

A principled solution to the problem of armchair knowledge

MARTIN DAVIES

The problem of armchair knowledge arises when there are armchair warrants for believing the premises of a palpably valid argument, yet it is implausible that the question whether or not the conclusion of the argument is true can be settled from the armchair. In the first lecture, I presented three instances of the problem, arising from an architecturalist argument, (LOT), an externalist argument, (WATER), and an argument about colour concepts, (RED). Other instances could be presented; I shall mention some later in this lecture.

My strategy for solving the problem is to say that, sometimes, I cannot settle the question whether or not the conclusion of an argument is true by deploying the warrants that I have for believing the premises, because I cannot rationally avail myself of those warrants within the context of the question-settling project. When this is so, I say that warrant is not transmitted from the premises to the conclusion.

If this strategy is not to appear *ad hoc*, then we need a general account of when warrant is not, in this sense, transmitted. In the second lecture, I explained and defended my hypothesis that non-transmission of warrant is the analogue in the epistemological domain of the dialectical phenomenon of begging the question. I began by reviewing Frank Jackson's (1987) account of begging the question and then refining and generalising it. Next, I described in some detail the analogy between the structure of propounded arguments that beg the question and the structure of certain epistemological situations. In the simplest situations of this type, a thinker has warrants for believing the premises of a palpably valid argument; but doubt about the truth of the conclusion would directly rationally require the adoption of a belief which, if it were warranted, would defeat the warrant for some premise. Such a belief, even if it is not warranted, cannot be rationally combined with taking as a warrant the element that putatively warrants that premise. That is, the belief, whether warranted or not, R-defeats the warrant for that premise.

To defeat a warrant, in the sense that is relevant here, is not to outweigh it but to remove it. So, when a warranted belief defeats the warrant for a premise, P_i , that is constituted by a putatively warranting element, w_i , it is not that w_i still weighs in favour of P_i but the warrant for the defeating belief weighs more heavily against P_i . Rather, when the warrant for P_i is defeated, w_i no longer weighs in favour of P_i at all. Similarly, when a belief R-defeats the warrant that is constituted by w_i , someone with that belief cannot rationally take w_i as weighing in favour of P_i at all.

The final step in the second lecture was to describe the epistemic project of settling a question. It is a project whose conduct is conditioned by the supposition that it is as likely as not that the proposition in question is false. There are restrictions on the warrants that I can rationally deploy within the context of such a project. For, just as a belief may rationally require that I not avail myself of a particular warrant, so the corresponding supposition rationally requires that I not avail myself of that warrant *within the context of a project whose conduct is conditioned by that supposition*. That is, just as a belief may R-defeat a warrant, so a supposition with the same content S-defeats that warrant.

After all this, my hypothesis came to appear not so much bold as virtually truistic. We finished with a limitation principle that codifies that hypothesis:

Limitation Principle (revised and generalised)

Epistemic warrant cannot be transmitted from the premises, P_1, \dots, P_n , of a valid argument to its conclusion Q if the warrants for believing the premises are constituted by putatively warranting elements, w_1, \dots, w_n , and, for one of the premises, P_i , someone who (i) accepted all the putatively warranting elements w_1, \dots, w_n in themselves, and (ii) accepted all the premises save P_i as so warranted, but (iii) doubted the conclusion Q would be directly rationally committed to a belief that would R-defeat the warrant for P_i that is constituted by w_i .

Since the notions of defeat, R-defeat, and S-defeat run in parallel, the limitation principle could be formulated in terms of defeat, or S-defeat, instead of R-defeat. The formulation in terms of defeat would say:

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if . . . for one of the premises, P_i , someone who met conditions (i), (ii), and (iii) would be directly rationally committed to a belief which, if it were warranted, would defeat the warrant for P_i that is constituted by w_i .

The formulation in terms of S-defeat would say:

Epistemic warrant cannot be transmitted from the premises of a valid argument to its conclusion if . . . for one of the premises, P_i , someone who met conditions (i) and (ii) and (iii) suppositionally doubted the conclusion Q would be directly rationally committed to a supposition that would S-defeat the warrant for P_i that is constituted by w_i .

I said at the beginning of the second lecture that there is more than one notion of begging the question in the vicinity. There is also more than one notion of transmission, and correspondingly of non-transmission, of warrant. I do not claim that my notion of non-transmission of warrant is more important for epistemology than other notions. I shall come to a comparison in the final section of this lecture. But my notion is tailored to the task of solving the problem of armchair knowledge.

When I say that warrant is not transmitted from the premises to the conclusion I mean that I cannot settle the question whether or not the conclusion is true by deploying the warrants that I have for believing the premises. This is not because the warrants somehow weigh less heavily within the context of the question-settling project. Rather, it is because, within the context of that project, I cannot rationally avail myself of all those warrants. I cannot rationally regard all the premises as being warranted by their respective putatively warranting elements.

1. Applying the limitation principle

To complete the main task of these lectures, I need to show how the limitation principle provides solutions to the three instances of the problem of armchair knowledge that I presented in the first lecture.

1.1 The limitation principle and Aunty's argument

The first instance arises from the following 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 warrant for believing the first premise is provided by the exemplar-based introduction rule, *TB-I*, which is plausibly a primitive or non-derived component in our conception of a thinking being. The warrant for believing the second premise is provided by the elimination rule, *TB-E*. This is not plausibly a primitive or non-derived component in our conception, but is supported by a battery of philosophical theory including Aunty's argument (which was briefly outlined in the first lecture).

The epistemological situation here is not analogous to the structure of propounded arguments that meet the *basic condition* for begging the question. Doubt about the conclusion, LOT(3), does *not* by itself directly rationally commit a thinker to a belief which, if warranted, would defeat the warrant for believing LOT(1), nor to a belief which, if warranted, would defeat the warrant for believing LOT(2).

But suppose that I were to believe that it is as likely as not that I am not an LOT being *and* to accept the putatively warranting elements for both LOT(1) and LOT(2). This latter would be to accept that our conception of a thinking being includes the introduction rule, *TB-I*, and to accept the philosophical argument to the conclusion that the elimination rule, *TB-E*, is also a component in our conception of a thinking being. Taken together with my belief that it is as likely as not that I am not an LOT being, this would rationally commit me to the belief that it is as likely as not that no set can be assigned to the (putative) concept of a thinking being as its extension and that no genuine concept corresponds to our conception of a thinking being.

The route from doubt about LOT(3) to this belief would not go via doubt about the truth of the premises. And, given this belief – that it is as likely as not that that no genuine concept corresponds to our conception of a thinking being – I could scarcely rationally avail myself of the warrant for believing LOT(1) that is constituted by the introduction rule's being a component in that conception. Thus, someone who doubted the conclusion, and accepted the putative warranting elements for LOT(1) and LOT(2), even without accepting either premise as so warranted, would be directly rationally committed to a belief which, if warranted, would defeat the warrant for believing LOT(1), that is, a belief which, whether warranted or not, would R-defeat the warrant for believing LOT(1).¹ As a result of this, the limitation principle is triggered and epistemic warrant is not transmitted from LOT(1) and LOT(2) to LOT(3). That is, we do not have the unwanted consequence that I could settle the question whether or not the LOT hypothesis is true of me by deploying my armchair warrants for believing LOT(1) and LOT(2).

We noted in the first lecture that an instance of the problem of armchair knowledge arises even if the warrant for believing LOT(1) is not the *a priori* warrant provided by *TB-I*, but a first-personal warrant provided by my awareness of my own mental states. Given only what I have just said, it is not clear how the limitation principle would be triggered in this case. For it is not clear that acceptance of the putative warranting elements for both premises would bring to light the potential tension in the conception of a thinking being. The reason is that what I have said gives the impression that acceptance of the putative warranting element for LOT(2) involves recognising *TB-E*, but not *TB-I*, as a component in our conception of a thinking being.

However, this impression is not correct. The introduction rule already figured in our defence of Aunty's argument against the worry about eliminativism. So, someone who doubted the conclusion LOT(3), and accepted just the putative warranting element for LOT(2), would already be directly rationally committed to a belief which, if warranted, would defeat any warrant for believing LOT(1). Thus, even if my warrant for believing LOT(1) is not provided by the introduction rule *TB-I*, we avoid the unwanted

consequence that I could, from the armchair, settle the question whether or not the LOT hypothesis is true of me. (Recall that I said in the first lecture that worrying about eliminativism was not just delaying the main event. It should now be clear that the details of our response to the worry about eliminativism play an important role in the solution to this instance of the problem of armchair knowledge.)

1.2 *The limitation principle and colour concepts*

The second instance of the problem of armchair knowledge arises from this argument:

RED(1) The Essendon stripe is red.

RED(2) If the Essendon stripe is red then there is a type of inner state that is typically caused in people who have it before their eyes.

Therefore:

RED(3) There is a type of inner state that is typically caused in people who have the Essendon stripe before their eyes.

Here, the warrant for believing the first premise is provided by the exemplar-based introduction rule, *Red-I*, while the warrant for believing the second premise is provided by the more theoretically based elimination rule, *Red-E*.

Given the structural similarity between this case and the case of (LOT), it is clear that someone who doubted the conclusion RED(3), and accepted the putative warranting elements for both RED(1) and RED(2), would be directly rationally committed to a belief which, if warranted, would defeat the warrant for believing RED(1). So the limitation principle is again triggered. Thus, we avoid the unwanted consequence that I could, from the armchair, settle the question whether or not the Essendon stripe typically causes the same type of inner state in different people.

We face a small complication when we recall that an instance of the problem of armchair knowledge arises even if the warrant for believing RED(1) is not the *a priori* warrant provided by *Red-I*, but a perceptual warrant provided by looking at an Essendon football jumper. In order to deal with the problem as it arises in this variant case, we need to maintain that acceptance of the philosophical-theoretical putative warranting element for RED(2) already involves recognising, not only *Red-E*, but also *Red-I*, as a component in our conception of red. Given the evident individuating inadequacy of the principles D1 and D2, this is not an implausible thing to maintain. Thus, even when my warrant for believing RED(1) is perceptual, we avoid the unwanted consequence that, without rising from the armchair save to look at an Essendon football jumper, I could settle the question whether or not the Essendon stripe typically causes the same type of inner state in different people.

1.3 *The limitation principle and externalism*

The third instance of the problem of armchair knowledge arises from the following argument:

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

WATER(2) If I am thinking that water is wet then I meet environmental condition E.

Therefore:

WATER(3) I meet environmental condition E.

I have a first-personal warrant for believing WATER(1) and a philosophical-theoretical warrant for believing WATER(2).

It is not plausible that doubt about the conclusion, WATER(3), by itself rationally commits a thinker to a belief which, if warranted, would defeat the warrant for believing WATER(1), or to a belief which, if warranted, would defeat the warrant for believing WATER(2). So once again it is important that the limitation principle makes use of the epistemological analogue of the *generalised condition* for begging the question, and not just the basic condition. As in the earlier cases, most of the work is done by the putative warranting element provided by philosophical theory – in this case, externalist philosophical theory.

The theory that provides a warrant for believing WATER(2) is general. It is not just about me, and it is not just about the thought that water is wet. The theory underwrites the following general and schematic externalist dependence thesis:

Necessarily (If x is thinking *that* ___ water ___ then x meets environmental condition E).

If x does not meet environmental condition E then x does not believe, doubt, confirm, disconfirm, or even entertain any proposition in which the concept of water is a constituent. So it does not seem overly dramatic to say that externalist theory underwrites the thesis:

Necessarily (If there is such a thing for x to think as *that* ___ water ___ then x meets environmental condition E).

Suppose, now, that I were to believe that it is as likely as not that I do not meet environmental condition E *and* I were to accept the putatively warranting element for WATER(2). This would rationally commit me to the belief that it is as likely as not that there is no such thing for me to think as *that water is wet*, or *that water falls from the clouds*, or *that I drink water*, or *that I am thinking that water is wet*. That belief, if it were warranted, would defeat any warrant for believing WATER(1). Given that belief, whether warranted or not, I could scarcely rationally avail myself of any warrant for believing WATER(1). In particular, I could scarcely rationally take my awareness of my own conscious mental states as a warrant for believing WATER(1) if I were to believe that it is as likely as not that I am subject to an illusion of content.

If this is right then someone who doubted the conclusion WATER(3), and accepted the putative warranting element for WATER(2), would be directly rationally committed to a belief that would R-defeat the warrant for WATER(1). So the limitation principle is triggered and epistemic warrant cannot be transmitted from the premises WATER(1) and WATER(2) to the conclusion WATER(3). Thus, we avoid the unwanted consequence that I could settle the question whether or not I meet environmental condition E by deploying my armchair warrants for believing WATER(1) and WATER(2).

Of course, in reality I do meet environmental condition E, and there is such a thing for me to think as WATER(1). I have no warranted doubt as to whether I meet condition E, and I have undefeated warrants for believing both WATER(1) and WATER(2). Furthermore, I do not even have an unwarranted doubt as to whether I meet condition E, and I can rationally avail myself of the warrants for believing WATER(1) and WATER(2).

But the question-settling project would begin with my regarding the question as open *pro tem*. The conduct of the project – though not the conduct of my thinking outside the context of that project – would then be conditioned by the supposition that it is as likely as not that I do not meet environmental condition E. Within the context of that project, as also outside it, I accept the battery of philosophical theory that is the putative warranting element for WATER(2). And this, with my initial supposition, directly rationally commits me to supposing that it is as likely as not that there is no such thing for me to

think as WATER(1). But then I cannot, *within the context of that project*, rationally take my awareness of my own conscious mental states as constituting a warrant for believing WATER(1) – even though it does, in reality, constitute such a warrant, and even though I can, outside the context of that project, rationally avail myself of that warrant. In short, I cannot deploy my warrants for believing WATER(1) and WATER(2) in order to settle the question whether or not WATER(3) is true. For settling the question involves undertaking this relatively ambitious epistemological project – a project that is more ambitious than the project of deciding what to believe about the answer to the question.

1.4 A corollary to the limitation principle

The ways in which the limitation principle has been applied to provide solutions to these three instances of the problem of armchair knowledge have a common theme. In the cases of (LOT) and of (RED), the warrant for a premise is R-defeated by a belief that it is as likely as not that a particular conception does not correspond to a genuine concept. If there were a gap where the concept in question was supposed to be then there would be no such proposition, no such thing to think, as the premise. In the case of (WATER), it is explicit that the R-defeating belief is the belief that it is as likely as not that there is no such thing for me to think as the premise *that I am thinking that water is wet*.

The common theme is that, for one of the premises, a belief that the conclusion is false – taken together with acceptance of putatively warranting elements – leads to a belief that there is no such thing to think as that premise. Similarly, a supposition that it is as likely as not that the conclusion is false leads to a supposition that it is as likely as not that there is no such thing to think as that premise.

Thus, these three applications of the limitation principle can be regarded as applications of the following plausible corollary:

Corollary

Epistemic warrant cannot be transmitted from the premises, P_1, \dots, P_n , of a valid argument to its conclusion Q if the warrants for believing the premises are constituted by putatively warranting elements, w_1, \dots, w_n , and, for one of the premises, P_i , someone who (i) accepted all the putatively warranting elements w_1, \dots, w_n in themselves, and (ii) accepted all the premises save P_i as so warranted, but (iii) doubted the conclusion Q would be directly rationally committed to a doubt as to whether there is any such proposition for the thinker to think as P_i .²

2. Presupposing the requirements for thought

The (LOT), (RED), and (WATER) arguments are all palpably valid. In each case, I have armchair warrants for believing their premises. But, according to the limitation principle and its corollary, this does not have the unwanted consequence that I can settle the question whether or not the conclusion is true by deploying those warrants. In this sense, epistemic warrant is not transmitted from the premises to the conclusion.

In each case, the structure of the epistemological situation that triggers the limitation principle is the same. Someone who accepted the putative warranting element for the second premise (in each case, a piece of philosophical theory) and doubted the conclusion (thinking that it is as likely as not that the conclusion is false) would be directly rationally committed to a belief that would R-defeat the warrant for the first premise. The rational commitment to that belief would be direct in the sense that it would not go via a rational commitment to doubt the first premise.

If we shift from belief and R-defeat to supposition and S-defeat, then we can say that someone who accepted the philosophical theory that supports the second premise and supposed that it is as likely as not that the conclusion is false would be directly rationally committed to a consequent supposition that would S-defeat the warrant for the first premise. The consequent supposition would be that it is as likely as not that there is no such thing to think as the first premise. Within the context of a question-settling project whose conduct was conditioned by the original supposition, and so also by the consequent supposition, the thinker could not rationally take the warrant for the first premise as a warrant, and so could not deploy that warrant to help settle the question whether or not the conclusion was true.

2.1 Empirical requirements and armchair warrants

From the fact that I cannot rationally avail myself of a warrant within the scope of a particular project, with its conduct conditioned by a particular supposition, it does not follow that I do not really have such a warrant or that I cannot avail myself of that warrant outside the context of that project. For each of the three arguments – (LOT), (RED), and (WATER) – I cannot rationally avail myself of a warrant for the first premise within the context of a question-settling project that concerns the conclusion. But, in each case, I do really have a warrant for believing the first premise, whether it be an *a priori* warrant, a first-personal warrant, or a perceptual warrant.

This warrant is defeasible. It is subject to S-defeat, to R-defeat, and, if doubt about the conclusion is warranted, to defeat proper. So we should ask whether my really having a warrant for believing the premise in the first place depends on my having a substantive antecedent warrant for believing that there *is* such a thing for me to think as that premise. In each case, the truth of the conclusion is, according to the philosophical theory that supports the second premise, a necessary condition for there being such a thing for me to think as the first premise. So we should also ask whether my having a warrant for believing the first premise depends on my having a substantive antecedent warrant for believing the conclusion. Is the direction of inference the opposite of the direction of justification?

In his seminal contribution to the topic of externalism and first-person authority, Tyler Burge says (1988, pp. 653–4):

Among the conditions that determine the contents of first-order empirical thoughts are some that can be known only by empirical means. To think of something as water, for example, one must be in some causal relation to water – or at least in some causal relation to other particular substances that enable one to theorize accurately about water. . . . To know that such conditions obtain, one must rely on empirical methods. To know that water exists, or that what one is touching is water, one cannot circumvent empirical procedures. But to *think* that water is a liquid, one need not *know* the complex conditions that must obtain if one is to think that thought.

I agree with this. In order to think that water is wet, and even to know that I am thinking that water is wet, I do not need to know anything of externalist philosophical theory, and I do not need to know that the conditions required by that theory actually obtain.

This, in essence, is why the problem of armchair knowledge arises from the (WATER) argument. For suppose that my having a warrant for believing that I am thinking that water is wet were to depend on my knowing that various environmental conditions were met. As Burge says, ‘To know that such conditions obtain, one must rely on empirical methods’. So, my warrant for believing the first premise, WATER(1), would

not be an armchair warrant. It would depend on my having conducted detailed empirical investigation of the world around me.

Similarly, suppose that my having a warrant for believing that I am a thinking being were to depend on my knowing that I am an LOT being. Then, my warrant for believing the first premise of the (LOT) argument would not be an armchair warrant. It would depend on my having conducted detailed empirical investigation of the world within me.

Burge's remarks help to explain why, even though externalist arguments and architecturalist arguments uncover empirical requirements for thoughts, the warrants for believing the first premises of the (WATER) and (LOT) arguments are still armchair warrants. My having a warrant for believing that I am thinking that water is wet, or that I am a thinking being, does not depend on my having a substantive antecedent empirical warrant for believing that environment conditions E are met or that the LOT hypothesis is true.

2.2 *Two notions of presupposing*

Burge also says (1988, p. 653; emphasis added): 'It is uncontroversial that the conditions for thinking a certain thought must be *presupposed* in the thinking.' There is more than one way to interpret this remark. On one possible interpretation, Burge is saying that, in thinking that *I am thinking that water is wet*, I assume that various conditions on the world around me are met. In the same spirit, we could say that, in thinking that *I am a thinking being*, I assume that various conditions on the world within me are met. But, in each case, I do not need to have any substantive antecedent positive warrant for the assumption. I make it without warrant or justification, but not without entitlement or right.

According to this first interpretation, Burge's account of authoritative self-knowledge makes crucial use of the distinction between earned and unearned assumptions. First-personal warrants require empirical background assumptions. But the warrants are still armchair warrants because the assumptions are unearned. However, as I suggested in the second lecture (section 3.2) this appeal to unearned assumptions still over-intellectualises the epistemological situation.

It is surely not very plausible that, just in thinking that *I am thinking that water is wet*, I assume that environmental conditions E obtain. And it is not very plausible that, just in *thinking that I am a thinking being*, I assume that I am an LOT being. In response to this worry, it might be said that, for a thinker who has not engaged in any philosophical theorising, the assumption is just that there is such a thing to think as that I am thinking that water is wet or that there is such a thing to think as that I am a thinking being. Perhaps it is only philosophical theorising that leads me from this basic assumption to the further assumption that environmental conditions E obtain or that I am an LOT being, and the entitlement to the basic assumption is then transmitted to the more specific assumption.

But even the basic assumption seems too sophisticated to be required of every thinker who enjoys authoritative self-knowledge. So, although the notion of an unearned assumption will surely have some application in epistemological theory, it is better to interpret Burge in a different way. He is not saying that I assume that the requirements for thought actually obtain. Rather, he is saying that I rely on the obtaining of the requirements for thought rather as I rely on the reliability of perceptual mechanisms in normal conditions. I rely on these things even though they may be beyond my conceiving.

In fact, in response to a recent paper of mine (Davies, 2003b), Burge himself is explicit about this (2003, p. 253):³

I do *not* assimilate this notion of presupposition to a notion of assumption by the individual . . . In order to think that water is wet, an individual need not have the concepts necessary to assume that the relevant conditions for thinking the thought are in place. A child can think that water is wet without having the concepts *condition, environment, causal relation between environment and individual subject, normal*, and so on. I did not intend presupposition to be a propositional attitude. It is an impersonal relation between the thinking and actual principles or conditions governing its possibility.

Such a presupposition plays no epistemic role in justifying . . . an individual's authoritative self-knowledge.

While Burge is particularly concerned with self-knowledge, the point applies more generally. If I am to be warranted in believing a proposition then I must be able to think or entertain that proposition. The requirements for thinking the thought are presupposed, whether they be requirements on the external environment, or on cognitive architecture, or on the inner states of people who look at the diagonal stripe on an Essendon football jumper. In the case of each argument – (LOT), (RED), or (WATER) – my having a warrant for believing the first premise does not depend on any substantive antecedent warrant for believing anything else, nor on any background assumptions, whether earned or unearned. But still, a warranted belief that certain conditions – presupposed conditions – do *not* obtain would defeat the warrant, and even an unwarranted belief that those conditions do not obtain, or that it is as likely as not that those conditions do not obtain, would R-defeat the warrant. This is all that is needed if we are to apply the limitation principle and solve the instance of the problem of armchair knowledge that is presented by each of these three arguments.

3. *A priori* and deep necessity

Because the conditions for thinking a thought must be presupposed in the thinking and the conditions are often empirical, many warrants, including many *a priori* warrants, are subject to a kind of empirical defeat. A warrant would be defeated by a warranted belief that it is as likely as not that the conditions for thinking the (putatively) warranted thought do not obtain.⁴

This fact is important as we consider the relationship between *a priori* and deep necessity – that is, truth on the diagonal, in the two-dimensional semantic framework (Evans, 1979; Davies and Humberstone, 1980; Davies, forthcoming). There is quite a powerful intuition that what can be established *a priori* cannot depend on any contingent feature of reality. So, suppose that we have an *a priori* warrant for believing a proposition P. Then there is a *prima facie* intuition that P should be deeply necessary – that is, true on the diagonal, or true 'no matter what'. But, if the conditions for thinking P are presupposed in the thinking, then this initial intuition has to be hedged. The proposition P should be true provided only that the presupposed conditions obtain.

3.1 'I am here now'

Thus, I believe that I am here now. It is plausible, and I shall not dispute, that I have an *a priori* warrant for believing that I am here now. It is, of course, contingent that I am here now – I could have been elsewhere. It is not true 'no matter what', it is not true in all possible worlds, that I am here now. But what we need to ask is whether it is true that I am here now in all the worlds where the presupposed conditions obtain. There are two

kinds of presupposed conditions that we need to consider. There are conditions for thinking any thought of the ‘I am here now’ type; and there are conditions for thinking the very thought, with the very truth conditions, that I am thinking here and now.

Gareth Evans argues that being able to think about a particular place as ‘here’ is not a trivial matter (1982, p. 161):

We are prepared to suppose that there is a determinate thought here – that the subject has a definite place in mind – because we know that subjects do have a capacity to select one position in egocentric space, and to maintain a stable dispositional connection with it. . . . If the subject . . . does know which place his thought concerns . . . this will be manifestable only in manifestations of that stable dispositional connection . . .

What this suggests is that someone who is unable, for a while, to maintain a stable dispositional connection with a position (and unable to keep track of his movement through space) is likewise unable, for that while, to have indexical thoughts about places. So conditions of the first kind – conditions for thinking any thought of the ‘I am here now’ type – include the condition that, at the time of the thinking, I should be related to *some* place in a stable way and should not, for example, be undetectably moving.

But appeal to conditions of the first kind is not sufficient. For it is not true that I am here now in all the worlds where I am stationary at *some* place – I could have been stationary elsewhere.

Conditions of the second kind include the condition that I, MKD, should be at this place *p* at this time *t*. For I cannot think this very thought, that I am here now, with its particular truth conditions, unless I think it at this very place at this very time. According to the hedged intuition about the one-way entailment relation from *a priori* to deep necessity or truth ‘no matter what’, what I have an *a priori* warrant for believing may not be deeply necessary, but it should be true provided that the presupposed conditions obtain. And, of course, it is true that I am here now in all the possible worlds where MKD is at place *p* at time *t*.

3.2 ‘Here’ thoughts and the problem of armchair knowledge

Reflection on the conditions for thinking any thought of the ‘I am here now’ type leads us to another instance of the problem of armchair knowledge.

Evans presents a vivid example of a thinker who fails to keep track of his movement through space (ibid., p. 201):

A person might lie in bed in hospital thinking repeatedly ‘How hot it was here yesterday’ – supposing himself to be stationary in the dark. But his bed might be very well oiled, and be pulled by strings, so that every time he has what he takes to be the same thought, he is in fact thinking of a different place, and having a different thought.

As Evans describes the case, this thinker has several instantaneous thoughts about different places. But we can adapt the example by imagining that the person thinks, slowly, carefully, not wanting to knock anything over in the dark, ‘There’s a bottle of whisky just here’. In general, 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 several yards 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. The subject has no determinate place in mind.

Suppose now that I am stationary in bed in the dark, thinking, ‘There’s a bottle of whisky just here’ – a thought that is true if there is indeed a bottle of whisky located at a particular position just next to the bed. Suppose, too, that it is correct, as a matter of philosophical theory, that someone who neither maintains a stable dispositional connection with a position nor keeps track of his movement through space is unable to have indexical thoughts about places. And now consider the following argument:

- BED(1) I am thinking that there’s a bottle of whisky just here.
 BED(2) If I am thinking that there’s a bottle of whisky just here, then I am stationary.
 Therefore:
 BED(3) I am stationary.

I have a first-personal warrant for believing BED(1) and a philosophical-theoretical warrant for believing BED(2). But it is implausible that I could, from the armchair or the bed, settle the question whether or not I am stationary – whether instead, perhaps, I am being moved silently along a darkened hospital corridor.

The solution to this instance of the problem of armchair knowledge follows the contours of the solution to the instance that arises from the (WATER) argument. The warrant for believing BED(2) is a piece of philosophical theory that underwrites a general and schematic thesis along the following lines:

Necessarily (If x is thinking *that* ___ *here* ___ then x is stationary).

The theory also underwrites some such thesis as:

Necessarily (If there is such a thing for x to think as *that* ___ *here* ___ then x is stationary).

Suppose, now, that I were to doubt that I am stationary *and* to accept the philosophical theory that supports BED(2). This would rationally commit me to a doubt as to whether there is any such thing for me to think as that I am thinking that there’s a bottle of whisky just here. So the corollary to the limitation principle is triggered.

3.3 Complex demonstratives

Similar issues arise when we consider thoughts that involve perceptual demonstrative concepts, such as *that dog*, *that crooked stick*, or *that man wearing a carnation*. Suppose that such thoughts are object-dependent thoughts (Evans, 1982; McDowell, 1984, 1986) and that the predicate concept within the complex demonstrative concept plays an individuating role in such a way that no object can be thought of as *that F* unless the object is F (Davies, 1982). Suppose, in particular, that I see a crooked stick and think of it under the complex demonstrative concept *that crooked stick*. Then it is plausible, and I shall not dispute, that I have an *a priori* warrant for believing that that crooked stick is crooked.

Now consider the argument:

- STICK(1) That crooked stick is crooked.
 STICK (2) If that crooked stick is crooked then some stick is crooked.
 Therefore:
 STICK (3) Some stick is crooked.

The conclusion of this argument is clearly contingent, and this makes it vivid that the first premise is not true ‘no matter what’. So, once again, we need to check whether STICK(1), and so also STICK(3), is at least true in all the worlds where the presupposed

conditions obtain. But clearly, STICK(3) does pass this test. For, in thinking STICK(1), it is presupposed that there is a crooked stick that I am visually attending to.

The (STICK) argument also gives rise to an instance of the problem of armchair knowledge. I can have *a priori* armchair warrants for believing both STICK(1) and STICK(2), yet it is intuitively implausible that I could, from the armchair, settle the question whether or not there are any crooked sticks.

The solution to this instance of the problem takes a form that is, by now, very familiar. Doubt about the truth of the conclusion, taken together with acceptance of the philosophical theory that provides the warrant for believing the conditional premise STICK(2), would directly rationally commit me to a doubt as to whether there is any such thing to think as the premise STICK(1). Once again the corollary is triggered.

4. The problem of easy knowledge

Stewart Cohen (2002) raises a problem for epistemological views that seek to avoid a particular sceptical challenge by allowing what he calls ‘basic knowledge’. The sceptical challenge is posed by the problem of the criterion (2002, p. 309):

[A] natural intuition (pretheoretically anyway) is that a potential knowledge source, e.g., sense perception, can not deliver knowledge unless we know the source is reliable. But surely our knowledge that sense perception is reliable will be based on knowledge we have about the workings of the world. And surely that knowledge will be acquired, in part, by sense perception. So it looks as if we are in the impossible situation of needing sensory knowledge prior to acquiring it. . . . Scepticism threatens.

This challenge is, of course, similar to the sceptical challenge that Wright (1985) develops from his reflections on the intuitive inadequacy of Moore’s argument and to the argument that Pryor develops on behalf of the sceptic in section 2 of ‘The sceptic and the dogmatist’ (2000).

One way to avoid the challenge is to deny the principle (ibid.):

KR A potential knowledge source K can yield knowledge for S, only if S knows K is reliable.

To allow that a belief source, such as sense perception, ‘can deliver knowledge prior to one’s knowing that the source is reliable’ is to allow *basic knowledge* (ibid., p. 310). But, Cohen argues (ibid., p. 311): ‘[O]nce we allow for basic knowledge, we can acquire reliability knowledge very easily – in fact, all too easily, from an intuitive perspective. . . . We can call this “The Problem of Easy Knowledge”.’

Cohen’s first instance of the problem of easy knowledge arises from the following palpably valid argument (ibid., pp. 312–3):

TABLE(1) This table is red.

TABLE(2) If this table is red then it is not the case that this table is white but illuminated by red lights.

Therefore:

TABLE(3) It is not the case that this table is white but illuminated by red lights.

According to someone who allows for basic knowledge, my warrant for believing TABLE(1) is constituted by my visual experience as of the table being red, in the absence of any warrant for doubting that the lighting conditions are normal or that my perceptual apparatus is functioning properly. My warrant for believing TABLE(1) does not depend

on my having any antecedent warrant for believing TABLE(3). I have a simple *a priori* warrant for believing TABLE(2); and it is straightforward to perform the *modus ponens* inference. But, as Cohen says (ibid., p. 313), ‘it seems very implausible to say that I could in this way come to know that I’m not seeing a white table illuminated by red lights’.⁵

Clearly, what we have here is another argument which, like Moore’s argument, we would expect to be grouped together with the arguments that give rise to instances of the problem of armchair knowledge. It is surely implausible that I could settle the question whether or not the table is white but illuminated by red lights by deploying my warrants for believing TABLE(1) and TABLE(2). But we can avoid this unwanted consequence.

For, suppose that I were to doubt the truth of the conclusion TABLE(3). Suppose that I were to believe that it is as likely as not that this is a white table illuminated by red lights. This belief – this doubt – would directly rationally commit me to a belief that would R-defeat the perceptual warrant for believing TABLE(1). For, if I were to believe that it was as likely as not that lighting conditions were abnormal, with coloured lights illuminating white objects, then I could not rationally take the putatively warranting element, my visual experience as of the table being red, as a warrant for believing TABLE(1). So the original limitation principle (not the corollary) is triggered.

Cohen’s second instance of the problem of easy knowledge arises from the thought that, once basic perceptual knowledge is allowed, we have an easy – too easy – route to knowledge that perceptual experience is reliable, or at least to a battery of evidence that seems to support the hypothesis that perceptual experience is reliable. Cohen calls this form of the problem of easy knowledge ‘the problem of easy evidence’ (ibid., pp. 317–8).

Consider the following argument (ibid., p. 318):

EVIDENCE(1) This table is red.

EVIDENCE(2) This table visually appears to be red.

Therefore:

EVIDENCE(3) On this occasion, at least, my colour vision operated correctly.

As before, I have a perceptual warrant for believing EVIDENCE(1). Furthermore, I have a first-personal warrant for believing EVIDENCE(2). But it is implausible that I could settle the question whether or not on this occasion my colour vision operated correctly just by deploying those warrants for believing EVIDENCE(1) and EVIDENCE(2).

Once again, the limitation principle allows us to avoid the unwanted consequence. For suppose that I were to doubt the truth of the conclusion EVIDENCE(3). Suppose that I were to believe that it is as likely as not that, on this occasion, my colour vision is not operating correctly. Then I could not rationally take my visual experience as of the table being red as providing any support at all for EVIDENCE(1). In short, my solution to the problem of armchair knowledge generalises to Cohen’s problem of easy knowledge, including his problem of easy evidence.⁶

5. Justified belief

In these lectures, I have been concerned to solve the problem of armchair knowledge. The notion of transmission of epistemic warrant that I have been using, and for which I have proposed a limitation principle, is tailored to that task. It has to do with the epistemic project of settling a question because what is problematic about armchair knowledge is the apparent consequence that substantive empirical questions can be settled from the armchair. I said at the beginning of this lecture that I am not claiming that this notion of transmission of warrant is of more importance for epistemology than other notions. In

this final section of the lecture, I shall undertake a brief comparison with one other notion of transmission, and correspondingly of non-transmission.

5.1 Two epistemic projects

The epistemic project of settling a question is, in a sense, the analogue of the dialectical phenomenon of propounding an argument with the purpose of convincing a doubter. At the beginning of the second lecture, the convincing purpose of arguing was contrasted with the teasing-out purpose, and there seems to be an analogue of this teasing-out function in the thinking of a single subject. Surely one aspect of the management of one's web of beliefs is to tease out the consequences of one's beliefs, checking for relationships of support, tension, or inconsistency. If I review my beliefs P_1, \dots, P_n and notice a valid argument from those premises to Q then I shall adopt the belief Q or, if other considerations argue against Q , then I shall reconsider my beliefs P_1, \dots, P_n . So, we might say, the epistemic project of deciding what to believe is the analogue of the dialectical phenomenon of propounding an argument with the purpose of teasing out the consequences of a hearer's beliefs. And, as we noted towards the end of the second lecture, the epistemic project of settling a question is more ambitious than the epistemic project of deciding what to believe about the answer to the question.

If I have done well doxastically in believing P_1, \dots, P_n , and if I notice a valid argument from these premises to Q then, other things being equal, I shall do well doxastically if I believe Q . It is natural to connect this last point with the notion of justification. But it is important to draw a distinction between what justifications there are for believing propositions and what justifications I avail myself of, or can rationally avail myself. Facts about justifications are not, in themselves, psychological facts. So we might speak of the abstract space of justifications.

If there is, in the abstract space of justifications, a justification for believing a proposition and I do believe that proposition, then I think the thing that is the thing to think. To that extent, I do well doxastically. If there is a justification for believing each of P_1, \dots, P_n and I therefore do well doxastically in believing those premises then there is also a justification for believing Q , and I shall do well doxastically if I believe that conclusion. But there are also richer notions of doing well doxastically. If there is a justification for believing a proposition and I believe it *because* there is that justification for doing so then I do well doxastically in a richer sense. If I believe P_1, \dots, P_n because there are those justifications for doing so then I do well doxastically in that richer sense. If I believe the conclusion Q because it follows from premises that I believe because there are those justifications for doing so, then once again I do well doxastically in that richer sense.

Corresponding to these and other senses of doing well doxastically there are notions of transmission of justification or warrant from the premises to the conclusion of a valid argument. But these notions, which are connected with deciding what to believe, are different from the notion of transmission of warrant with which I have been primarily concerned, the notion that is connected with settling a question.

5.2 Another notion of non-transmission of warrant

The most flat-footed notion of transmission of warrant that is connected with deciding what to believe, and with doing well doxastically, will have it that warrant is always transmitted from the premises to the conclusion of a valid argument. But suppose that there is a normative requirement that we should conform the structure of our web of beliefs to the structure of the abstract space of justifications. In that case, I shall, in a

sense, not do well doxastically if the direction of inference is the opposite of the direction of justification. There will be something amiss if I argue from P to Q when the justification for believing P depends on an antecedent justification for believing Q. A diagnostic sign that there is something amiss in this way is that I cannot, by performing the inference from P to Q, arrive at a *first* justification for believing Q. According to the alternative notion of non-transmission of epistemic warrant – though not according to my notion – this is the kind of case in which warrant is not transmitted from the premises to the conclusion of a valid argument.

Thus, for example, Wright says (2003, p. 57):

[A *cogent* argument] is an argument, roughly, whereby someone could/should be moved to rational conviction of the truth of its conclusion – a case where it is possible to *learn* of the truth of the conclusion by getting warrant for the premises and then reasoning to it by the steps involved in the argument in question. Thus a valid argument with warranted premises cannot be cogent if the route to warrant for its premises goes – of necessity, or under the particular constraints of a given epistemic context – via a prior warrant for its conclusion. Such arguments, as we like to say, ‘beg the question’.

Say that a particular warrant, *w*, *transmits* across a valid argument just in case the argument is cogent when *w* is the warrant for its premises.

Wright here connects non-transmission of warrant with begging the question. But the notion of begging the question that goes with his notion of non-transmission is not – or not clearly – Jackson’s notion. It is plausible that it is the notion that is connected with circularity. For example, Irving Copi (1961, pp. 65–6), says that someone who ‘assumes as a premise the very conclusion he intends to prove’ commits the fallacy of begging the question.

In these lectures, I have laboured to make the point that non-transmission, in the sense that is connected to the epistemic project of settling the question, does not require that the warrant for one of the premises should depend on an antecedent warrant for believing the conclusion. In the case of Moore’s argument, the sceptic says that the perceptual experience described in MOORE(0) supports MOORE(1) only if there is an antecedent warrant for believing MOORE(3). Pryor (2000, forthcoming) says that the sceptic is wrong about this. But, on my account, Moore’s argument is an example of transmission failure even though, as I suppose, Pryor is right and the sceptic is wrong. So cases of non-transmission in my sense will not always be cases of non-transmission in Wright’s sense.

Conversely, there are examples that are not cases of non-transmission in my sense, but are cases in which the argument could not be a route to a *first* warrant for believing the conclusion. Thus, consider the following frankly circular variation on Moore’s theme:

CIRCULAR(0) I am having an experience as of one hand [here] and another [here].

CIRCULAR(1) I have hands.

Therefore:

CIRCULAR(2) I have hands.

Clearly, I cannot learn for the first time of the truth of the conclusion CIRCULAR(2) by having a perceptual warrant for believing CIRCULAR(1) and then reasoning from there to CIRCULAR(2). And the situation is no better if we insert a needless detour between premise and conclusion.

But my limitation principle does not classify this as a case of non-transmission. For someone who doubted the conclusion, who believed that it was as likely as not that he did

not have hands, would not be directly rationally committed to any belief that would R-defeat the perceptual warrant for believing CIRCULAR(1) – aka MOORE(1). I can deploy the perceptual warrant for believing the premise to settle the question whether or not the conclusion is true.

The two notions of non-transmission of warrant dissociate in both directions. But they are not mutually exclusive. Some arguments are examples of non-transmission in both senses. Thus, for example, if the justificatory situation in Moore's argument were as the sceptic says that it is (Wright, 1985, p. 437), then Moore's argument would be an example of non-transmission in Wright's sense, and it would still be an example of non-transmission in my sense (provided, of course, that there were to be a warrant for believing the premise, MOORE(1), in the first place). Similarly, if the justificatory situation in the (ELECTION) argument is as the sceptic says that it is (ibid., p. 436), then that argument would be an example of non-transmission in both senses.

However, just because I have laboured to make the point that non-transmission, in my sense, does not require that the warrant for one of the premises should depend on an antecedent warrant for believing the conclusion, I must allow that many of the arguments that I have considered do offer routes to a first warrant for believing their conclusions. If the warrants for believing the premises are *a priori*, then I am *a priori* rationally committed to believing the conclusion. More generally, if there are armchair warrants for believing the premises then I am rationally committed, in the armchair, to believing the conclusion.

In these lectures, I have provided a solution to the problem of armchair knowledge. I hold that, for each of the problematic arguments that I have considered, I cannot settle the question whether or not the conclusion is true by deploying my armchair warrants for believing the premises. But it is no part of my project to deny that the armchair methodology of philosophy can yield rational commitments to believe, and first justifications for believing, substantive empirical claims about the world.

Conclusion

My solution to the problem of armchair knowledge is to say that, in every problematic instance, I cannot rationally avail myself of the warrants for believing the premises within the context of the epistemic project of settling the question whether or not the conclusion is true. But I have also allowed that I do have a warrant for believing the conclusion. Each problematic argument offers an armchair route to a warrant – even to a first warrant – for believing its conclusion. So, do I or do I not have armchair knowledge of the conclusion? This depends on whether our ordinary knowledge claims and knowledge attributions go in step with the status of a belief as warranted *simpliciter* or with the status of a belief as having a question-settling warrant. But the question of the truth conditions of knowledge attributions will have to await another occasion.

Notes

¹ If doubt as to whether our conception corresponds to a genuine concept, and so doubt as to whether there is any such thing to think as that I am a thinking being, would R-defeat the warrant for LOT(1), then presumably that doubt would R-defeat the warrant for LOT(2) as well.

² Crispin Wright (2003) raises a problem for an earlier version of the corollary that is less general and also, unfortunately, less clearly expressed (Davies, 2003, p. 45):

Corollary (old version)

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

Putting this principle into the terminology that I am now using, it is intended to say that warrant cannot be transmitted if, for some premise, P_i , someone who accepted the putatively warranting elements for all the premises save P_i , and who doubted the conclusion, would be rationally committed to a doubt as to whether there is any such proposition for the thinker to think as P_i .

Wright (2003, p. 73) notes that this principle has consequences that differ from those of his own favoured way of limiting the transmission of warrant. He then objects to the principle as follows (ibid., p. 74):

[The principle] is going to block transmission in [(WATER)] *whatever* the nature of the warrant accepted for its conditional premise [WATER(2)]. And that is clearly a bridge too far. For [(WATER)] is fine as a vehicle for warrant transmission provided the warrant for its major premise is *empirical*.

My reply is that the objection turns on not distinguishing between acceptance of warrants – what I now call putative warranting elements – in themselves and acceptance of premises as warranted. For, where I speak of acceptance of the warrant for a premise, Wright speaks of a belief ‘that there is (all things considered) warrant to accept [that premise]’ (ibid.).

But we also need to assure ourselves that the corollary is properly sensitive to the nature of the warrants for the premises. Does it allow that warrant could be transmitted from the premises of the (WATER) argument to its conclusion if the warrant for believing WATER(2) were empirical in character? Here is the example that Wright describes (2003, p. 74):

Suppose, for instance, I am a dyed-in-the-wool anti-externalist who – perhaps as result of brain injury – lacks all recollection of (testimony of) encounters by myself or anyone in my speech community with water. Interested in what are the actual empirical preconditions for possession of the concept water, I do some interactive anthropology. An electronic interrogation of members of a variety of societies – not extending to my own – who possess the concept discloses that they all have a history of interaction with the stuff. I therefore propose on inductive grounds that this probably goes without exception, and so reason my way to the prediction that this will also prove true of my own speech community.

Abstract away from the community aspect of this example, and suppose that I doubt the truth of WATER(3). I believe that it is as likely as not that I have had no contact with water. Suppose that I accept the putative warranting elements for both premises. I accept

that I am aware of a conscious mental state that is at least subjectively indistinguishable from thinking that water is wet. And I accept the battery of empirical evidence that provides inductive support for a generalisation of which WATER(2) is an instance. There is no evident direct route from this belief and this acceptance to a belief that would R-defeat the warrant for believing WATER(1).

Suppose that I actually believe that it is rather *more* likely than not that I have had no contact with water. And suppose that, in addition to accepting the putative warranting elements for both premises, I accept WATER(2) as so warranted. Then I shall believe that it is more likely than not that I am not, after all, thinking that water is wet. Since it certainly seems to me that I am thinking that water is wet, I may be driven to the belief that I may be subject to some kind of cognitive illusion. But, even if that is a belief that would R-defeat the warrant for believing the premise WATER(1), the route to it is not direct, but goes via doubt about the truth of that very premise.

Or suppose that I doubt the truth of the conclusion, accept the putative warranting elements for both premises, and accept WATER(1) as so warranted. This does not lead directly to any R-defeater for inductive warrant but only to the belief that WATER(2) may well be false, and so that it may well be false despite the inductive support for a generalisation of which it is an instance.

In short, Wright is right that any limitation principle on the transmission of epistemic warrant from premises to conclusion must be sensitive to the nature of the warrants for believing the premises. But this is no objection since the principles that I am proposing are so sensitive.

³ I say (Davies, 2003b, p. 111):

Perhaps there is a distinction to be drawn between assumptions and presuppositions. But, on the face of it, Burge is allowing that in thinking that water is wet, or in thinking that I am thinking that water is wet, I presuppose or assume that the conditions necessary for me to think that thought do obtain.

In response, Burge says (2003, p. 253): 'I do *not* assimilate this notion of presupposition to a notion of assumption by the individual, as Davies conjectures.'

⁴ In the first lecture (section 2.4), I mentioned two reasons why it would not be right to say that philosophy has a purely *a priori* methodology. The first reason was that, in philosophy, we presume upon commonplace assumptions that are empirical in character. The second reason was that conceptual negotiation in a disobliging world has to take account of the ways in which the world has turned out to be disobliging. Now we have a third reason not to say that the methodology of philosophy is *purely a priori*. Even philosophical claims for which there are *a priori* warrants may be subject to empirical defeat. See Field, 1996.

⁵ Wright (2003) also discusses this example (with a wall instead of a table); see pp. 60–3.

⁶ Pryor (forthcoming) also discusses Cohen's (2002) problem of easy knowledge.