

ON EPISTEMIC ENTITLEMENT

By Crispin Wright and Martin Davies

II—Martin Davies

EPISTEMIC ENTITLEMENT, WARRANT TRANSMISSION, AND EASY KNOWLEDGE

ABSTRACT Wright's account of sceptical arguments and his use of the idea of epistemic entitlement are reviewed. His notion of non-transmission of epistemic warrant is explained and a concern about his notion of entitlement is developed. An epistemological framework different from Wright's is described and several notions of entitlement are introduced. One of these, negative entitlement, is selected for more detailed comparison with Wright's notion. Thereafter, the paper shows how the two notions of entitlement have contrasting consequences for non-transmission of warrant and how they go naturally with two conceptions of the presuppositions of epistemic projects. Problems for negative entitlement are explained and solutions are proposed.

In the first section of his paper, Crispin Wright distinguishes two kinds of sceptical paradox – the Cartesian and the Humean – which, he suggests, ‘capture, in essentials, all that we have to worry about’. Sceptical arguments of both kinds are supposed to show that we lack a warrant for crucial propositions that Wright calls ‘cornerstones’. But in each case there is a gap in the argument. The sceptic needs to move from our having no evidential justification to our having no warrant at all. So one strategy for responding to the sceptic is to appeal to a kind of warrant that is not a matter of evidential support – a kind of warrant that ‘we do not have to *do any specific evidential work* to earn’. This is the kind of warrant that Wright calls ‘entitlement’. The main business of the central sections (3–8) of Wright's paper is then to make a start on the major philosophical project of providing a substantive account of our epistemic entitlements. So he considers the prospects for, and the limitations of, strategic entitlements and entitlements of cognitive project, of rational deliberation, and of substance.

In my paper, I focus on structural features of Wright's notion of entitlement. I begin (section 1) with an earlier discussion of sceptical arguments that covers some of the same ground as the first section of his paper in this symposium. Then, in section 2, I turn to the question that Wright addresses in his second section, namely, the question of what epistemic entitlement is an entitlement to do (to believe, to accept, or something else). I suggest that Wright's answer to this question, when taken together with some ideas about transmission of epistemic warrant, may impose a strain on our ordinary thinking about the proper management of our web of belief.

While I do not press that concern about Wright's notion of entitlement, I go on to describe a different notion, negative entitlement (section 3). I show how the two notions yield different answers to questions about transmission of epistemic warrant and how the negative notion does not impose the just-mentioned strain on our ordinary thinking about belief (section 4). And I connect the two notions of entitlement with two ways of thinking about the presuppositions of a cognitive or epistemic project (section 5). In the final two

sections, I explain a problem that is faced by the epistemological framework into which the notion of negative entitlement naturally fits and then propose a solution.

I

Two patterns of sceptical argument: Dreaming and I-II-III. In his British Academy Lecture, ‘Facts and certainty’ (1985), Wright begins from ‘two simple patterns of argument which can be brought to bear upon a variety of large regions of discourse so as to generate what seem to be genuine sceptical paradoxes’ (1985, p. 430). Arguments that exhibit the first pattern make use of sceptical possibilities that I am dreaming, that I am a brain in a vat, or that I am hallucinating. Thus, for example, if the sceptic can argue for the principle (C) that ‘at no time *t* do I have sufficient reason to believe that I am not dreaming at *t*’ (ibid., p. 432) then we seem to be led to the conclusion that perception does not provide a basis for knowledge or even for reasonable belief. And the sceptic does appear to be well placed to argue for that principle, to the extent that neither empirical evidence nor *a priori* considerations can furnish me with a reason to believe that I am not dreaming.

The second pattern of sceptical argument is to be appreciated by ‘reflecting on the intuitive inadequacy of G.E. Moore’s [1959] “proof” of the existence of the external world’ (ibid., p. 434). Moore’s argument can be set out as follows:

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

MOORE(II) I have hands.

 If I have hands then an external world exists.

Therefore:

MOORE(III) An external world exists.

MOORE(I) is a proposition about Moore’s experience and, on one conception of perceptual warrant, MOORE(II) is arrived at by inference from this proposition. On another conception, it is the experience itself, rather than a belief about the experience, that provides the defeasible warrant for believing the proposition about hands. But, whichever conception of perceptual warrant is adopted, the key question at this point in Wright’s account is whether the support for MOORE(II) is transmitted to MOORE(III) across the *modus ponens* inference in which the conditional premise is supported by an elementary piece of philosophical theorising.

On behalf of the sceptic, Wright (ibid, pp. 435–6) asks us to compare Moore’s argument with the following:

ELECTION(I) Jones has just written an ‘X’ on that piece of paper.

ELECTION (II) Jones has just voted.

 If Jones has just voted then an election is taking place.

Therefore:

ELECTION (III) An election is taking place.

Here, the evidence summarised in ELECTION(I) provides defeasible support for ELECTION(II); and this premise, together with the conditional premise that is warranted by a conceptual connection between voting and elections, clearly entails ELECTION(III).

Given this relationship between ELECTION(II) and ELECTION(III), we might expect that empirical evidence against ELECTION(III) would count against ELECTION(II) by going into the scales on the opposite side from the evidence summarised in ELECTION(I). In particular, we might expect that evidence that there is no election taking place would leave intact the status of Jones's writing an 'X' on the paper as evidence supporting the belief that Jones has just voted. It would simply outweigh that evidence. But, Wright stresses, this is not, in general, the correct picture (ibid., p. 436):

Imagine . . . that you live in a society which holds electoral 'drills' as often as we hold fire drills, so that the scene you witness of itself provides no clue whether a genuine election is going on or not. In that case, unless you have further information, the knowledge that Jones has placed an 'X' on what looks like a ballot paper has no tendency whatever to support the claim that he has just voted.

In a situation where I have reason to believe that what I am watching is a drill rather than an election, the support ordinarily provided for ELECTION(II) by ELECTION(I) is not outweighed but removed.¹ So, Wright says (p. 436):

the evidential support afforded by [ELECTION(I)] for [ELECTION(II)] is itself conditional on the *prior* reasonableness of accepting [ELECTION(III)]. . . . [K]nowledge of the first does not begin to provide support for the second unless it is *antecedently* reasonable to accept the third.

The imagined sceptic then says that Moore's argument is relevantly similar (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.

As a result (ibid.; emphasis added): 'Only if Moore *already* has grounds for [MOORE(III)] does [MOORE(I)] tend to support [MOORE(II)].'

The sceptic's point is that ELECTION(III) cannot be supported by inference from ELECTION(II) when this is supported in turn by evidence of the kind described in ELECTION(I). Similarly, MOORE(III) cannot be supported by inference from particular claims like MOORE(II) when these are supported in turn by evidence of the kind described in MOORE(I), that is, by the evidence provided by putatively perceptual experiences. According to the sceptic, independent and antecedent support for ELECTION(III) and

¹ We are here in the vicinity of John Pollock's (1974) distinction between rebutting and undercutting defeaters. Wright does not commit himself to any specific analysis of the distinction between outweighing and removing evidential support, but that there is some such distinction is intuitively very plausible and I shall follow Wright in presuming upon it.

When Wright speaks of the scene 'provid[ing] no clue whether a genuine election is going on or not', should we think of this as the background information making it rational to assign equal probabilities to the genuine election possibility and the election drill possibility? If that were the situation, then the probability of the election possibility given Jones's writing an 'X' would be 0.5 (supposing that Jones would write an 'X' only if there were either a genuine election or a drill). The probability of Jones's having just voted given Jones's having just written an 'X' would also be about 0.5 and, presumably, significantly higher than its prior probability. So it would be difficult to maintain that Jones's writing an 'X' 'has no tendency whatever to support the claim that he has just voted'. The main point is that Jones's writing an 'X' does nothing to support the claim that he has just voted as against the claim that he has just taken part in an election drill.

MOORE(III) is what is needed. But while independent evidence in support of ELECTION(III) might be gathered, there is no prospect of such support for MOORE(III). If this pattern of sceptical argument is accepted then, as Wright says (*ibid.*, p. 438):

we seem bound to recognize that all our evidential commerce is founded upon assumptions for which we have no reason whatever, can get no reason whatever, and which may yet involve the very grossest misrepresentation of reality.

Of course, Wright himself is by no means committed to the sceptical conclusion. The question, though, is how to avoid it. In the second half of ‘Facts and certainty’, he notes that we could escape the sceptical bind ‘if it could be reasonable to accept a group III proposition *without reason*; that is, without evidence’ (*ibid.*, p. 450). This leads Wright to consider the possibility that there are propositions that lie outside the domain of *cognitive achievement*. These propositions would not be known in a narrow sense; but they might still be known in a more inclusive sense. As Wittgenstein says in *On Certainty* (1969), 357–9:

357. One might say: “‘I know’ expresses *comfortable* certainty, not the certainty that is still struggling.”

358. Now I would like to regard this certainty, not as something akin to hastiness or superficiality, but as a form of life. . . .

359. But that means I want to conceive it as something that lies beyond being justified; as it were, as something animal.

Wright actually explores the idea that these propositions lie outside the domain of cognitive or epistemic achievement because they lie outside the domain of truth-evaluability – they are *not fact-stating*. But it seems that the general structure of Wright’s proposal as involving narrower and more inclusive notions of knowledge or warrant could be retained even if we were not to go so far as to deny the fact-stating status of the propositions to which only the more inclusive notion (‘comfortable certainty’) applied. We might distinguish between a narrower notion of knowledge or warrant that is an *achievement* and a more inclusive notion that embraces assumptions that we are epistemically *entitled* to make. As Wright says at the very end of the paper (1985, p. 471; second emphasis added):

If . . . the concept of reasonable belief . . . embraces *certainty* in Wittgenstein’s more inclusive sense, then the argument for C [the principle mentioned in the first paragraph of this section] fails: it does not follow from the impossibility of my *achieving* cognition that I am not dreaming at *t* that I cannot be legitimately certain that I am not.

Before turning from Wright’s 1985 lecture to his paper in the present symposium, I want to draw attention to one aspect of his account of the sceptic’s arguments. In the discussion of both the (ELECTION) argument and the (MOORE) argument, Wright’s sceptic appears to take a small but important step. From the agreed point that if it were antecedently reasonable to *reject* the type-III proposition then the putative support for the type-II proposition would be removed, the sceptic moves to the apparently different claim that it is only if it is antecedently reasonable to *accept* the type-III proposition that the type-II proposition is really supported. Wright himself does not dispute this latter claim. His appeal to the idea of epistemic entitlement is intended as a way of providing the antecedent warrant that the sceptic demands. But we could dispute the claim and, in what follows, I consider epistemological accounts that do dispute it.

II

Cornerstones, entitlement, and non-transmission of warrant. In his paper in this symposium, Wright returns to the two kinds of sceptical paradox, organising his discussion around the notion of a proposition being a *cornerstone* for a given region of thought. If we were to lack warrant for the cornerstone proposition then we could not rationally claim to have warrant for any belief in the region.² In line with the first pattern of sceptical argument, the negation of the dreaming hypothesis, or of the brain-in-a-vat hypothesis, is taken to be a cornerstone for a large class of beliefs including, we may suppose, many perceptually based beliefs.³ In line with the second pattern of sceptical argument, type-III propositions are taken to be cornerstones for corresponding regions of type-II beliefs.

As in the earlier account, it is allowed that the project of assembling evidential support for the cornerstone would be futile. In the case of the first pattern (p. 168): ‘So it appears that my acquiring a warrant by empirical means for the proposition that I am not now dreaming requires that I *already have* a warrant for that same proposition. So I cannot ever acquire such a warrant (for the first time).’ In the case of the second pattern (p. 170): ‘So, again, there is a vicious circle.’ The sceptical conclusion threatens. But, as before, there appears to be a way out of the sceptical bind (p. 174):

Suppose there is a type of rational warrant which one does not have to *do any specific evidential work* to earn . . . Call it *entitlement*. If I am entitled to accept P, then my doing so is beyond rational reproach even though I can point to no cognitive accomplishment in my life . . . whose upshot could reasonably be contended to be that I had come to know that P, or had succeeded in getting evidence justifying P.

This appeal to antecedent entitlement – entitlement to our cornerstones – is the beginning of the *unified strategy* for responding to both patterns of sceptical argument.

We should ask, as Wright does in the second section of his paper, what the nature of this entitlement is. What is it an entitlement *to do*? It is initially introduced as an entitlement to ‘accept’ a proposition P, and Wright goes on to suggest that it is *not* a non-evidential warrant to *believe* a proposition P but something like a warrant to act on the assumption that P, take it for granted that P, or trust that P (p. 176). Later, he settles on the idea that entitlement is rational trust (p. 194). So the overall picture is that, against the background of our rational trust that the type-III proposition is true, the type-II proposition is supported by evidence and we have a warrant to believe it.

Someone might query whether something *less* than an antecedent warrant to *believe* the type-III proposition can really secure this favourable outcome for the type-II proposition. Wright addresses this question in the concluding section of his paper.⁴ But suppose, for the moment, that we do indeed have a warrant to believe the type-II proposition. Then a second question arises. Given that the type-II proposition obviously entails the type-III proposition,

² We should note that, in the definition of a cornerstone, Wright says that from a lack of warrant for the cornerstone proposition it would follow that one *could not rationally claim* warrant for any belief in the region – not that one *would not have* warrant for any belief in the region. At the outset, he is not explicit about the significance that he attaches to the distinction; but it looms large in the concluding section of his paper.

³ Wright, 1985, p. 431.

⁴ See his discussion of the ‘leaching’ problem, pp. 207–9.

do we end up with a warrant to believe the type-III proposition as well? Does the I-II-III argument serve to transform the lead of rational trust into the gold of justified belief? In order to address this question, we need first to sketch Wright's ideas about non-transmission of epistemic warrant – going back, once again, to his 1985 lecture.

Within his discussion of the second pattern of sceptical argument – the I-II-III pattern – Wright introduces the idea of non-transmission of epistemic warrant in the specific form of non-transmission of evidential support (1985, p. 436–7):

It simply is not true that whenever evidence supports a hypothesis, it will also support each proposition which follows from it. The important class of exceptions illustrated are cases where the support offered to the hypothesis is conditional upon its being independently reasonable to accept one in particular of its consequences.

So, for example, it may be that the evidence described in ELECTION(I) supports ELECTION(II). But this support is not transmitted to ELECTION(III) because (at least according to Wright's sceptic) the support offered for ELECTION(II) is already conditional upon its being antecedently reasonable to accept ELECTION(III). Similarly, even if the evidence described in MOORE(I) were to support MOORE(II), this support would not be transmitted to MOORE(III). For the support offered for MOORE(II) would be conditional upon its being (*per impossibile*, according to the sceptic) antecedently reasonable to accept MOORE(III).

This suggests a first shot at a general principle limiting transmission of epistemic warrant⁵ – something along the following lines:

Non-transmission of warrant

Epistemic warrant is not transmitted from the premises of a valid argument to its conclusion if the putative support offered for one of the premises is conditional on its being antecedently and independently reasonable to accept the conclusion.⁶

In his new paper, the issue of non-transmission arises again when Wright says (p. 172):

[T]ype-III propositions cannot be warranted by transmission of evidence provided by type-I propositions for type-II propositions across a type-II to type-III entailment – rather it's only if one already has warrant for the type-III proposition that any type-II propositions can be justified in the first place.

The point here is not that evidential support cannot be transmitted across the type-II to type-III entailment because, given the sceptical argument, there is no evidential support for the type-II proposition in the first place. Rather, even supposing that there is evidential support for the type-II proposition and that there is some kind of antecedent warrant, perhaps entitlement, for the type-III proposition (since it is a cornerstone), the evidential support for the type-II proposition still cannot be transmitted to the type-III proposition. It provides no additional support for the type-III proposition. As Wright puts it in another recent paper, the I-II-III arguments are not *cogent* (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

⁵ See also Wright 2000, 2002, 2003.

⁶ We can allow that the putative warrant for believing the premise of an argument might be provided by some further warranted belief or by something other than a belief such as, for example, a perceptual experience.

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.

With this much about non-transmission of warrant by way of background, we can return to the question whether, at the end of a I-II-III argument, the entailed type-III proposition ends up with any more warrant than it started out with – whether, by following through a I-II-III argument, we can perform a kind of epistemic alchemy. It seems to me that the principles governing transmission of warrant dictate a negative answer to this question. The direction of the inferential step from type-II to type-III is opposite to the direction in the space of warrants – for, according to Wright, the warrant for the type-III proposition is *antecedent* to the warrant for the type-II proposition. But this negative answer, taken together with the idea that we do indeed have a warrant to believe the type-II proposition, seems to impose some strain on our ordinary thinking about the proper management of our web of belief.

Ordinarily, we think that, if I review some of my beliefs, P_1, \dots, P_n , and notice a valid argument from those premises to *Q* then I should adopt the belief *Q* or, if other considerations argue against *Q*, then I should reconsider my beliefs P_1, \dots, P_n . If there are warrants for me to believe P_1, \dots, P_n then, if I also believe *Q*, I shall again believe something for which there is a warrant. I shall think the thing that is the thing to think. But there is a distinction between believing something that is, as it happens, the thing to think and believing something *because* it is the thing to think. If I believe P_1, \dots, P_n because there are warrants for doing so, then I do well doxastically. If I start out believing P_1, \dots, P_n because there are warrants for doing so, and I go on to believe *Q* precisely because it follows from those premises, once again I do well doxastically.

These familiar thoughts suggest that, given the obvious entailment in the I-II-III argument, if we believe the type-II proposition that is supported by the evidence described in the type-I proposition, then we should also believe the type-III proposition that is the argument’s conclusion. If considerations about non-transmission argue for going no further than the antecedent trust in the type-III proposition then we should reconsider whether belief is the proper attitude towards the type-II proposition. At the end of the second section of his paper, Wright offers a quick response to what is, I think, nearly enough this concern. But the response does not quite address head-on the question whether, if we start out with rational trust and then consider the I-II-III argument, we should, in the end, believe the type-III proposition.⁷

⁷ The concern is attributed to Stephen Schiffer; p. 177, n. 8. The response has two components. One is that closure principles for specific kinds of warrant are liable to be subject to restrictions. Wright’s example is evidential warrant; an even more obvious case is non-inferential warrant. The other component is that warrant, construed inclusively so as to encompass both evidential justification and entitlement, is subject to less restricted, or even unrestricted, closure principles.

One way of responding to the concern would be to allow that entitlement is, after all, entitlement to believe. With just that change, Moore’s argument, for example, would remain a case of non-transmission of warrant. But despite the non-transmission, there would be closure of warrant to believe. We would have a warrant for

I am not committed to the view that this concern poses a serious threat to Wright's account of epistemic entitlement. But I shall go on to describe a different notion of entitlement for which the concern does not arise.

III

The structure of entitlement. As I noted at the end of section 1, there is a striking feature of Wright's appeal to entitlement as a way of escaping the sceptical bind. He does not challenge the sceptic's claim that, in order to have a warrant for the belief that there is, say, a computer in front of me, I need an antecedent warrant for ruling out the dreaming hypothesis, the brain-in-a-vat hypothesis, and so on (the first pattern of sceptical argument). Nor does he challenge the sceptic's claim that, in order to have a warrant for believing a type-II proposition such as the proposition that I have hands, I need an antecedent warrant for believing (or at least for trusting in the truth of) a type-III proposition such as the proposition that an external world exists (the second pattern of sceptical argument). The appeal to entitlement is supposed to make good the accepted need for an antecedent warrant – despite the fact that there seems to be no way to earn such a warrant.

An alternative strategy would be to challenge the sceptic's claim. We could deny that, in general, we need all these antecedent warrants. We could do this, even while allowing that both my warrant for believing that there is a computer in front of me and my warrant for believing that I have hands are defeasible. We could allow that there are propositions that are rather like cornerstones in that a warranted doubt about such a proposition defeats a putative warrant for any belief in the corresponding region. Furthermore, we could allow that, for such a cornerstone-like proposition P, even an unwarranted doubt about P would make it impossible for one rationally to avail oneself of a warrant for any belief in the corresponding region.⁸ But a cornerstone-like proposition is not yet a cornerstone. According to the alternative strategy, as according to Wright's strategy, doubt about P would be epistemically damaging. But the alternative strategy would not, in general, allow that, in order to avoid epistemic damage, we need a positive warrant – earned or unearned – for some attitude towards P that excludes doubt. For there may be no doubt and no reason to doubt even though doubt is not excluded by a competing attitude towards the proposition P. Indeed,

believing the type-II proposition; and at the conclusion of the argument we would still have what we had at the outset, namely, a warrant for believing the type-III proposition. The warrants would be of different types. The first would be an achieved evidential warrant; the second would be a warrant of entitlement. But we should not expect unrestricted closure for specific kinds of warrant.

A more radical way of responding to the concern about the proper attitude that should be taken towards type-III propositions would be to regard those propositions as being outside the domain of our ordinary thinking about doing well doxastically. Lying outside the domain of cognitive achievement, those propositions would also be 'outside the domain of what may be known, reasonably believed, or doubted' (1985, pp. 470–1). So, even if belief were the proper attitude towards a type-II proposition, belief could not be the proper attitude towards the entailed type-III proposition.

⁸ Here we have a distinction between the conditions for having a warrant and the conditions for rationally claiming (or rationally availing oneself of) a warrant. Someone with an unwarranted doubt about a cornerstone-like proposition P may still have an undefeated warrant for believing that he has hands, for example. But he cannot rationally combine that doubt with a claim to have such a warrant.

someone might have no doubt about P and no reason to doubt P without even being able to grasp the proposition P.⁹

This, in barest outline, is the strategy that James Pryor (2000, 2004) adopts in the case of the justification of perceptual beliefs. Pryor develops an argument on behalf of the sceptic and ends with (2000, p. 532):

[The skeptic about perceptual justification] says that if you're to be justified in believing that things are as they perceptually seem to you, you need to have antecedent reason to believe that you're not in certain sceptical scenarios.

But while Wright grants the sceptic this point, Pryor disputes it (ibid.):

According to the dogmatist [about perceptual justification], when you have an experience as of *p*'s being the case, you have a kind of justification for believing *p* that does not presuppose or rest on any other evidence or justification you may have. [T]o be justified in believing *p*, you do *not* need to have the antecedent justification the sceptic demands.

Pryor does not make explicit use of any notion of entitlement and he certainly has no need for the idea that one is entitled to adopt some attitude towards the propositions for which the sceptic demands antecedent warrant – that there is an external world, that one is not dreaming, or that one's perceptual apparatus is working properly. But we could introduce a notion of entitlement into the dogmatist account in one of at least two ways.

First, since warranted doubt about one's perceptual apparatus would defeat the warrant for a perceptual belief and even unwarranted doubt would prevent one from availing oneself of that warrant, we could say that one has an entitlement not to doubt, not to call in question, or not to bother about, various things – unless there is some reason to doubt. This would be an entitlement *not to adopt the attitude of doubt* where Wright has an entitlement *to adopt the attitude of trust*. If entitlement is introduced into the dogmatist account in this first way, then it is a *negative* notion that operates at the *same* point – the cornerstone or cornerstone-like proposition – as does Wright's positive notion. It is this notion of negative entitlement that I shall compare with Wright's notion in the sections that follow.¹⁰

⁹ We should note, once again, that Wright's definition of a cornerstone has it that without a warrant for a cornerstone proposition one cannot *rationally claim* to have a warrant for any belief in the region. On the alternative conception, as presumably on Wright's conception, the requirements for rationally claiming a warrant are stricter than the requirements for having a warrant. But, on the alternative conception, in the absence of doubt (warranted or unwarranted) about a cornerstone-like proposition, one can both have a warrant, and rationally claim to have a warrant, for believing that one has hands, for example.

¹⁰ The difference between negative entitlement and Wright's notion of entitlement is not adequately captured by the distinction between not adopting the attitude of doubt and adopting the attitude of trust. Indeed, it might be said that this latter distinction marks no significant difference in cases where the question whether one's perceptual apparatus is working properly is allowed to arise. For, if one is entitled not to adopt the attitude of doubt and the question arises then, presumably, one should adopt some such attitude as belief or trust. The difference between negative entitlement and Wright's notion must be understood, rather, against the background of the difference between two epistemological frameworks. The dogmatist disputes the sceptic's demand and says that, in the absence of any reason to doubt that one's perceptual apparatus is working properly, a perceptual experience itself provides an epistemically adequate warrant for belief. This warranted perceptual belief may then figure in inferential warrants for other beliefs such as the belief that an external world exists or the belief that one's perceptual apparatus is working properly. The difference between the two epistemological frameworks is, in large part, a difference over the proper justificatory order. See below, especially sections 6 and 7. (I am indebted, here and elsewhere, to discussion with Paul Horwich.)

Second, while the sceptic says that we have no warrant for our perceptual beliefs without an antecedent warrant for a cornerstone proposition, we could say that one has an entitlement to one's perceptual beliefs provided only that there is no warrant for doubt about the cornerstone-like proposition. If entitlement is introduced into the dogmatist account in this second way, then it is a *positive* notion; indeed, it is a more positive notion than Wright's since it is entitlement *to adopt the attitude of belief* rather than just entitlement *to adopt the attitude of trust*. But it operates at a *different* point from Wright's notion. Roughly, it applies to type-II propositions rather than type-III propositions.

It is this second notion of entitlement that we find in the work of Tyler Burge. He says, for example (1993, p. 458–9):

The distinction between justification and entitlement is this: Although both have positive force in rationally supporting a propositional attitude or cognitive practice, and in constituting an epistemic right to it, entitlements are epistemic rights or warrants that need not be understood by or even accessible to the subject. We are entitled to rely, other things equal, on perception, memory, deductive and inductive reasoning, and on – I will claim – the word of others. The unsophisticated are entitled to rely on their perceptual beliefs. Philosophers may articulate these entitlements. But being entitled does not require being able to justify reliance on these resources, or even to conceive such a justification.

This passage might suggest a third way to introduce entitlement into the dogmatist account. Perhaps we could introduce a notion of entitlement *to rely on* various cognitive capacities or faculties. We could say that one is entitled to rely on the proper operation of one's perceptual apparatus unless there is some reason to think that it is not working properly. This does not, so far, sound like an entitlement to adopt any attitude towards a proposition. But, of course, reliance on one's perceptual apparatus would play a role in one's doing something, just as reliance on the proper operation of a power drill might play a role in one's putting together a piece of furniture. In particular, we should say that one is entitled to rely on the proper operation of one's perceptual apparatus *in forming beliefs* about one's perceptible environment. So, the third notion of entitlement is, after all, closely involved with the adoption of attitudes, in particular, with the formation of beliefs. Indeed, Burge says (2003a, p. 531): 'An epistemic entitlement to rely on a perceptual state or a perceptual system just *is* an entitlement to hold appropriately associated perceptual beliefs.'

IV

Unearned assumptions and negative entitlement. If the third notion of entitlement is not really separate from the second notion, then we only need to consider two notions of entitlement that might be introduced into the dogmatist account. One is the negative notion of entitlement not to doubt, not to call in question, or not to bother about, Wright's type-III propositions – unless there is a reason to doubt. The other is the notion that Burge uses: entitlement is a species of warrant for beliefs, and it applies to Wright's type-II propositions. Each notion is legitimate – though, of course, if both are in play then we need to mark the difference terminologically. But it is vital that the negative notion of entitlement not to call cornerstone-like propositions in question should be distinguished from Wright's own notion of an unearned warrant to assume, whether this is to believe or to trust, that cornerstone

propositions are true. As we shall see at the end of this section, these two notions have quite different consequences for non-transmission of warrant.

Consider again the argument (ELECTION). We are asked to imagine ‘that you live in a society which holds electoral “drills” as often as we hold fire drills, so that the scene you witness of itself provides no clue whether a genuine election is going on or not’ (Wright, 1985, p. 436). It is plausible that, if I lived in such a society, then the evidence of Jones writing an ‘X’ on what looked like a ballot paper would not constitute a warrant for my believing that Jones had just voted. From that starting point, I would need additional information, some positive reason to think that this was not a drill but a genuine election, before that evidence could provide a warrant for believing that Jones had just voted. In short, I would need to earn the assumption that an election is taking place [ELECTION(III)].

But there is another kind of case. Suppose that I live in a society where there are elections every few years and no election drills, nor even rumours of election drills. In this case it is not nearly as plausible that the evidence of Jones writing an ‘X’ on what looks like a ballot paper could not constitute a warrant for believing that Jones had just voted unless I had an antecedent positive reason for ruling out the election drill possibility.¹¹

It might be suggested that we can acknowledge the importance of the difference between these two cases even while agreeing that, in both, the evidence constitutes a warrant only against the background of an assumption that this is a genuine election and not just a drill. For, it might be said, the important distinction is between an assumption that is earned and an assumption that is unearned – perhaps a default assumption. In the first case, the background assumption has to be earned. I need to do something substantive to rule out the election drill possibility. In the second case, it might be said, the same background assumption is unearned. I just assume, or take it for granted, that an election, rather than a drill, is taking place. And I have the epistemic right to take this for granted in the absence of any reason to doubt that an election is taking place or to think that it might just be a drill. I have a warrant for my assumption even though I have undertaken no ‘specific evidential work’ to earn that warrant.¹²

This is Wright’s notion of entitlement and, whether or not he would make use of it in this particular case, the notion surely holds some appeal. It promises, not only some relief from the sceptic’s challenge, but also a fairly plausible account of the way in which, when there are no election drills nor even rumours of drills, the evidence of Jones writing an ‘X’ on what looks like a ballot paper could provide a warrant for believing that Jones had just voted.

However, it is not clear that, in every case where earning the right to a background assumption is intuitively not required, we should postulate an unearned or default

¹¹ As described, the two cases differ in whether or not there really are election drills. If there is an intuitive difference between the cases just as described, then presumably this reflects some externalist element in our conception of warrant. But we can take it that, in the first case, I have a belief – indeed, a warranted belief – that what looks like an election is quite likely to be an election drill. In the second case, I have no such belief and no warrant for such a belief.

¹² This account would, of course, face the question why, in the absence of empirical investigation, the assumption that a genuine election is taking place enjoys default status, not only descriptively – this is what I do assume – but normatively – this is what I have the epistemic right to assume. This question would have to be answered by a substantive philosophical account of our epistemic entitlements, of the kind offered by Wright in his paper. (This is not to say that Wright himself is committed to our having an entitlement to the assumption that a genuine election is taking place.)

background assumption. A thinker to whom we would, on this proposal, credit a default background assumption may not have any such assumption in mind. Even in a situation where there are regular elections and no drills, I might not have in mind any assumption to the effect that this is a genuine election and not an election drill. I might not even have the concept of an election drill.

In the case of a perceptual warrant, a thinker may simply *take* the deliverances of perceptual experience as veridical, without having in mind the assumption *that* the deliverances of perceptual experience are veridical. Certainly a thinker need not have in mind any assumptions about lighting conditions being normal, about perceptual apparatus working properly, or about not being the envatted victim of a powerful but deceptive scientist. A thinker could have a perceptual warrant for a belief while lacking the intellectual resources even to formulate such assumptions. As Burge (1993, 2003a, 2003b) argues, retaining the notion of an assumption – a kind of propositional attitude – in all such cases is an over-intellectualisation of the epistemological situation.

Consistently with Burge's or Pryor's account of the epistemological situation, and using Burge's notion of entitlement, we could say that the thinker is entitled to his perceptual belief that he has hands and that he is entitled to rely on the proper operation of his perceptual apparatus. Switching to the negative notion of entitlement, we could add that the thinker is entitled not to bother about, nor even to consider, the possibility that his perceptual apparatus might not be operating properly. But we must not slide from this to the idea that, since the thinker does not doubt that his perceptual apparatus is operating properly, he *assumes* this. For the thinker need not be capable of adopting any attitudes towards that proposition.

We might say that the negative notion of entitlement not to doubt a type-III proposition takes account of the Wittgensteinian idea that we are dealing here with 'something animal' (*OC*, 359) more fully than does Wright's notion of entitlement to trust that the proposition is true. But however that may be, there is a clear difference between the two notions in respect of their consequences for transmission of warrant.

If, as Wright does, we accept the sceptic's claim that we need an antecedent warrant for adopting some propositional attitude towards type-III propositions, then Moore's argument, for example, is a case of non-transmission of warrant. The argument involves a vicious circularity. Warrant is not transmitted from MOORE(II) to MOORE(III) because the putative support offered for MOORE(II) by the experience described in MOORE(I) is conditional on its being antecedently reasonable to accept MOORE(III). Thus, as we saw towards the end of section 2, if the antecedent warrant – entitlement – to accept MOORE(III) is an entitlement to trust, rather than an entitlement to believe, then it seems that the inference from MOORE(II) to MOORE(III) cannot provide a warrant for believing MOORE(III). And this is so, even if we do have a warrant for believing MOORE(II).

But suppose that, with Burge and Pryor, we say that there is no need for an antecedent warrant – not even for an antecedent unearned warrant. Suppose we say that the evidence described in a type-I proposition by itself supports the type-II proposition. Then there is no vicious circularity. The direction of the inferential step coincides with the direction in the space of warrants and the I-II-III argument could be a route to a first warrant for believing that there is an external world [MOORE(III)]. In short, if we reject the sceptic's claim and employ only the negative notion of entitlement for type-III propositions, then Moore's

argument is *not* a case of transmission-failure.¹³ It is for this reason that I said, at the end of section 2, that for the alternative notion of entitlement – negative entitlement – the concern about whether following through a I-II-III argument can yield a warrant to believe the type-III proposition does not arise.

In sections 6 and 7, we shall return to the consequences of the two notions of entitlement – Wright’s notion and the negative notion – for issues about warrant transmission. But first, I want to connect that difference with the idea of the presuppositions of a cognitive or epistemic project.

V

Presuppositions and entitlements of cognitive project. When, in section 5 of his paper, Wright turns to entitlements of cognitive project, he says (p. 188):

To take it that one has acquired a justification for a particular proposition by the appropriate exercise of appropriate cognitive capacities – perception, introspection, memory, or intellection, for instance – always involves various kinds of presupposition. These presuppositions will include the proper functioning of the relevant cognitive capacities, the suitability of the occasion and circumstances for their effective function, and indeed the integrity of the very concepts involved in the formulation of the issue in question.

A presupposition of a cognitive project is defined as a proposition P for which: ‘to doubt P (in advance) would rationally commit one to doubting the significance or competence of the project’ (p. 191).

Some presuppositions have the feature that, although it would be possible to undertake an investigation as to their truth or falsity, an attempt to provide evidence of their truth would involve a further cognitive project with its own presuppositions ‘of no more secure a prior standing’, and an attempt to provide evidence of the truth of these presuppositions in turn would involve yet a further cognitive project, and so on. In short, a presupposition P may have the feature that to accept ‘an onus to justify P’ would be to ‘undertake a commitment to an infinite regress of justificatory projects’ (p. 191). The key idea about entitlements of project is that such a presupposition is an entitlement provided only that (ibid.): ‘We have no sufficient reason to believe that P is untrue.’ And to say that a presupposition is an entitlement is to say that ‘we should – are rationally entitled to – just go ahead and trust’ that the presupposition is met (p. 192).

According to the definition of a presupposition, doubt about the presuppositions of an epistemic project would be epistemically damaging. It is consistent with the definition to suppose that warranted doubt, or a warrant for doubt, about a presupposition would defeat any putative warrant that an epistemic project might yield and that even unwarranted doubt would make it impossible for one rationally to avail oneself of any warrant that the project might yield. So, rational pursuit of an epistemic project requires negative entitlement to the project’s presuppositions. But Wright’s account of entitlements of cognitive project goes beyond this. For it suggests that, in order rationally to carry out an epistemic project, and in order rationally to take oneself to have arrived at a reason to believe a particular conclusion, one needs to adopt an attitude – trust – towards the presuppositions of that project. So, do we

¹³ See Pryor, 2004.

need to assume, trust, or believe that the presuppositions of our epistemic projects are true? Or is it enough that we should not doubt those presuppositions or call them into question? Do we need entitlement as Wright conceives it or only negative entitlement?

In earlier work,¹⁴ I have considered the presuppositions of epistemic projects in the context of arguments about self-knowledge and externalism about content.¹⁵ I have been particularly concerned with the presupposition that there is such a proposition to think as the purported proposition that formulates the issue on which the project is focused. I take it that this is encompassed in one of the presuppositions that Wright mentions, namely, the presupposition of ‘the integrity of the very concepts involved in the formulation of the issue in question’.

In his seminal contribution to the topic of externalism and self-knowledge, 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.

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 there seems to be a problem with combining self-knowledge and externalism. For it seems that self-knowledge and philosophical theorising together provide a route to too-easy knowledge that certain environmental conditions obtain (McKinsey, 1991).

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.’ In my view, this is the heart of the solution to the apparent problem posed by self-knowledge and externalism.¹⁶ But my present concern is with two different ways of interpreting Burge’s remark about presupposition.

On one possible interpretation – *not* the one that Burge intends – 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. I do not need to earn an antecedent positive warrant for the assumption; rather, I am entitled to make the assumption. According to this first interpretation, Burge’s account of authoritative self-knowledge makes use of something like Wright’s notion of entitlement to empirical background assumptions.

¹⁴ Davies, 1998, 2000, 2003a, 2003b.

¹⁵ See, for example, Davies (1998, p. 354): ‘In any given epistemic project, some propositions will have a presuppositional status. Suppose that the focus of the project P is the proposition A, and that the investigation is carried out using method N. Then within P it is presupposed, for example, that A is a hypothesis that can be coherently entertained (can be believed, doubted, confirmed, disconfirmed); and it is also presupposed that N is a method that can yield knowledge, at least with respect to A.’

¹⁶ See especially Davies, 2003b.

This appeal to unearned assumptions seems once again to over-intellectualise the epistemological situation. For it is not especially plausible that, just in thinking that *I am thinking that water is wet*, I assume that some particular environmental conditions E obtain. 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 proposition to think as that *I am thinking that water is wet*. Perhaps it is only philosophical theorising that leads me from this basic assumption to the further assumption that environmental conditions E obtain. 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 – or trust or believe – 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 (2003b, p. 264):¹⁷

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 Q then I must be able to think or entertain that proposition. The requirements for thinking the thought are presupposed in the thinking. This does not mean that, in order to be warranted in believing Q, I must have either an earned or an unearned warrant for assuming that those presupposed conditions C are met. But still, a warranted belief that the conditions C are *not* met would defeat the warrant for Q. And even an unwarranted belief that those conditions are not met would make it impossible for me rationally to avail myself of the warrant for Q.¹⁸

If this is right then, at least in the case of one of the presuppositions of an epistemic project, it is just negative entitlement that is required.

¹⁷ I say (Davies, 2003b, p. 117–8): ‘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 (2003b, p. 264): ‘I do *not* assimilate this notion of presupposition to a notion of assumption by the individual, as Davies conjectures.’

¹⁸ If I am to have a warrant for believing Q then I must not have a warranted doubt about whether the presupposed conditions C are met. If I am rationally to claim to have a warrant for believing Q then I must not have a doubt – warranted or not – that the conditions C are met.

VI

The problem of easy knowledge. The title of Wright’s paper implicitly suggests that one problem that an account of entitlement might face is that it should make knowledge or warrant too easy. Indeed, this problem arises for the dogmatist epistemology of perceptual beliefs into which the negative notion of entitlement fits so naturally. I explain the problem in this section and then, in the next section, propose a solution.

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 similar to the one that Wright (1985) develops from his reflections on the intuitive inadequacy of Moore’s argument.¹⁹ In Cohen’s exposition, it 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. . . . Skepticism threatens.

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

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

To deny that principle (KR) and accept 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 example (ibid., pp. 312–3):

TABLE(I) I am having an experience as of this table being red.
TABLE(II) This table is red.
 If this table is red then it is not the case that this table is white but illuminated by red lights.

Therefore:

TABLE(III) 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(II) is constituted by the visual experience described in TABLE(I), provided only that I do not have any reason to doubt that the lighting conditions are normal or that my

¹⁹ It is also similar to the argument that Pryor develops on behalf of the sceptic in section 2 of ‘The skeptic and the dogmatist’ (2000). See Pryor (2004) for his way of responding to Cohen’s (2002) problem of easy knowledge.

perceptual apparatus is working properly. In particular, my warrant for believing TABLE(II) does not depend on my having any antecedent warrant for believing TABLE(III). I have a simple *a priori* warrant for believing the conditional premise; 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’.

Wright (2003, pp. 60–3) discusses this same example – though with a wall instead of a table. His view, like Cohen’s, is that one could not, just by looking at a surface, acquire a warrant for believing that it is not illuminated by red lights. And, according to Wright, the principles governing transmission of warrant have the desired result that the evidential warrant for TABLE(II) is not transmitted to TABLE(III) (ibid., p. 61):

While you have – no doubt quite justifiably – taken it for granted that the conditions were generally suitable for the acquisition of reliable information by casual-perceptual means, it would be absurd to pretend that you had *gained a reason for thinking so . . .* just by dint of the fact that those specific possibilities [such as deceptive lighting] are logically excluded by the beliefs which, courtesy of your background assumption, you have now confirmed.

The warrant for TABLE(II) is not transmitted to TABLE(III) because it only counts as a warrant in the first place because of my antecedent warrant for assuming, or taking for granted, TABLE(III).

However, suppose that, as against Wright’s view, there is no need for an antecedent warrant – not even an antecedent unearned warrant – for assuming, trusting, or believing that TABLE(III) is true. Suppose that the evidence described in TABLE(I) by itself supports TABLE(II). Then, not only do I have an evidential warrant for believing TABLE(II), but also, by following through the *modus ponens* argument, I can gain – perhaps for the first time – a warrant for believing TABLE(III). When the evidence described in TABLE(I) is the warrant for TABLE(II), the argument from TABLE(II) and the conditional premise to TABLE(III) is *cogent* and the principle limiting transmission of epistemic warrant (section 1) is not triggered. So the epistemological view that goes naturally with the notion of negative entitlement faces the problem of easy knowledge in a way that Wright’s view does not.

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(0) I am having an experience as of this table being red.

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), constituted by the visual experience described in EVIDENCE(0). Furthermore, I have a first-personal warrant for believing EVIDENCE(2). From EVIDENCE(1) and EVIDENCE(2), it follows that the table

is the way it visually appears to be. But, according to Cohen, it is implausible that I could in this way come to know that, at least on this occasion, my colour vision operated correctly or that I could in this kind of way amass evidence to support the claim that my colour vision is generally reliable.

VII

Limiting the rational deployment of warrants. If we adopt the notion of negative entitlement and the epistemological framework into which it naturally fits, then Moore's argument turns out not to be an example of transmission-failure (section 4). Furthermore, as we have just seen, looking at a table turns out to be a way of gaining a warrant for believing, first, that it is not a white table illuminated by red lights and, second, that one's colour vision is operating correctly. Earlier, we noted that, if warrant is transmitted from premises to conclusion in Moore's argument, then a concern (developed towards the end of section 2) does not arise. That is a good result. But there is also a cost. For if Moore's argument is not an example of transmission-failure then we need something else to say about why it is unsatisfying as a response to the sceptic.

That is one problem for negative entitlement; and Cohen's problem of easy knowledge is another. Both problems have the same source, namely, the shift from Wright's notion of entitlement and his acceptance of the sceptic's demand to the negative notion of entitlement and rejection of the sceptic's demand. For, once that shift is made, the (MOORE) argument, and the (TABLE) and (EVIDENCE) arguments, no longer trigger the principle that limits transmission of warrant as Wright defines that notion. Thus in the case of the (EVIDENCE) argument, for example, Wright will claim that the experience described in EVIDENCE(0) counts as a warrant for believing EVIDENCE(1) only because we have an antecedent warrant – entitlement – for assuming the truth of the conclusion EVIDENCE(3). But, after the shift, we have to deny this. In this section, I want to sketch an idea that offers an alternative response to the problems posed by transmission of epistemic warrant in these arguments.

If we replace Wright's notion of entitlement to trust with the purely negative notion, then we deny that a lack of positive warrant for a cornerstone-like proposition P is, by itself, epistemically damaging. But warranted doubt about P is still damaging. For it defeats or removes a putative warrant for any belief in the corresponding region. And if that is so then even an unwarranted doubt is damaging to our pursuit of epistemic projects. For a doubt about P rationally commits one to not availing oneself of a warrant for any belief in the corresponding region.²⁰

If this much is right then we can, it seems to me, take one further step. For suppose that we undertake an epistemic project whose conduct is conditioned by an initial supposition that P is, or may very well be, false. Suppose, in short, that the conduct of the project is conditioned by suppositional doubt about P. Then, within the context of that project it is not

²⁰ Pryor (2004) draws a distinction between 'what you have *justification* for believing, and what you are *rationally committed to believing*, given the beliefs and doubts you in fact already have'. He then says (ibid.): '[S]uppose you suspect that your color vision might not be working properly. This doubt is in fact unjustified, but you have not realized that. In any case, I'm inclined to say that your doubt would make it irrational for you to form any beliefs about color, on the basis of your visual experiences. Even though your experiences might very well be giving you justification for those beliefs.'

possible rationally to avail oneself of those warrants that would, in reality, be defeated by a warranted doubt about P.

Now there is a kind of epistemic project whose conduct is conditioned by suppositional doubt, namely, the project of *settling the question* whether a particular proposition is true. As I understand this idea, the epistemic project of settling a question begins with my regarding that question as open *pro tem*. By this, I do not mean that I work myself into a state of really doubting that the proposition is true. When I begin by regarding a question as open, I suppose, for the purposes of the project, that the proposition is, or may very well be, false. This initial supposition then conditions my conduct of the project. In particular, it conditions my rational deployment of epistemic warrants as I try to answer – to settle – the question that I have begun by regarding as open. If I try to settle the question whether or not some cornerstone-like proposition P is true then I begin by suppositionally doubting that P is true and my conduct of the question-settling project is conditioned by that suppositional doubt. But then, within the context of that project, I cannot rationally avail myself of a warrant for any belief in the corresponding region. So, in particular, I cannot deploy those warrants to settle the question in favour of P. This is so even though in reality – outside the context of the question-settling project – I do have warrants for those beliefs. And it is so whether in reality I have a positive entitlement to assume P (as on Wright's account) or just a negative entitlement not to call P into question (as on the alternative account).

To see how a suppositional doubt conditions the conduct of a question-settling project, imagine that I undertake the project of settling the question whether or not there is an external world as ordinarily conceived – whether instead, perhaps, I am the envatted victim of a powerful but deceptive scientist. In particular, imagine that I attempt to settle that question in favour of an external world by deploying the warrants that I have for believing the premises of Moore's argument. I begin the project by regarding the question of the truth of the conclusion of Moore's argument as open *pro tem*. So my conduct of the question-settling project is conditioned by the initial supposition that Moore's conclusion is, or may very well be, false.

As the sceptic points out, if I really believed what I now suppose – if I really doubted Moore's conclusion – then I could not rationally regard my experience as constituting a warrant for believing MOORE(1). Just so, within the context of a project whose conduct is conditioned by the suppositional doubt, I cannot rationally avail myself of the warrant for believing MOORE(1). In short, I cannot settle the question whether or not the conclusion of Moore's argument is true – and, in particular, cannot settle it in favour of Moore's conclusion – by deploying the epistemic warrants that I have for believing the premises.

The idea of the epistemic project of settling a question offers – I hope – an alternative response to the problems posed by transmission of epistemic warrant. In a fuller treatment, this idea would be set against the background of a distinction between two kinds of epistemic project – deciding what to believe and settling a question – and two notions of transmission of epistemic warrant. The two kinds of epistemic project correspond to two dialectical purposes of arguing that Frank Jackson calls the teasing-out purpose and the convincing purpose.²¹ Principles limiting the two kinds of transmission of warrant correspond to limitations on arguments that can be effectively propounded for each of the two purposes and so to two notions of how an argument can beg the question. But all of that is for another

²¹ Jackson, 1987, chapter 6.

paper.²² Here, I only want to indicate how the idea of settling a question gives us something to say about the intuitive inadequacy of Moore's argument and about Cohen's problem of easy knowledge.

In the concluding section of his paper, Wright says (p. 206):

In general, it has to be recognised that the unified strategy can at most deliver a *sceptical solution* . . . The unified strategy . . . concedes the basic point of the sceptical arguments to which it reacts, namely that we do indeed have no claim to know, in any sense involving possession of evidence for their likely truth, that certain cornerstones of what we take to be procedures yielding knowledge and justified belief hold good.

In the alternative epistemological framework that I have been exploring, the warrants for the premises of Moore's argument provide us with a reason to believe the conclusion. Warrant is transmitted, in Wright's sense: MOORE(III) is what we should believe. But we cannot rationally avail ourselves of the warrant for MOORE(II) within the epistemic project of settling the question whether or not MOORE(III) is true – even though, in reality, it is an epistemically adequate warrant. This is the point at which we concede something to the sceptic. Moore's argument furnishes a warrant, but not a question-settling warrant, for MOORE(III).²³

Turning now to the first of Cohen's two examples of the problem of easy knowledge, we can say that it is 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(II) and the conditional premise. But the unwanted result is not a consequence of allowing basic knowledge. Suppose that I were to doubt the truth of the conclusion TABLE(III). Suppose that I were to believe that this is, or may very well be, a white table illuminated by red lights; in short, that the lighting is, or may very well be, deceptive. This belief – this doubt – would rationally commit me to not availing myself of the perceptual warrant for believing TABLE(II). Just so, if I were to begin an epistemic project by suppositionally doubting TABLE(III) then, within the context of that project, I could not rationally avail myself of the perceptual warrant for believing TABLE(II). I have an epistemically adequate warrant for believing TABLE(II) and warrant is transmitted, in Wright's sense: TABLE(III) is what I should believe. But I cannot deploy my warrant for TABLE(II) in order to settle the question whether or not TABLE(III) is true.

Finally, the situation is similar with the second of Cohen's examples. If I were to doubt the truth of EVIDENCE(3) then I could not rationally take my visual experience as of the table being red as providing any support at all for EVIDENCE(1). Just so, in the context of the project of settling the question whether or not EVIDENCE(3) is true, I cannot rationally avail myself of the perceptual warrant for EVIDENCE(1).

VIII

Conclusion. Wright accepts the sceptic's demand that there should be antecedent epistemic warrants for the propositions that Wright calls 'cornerstones' and for the presuppositions of

²² 'Two purposes of arguing and two epistemic projects', forthcoming.

²³ It remains to connect this point with Wright's distinction between knowing and having a *claim* to know (p. 206, n. 22).

cognitive and epistemic projects. Epistemic entitlement, as Wright conceives it, is a kind of epistemic warrant, though not a warrant that is earned.

In this paper, I have been exploring the prospects for a different notion of entitlement – I have called it ‘negative entitlement’. This is not a kind of epistemic warrant, for negative entitlement is not an entitlement to assume, trust, or believe any proposition. Since we can scarcely hope to earn an antecedent warrant for Wright’s cornerstone propositions, and since it is not coherent to aim to earn a warrant for the presuppositions of every epistemic project in which we engage, negative entitlement fits naturally into an epistemological framework in which the sceptic’s demand is disputed. Thus, for example, negative entitlement goes naturally with the dogmatist account of perceptual justification that Pryor favours.

If the sceptic’s demand is accepted then there are examples of non-transmission of epistemic warrant, as Wright defines that notion. These examples include Moore’s argument and the two problematic arguments that Cohen considers. So we have a ready account of the intuitive inadequacy of Moore’s argument as a response to the sceptic; and we avoid Cohen’s problem of easy knowledge. But if the sceptic’s demand is disputed then Moore’s argument is not an example of transmission-failure and neither are Cohen’s two arguments. So we need to find something else to say about why Moore’s argument is unsatisfying; and we need a solution to the problem of easy knowledge. I have tried to meet these needs by appealing to the idea that, even where warrant is transmitted from the premises to the conclusion of an argument, still it may not be possible rationally to deploy the warrants for the premises in order to settle the question whether or not the conclusion is true.

I have defended negative entitlement and the epistemological framework into which it fits by responding to the objections that arise out of the differences from Wright’s account over warrant transmission. But I have been relatively cautious in arguing for negative entitlement and against Wright’s notion. Still less have I ventured any substantive account of our epistemic entitlements as Wright does in the central sections of his paper. At the outset, I contrasted epistemic entitlement with epistemic achievement. But while entitlement is not itself an epistemic achievement, understanding the nature and source of entitlement would certainly be an epistemological achievement of some magnitude.²⁴

*Philosophy Program, Research School of Social Sciences
Australian National University
Canberra, ACT 0200
Australia
Email: martin.davies@anu.edu.au*

²⁴ My first thoughts about warrant transmission and about transmission-failure were prompted by reading Crispin Wright’s British Academy Lecture. I have continued to learn from that lecture and from Wright’s subsequent writings on the topic. I am much indebted to Tyler Burge and Jim Pryor – to their writings on epistemology and, in recent years, to conversations with each of them. Special thanks to Mark Greenberg for comments and conversations over many years and for his detailed comments on a penultimate draft.

REFERENCES

- Burge, Tyler 1988: Individualism and self-knowledge. *Journal of Philosophy*, 85, 649–63.
- Burge, Tyler 1993: Content preservation. *Philosophical Review*, 102, 457–88.
- Burge, Tyler 2003a: Perceptual entitlement. *Philosophy and Phenomenological Research*, 67, 503–48.
- Burge, Tyler 2003b: Reply to Martin Davies. In M.J. Frápolli and E. Romero (eds), *Meaning, Basic Self-Knowledge, and Mind: Essays on Tyler Burge*. Stanford: CSLI Publications, 250–7.
- Cohen, Stewart 2002: Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65, 309–29.
- Davies, Martin 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, Martin 2000: Externalism and armchair knowledge. In P. Boghossian and C. Peacocke (eds), *New Essays on the A Priori*. Oxford: Oxford University Press, 384–414.
- Davies, Martin 2003a: The problem of armchair knowledge. In S. Nuccetelli (ed.), *New Essays on Semantic Externalism and Self-Knowledge*. Cambridge, MA: MIT Press, 23–55.
- Davies, Martin 2003b: Externalism, self-knowledge and transmission of warrant. In M.J. Frápolli and E. Romero (eds), *Meaning, Basic Self-Knowledge, and Mind: Essays on Tyler Burge*. Stanford: CSLI Publications, 105–130.
- Jackson, Frank 1987: *Conditionals*. Oxford: Blackwell Publishers.
- McKinsey, Michael 1991: Anti-individualism and privileged access. *Analysis*, 51, 9–16.
- Moore, G.E. 1959: Proof of an external world. In *Philosophical Papers*. London: Allen and Unwin, 127–50.
- Pollock, John 1974: *Knowledge and Justification*. Princeton, NJ: Princeton University Press.
- Pryor, James 2000: The skeptic and the dogmatist. *Noûs*, 34, 517–49.
- Pryor, James 2004: Is Moore’s argument an example of transmission-failure? *Philosophical Perspectives*, 18.
- Wittgenstein, Ludwig 1969: *On Certainty*. Oxford: Basil Blackwell.
- Wright, Crispin 1985: Facts and certainty. *Proceedings of the British Academy*, 71, 429–72.
- Wright, Crispin 2000: Cogency and question-begging: Some reflections on McKinsey’s paradox and Putnam’s proof. *Philosophical Issues*, 10, 140–63.
- Wright, Crispin 2002: (Anti-)sceptics simple and subtle: G.E. Moore and John McDowell. *Philosophy and Phenomenological Research*, 65, 331–49.
- Wright, Crispin 2003: Some reflections on the acquisition of warrant by inference. In S. Nuccetelli (ed.), *New Essays on Semantic Externalism and Self-Knowledge*. Cambridge, MA: MIT Press, 57–77.