language.

Therefore:

MEANING(3)   I have tacit knowledge of a compositional semantic theory for my language.

Suppose that *s* is a hitherto unused and unconsidered sentence built from words and constructions that occur in other sentences that I have used. When I hear or consider the sentence *s* for the first time, I am able to assign it a meaning, say the meaning that *p*. I may know that this is what *s* means. I may know that my judgement that *s* means that *p* is correct – that this is what *s* does and did and would mean whether or not I used it. Furthermore, I may know this without engaging in any empirical investigation of my language processing system. So I have armchair knowledge of the first premise. Then, if the development of the arguments advanced by Loar and Schiffer is correct, I also have armchair knowledge of the conditional premise. But the conclusion, which follows so obviously from these premises, concerns the structure of the language-processing system and surely cannot be known without a substantial programme of empirical research. Armchair methodologies suffice for knowledge of the premises, but knowledge of the conclusion requires an investigative methodology.

Example 3: Environmental requirements for thought

In Lecture 2, we considered externalist dependence theses. Such a thesis might support the second premise of the argument:

WATER(1) I am thinking that water is wet.

WATER(2) If I am thinking that water is wet then I am (or have been) embedded in an environment that contains samples of water.

Therefore:

WATER(3) I am (or have been) embedded in an environment that contains samples of water.

The instance of the problem of armchair knowledge that is raised by this argument reveals a prima facie tension between externalism about content and first-person authority about the contents of our own thoughts.

I distinguished the achievement problem for first-person authority given externalism from the consequence problem. Approaches to the achievement problem typically make some use of the fact that the content of my second-order thought that I am thinking that water is wet, for example, embeds the content of my first-order thought that water is wet. This is used to show that no special problem for the achievement of self-knowledge flows from the fact that the content of a first-order thought is, in part, externalistically determined. But to the extent that the truth of externalist dependence theses is no bar to the achievement of self-knowledge, the consequence problem for first-person authority given externalism becomes more acute. I suggested that a pointer towards a solution to the consequence problem might be seen in Burge's remark that 'the conditions for thinking a certain thought must be presupposed in the thinking'.[34] But, in the second lecture, I did nothing to develop that suggestion.

Example 4: Teleological theories of content

The instance of the problem of armchair knowledge that results from externalist philosophical theorising about thoughts involving natural kind concepts clearly belongs in a larger category. There are, for example, other varieties of externalism, including the externalism about so-called 'object-dependent thoughts' that is familiar from the work of Gareth Evans and John McDowell.[35] It is also possible to adapt material from Evans (1982, p. 201)

so as to generate externalist examples that involve indexical concepts. Thus, it is plausible that a thinker who essays a 'here'-thought, but who is moving through space even as he thinks, fails to think any determinate thought at all. If the thought that he essays as he moves is, 'There's a bottle of whisky just here', then there is no place such that the correctness of the putative thought would turn on whether there is a bottle of whisky at that place. This piece of philosophical theory, when combined with a doctrine of first-person authority for someone who thinks a 'here'-thought while stationary, seems to yield an unexpected route to knowledge that one is stationary (rather than, say, being moved silently along a darkened hospital corridor, in Evans's example).[36]

But, more generally, all these externalist examples belong with other substantive requirements for thought that issue from philosophical theories about thought content.[37] Thus consider the argument:

TELEO(1) I think many things.

TELEO(2) If I think many things then I have such-and-such a kind of evolutionary history and did not come into existence just a few minutes ago.

Therefore:

TELEO(3) I have such-and-such a kind of evolutionary history and did not come into existence just a few minutes ago.

By the assumption of first-person authority, I can have armchair knowledge of TELEO(1). If teleological theories of content are correct, then I can have armchair knowledge of TELEO(2) by engaging in philosophical theorising. But, while TELEO(3) is doubtless true and even something that I do know, it seems to fall outside the scope of armchair knowledge.

1.2 Other problem cases for transmission of epistemic warrant

In this sub-section, I consider two cases which, while not strictly instances of the problem of armchair knowledge, nevertheless illustrate problems about transmission of epistemic warrant.

Example 5: Moore's anti-sceptical argument

Some restriction on the transmission of epistemic warrant seems already to be needed for an adequate account of putative anti-sceptical arguments like Moore's:

MOORE(1) Here is one hand and here is another.

MOORE(2) If here is one hand and here is another then an external world exists.

Therefore:

MOORE(3) An external world exists.

Of this purported anti-sceptical argument, Barry Stroud says (1984, p. 86):

Once we are familiar with the philosophical problem of our knowledge of the external world, I think we immediately feel that Moore's proof is inadequate.

The sense that what Moore offered is inadequate as an anti-sceptical argument does, indeed, seem to be compelling. We want to allow, with Moore, that the conditional premise can be known by a very elementary piece of philosophical theorising (or conceptual analysis); and we want to allow that the first premise can be known just by looking at one hand and then the other. The argument is palpably valid; but it seems wrong to suppose that our knowledge of the premises could provide us with a route to knowledge of the anti-sceptical conclusion, or that the conclusion has epistemic warrant transmitted to it from the premises.

When I say that epistemic warrant is not transmitted from the premises to the conclusion of Moore's argument, I am not suggesting that someone who confidently believes MOORE(1) and MOORE(2) would be wrong to believe MOORE(3). I am not even suggesting that the belief in MOORE(3)

is epistemologically out of order. The proposal is, rather, that the epistemically adequate warrants for believing the premises may not themselves add up to an epistemically adequate warrant for believing the conclusion.

We observed (in Section 2) that it is possible to deny transmission of warrant without denying closure of knowledge under known entailment. Even if warrant is not transmitted from the premises to the conclusion of a particular valid argument, it might be that anyone who knows the premises also knows the conclusion because it is possible to have knowledge of the conclusion without any warrant at all.

Alternatively, it might be that anyone who has a warrant for the premises inevitably has some other, prior and independent, warrant for the conclusion. In principle, it seems that there are two significantly different ways in which this could happen. In one kind of case, anyone who has a warrant for the premises inevitably already has a warrant for the conclusion, and this warrant for the conclusion does not contribute to the warrant for the premises. This pattern of warrants would be consistent with having closure of knowledge without transmission of warrant. But it is important to notice that the pattern does not actually guarantee non-transmission. It might be that the warrant is indeed transmitted from premises to conclusion, but that this fact is obscured from us by the fact that the transmitted warrant for the conclusion is not needed, given the availability of the prior and independent warrant.

In fact, it might be suggested that this is the situation with Moore's argument. On this view, anyone who looks at one hand and then the other and thereby achieves a perceptual warrant for MOORE(1) already has a more than adequate warrant for MOORE(3) provided by perception of countless other objects. So (continuing with the suggestion), the warrant for MOORE(1), provided by perception of two hands, is indeed transmitted to MOORE(3). But it is not needed, as the conclusion MOORE(3) is already as warranted as can be. According to this suggestion, an illusion of non-transmission is created by a kind of epistemic 'ceiling effect'. It is not easy definitively to rebut

this suggestion. But anyone who is convinced, with Stroud, that Moore's premise, 'Here is a hand and here is another', does not warrant the anti-sceptical conclusion is likely to retain their conviction if we add to the premise 'and another and another . . .'. What seems implausible is not just that one little piece of perceptual evidence could defeat the sceptic, but rather that any amount of perceptual evidence could perform that epistemic service.

There is a second kind of case in which anyone who has a warrant for the premises of an argument inevitably has some prior warrant for the conclusion. In this kind of case, the prior warrant for the conclusion figures as a component in the warrant for the premises. This pattern of warrants involves a kind of epistemic circularity, but it is not obvious that it constitutes non-transmission of warrant. Consider, for example, a very simple case of this kind. The premises of the argument are A and B, and the argument proceeds through the conjunction A&B to the first conjunct A. Clearly, anyone who has a warrant for the premises already has a warrant for the conclusion. But it is far from clear that the warrant for the premises is not transmitted to the conclusion. In fact, the most natural account of this case would be that the warrants for the two premises A and B add up to a warrant for the conjunction A&B and that this amounts to a warrant for the conclusion A. But the transmitted warrant for the conclusion A is no more epistemically adequate than the warrant for A&B and this, in turn, is no more epistemically adequate than the warrant for B. So the pointlessly indirect inferential route to the conclusion A results in transmission of the warrant for the premise A; but the transmitted warrant is liable to be contaminated and weakened along the way.

We began from the idea that non-transmission of epistemic warrant is consistent, in principle, with closure of knowledge under known entailment. It might be that knowledge of the conclusion does not require a warrant; or it might be that knowledge of the conclusion is provided by a prior warrant. We have considered two kinds of case in which there is a prior warrant for the conclusion. The first kind of case was consistent with, but did not guarantee, non-transmission; the second kind of case seemed to amount to pointlessly

circular transmission of warrant. So, although it remains plausible that Moore's argument is a case of non-transmission of warrant, we do not yet have any account of how this non-transmission arises. We turn now to consider the possibility that knowledge of the conclusion of Moore's argument, MOORE(3), does not require a justification or warrant.

I want to focus on the role of MOORE(3) as a background assumption for ordinary perceptual knowledge. Crispin Wright says (1985, p. 437):

Once the hypothesis is seriously entertained that it is as likely as not, for all I know, that there is no material world as ordinarily conceived, my experience will lose all tendency to corroborate the particular propositions about the material world which I normally take to be certain.

Perceptual knowledge does seem to depend on the assumption (or at least the not calling into question) that an external world exists. But we need to be explicit about the relationship between knowledge and justification or warrant. On one conception of this relationship, any proposition that figures as a background assumption in a project of warranting must itself be warranted, or at least be susceptible of being warranted. According to this first conception, unless the conclusion of Moore's argument can somehow be warranted, what we ordinarily take for perceptual knowledge lacks a proper warrant. But there is an alternative conception of the relationship between knowledge and warrant. According to this second conception, the fact that an unwarranted proposition is a background assumption in a warranting project does not always reflect negatively on the epistemic status of propositions that are warranted by the project. There are assumptions that we make without warrant or justification, but not without entitlement or right. It is a substantive task for epistemology to provide an account of the nature of this entitlement. But for present purposes I shall, without giving such an account, simply proceed on the basis of this second conception of the relationship between knowledge and warrant. Some background assumptions are epistemically in good order; and if they are epistemically in good order then we can regard them as cases of knowledge without warrant.

If this conception is correct, then we can indeed offer Moore's argument as a case of non-transmission of epistemic warrant without being committed to saying that it is a counterexample to the closure of knowledge under known entailment. Non-transmission arises because the conclusion of the argument plays a crucial role as a background assumption in the project of warranting the premises. It is only against the background of the assumption of MOORE(3) that the perceptual warrant for MOORE(1) counts as a warrant. There is a kind of epistemic circularity here, but it is a more subtle kind than we saw in the case of the pointlessly indirect inferential route to a conclusion.

Example 6: Colour concepts

In 'Naming the colours', David Lewis begins from the thought that our folk theory of colours contains principles linking colours and colour experiences, such as:[38]

When a red thing is before someone's eyes, it typically causes in him an experience of redness.

If our concepts of colours and of colour experiences are concepts of properties of objects and of inner states that are implicitly defined by our folk theory, then conceptual analysis is liable to lead us to such 'definitions' as these:[39]

D1   *Red* is the surface property of things which typically causes experience of red in people who have such things before their eyes.

D2   *Experience of red* is the inner state of people which is the typical effect of having red things before the eyes.

The problem with D1 and D2 is that what they say, while true, does not distinguish the pair <red, experience of red> from other similar pairs such as <green, experience of green>. A further chapter must be added to the folk theory of colour in order to individuate specific colours, and Lewis suggests

that this further chapter can come in different versions, each specifying relatively parochial examples that serve well enough the needs of some subcommunity of the population. Thus, amongst followers of Australian Rules football, it will suffice to say 'that red is the colour of the diagonal stripe on an Essendon Football Club jumper'.[40]

With this much by way of background, we can consider the following *modus ponens* inference:

RED(1)   This [pointing at the diagonal stripe on an Essendon jumper] is red.

RED(2)   If this is red then there is a type of colour experience and a type of inner state that is typically caused in people who have this before their eyes.

Therefore:

RED(3)   There is a type of colour experience and a type of inner state that is typically caused in people who have this before their eyes.

By relying on my mastery of the exemplar component of the concept of red (the parochial exemplar component that applies to my group), I can know that this Essendon stripe is red. Indeed, I have more than one way of knowing this since I can often know what colour something is just by looking at it. Having seen many Essendon jumpers, I can recognise this item as being the colour of the Essendon diagonal stripe. Either way, knowledge of RED(1) is available to me ahead of any investigation of other people's colour experiences or inner states. By relying on my grasp of the theoretical component of the concept of red (including the principles D1 and D2), I can know that if something is red then there is a type of colour experience and a type of inner state that is typically caused in people who have that thing before their eyes. So I can know RED(2). But it is implausible that, without rising from the armchair save perhaps to look at an Essendon football jumper, I can know the conclusion RED(3).

At the beginning of 'Naming the colours', Lewis says that 'it is a Moorean fact that there are colours rightly so-called'.[41] This remark suggests that certain claims about colours and colour experiences have the status of presuppositions or unquestioned background assumptions in our everyday use of colour concepts to classify the things that we see. It also suggests that these claims, like Moore's conclusion, cannot have epistemic warrant transmitted to them from premises that acquire their warrant in our everyday epistemic projects. That is just what I shall be claiming.

2.   Limitation principles and begging the question

In an earlier paper,[42] I used Wright's discussion of Moore's argument,[43] along with some remarks of Wittgenstein's in *On Certainty*, to motivate a limitation on transmission of epistemic warrant along the following lines:

First Limitation Principle (early version):

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if, for one of the premises, the truth of the conclusion is a precondition of our warrant for that premise counting as a warrant.

This principle appears to have the desired consequence, in respect of Example 5, that epistemic warrant cannot be transmitted from the premises to the conclusion of Moore's argument. But, in this initial formulation, the principle is problematic in a number of respects. It makes use of the unexplained notion of a precondition; and if this is interpreted simply as a necessary condition then the principle is certainly open to counterexamples.

In order to improve the formulation of the principle and to provide it with some clearer motivation, we need to make use of the idea that failure of transmission of epistemic warrant is the analogue, within the thought of a single subject, of the dialectical phenomenon of begging the question.[44] It is often said that Moore's argument begs the question against the sceptic, but what we need is an explicit account of what makes an argument question-begging and for this I rely on Frank Jackson.[45] He says that an argument begs

the question when 'anyone – or anyone sane – who doubted the conclusion would have background beliefs relative to which the evidence for the premises would be no evidence'.[46]

According to Jackson's view of what is achieved by advancing an argument for a conclusion, the speaker invites the hearer to borrow evidence, or other considerations, in favour of the premises of the argument. By her choice of premises she provides an indication as to what kind of considerations these are. Typically, evidence counts in favour of a proposition only relative to particular background assumptions and often the relevant background assumptions are shared between speaker and hearer. But when background assumptions are not shared, it is possible that the considerations that count in favour of the premises relative to the speaker's background assumptions do not count in favour of the premises relative to the hearer's background assumptions. Suppose that a speaker sets out to convince a doubting hearer of the truth of some conclusion. The speaker begs the question against the hearer if the hearer's doubt rationally requires him to adopt background assumptions relative to which the considerations that are supposed to support the speaker's premises no longer provide that support. A question-begging argument 'could be of no use in convincing doubters'.[47]

Convincing a doubter and providing an epistemic warrant both involve ruling out various ways in which a proposition could have been false. A speaker's evidence for her premises rules out various ways in which those premises could have been false, ways left open by the speaker's background assumptions. A hearer who doubts the conclusion of the argument may have background assumptions that leave a wider range of possibilities open, and the speaker's evidence for the premises may not rule all those possibilities out. Indeed, the speaker's evidence may leave untouched the ways in which, according to the hearer, the conclusion could be false.

In a similar way, a thinker's epistemic warrants for believing the premises of an argument rule out various ways in which those premises could have been false, ways left open by background assumptions that the thinker is, in that context, epistemically entitled to make. But it does not follow that those same warrants rule out all the ways in which the conclusion could have been false that are left open by assumptions that the thinker is entitled to make. It may be, for example, that in the context of considering the conclusion the thinker is entitled to make fewer assumptions than when only the premises were under consideration. In that case, epistemically adequate warrants for the premises may be inadequate to rule out the possibilities that now need ruling out. On the other hand, it may be that the assumptions that the thinker is entitled to make already leave open no ways in which the conclusion could have been false. In that case, the thinker is epistemically entitled to believe the conclusion but the epistemic warrants for believing the premises contribute nothing to that entitlement. In either case, the thinker's epistemic warrants for believing the premises do not add up to an epistemically adequate warrant for believing the conclusion.

The analogy between convincing a doubter and providing an epistemic warrant motivates the following principle:

First Limitation Principle (revised version)

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if, for one of the premises, the warrant for that premise counts as a warrant only against the background of certain assumptions and acceptance of those assumptions cannot be rationally combined with doubt about the truth of the conclusion.

In order to apply this principle to any particular argument, we need to identify assumptions such that, for one of the premises, it is only against the background of those assumptions that the epistemic warrant for the premise counts as a warrant. Then we need to show that acceptance of those assumptions cannot be combined with doubt about the truth of the conclusion. Moore's argument does seem to fit this pattern.[48] But it is not clear how we can apply the principle to explain the failure of transmission of warrant from LOT(1)

and LOT(2) to LOT(3), for example.

2.1 Limitation principles and subpersonal requirements for thought

In the case of the conditional premise LOT(2), it might reasonably be said that it is only against the background of the assumption of intentional realism that the premise is supported by the neo-Fregean philosophical theory on which Aunty's argument draws. But it is surely not true that acceptance of this assumption cannot be rationally combined with doubt about the conclusion LOT(3). Many philosophers sanely believe that intentional realism is true but the language of thought hypothesis is false. In the case of the first premise, LOT(1), the epistemic warrant is constituted either by grasp of the exemplar component of the concept of a thinking being or else by awareness of my own conscious mental states. But in neither case is there an obvious candidate for the role of background assumption without which the epistemic warrant would not count as a warrant.

There is, however, a very basic assumption that lies in the background of any epistemic project; namely, the assumption that there is such a proposition as the proposition for which one is attempting to provide evidence, justification or warrant. The notion of a proposition that figures in this assumption is not to be construed in a metaphysically committed way. If a thinker is attempting to provide a warrant for believing A then the basic background assumption is simply that there is such a thing to think as A. If there were no such thing to think as A then there could be no question of anything constituting an epistemically adequate warrant for believing A. So we can make explicit a second principle that is arguably a consequence of the first:

Second Limitation Principle (revised version)[49]

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if, for one of the premises, acceptance of the assumption that there is such a proposition for the knower to think as that premise cannot be rationally combined with doubt about the truth of the conclusion.

One way in which the assumption that figures in this principle could turn out to be false would be if one of the purported conceptual constituents in the premise were revealed to be internally incoherent, dictating contradictory answers to the question whether some particular item falls under the concept. To that extent, the principle holds some promise of providing a solution to the instance of the problem of armchair knowledge that is posed by Aunty's argument. But, on the other hand, it is clear that acceptance of the assumption that there is such a thing to think as that I am a thinking being – and, in particular, acceptance of the assumption that the concept of a thinking being is in good order – can be rationally combined with doubt about the truth of the LOT hypothesis. It is only the acceptance of Aunty's argument that generates rational tension between acceptance of the background assumption and doubt about the conclusion.

It is clear what kind of modification of the principle is required if it is to provide a solution to the problem of armchair knowledge that arises from Aunty's argument and the required modification is not merely opportunistic or *ad hoc*. To see this, we need to return to begging the question and focus on the fact that arguments may have several premises. Suppose that a speaker advances a multi-premise argument in an attempt to convince a hearer who doubts that argument's conclusion. The speaker offers various considerations for borrowing; they are considerations that count in favour of the premises relative to the speaker's background assumptions. If the hearer's doubt by itself rationally requires him to adopt background assumptions relative to which one of the speaker's premises is no longer supported by the considerations that she offers for borrowing then the speaker begs the question against the hearer. That is the kind of case that Jackson describes.

But there is a more complicated scenario in which it is no less true that the argument, as advanced by the speaker, will be of no use in convincing the doubting hearer. If the hearer is to be convinced then he must accept the

considerations that the speaker offers in support of her premises. In addition, he must not differ from the speaker in his background assumptions in such a way that the premises are no longer supported by those considerations. Suppose that the hearer's doubt about the conclusion, when put together with acceptance of the considerations that the speaker offers in support of the premises, rationally requires him to adopt background assumptions relative to which one of the premises is no longer supported by those considerations. That is enough to ensure that the argument, as advanced by the speaker, will be of no use in convincing the hearer. So, if failure of transmission of epistemic warrant is the analogue, within the thought of a single subject, of the dialectical phenomenon of begging the question then we should expect the following pair of limitation principles, of which the second is arguably a consequence of the first:[50]

First Limitation Principle (generalised version)

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if, for one of the premises, the warrant for that premise counts as a warrant only against the background of certain assumptions and acceptance of those assumptions together with the warrants for all the premises cannot be rationally combined with doubt about the truth of the conclusion.

Second Limitation Principle (generalised version)

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if, for one of the premises, acceptance of the assumption that there is such a proposition for the knower to think as that premise together with the warrants for all the premises cannot be rationally combined with doubt about the truth of the conclusion.

It is this last principle that provides a solution to the instance of the problem of armchair knowledge that is posed by Aunty's argument.

Suppose that a thinker accepts that there is such a thing to think as the premise LOT(1), that he himself is a thinking being. Suppose, in particular, that he accepts that there is no internal incoherence, no source of contradictions, in the concept of a thinking being. In that case, the thinker must accept the assumption that the items, such as himself, that meet the sufficient condition for falling under the concept also meet the necessary condition. Acceptance of that assumption does not, by itself, rationally preclude doubt about whether he himself is an LOT being. But suppose, in addition, that the thinker accepts the epistemic warrants for the premises LOT(1) and LOT(2). The epistemic warrant for believing LOT(1) is provided either by the exemplar component of the concept of a thinking being or else by his awareness of his own conscious mental states. But it is the warrant for believing the conditional premise LOT(2) that figures crucially in the solution to the problem of armchair knowledge. That warrant is provided by a battery of philosophical theory and by Aunty's argument. Acceptance of the assumption that the items that meet the sufficient conditions for falling under the concept of a thinking being also meet the necessary conditions *and of the warrant for LOT(2)* cannot be rationally combined with doubt about whether the thinker himself is an LOT being. So the limitation principle is triggered and epistemic warrant cannot be transmitted from LOT(1) and LOT(2) to the conclusion LOT(3). The warrants for believing the premises do not add up to an epistemically adequate warrant for the conclusion.

Example 6: Colour concepts

The problem about transmission of epistemic warrant that is presented by Lewis's account of colour concepts has a solution that is very similar to the solution to the problem presented by Aunty's argument. For the concept of red, like the concept of a thinking being, has an exemplar-based sufficient-conditions component and a theory-based necessary-conditions component.

According to the (parochial) exemplar component of the concept of red, being the colour of the Essendon stripe is sufficient for being red: Essendon stripes (at least) are red things. From the theoretical component we can de-

rive a necessary condition for being red: If something is red then there is a type of colour experience and a type of inner state that is typically caused in people who have that thing before their eyes. But there is no logical guarantee that there is a single type of colour experience and a single type of inner state that is typically produced in people by the diagonal stripe on an Essendon jumper. If there is not, then the two components of the concept yield contradictory pronouncements. If the world turns out to be disobliging in this respect then our current colour concepts will be of no use to us and we must negotiate our way to revised, presumably relativised, colour concepts.

Acceptance of the assumption that there is such a thing to think as the premise RED(1) involves acceptance of the assumption that the items that meet the sufficient conditions for falling under the concept of red also meet the necessary conditions. But acceptance of this *and of the warrant for RED(2)* cannot be rationally combined with doubt about the truth of RED(3). So the limitation principle is triggered and epistemic warrant cannot be transmitted from RED(1) and RED(2) to the conclusion RED(3).

2.2 Limitation principles and environmental requirements for thought

In the earlier paper that I have already mentioned,[51] I proposed a principle that was intended to apply to instances of the problem of armchair knowledge generated by externalist dependence theses:

Second Limitation Principle (early version):

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if, for one of the premises, the truth of the conclusion is a precondition of the knower even being able to believe that premise.

This principle certainly has the desired result. According to externalist philosophical theorising, my being embedded in an environment that contains water is a necessary condition for my believing or even thinking that water is wet. It is also a necessary condition for my thinking any other thought in which the concept of water is deployed; in particular, for my thinking that I am thinking that water is wet. So the truth of WATER(3) is a necessary condition of my even being able to think WATER(1), and this triggers the Second Limitation Principle.[52]

The early version of the Second Limitation Principle has the desired result; but in other respects it is far from satisfactory. The worry is not, primarily, that the principle is open to counterexamples, but rather that no independent motivation for the principle was provided. In short, the early version of the Second Limitation Principle appears to be completely *ad hoc*. Fortunately, we do not need to rely on the early version. We now have properly motivated versions of the Second Limitation Principle, and it only remains to confirm that one or the other of these provides a solution to the instance of the problem of armchair knowledge that arises from externalism and self-knowledge.

The revised version of the principle is not adequate to this task. There is no immediately obvious incompatibility between, on the one hand, acceptance of the assumption that there is such a thing for me to think as that I am thinking that water is wet and, on the other hand, doubt as to whether I am (or have been) embedded in an environment that contains samples of water. It is only in the context of a philosophical theory of externalism that there is a tension between this acceptance and this doubt.

But the generalised version of the principle does have the consequence that epistemic warrant is not transmitted from the premises of the argument to the conclusion. The warrant for the conditional premise, WATER(2), is a piece of philosophical theory that supports both the externalist dependence thesis:

Necessarily(If I am thinking *that water is wet* then I am (or have been) embedded in an environment that contains samples of water).

and the thesis:

Necessarily(If I am thinking *that I am thinking that water is wet* then I am (or have been) embedded in an environment that contains samples of water).

The theory supports the first externalist dependence thesis because it supports the claim:

Necessarily(If there is such a thing for me to think as *that water is wet* then I am (or have been) embedded in an environment that contains samples of water).

Equally, it supports the claim:

Necessarily(If there is such a thing for me to think as *that I am thinking that water is wet* then I am (or have been) embedded in an environment that contains samples of water).

So acceptance (i) of the assumption that there is such a thing for me to think as WATER(1) and (ii) of the warrant for WATER(2) cannot be rationally combined with doubt about the truth of WATER(3). According to the generalised version of the Second Limitation Principle, then, epistemic warrant cannot be transmitted from the premises WATER(1) and WATER(2) to the conclusion WATER(3).

Example 4: Teleological theories of content

The solution to the instance of the problem of armchair knowledge that is presented by teleological theories of content follows the contours of either Example 1 or Example 3, depending on the exact construal of the first premise:

TELEO(1) I think many things.

If this premise is tantamount to 'I am a thinking being', then the solution is just as for Example 1, the problem that arises from Aunty's argument. Acceptance of the assumption that the items that meet the sufficient condition for falling under the concept of a thinking being also meet the necessary condition, when it is coupled with acceptance of the theory of content that constitutes the warrant for TELEO(2), cannot be rationally combined with doubt about truth of TELEO(3).

If, on the other hand, TELEO(1) is tantamount to 'I think many things, including that water is wet, that Essendon stripes are red, that penguins waddle, and more', then the solution is as for Example 3, the problem generated by externalist dependence theses. The conditional premise, TELEO(2), is now construed as: 'If I think many things, including that water is wet, that Essendon stripes are red, that penguins waddle, and more, then I have such-and-such a kind of evolutionary history and did not come into existence just a few minutes ago.' The warrant for this premise is a piece of philosophical theory that supports a battery of teleological dependence theses, including:

Necessarily(If I think *that Essendon stripes are red* then I have such-and-such a kind of evolutionary history and did not come into existence just a few minutes ago.

Necessarily(If I think *that penguins waddle* then I have such-and-such a kind of evolutionary history and did not come into existence just a few minutes ago.

The teleological theory of content supports these theses about thinking because it supports the corresponding claims about content, including:

Necessarily(If there is such a thing for me to think as *that penguins waddle* then I have such-and-such a kind of evolutionary history and did not come into existence just a few minutes ago.

Crucially, this same teleological theory of content also supports claims such as:

Necessarily(If there is such a thing for me to think as *that I think that penguins waddle* then I have such-and-such a kind of evolutionary history and did not come into existence just a few minutes ago.

and so also:

Necessarily(If there is such a thing for me to think as *that I think many things, including . . .* then I have such-and-such a kind of evolutionary history and did not come into existence just a few minutes ago.

So acceptance (i) of the assumption that there is such a thing for me to think as TELEO(1) and (ii) of the warrant for TELEO(2) cannot be rationally combined with doubt about the truth of TELEO(3), and this triggers the generalised version of the Second Limitation Principle.

3.   Limitation principles and the objectivity of meaning

In the first section of this lecture, I provided six examples to substantiate the claim that problems about transmission of epistemic warrant are relatively widespread. In the second section, I argued that limitation principles on transmission of epistemic warrant can be motivated by an analogy between providing a warrant and convincing a doubter. Failure of transmission is the analogue of begging the question. I showed how five of the six problems (three instances of the problem of armchair knowledge and two closely related problems) can be solved by appeal to versions of the First and Second Limitation Principles.[53] It remains to say something about Example 2: Meaning and tacit knowledge.

According to the philosophical theory that supports the conditional premise, MEANING(2), if sentences that are never used or even considered are to have determinate meanings then the language user must have tacit knowledge of a compositional semantic theory. If a speaker had only phrasebook knowledge of the meanings of a finite set of sentences then there would be no basis for crediting her with speaking a language in which sentences outside that set had determinate meanings.[54] In the absence of tacit knowledge of a compositional semantic theory, the application of the concept of meaning to an unused sentence, *s*, would be indeterminate. Any specific judgement about the meaning of *s* in this speaker's language would be incorrect. But it does not appear to follow from this philosophical theory about the objectivity of

meaning that if MEANING(3) were false then there would be no such thing for me to think as MEANING(1). Rather, if MEANING(3) were false because I did not have tacit knowledge of a compositional semantic theory then MEANING(1) would be thinkable but false. So it is not very plausible that the Second Limitation Principle will be applicable to this example.

The solution to the instance of the problem of armchair knowledge that is posed by the argument about meaning and tacit knowledge must lie with the First Limitation Principle. What we need to show is that the warrant for MEANING(1) counts as a warrant only against the background of certain assumptions and that acceptance of those assumptions cannot be combined with doubt about the truth of MEANING(3) – or at least that acceptance of those assumptions together with the warrant for MEANING(2) cannot be combined with doubt about MEANING(3).[55] A fully satisfying account of the issues surrounding the warrant for MEANING(1) would require nothing less than an adequate epistemology of understanding. But perhaps it is sufficient for present purposes to suggest that one route to knowledge of meaning is, under appropriate conditions, to take an impression of meaning at face value.

Suppose, for a moment, that the philosophical theory about the objectivity of meaning is correct and that things are as that theory says they need to be. In particular, I have tacit knowledge of a compositional semantic theory and this tacit knowledge underwrites the meanings of sentences that belong to my language even though I never get around to using them. Meaning in my language is not constituted by my having an impression of meaning, both because unconsidered sentences have meanings and because impressions of meaning can, in principle, be misleading or illusory.[56]

Suppose that *s* is a hitherto unused and unconsidered sentence built from words and constructions that occur in other sentences that I have used. And suppose that, in virtue of my having tacit knowledge of compositional meaning rules for those words and constructions, *s* determinately means that

p in my language. If I now hear or consider *s* for the first time (hearing it in reality or in my mind's ear, as it were) then I may hear it *as* meaning that *p* and, taking that impression of meaning at face value, I may judge that *s* does mean that *p*. The suggestion is that, under appropriate conditions, this judgement amounts to knowledge.

We do not have to be in the grip of a purely reliabilist epistemology to find it plausible that one condition that is necessary, if this judgement is to be knowledge, is that the same states of tacit knowledge that contribute to the constitution of *s* as meaning that *p* should figure in the causal explanation of *s*'s being heard as meaning that *p*. If taking an impression of meaning at face value is to be a route to knowledge, then the mechanisms that generate the impression of meaning should be mechanisms that reliably track the truth about meaning. It would be too restrictive to insist that every knower should be able to conceptualise this requirement and explicitly assume that it holds. Language users with no conception of mechanisms that embody tacit knowledge of semantic rules, or even with no conception of mechanisms that generate impressions of meaning, can surely come to know what sentences mean by taking impressions of meaning at face value. But, on the other hand, if a language user has the conceptual sophistication to consider this requirement and actually doubts that it holds then this seems to rule out the possibility of gaining knowledge of meaning simply by taking impressions of meaning at face value.[57]

When, as in this case, there is a logical gap between having an impression and that impression's being veridical, one is justified in taking the impression at face value only against the background of an assumption (a not calling into question) that certain reliabilist conditions related to the production of that impression are met. The impression furnishes an epistemic warrant for the judgement that things are as they seem to be only against the background of that assumption. The assumption against the background of which an impression of meaning furnishes a warrant when it is simply taken at face value may not be very specific; it may speak of reliability in general

rather than of mechanisms that embody tacit knowledge in particular. But, given the philosophical theory that provides the warrant for MEANING(2), a general assumption of reliability can be elaborated into the particular assumption about impressions of meaning being generated by mechanisms that embody tacit knowledge of semantic rules. So, it is not possible rationally to combine acceptance of (i) the assumption of reliability against the background of which the warrant for MEANING(1) counts as a warrant and (ii) the philosophical theory that provides the warrant for MEANING(2) with doubt about the truth of MEANING(3). This is what we needed to show in order to trigger the First Limitation Principle.

If I were to doubt that I have tacit knowledge of a compositional semantic theory for my language then I could not resolve that doubt by reviewing the considerations that would ordinarily count in favour of MEANING(1) and MEANING(2). For, in the presence of that doubt, and given the considerations in favour of MEANING(2), the consideration that would ordinarily count in favour of MEANING(1) would no longer justify that belief. Analogously, if *you* were to doubt that I have tacit knowledge of a compositional semantic theory for my language then I would be begging the question against you if I tried to convince you by offering those considerations for borrowing.

In ordinary circumstances, it does not occur to me to doubt that the reliabilist conditions for gaining knowledge by taking an impression of meaning at face value are met. Against the background of that assumption (that not calling into question), the impression of meaning provides knowledge that *s* means that *p* by ruling out various relevant alternatives to MEANING(1), such as that *s* means that *q* or that *s* means that *r*.[58] But, even taken together with the philosophical theory that supports MEANING(2), the impression that *s* means that *p* does nothing to rule out the most obviously salient alternative to MEANING(3), namely that I do not have tacit knowledge of a compositional semantic theory and that my impressions of objective meaning are illusory. My epistemic warrants for the two premises of the *modus ponens* inference do not add up to an epistemic warrant for the conclusion.

Warrant is not transmitted from premises to conclusion.[59]

Conclusion

In the first two lectures, I developed two instances of the problem of armchair knowledge. In this final lecture, I have shown that it would not be right to suppose that instances of the problem arise only from a couple of idiosyncratic philosophical arguments (Section 1). I have then motivated some principled limitations on transmission of epistemic warrant and shown how these provide solutions to various instances of the problem of armchair knowledge and to some closely related problems (Section 2). In Section 3, I have considered one instance of the problem at greater length. There are many difficult questions concerning the epistemology of understanding. But I am reasonably confident that even this last instance of the problem of armchair knowledge can be solved in a well motivated way. Being in the armchair, down and out, still seems like an attractive philosophical position.

References

Block, N. 1981: Psychologism and behaviorism. *Philosophical Review*, 90, 5–43.

Block, N. 1995: The mind as the software of the brain. In D.N. Osherson and E.E. Smith (eds), *An Invitation to Cognitive Science, Volume 3: Thinking*. Cambridge, MA: MIT Press, 377–425.

Boghossian, P.A. 1997: What the externalist can know a priori. *Proceedings of the Aristotelian Society*, 97, 161–75.

Brewer, B. 1995: Compulsion by reason. *Proceedings of the Aristotelian Society, Supplementary Volume* 64, 237–53.

Brewer, B. 1996: Internalism and perceptual knowledge. *European Journal of Philosophy*, 4, 259–75.

Brewer, B. 1999: *Perception and Reason*. Oxford: Oxford University Press.

Brown, J. 1995: The incompatibility of anti-individualism and privileged access. *Analysis*, 55, 149–56.

Brown, J. 1998: Natural kind terms and recognitional capacities. *Mind*, 107, 275–303.

Burge, T. 1979: Individualism and the mental. In P.A. French, T.E. Uehling and H.K. Wettstein (eds), *Midwest Studies in Philosophy, Volume 4: Studies in Metaphysics*. Minneapolis: University of Minnesota Press, 73–121.

Burge, T. 1988: Individualism and self-knowledge. *Journal of Philosophy*, 85, 649–63.

Burge, T. 1996: Our entitlement to self-knowledge. *Proceedings of the Aristotelian Society*, 96, 91–116.

Burge, T. 1998: Reason and the first person. In C. Wright, B.C. Smith and C. Macdonald (eds), *Knowing Our Own Minds*. Oxford: Oxford University Press, 243–70.

Chomsky, N. 1965: *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Davies, M. 1987: Tacit knowledge and semantic theory. Can a five per cent difference matter? *Mind*, 96, 441–62.

Davies, M. 1991a: Concepts, connectionism, and the language of thought. In W. Ramsey, S. Stich and D. Rumelhart (eds), *Philosophy and Connectionist Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates, 229–57.

Davies, M. 1991b: Individualism and perceptual content. *Mind,* 100, 461–84.

Davies, M. 1992a: Aunty's own argument for the language of thought. In J. Ezquerro and J.M. Larrazabal (eds), *Cognition, Semantics and Philoso-*

phy: Proceedings of the First International Colloquium on Cognitive Science. Dordrecht: Kluwer Academic Publishers, 235–71.

Davies, M. 1992b: Perceptual content and local supervenience'. *Proceedings of the Aristotelian Society,* 92, 21–45.

Davies, M. 1993: Aims and claims of externalist arguments. In E. Villanueva (ed.) *Philosophical Issues Volume 4: Naturalism and Normativity*. Atascadero, CA: Ridgeview Publishing Company, 227–49.

Davies, M. 1995: Two notions of implicit rules. In J.E. Tomberlin (ed.), *Philosophical Perspectives 9: AI, Connectionism and Philosophical Psychology*. Atascadero, CA: Ridgeview Publishing Company, 153–83.

Davies, M. 1996: Externalism and experience. In A. Clark, J. Ezquerro and J.M. Larrazabal (eds.), *Philosophy and Cognitive Science: Categories, Consciousness and Reasoning*. Dordrecht: Kluwer Academic Publishers, 1–33. Reprinted in N. Block, O. Flanagan and G. Güzeldere (eds.), *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press, 1997, 309–28.

Davies, M. 1998: Externalism, architecturalism, and epistemic warrant. In C. Wright, B.C. Smith and C. Macdonald (eds), *Knowing Our Own Minds*. Oxford: Oxford University Press, 321–61.

Davies, M. 2000a: Persons and their underpinnings. *Philosophical Explorations*, 3, 43–62.

Davies, M. 2000b: Externalism and armchair knowledge. In P. Boghossian and C. Peacocke (eds), *New Essays on the A Priori*. Oxford: Oxford University Press, 384–414.

DeRose, K. 1999: Contextualism: An explanation and defense. In J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*. Oxford: Blackwell Publishers, 187–205.

Evans, G. 1981: Semantic theory and tacit knowledge. In S. Holtzman and C. Leich (eds), *Wittgenstein: To Follow a Rule*. London: Routledge and Kegan Paul, 118–37. Reprinted in *Collected Papers*. Oxford: Oxford University Press, 1985.

Evans, G. 1982: *The Varieties of Reference*. Oxford: Oxford University Press.

Field, H. 1996: The apriority of logic. *Proceedings of the Aristotelian Society*, 96, 359–79.

Fodor, J. 1985: Fodor's guide to mental representation. *Mind*, 94, 55–97. Reprinted in *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press, 1990.

Fodor, J. 1987: *Psychosemantics*. Cambridge, MA: MIT Press.

Heil, J. 1988: Privileged access. *Mind*, 97, 238–51.

Jackson, F.C. 1987: *Conditionals*. Oxford: Blackwell Publishers.

Lewis, D. 1992: Meaning without use: Reply to Hawthorne. *Australasian Journal of Philosophy*, 70, 106–10.

Lewis, D. 1997: Naming the colours. *Australasian Journal of Philosophy*, 75, 323–42.

Loar, B. 1981: *Mind and Meaning*. Cambridge: Cambridge University Press.

McDowell, J. 1984: De re senses. *Philosophical Quarterly*, 34, 283–94.

McDowell, J. 1986: Singular thought and the extent of inner space'. In P. Pettit and J. McDowell (eds), *Subject, Thought, and Context*. Oxford: Oxford University Press, 137–68.

McKinsey, M. 1991: Anti-Individualism and privileged access. *Analysis*, 51, 9–16.

Atascadero, CA: Ridgeview Publishing Company, 499–533.

Rey, G. 1995: A not 'merely empirical' argument for a language of thought. In J.E. Tomberlin (ed.), *Philosophical Perspectives 9: AI, Connectionism and Philosophical Psychology*. Atascadero, CA: Ridgeview Publishing Company, 201–22.