

Published in *Synthese* 130 (2002), 279-302.

## MCKINSEY PARADOXES, RADICAL SCEPTICISM, AND THE TRANSMISSION OF KNOWLEDGE ACROSS KNOWN ENTAILMENTS

ABSTRACT. A great deal of discussion in the recent literature has been devoted to the so-called ‘McKinsey’ paradox which purports to show that semantic externalism is incompatible with the sort of authoritative knowledge that we take ourselves to have of our own thought contents. In this paper I examine one influential epistemological response to this paradox which is due to Crispin Wright and Martin Davies. I argue that it fails to meet the challenge posed by McKinsey but that, if it is set within an externalist epistemology, it may have application to a related paradox that concerns the problem of radical scepticism.

### 1. THE MCKINSEY PARADOX

According to semantic externalism, some of a subject’s thought contents are individuated, at least in part, by factors within that subject’s environment.<sup>1</sup> Seemingly, however, such a theory stands in direct opposition to the highly intuitive thesis—what I shall call the ‘privileged access’ thesis—that one can have *a priori* knowledge of one’s own thought contents, where this is understood as the claim that one can know them purely by reflection (i.e., without an empirical investigation). Crudely put, the concern is this: If semantic externalism and the privileged access thesis are both true, then it would seem to follow that one could gain knowledge of one’s ‘wide’ thought contents—and thus of non-mental *empirical* facts—by conducting a purely *a priori* investigation.<sup>2</sup>

Much of the recent debate in this respect has tended to cluster around a neat formulation of the tension due to Michael McKinsey (1991) which we shall represent as follows, where ‘K<sub>AP</sub>’ refers to *a priori* knowledge:

The McKinsey Argument

- (1) K<sub>AP</sub> [(MK1) I have a mental property, M].
- (2) K<sub>AP</sub> [(MK2) If I have mental property M, then I meet external condition E].
- (3) K<sub>AP</sub> [(MK3) Therefore, I meet external condition E].

The first line of the argument simply captures a paradigm case of an agent’s *a priori* knowledge, gained via privileged access, that he has a certain mental property. The second line concerns that same agent’s *a priori* knowledge, via the philosophical thesis of semantic

externalism, that the possession of this mental property entails that a certain external condition has been met—an empirical non-mental truth about that subject’s environment. Finally, the third line puts these two claims together to conclude that the subject must, therefore, have *a priori* knowledge that this empirical condition has obtained. Intuitively, since the external condition in question seems precisely the sort of thing that cannot be known *a priori*, this represents a *reductio* of the supposed compatibility of semantic externalism and privileged access. One of them must go.

In order to see this incompatibility in more detail, consider the following influential argument—due to Hilary Putnam (1975)—which has been used to motivate a form of semantic externalism. Let twin earth be like earth in every respect except that there is no water, just a qualitatively indistinguishable substance—‘twater’—that has the abbreviated microstructure XYZ. Now suppose that I, and my twin on twin earth, both have exactly the same narrow mental states,<sup>3</sup> and are equally as ignorant of chemistry. According to Putnam, even despite our shared narrow mental states, our respective utterances of sentences containing the word ‘water’ will have different semantic properties. On earth, my use of this word will refer to H<sub>2</sub>O, whereas on twin earth, my twin’s use of this word will refer to XYZ. On the assumption that meaning determines reference (so that a difference in reference suffices for a difference in meaning), it follows that my use of this word will mean water while my twin’s use of the word will mean twater. Hence, as Putnam famously put it, “meanings ain’t in the head”.

According to Paul Boghossian (1997), however, the combination of this proposal and privileged access leads to *a priori* knowledge of empirical truths. Boghossian takes Putnam to be arguing that the possession of water thoughts *entails* the existence of water (of H<sub>2</sub>O). As a result, it would appear to follow that *a priori* knowledge of one’s thought contents in this instance, coupled with *a priori* knowledge of the truth of semantic externalism (and its application to this particular concept), would entail *a priori* knowledge of the fact that water exists, which is clearly an empirical truth about the world. We thus get the following McKinsey-style argument:

#### The Putnam-McKinsey Argument

- (1) K<sub>AP</sub> [(P1) I have the concept water].
- (2) K<sub>AP</sub> [(P2) If I have the concept water, then water exists].
- (3) K<sub>AP</sub> [(P3) Water exists].

And since the existence of water seems to be the very sort of thing that is settled by empirical investigation and not by armchair reflection, it follows that we have an instance of the

McKinsey style *reductio* of the conjunction of privileged access and semantic externalism that we saw above.

One natural way to respond to a paradox of this sort is to decline to take the puzzle at face value and instead undertake a detailed investigation into the nature of the theses of semantic externalism and privileged access in order to see whether the argument can be blocked or, despite the odds, the conclusion accepted. In contrast, the more unorthodox approach to the problem that I shall be considering in this paper involves granting the truth of something like the above formulations of these two theses and investigating instead the *epistemological* intuitions that are at work here. It is important to note that this approach runs orthogonal to the standard line because, in principle at least, they could *both* offer plausible solutions to this puzzle. It could be, for instance, that not all construals of semantic externalism and privileged access license this problematic conclusion *and* that certain epistemological considerations show that, in any case, the inference to *a priori* knowledge does not go through. Accordingly, the epistemological account of the problem need not necessarily be viewed as being in *competition* with any of the current non-epistemological responses on offer.

## 2. THE WRIGHT-DAVIES PROPOSAL

Of the epistemological attempts to dissolve the McKinsey paradox, the most influential has been that put forward by Martin Davies (1998) and developed by Crispin Wright (2000a; cf. Wright 1997). Essentially, they argue that what arguments of this form expose is the fact that even valid arguments can have premises that are true and known and yet the conclusion be such that this knowledge will not *transmit* across the entailment to the conclusion.<sup>4</sup>

The best way to understand this notion of ‘Transmission’ is in contrast to the related (and weaker) epistemic principle of ‘Closure’ that Davies and Wright both accept. This can be (roughly) expressed as follows:

$$K_C: \quad \{K[\varphi] \ \& \ K[\varphi \rightarrow \psi]\} \rightarrow K[\psi]$$

In words, if an agent knows a proposition,  $\varphi$ , and knows that  $\varphi$  entails a second proposition,  $\psi$ , then that agent knows  $\psi$  also.<sup>5</sup> More informally still, one might express this principle as the claim that knowledge always transfers across known entailments. Note that the McKinsey template that we have just considered pivots upon an instance of this principle (where the ‘K’ operator picks out *a priori* knowledge throughout). Denying Closure could thus be one way in

which one might respond to this puzzle.<sup>6</sup> This is not, however, the course of action taken by Davies and Wright because they believe that there is a stronger sister principle to Closure in play here, which they call ‘Transmission’, and that it is this principle that is essential to the formulation of the paradox.

Transmission is held to be a stronger principle than Closure because it makes the extra demand that the knowledge that transfers across the known entailment should also preserve what Wright refers to as the “cogency” of the argument, which is its aptitude to produce *rational conviction*. Here is Wright:

A *cogent* argument is one whereby someone could be moved to rational conviction of the truth of its conclusion. (Wright 2000a, 140)

In essence, the idea is that the McKinsey argument, though not a counterexample to Closure, is nevertheless a case in which Transmission fails. That is, although the knowledge at issue in the premises will *transfer* to the conclusion (in the sense that if the premises are known then the conclusion must be known also), that knowledge will not *transmit* to the conclusion because the argument lacks cogency.

We shall examine the reasons that Davies and Wright offer for this contention in a moment. First, however, it is worthwhile outlining what dialectical consequences it has. To begin with, it is important to note that since Davies and Wright retain Closure it follows that if the agent does know the premises ((MK1) and (MK2)) of the McKinsey argument, then he must also know the conclusion, (MK3). What the Davis-Wright line denies is thus not that the agent might know the conclusion, but rather the more problematic contention that the agent knows the conclusion *on the basis of* legitimately inferring that conclusion from his knowledge of the premises. This point is important because what is problematic about the argument is not that the agent knows the conclusion, but rather that this knowledge has been derived in a purely *a priori* fashion by inferring it from the premises. If the Davis-Wright manoeuvre holds water, then this latter claim will be blocked because the agent will be unable to legitimately infer knowledge of the conclusion from his knowledge of the premises, and thus the argument will offer us no reason for thinking that the agent knows the conclusion on *a priori* grounds.

According to Davies and Wright, the reason why the McKinsey argument lacks cogency is because it is *question-begging* in the sense that an epistemic entitlement to believe the conclusion is already presupposed, without independent motivation, in the epistemic entitlement that the agent takes himself to have for believing the premises. In particular, what the McKinsey argument illustrates, according to Davies and Wright, is that an epistemic

entitlement to believe that certain external conditions have obtained is already taken for granted in one's putative *a priori* knowledge of the wide contents of one's own mental states. Given that this is so, it should come as no surprise to find that Closure holds since if one has knowledge of the premises and such knowledge presupposes an epistemic entitlement to believe the (true) conclusion, then how could it not be that the conclusion is not also known? Nevertheless, such an argument would provide no reason whatsoever to believe the conclusion of the argument, and thus Transmission would fail. One who doubted the conclusion would, for instance, be offered no reason to change his mind upon being presented with such a question-begging argument. As Wright expresses the matter:

Here is the point in a small nutshell. Suppose that I want to describe myself—truly as it happens—by using the words: “I believe that water is wet”. If some form of externalism is true which would be strong enough to sustain the second premise of the McKinsey argument, then there are external preconditions of my expressing a true belief by those words whose satisfaction I may nevertheless, without compromise of the warrant for my claim, have done nothing special to ensure. Can that warrant now licitly be extended to the claim that those preconditions *are* met? It should seem obvious that it cannot, for the simple reason that the warrant is in the first place conditional on the concession, as it were, that unless there is extant positive reason to doubt that the external conditions are met, the possibility that they are not can be ignored. Given that the warrant I start out with has this concession-dependent character, it naturally cannot be massaged by inference into a reason for a positive view about the issue which the concession was precisely a concession to take for granted. (Wright 2000a, 156)

In a still smaller nutshell, we might put the point by saying that one only has *a priori* access to one's wide thought contents on the assumption that certain empirical conditions are legitimately disregarded from the outset—such as the possibility that, appearances notwithstanding, one in fact lives on a ‘dry’ earth on which there is no water. Accordingly, one cannot, *contra* the McKinsey argument, reason one's way to *a priori* knowledge that such presuppositional empirical conditions obtain.

Of course, unlike other arguments where it is clear that the question is being begged in this way, it is far from transparent that the McKinsey argument is question-begging. The guiding diagnostic thought that underpins the Davies-Wright proposal, however, is that in every discourse there are propositions which play a presuppositional role—and hence which one cannot know via participation in that discourse—but that sometimes these presuppositions are not recognised as such by the participants of that discourse. By making explicit which propositions perform this ‘presuppositional’ function in any given discourse we can thus evade a great deal of philosophical confusion. As Davies puts it:

The intuitive idea [...] is something like this. In any given epistemic project, some propositions will have a presuppositional status. Suppose that the focus of project P is the proposition A, [... *and*] suppose that B is some proposition that has this presuppositional status in project P. Then P cannot itself yield

knowledge that B; nor can P play an essential role in yielding knowledge that B. (Davies 1998, 354)

Or, as Wright more succinctly expresses the matter:

Every epistemic project incorporates certain presuppositions which, whatever its upshot, it cannot provide warrant for. (Wright 2000a, 150)<sup>7</sup>

Note, however, that to say that knowledge of a presupposition of a certain epistemic project cannot be gained via that project is not to say that the presupposition is unknowable, only that it cannot be known via that route. This point is important because the truth of the presupposition will no doubt be entailed by many of the truths believed in that epistemic project. Insofar as there is knowledge in that discourse (and insofar as Closure holds), then one would expect the (known) presuppositions to be known as well. The point is thus not that the (known) presuppositions of a certain epistemic project cannot be known (which would almost certainly constitute a counterexample to Closure), but rather that they are not known via the epistemic project for which they play the presuppositional role. In particular, in the specific case of the McKinsey paradox they are not known via that particular *a priori* argument.

Davies and Wright attempt to make this line of argument compelling by showing how it serves to explain the inadequacies of other arguments that are meant to be analogous, in the relevant respects, to the McKinsey argument. One example that they both consider is G. E. Moore's (1925; 1939) infamous 'proof' of an external world on the basis of his putative knowledge that he has two hands,<sup>8</sup> an argument, which, absent the stipulation that the knowledge that is at issue always be *a priori*, can be expressed in terms of the familiar McKinsey template:

The Moorean Argument

- (1) K [(M1) I have two hands].
- (2) K<sub>AP</sub> [(M2) If I have two hands, then there exists an external world].
- (3) K [(M3) There exists an external world].

One who already doubted the existence of an external world would not be moved to change his mind by an argument of this form. Rather, he would simply retort that in claiming knowledge of (M1) Moore is thereby already *presupposing* that he has an epistemic entitlement for his belief in the conclusion of the argument, since it is only on this presupposition that he is in a position to coherently claim knowledge of (M1) in the first place. That is, the sceptic we have in mind here will argue as follows. Given that there is no *a priori* defence of his belief in (M3),<sup>9</sup> Moore's knowledge of (M1) could only be supported by empirical evidence, evidence which, moreover, could only supply sufficient epistemic support for (M1) on the supposition

that Moore was *already* epistemically entitled to believe (M3). Without this assurance, the evidence would be necessarily insufficient, and so not apt to support the kind of epistemic support for belief in (M1) necessary to buttress a knowledge claim. Hence, Moore's claim to know (M1) can only provide, at best, question-begging support for his claim to know the conclusion of the argument, (M3).

Moreover, as before, note that, at least as it stands, this feature of the Moorean argument does *not* suffice to indicate that it is a case in which Closure fails. Intuitively, if Moore does know that he has two hands, then he must know that there is an external world. Nevertheless, the point is that such an argument to this conclusion lacks cogency. If Moore does know this conclusion, then it is not as a result of a legitimate inference from his putative knowledge of (M1) and (M2) because such knowledge already presupposes an epistemic entitlement to believe the conclusion. The Moorean argument is thus a case in which, on closer analysis, we discover that Moore is simply presupposing that which he sets out to establish. No wonder, then, that the Moorean argument can seem to completely by-pass the sceptical issue.

Davies and Wright thus support their diagnosis of the McKinsey argument by showing how such an argument has an analogous form to another argument, the Moorean argument, which is more obviously question-begging and so lacking in cogency. Further support for this analogy can be offered by considering how one could just as well extend the Putnam version of the McKinsey paradox that we saw above along similar anti-sceptical lines.

Take the 'brain-in-a-vat\*' sceptical hypothesis to mean that hypothesis discussed by Putnam (1981, chapter 1) which stipulates that *everyone* has *always* existed as brains in vats coupled with the added flourish that such vats will not countenance (on pain of death for the inhabitant) the presence of H<sub>2</sub>O. The motivation behind these additions to the standard brain-in-a-vat hypothesis is that they ensure that no-one has ever had causal interaction either with water itself, or with someone else who has causally interacted with water. In this way, we can guarantee that, at least by the lights of Putnam's account at any rate, the possession of the concept of water actually entails the falsity of the brain-in-a-vat\* hypothesis. As a result, if an agent has *a priori* knowledge both that he possesses such a concept and of the truth of Putnam's theory, he can therefore infer, on *a priori* grounds, that he knows that he is not a brain-in-a-vat\*. We can thus extend the Putnam-McKinsey argument above as follows:

The Extended Putnam-McKinsey Argument

- (1)  $K_{AP} [(P1) \text{ I have the concept water}]$ .
- (2)  $K_{AP} [(P2) \text{ If I have the concept water, then water exists}]$ .

- (3)  $K_{AP}$  [(P3) Water exists].
- (4)  $K_{AP}$  [(P4) If I have the concept water, then the brain-in-a-vat\* sceptical hypothesis is false].
- (5)  $K_{AP}$  [(P5) The brain-in-a-vat\* sceptical hypothesis is false].

If one regards the gaining of *a priori* knowledge of (P3) as paradoxical, then one should likewise have reservations about gaining *a priori* knowledge of (P5) in this way. Indeed, Wright himself draws a similar anti-sceptical conclusion from the Putnam argument, arguing that it would follow from the Putnam account that one could know *a priori* that either oneself or members of one's speech community had had encounters with water. He writes:

[...] the idea that I can come to a warranted belief that my (community's) history includes interaction with water, or any other particular stuff, purely reflectively, without any reliance on empirical investigation, seems quite preposterous. (Wright 2000, 144).

McKinsey-style reasoning thus leads to some unpalatable conclusions. The Davies-Wright diagnosis attempts to disarm such reasoning by showing how empirical presuppositions are in fact implicit in even everyday *a priori* attributions of knowledge, and thus that, despite appearances, one cannot cogently infer, on purely *a priori* grounds, that one knows these presuppositions. It is thus with these presuppositions, and not the premises, that the radical (and otherwise surprising) conclusions of the McKinsey-type paradoxes are generated.<sup>10</sup>

### 3. PROBLEMS WITH THE WRIGHT-DAVIES PROPOSAL

Although this is a compelling diagnosis of the McKinsey paradox, it is unable to offer the solution that it advertises. This is because the success of the argument depends upon the McKinsey paradox being sufficiently analogous to the Moorean and extended (brain-in-a-vat\*) Putnam arguments and, unfortunately, this is not the case. For although it might be clear that the Moorean argument and the extended Putnam argument are question-begging, it is far from transparent that the basic McKinsey argument—such as the Boghossian formulation of Putnam's semantic externalism (P1-P3) that we saw above—is likewise question-begging.

In order to see this we need to make explicit exactly what the presuppositions which guide the Moorean and extended Putnam arguments are, and how they differ from the supposed presuppositions of the basic McKinsey template.

Recall that the objection to the Moorean argument was that Moore's entitlement to claim knowledge that he has two hands already presupposes an entitlement to believe that there is an external world. Why should that be the case? Why, for example, doesn't Moore's

empirical evidence for believing that he has two hands provide excellent grounds to support his knowledge of an external world? The answer is, of course, an old one; a perennial component of the standard sceptical debate. The reason why Moore's 'evidence' for believing that he has two hands is unable to support the hypothesis that there exists an external world, or that one of a number of other sceptical hypotheses can be discounted, is precisely because such sceptical hypotheses are *defined* such that although they are inconsistent with the truth of the bulk of our beliefs, were they to be true our evidence would be exactly the same.<sup>11</sup> Accordingly, although we might refer to these hypotheses as, broadly speaking, *empirical* hypotheses, this does not mean that they can be ruled-out on empirical grounds. Inevitably, then, if we are to claim empirical evidence for anything concrete about the world then we must first presuppose an entitlement to believe that the members of the class of such radical sceptical hypotheses that are able to defeat that claim are false.

Such are the presuppositions involved in the standard empirical case, and which act as a spur to sceptical doubt. The insight that Davies and Wright have is to recognise that the mere fact that an entitlement to believe these anti-sceptical proposition is presupposed in our epistemic practices does not thereby indicate that they are not known (as both the sceptic and the proponents of non-Closure charge), but only that we cannot cogently argue our way to knowledge of them. Moreover, since an entitlement to believe the presupposition that we are not radically in error might be thought to be implicit in all serious claims to know, so a similar diagnosis ought to be available to explain why the conclusion of the extended Putnam argument can be blocked. So far so good then (we shall consider these claims in more detail below).

If the basic McKinsey argument is to be dissolved in an analogous fashion to the Moorean and extended Putnam arguments then we would expect the presuppositions to be broadly similar. It is at this point that the Davies-Wright analysis goes awry. To begin with, note that the presupposition here *cannot* be of the same form as those that appear in the Moorean or extended Putnam arguments. That is, it cannot concern the denials of a set of relevant radical sceptical hypotheses. The reason for this is that to presuppose an entitlement to believe such propositions will not account for why an argument to the conclusion that, for example, one has *a priori* knowledge that water exists is question-begging. After all, that one is entitled to believe that there is an external world, or that one is not a brain-in-a-vat\*, does not itself entail that water exists. It could just be that one is in an external world or an 'envatted' environment that is completely dry. Why should it be, therefore, that simply presupposing an

epistemic entitlement to believe these propositions will serve to determine that one is thereby merely presupposing an epistemic entitlement to believe that water exists?

Of course, the Davies-Wright response would no doubt be to maintain that the analogy to the radical sceptical case should not be taken *that* seriously.<sup>12</sup> The problem, however, is that insofar as we do not take the analogy this far then the proposal itself starts to falter. For one thing, note that what made the idea that the anti-sceptical presuppositions in play in the Moorean case were in some sense necessary presuppositions of an empirical discourse was the fact that one could never, in principle, gain a non-question-begging warrant on empirical grounds to believe them. This feature is, however, peculiar to sceptical hypotheses that are defined in these terms. It will not so easily apply to other propositions where empirical grounds *are* available, but which are just not adduced.

Indeed, if the point were to allow the denials of non-sceptical scenarios to play this presuppositional role then this would appear to call the entire proposal into question. After all, if there are empirical presuppositions of one's entitlement to believe that water is wet, and if one is able, in principle, to gain empirical grounds to independently support those presuppositions, then, intuitively, one ought to do so. If this is the case, however, then it merely serves to undermine the putatively *epistemological* diagnosis of the McKinsey puzzle that Davies and Wright advertise by highlighting how the agent *doesn't* have genuine *a priori* knowledge, via privileged access, that water is wet after all. Recall that the starting-point of the Davies-Wright proposal was to accept the standard construal of the doctrines of privileged access and content externalism in order to show how certain epistemological considerations reveal that, appearances notwithstanding, they are consistent. By denying privileged access in these cases, however, the Davies-Wright line ends up with a completely different diagnosis of the puzzle which shows not that these theses are consistent, but rather that the truth of content externalism excludes the possibility of genuine privileged access as regards certain contents.

We can see this point more clearly by considering what possible candidates for the presuppositions in the Putnam McKinsey case might be if they are not to be the denials of radical sceptical hypotheses. For example, the thought might simply be that acquiring *a priori* knowledge that one has the concept water involves disregarding certain non-sceptical empirical possibilities such as that one does not have the concept of some ersatz substance like XYZ. The issue, however, is why it is that such possibilities can be disregarded without undermining the possession of *a priori* knowledge in the first place. After all, it is not as if such possibilities are in principle undetectable *sceptical* scenarios.<sup>13</sup> As a result, if we are to allow genuine *a priori*

knowledge of ‘wide’ thought contents, then it ought to be the case that knowledge of these possibilities gained by privileged access will indeed serve to rule-out such empirical scenarios. This may be a reason to restrict the thesis of privileged access, of course, but if Davies and Wright take this line then they cease to be proposing a purely epistemological diagnosis. Instead, they would merely be offering a rejection of at least the standard construal of privileged access insofar as it applies to wide contents. The detour via a discussion of Transmission would thus be irrelevant.<sup>14</sup>

It therefore seems that the Davies-Wright diagnosis at best only highlights certain anti-sceptical presuppositions that we make and thus only applies, if at all, to our putative knowledge of the denials of these hypotheses, not to knowledge of more mundane empirical propositions such as (P3). If, therefore, one is impressed by the paradoxical force of the original McKinsey arguments, then one will not find any comfort in the sort of epistemological diagnosis that the Davies-Wright analysis offers. Instead, a traditional approach that examines the truth of the two central theses involved—semantic externalism and privileged access—will be required.<sup>15</sup>

#### 4. EPISTEMOLOGICAL EXTERNALISM AND TRANSMISSION

Nevertheless, one might think that although the Davies-Wright line is unable to provide a compelling resolution to the McKinsey argument it could still be utilised to undermine a related McKinsey argument that involves anti-sceptical conclusions (what I shall henceforth refer to as the *anti-sceptical* McKinsey argument). For as we saw above with the extended Putnam version of the McKinsey argument, although the Davies-Wright line cannot account for the agent’s *a priori* knowledge of (P3), it does appear to be able to undermine the paradoxical force of the agent’s putative *a priori* knowledge of the anti-sceptical (P5) by highlighting how the agent is merely presupposing an epistemic entitlement to believe that conclusion at the start of his argument. (Similar considerations apply to the Moorean argument). As we shall see, I have some sympathy with this contention. Crucially, however, as I now hope to show, such a proposal is only plausible if set within an *externalist* epistemology since otherwise it will be in danger of collapsing into scepticism itself.

For the purposes of this paper I shall understand the internalist thesis to consist, at least minimally, in the following claim:

*Epistemological Internalism*

A necessary condition of an agent,  $a$ 's, knowledge of a proposition,  $\varphi$ , is that  $a$  has sufficient reflective access to the factors that make it such that  $a$  knows  $\varphi$  (rather than just truly believes  $\varphi$ ).

Although short on details, this characterisation should suffice for our purposes here since it captures the key internalist intuition that the satisfaction of an internal epistemic condition is *essential* for knowledge possession. In this way, internalists can allow external conditions (over and above the truth condition) to be important to knowledge possession without thereby endorsing an externalist epistemology which allows knowledge even in the absence of the satisfaction of such internal conditions.<sup>16</sup> Moreover, such a formulation is common to both proponents of the internalist thesis—such as Roderick Chisholm (1989, chapter 1)—and detractors—such as Alvin Plantinga (1993, chapter 1)—alike. Indeed, in a recent survey article on epistemology, James Pryor (2000) argues that this conception of the internalist thesis constitutes the “core internalist position”.

I shall therefore take externalism to be the denial of this thesis, and thus to consist in the claim that it is at least *possible* to have knowledge of a proposition whilst lacking sufficient reflective access to the factors that make it such that one knows.<sup>17</sup> Accordingly, the externalist thesis is consistent with a weak interpretation of that view which allows internal epistemic properties to play a pivotal role. All the externalist need deny, I take it, is that such internal epistemic properties are in any way *essential* to knowledge.<sup>18</sup>

Suppose, to begin with, that Davies and Wright are to be understood as working within an internalist conception of knowledge such that all knowledge fulfils the internalist rubric set out above. If this were so, then the knowledge that transfers to the anti-sceptical conclusion of the extended Putnam McKinsey argument (or, for that matter, the Moorean argument), though not *a priori*, must be internalist. Herein lies the tension, because, however one is to fill-in the details of the internalist account, it does seem that by the lights of such an account one cannot know the denials of radical sceptical hypotheses. After all, the very nature of radical sceptical hypotheses makes them such that there is no phenomenological difference between, say, sitting in one's chair in front of one's computer, and being a brain-in-a-vat being ‘fed’ such ‘experiences’ by a neuroscientist. And if there is no phenomenological difference then there can be no reflectively accessible difference either, at least insofar as we assume, as we have been doing here, that there is no *a priori* reason available for dismissing radical scepticism. But now the difficulty should be manifest. For insofar as Davies and Wright retain Closure, then the lack of knowledge of the denial of the radical sceptical hypothesis will translate into a lack

of knowledge of the instance of privileged access that we started with. Indeed, it would translate into a lack of knowledge of *any* everyday proposition that was known to be inconsistent with the sceptical hypothesis under consideration. Furthermore, there is the *extra* problem in this case that the subject's belief in this proposition is held to be *merely* presuppositional, and therefore, as Davies and Wright recognise, is not believed as a result of an empirical inquiry. On such a construal, it seems that the subject has no reflectively accessible *epistemic* grounds that would support his belief,<sup>19</sup> and thus that he cannot be in the possession of internalist knowledge of this proposition.<sup>20</sup>

On the one hand, then, Davies and Wright, construed as internalists, are apparently committed to arguing that the subject *could not* have (internalist) knowledge of one of the known entailments of the conclusion of the Putnam argument (namely, (P5)). On the other, the very same Davies and Wright, as adherents to the Closure principle, are committed to arguing that such a known entailment must be itself known. The only two coherent ways out of this contradiction (short of abandoning the strategy altogether) are to either deny the highly intuitive principle of Closure<sup>21</sup> or else endorse radical scepticism. Neither is particularly attractive. The latter is unappealing for very obvious reasons. The former is problematic because it would appear to undermine what was most interesting about the proposal, which was its avowed ability to resolve this issue with Closure intact. Indeed, if the answer is to simply deny Closure, then we could do *that* without entertaining this detour via Transmission at all. Moreover, since the standard non-Closure line incorporates an externalist epistemology,<sup>22</sup> this move may itself undermine any avowed commitment to the very internalist paradigm that created the problem in the first place.

Before countenancing such unpalatable alternatives it is thus germane to consider the possibility of rendering this thesis along externalist lines. Moreover, since this proposal leaves Closure intact, it may well be that we could further regard the externalist denial of the Transmission principle as posing a direct challenge to any argument for non-Closure.<sup>23</sup>

## 5. TRANSMISSION, CLOSURE, AND RADICAL SCEPTICISM

Potentially, there are two main advantages to presenting the Davies-Wright line in terms of an externalist epistemology. The first is that an externalist epistemology can make sense of the idea that an agent can know a 'presuppositional' anti-sceptical proposition even though he

lacks sufficient reflectively accessible grounds to support his belief in that proposition. The second is that we may be able to give a more feasible account of how to understand the denial of Transmission in terms of such an epistemology.

In what follows, I shall take the success of the first point for granted. Though a somewhat controversial assumption to make, it rests upon uncontroversial motivations. For although it may be a matter of much controversy *which* externalist account of knowledge should be given—and thus whether the correct externalist account would allow an agent to know these sorts of propositions—it is an entirely mundane contention that on the externalist account whether or not an agent knows a proposition is not decided by a lack of reflectively accessible grounds alone. Instead, the possession of such knowledge will be dependent upon other salient features of the subject's belief, such as, for example, the cognitive process via which it was acquired or the subjunctive relationships that it bears to the truth in near-by possible worlds.<sup>24</sup> Of course, were it to transpire, once the discursive dust has settled, that the most plausible accounts of externalist knowledge do not allow such knowledge, then this would suffice to bury the Davies-Wright proposal on *any* account of knowledge. Nevertheless, absent such a dramatic (and, I think, unlikely) eventuality, it is worthwhile exploring what an externalist account of Transmission would look like.

Recall that Transmission was defined in terms of Closure plus 'cogency', where the latter involved a capacity to produce rational conviction in the consequent proposition. One way of capturing this distinction—and thus the idea that these two notions come apart—is in terms of the nature of the grounds that are adduced to support the claims in question. After all, the problem with those cases where Transmission (but not Closure) fails—such as the sort of anti-sceptical question-begging argument that Moore proposes—is primarily that the kind of 'internal' reflectively accessible grounds that are offered for the premises are not able to support the conclusion because such grounds only support the premises in the first place on the presupposition that one *already* has an epistemic entitlement to believe the conclusion. Accordingly, what we are looking for is an epistemic distinction which allows that knowledge can transfer across known entailments in such scenarios without it thereby also being the case that the internal grounds that one offers for the premises must, perforce, follow. By the lights of an internalist epistemology which posits a close connection between the possession of knowledge and the presence of such internal grounds this distinction is unavailable, especially in the anti-sceptical case where there are no adequate reflectively accessible grounds to be had for the proposition in question at all. That this is so, however, need not prejudice the prospects

for an externalist epistemology—which draws no such essential connection between knowledge possession and the possession of internal grounds—accounting for this distinction.

One way of making this point explicit is to re-cast the distinction between Transmission and Closure as regards anti-sceptical question-begging arguments into a distinction between two construals of the Closure principle.<sup>25</sup> The first, what we shall call Closure *simpliciter*, we saw above. This simply demands that knowledge should transfer across known entailments without specifying the type of knowledge at issue. Since this makes no demand that the inference involved should be cogent, we would expect  $K_C$  to hold even in manifestly question-begging arguments where the premises are known. Contrast this principle with a related principle that holds that *internalist* knowledge—knowledge that is sufficiently supported by reflectively accessible grounds—should transfer across known entailments; in effect, that Closure should preserve the *type* of knowledge at issue. We shall formulate this internalist version of Closure as follows, where ‘ $K_I [\varphi]$ ’ refers to an agent’s knowledge of a proposition,  $\varphi$ , which meets the sort of internal conditions demanded by an internalist epistemology.<sup>26</sup>

$$K_{CI}: \quad \{K_I [\varphi] \ \& \ K_I [\varphi \rightarrow \psi]\} \rightarrow K_I [\psi]$$

In words, where one has internalist knowledge of a proposition,  $\varphi$ , and internalist knowledge that  $\varphi$  entails a second proposition,  $\psi$ , then one has internalist knowledge of  $\psi$ .

Clearly, although Closure may itself be a highly intuitive epistemic principle, if one is an externalist then it will not follow that this  $K_{CI}$  principle should also be intuitive. Indeed, it would seem that denying  $K_{CI}$  captures the underlying motivation for denying the Transmission principle. The reason for this is that just as Transmission was expressed in terms of the transfer of knowledge plus cogency across known entailments, so  $K_{CI}$  attempts to encapsulate the idea that a certain sort of knowledge that involves sufficient reflectively accessible grounds should transfer across known entailments. The parallels are striking. Both retain the idea of ‘mere’ knowledge—that an agent might know a presupposition of a certain epistemic project even though he is unable, in principle, to adduce sufficient reflectively accessible grounds to support that knowledge—an idea that, as we saw above, is only coherent within an externalist epistemology. Moreover, since cogency is characterised in terms of the adducement of grounds—indeed, in terms of a particular form of reasoning from grounded premises to a grounded conclusion—then an internalist construal of this notion seems entirely natural. If internalist knowledge is ‘closed’ under known entailments, then it ought to be true that the reflectively accessible grounds that support the antecedent knowledge should *transmit* across

the entailment to (cogently) support the consequent knowledge.

Accordingly, one can explain the intuition that it is Transmission (rather than Closure) that fails when it comes to anti-sceptical (and other question-begging) arguments such as Moore's in terms of the failure of internalist Closure in such cases. That is, one accounts for the distinction in terms of the intuition that, although if Moore knows the premises he must know the conclusion (in line with Closure), nevertheless, he cannot know this conclusion on the basis of whatever reflectively accessible grounds he adduces for the premises, and thus he lacks internalist knowledge of the conclusion.<sup>27</sup> We thus retain the basic Closure thought that if Moore has knowledge at all, then it rests upon further knowledge of the presupposed anti-sceptical propositions; whilst denying the more demanding internalist intuition that all internalistically construed knowledge must rest on further internalist knowledge of the presupposed propositions.

In order to see this point in more detail, consider again the Moorean argument that we saw above. Moore's stated aim was to *show* something with his proof, and, therefore, we might reasonably assume that it was a particular sort of knowledge that he was seeking—knowledge that was internalistically grounded; internalist *knowledge*. But given that the grounds he offers to support his knowledge of the premise that he has two hands presuppose an epistemic entitlement to believe that there is an external world, such grounds cannot transmit across the known entailment to give him internalist knowledge of the existence of an external world. If he knows that he has two hands, then he must know that there is an external world, but this does not mean that the grounds that he proffers to support his (internalist) knowledge that he has two hands transfer across the known entailment to provide him with internalist knowledge of his anti-sceptical conclusion. We can represent this reconfiguration of the Moorean argument as follows, where 'K<sub>AP&I</sub>' refers to knowledge that is both internalist and *a priori*, and 'K<sub>E</sub>' refers to merely externalist knowledge that does not fulfil the internalist rubric:

- (1\*) K<sub>I</sub> [(M1) I have two hands].
- (2\*) K<sub>AP&I</sub> [(M2) If I have two hands, then there exists an external world].
- (3\*) K<sub>E</sub> [(M3) There exists an external world].<sup>28</sup>

So understood, one can account for the fact that if Moore does know that he has two hands then he must know that there exists an external world without thereby being committed to arguing that Moore is able to pull himself up by his bootstraps such that he is able to reason his way to internalist knowledge of what is presupposed in that internalist knowledge.

Of course, the radical sceptic will no doubt object to this line of reasoning on the

grounds that although it captures the intuition which drives Closure that knowledge must rest upon knowledge, it does not do justice to the sceptical intuition that knowledge must rest upon knowledge of a particular internalist sort. This is certainly an important charge and one that any development of this account must confront. Nevertheless, it is important not to concede too much to the sceptic here. For what this line of argument offers us is a compelling way of responding to scepticism without denying Closure by making a distinction which, within an externalist epistemology, one has *prima facie* grounds to make. This should give us cause to press the sceptic on this matter and demand why it should be that all knowledge should rest upon internalist knowledge, rather than just knowledge *simpliciter*. Indeed, one could go further and argue that this way of looking at radical scepticism presents us with an intriguing diagnosis of why sceptical problems have seemed so intractable. The truth in scepticism is that we cannot internalistically know what is presupposed in our empirical inquiries—that we are not the victims of radical sceptical hypotheses. Nevertheless, the radical sceptical line of argument becomes questionable at the point at which it converts this contention into the claim that a lack of such internalist knowledge should impact upon the epistemic credentials of our beliefs in general. If the line of argument sketched here is sound, then we have good cause to dispute this move. We can thus account for the intuitions driving both radical scepticism and arguments for non-Closure without either endorsing radical scepticism or denying Closure.

Finally, it is worth briefly outlining how this strategy impacts upon the Extended Putnam-McKinsey Argument that we saw above and which also had an anti-sceptical conclusion. We begin, as before, by noting how even the *a priori* knowledge of the first premise that is gained via privileged access rests upon the assumption that one is not the victim of a radical sceptical hypothesis. What is different in this case is that because the knowledge is based upon evidence which is gained by *a priori* means (i.e., privileged access) this knowledge will, at least in the paradigm case, be internalist *and a priori*.<sup>29</sup> The rest of the argument will then unfold pretty much as before:

- (1\*)  $K_{AP\&I}$  [(P1) I have the concept water].
- (2\*)  $K_{AP\&I}$  [(P2) If I have the concept water, then water exists].
- (3\*)  $K_{AP\&I}$  [(P3) Water exists].
- (4\*)  $K_{AP\&I}$  [(P4) If I have the concept water, then the brain-in-a-vat\* sceptical hypothesis is false].

As we saw above, since what is presupposed in order to get to line (1\*) in the first place does not beg the question against the conclusion reached at (3\*), hence the basic McKinsey argument *will* go through. Nevertheless, since what is presupposed does beg the question against the anti-sceptical line (5\*), the only conclusion that can be drawn from this extended

McKinsey argument is the following:

(5\*)  $K_E [(P5) \text{ The brain-in-a-vat* sceptical hypothesis is false}]$ .

As with the Moorean case, given that an epistemic entitlement to believe this proposition was presupposed from the outset, so this conclusion should not come as any surprise. Moreover, since the argument to this conclusion lacks cogency, it therefore offers us no *a priori* grounds to believe this conclusion, and thus we have no reason to regard the agent's knowledge of (P5) as being itself *a priori*. All that we may conclude is that if it is indeed true that the agent has knowledge of the premises then that agent must also have externalist knowledge of the presupposition that drives those premises, nothing more.

## 6. CONCLUDING REMARKS

Two important conclusions follow from the examination of the Davies-Wright line undertaken here. The first is that the prospects for an epistemological diagnosis of the McKinsey paradox are dim indeed. This result therefore puts the issue back into the arena of the philosophy of language and mind. The second is that although the Davies-Wright argument is unable to offer a persuasive dissolution of the McKinsey paradox, when presented in terms of an externalist epistemology it does translate into a plausible account of why it is that certain *anti-sceptical* arguments are unconvincing. There are thus strong grounds for taking the general approach to the McKinsey paradox that Davies and Wright advocate out of the McKinsey debate and reconfiguring it in terms of the wider issue of radical scepticism.<sup>30</sup>

## NOTES

<sup>1</sup> For the main proponents of theses of this sort, see Putnam (1975; cf. Putnam 1981); Burge (1979); and Evans (1982).

<sup>2</sup> For the core discussions of this supposed incompatibilism, see Davidson (1987); Burge (1988); Boghossian (1989; 1997); McKinsey (1991); Brueckner (1992b); Brown (1995); and Davies (1998).

<sup>3</sup> That is, roughly, mental states which supervene upon one's intrinsic properties.

<sup>4</sup> Davies and Wright actually tend to express the point in terms of *warrant*, but, in order to bring this discussion into line with the standard formulation of the McKinsey paradoxes, I shall reformulate their point throughout in terms of knowledge. In any case, this difference of exposition should be philosophically harmless. Insofar as warrant is the sort of robust epistemic notion that has relevance for this debate—such that it is sufficient, with true belief, for knowledge—then we can simply take it for granted that the premises of the argument are, in fact, true, and thus that it is knowledge that is at issue here.

<sup>5</sup> So as to avoid philosophically uninteresting counterexamples to the epistemic principles that I discuss here, the reader should henceforth take it for granted (where applicable) that the agent does have the requisite belief in the consequent proposition and that all occurrences of the ‘K’ operator are indexed to a particular time and a particular agent.

<sup>6</sup> There are, of course, those who, for reasons independent of the issue in hand, do deny the Closure principle. For the main texts in this regard, see Dretske (1970; 1971), and Nozick (1981).

<sup>7</sup> Both Davies and Wright regard this ‘presupposition’ thesis as being implicit in Wittgenstein’s final notebooks (published as *On Certainty* (1969)), where he argues that there are propositions which perform a hidden “hinge” role in a discourse such that one cannot acquire a warrant for that proposition (*a fortiori*, come to know that proposition) via any activity pursued within that discourse. As Wittgenstein expresses the matter: “[...] the questions that we raise and our doubts depend upon the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn. That is to say, it belongs to logic of our scientific investigations that certain things are *in deed* not doubted. But it isn’t that the situation is like this: We just *can’t* investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put.” (Wittgenstein 1969, §§341-3) Whether this is an accurate rendering of the Wittgensteinian position in this regard is, of course, another matter. For discussion of this point, see Pritchard (2000d). Wright has endorsed a related reading of the “hinge” proposition thesis in other works, see Wright (1985; 1991). For discussion, see Pritchard (2000b; 2000c; 2001).

<sup>8</sup> Note that I do not mean to claim that this is an accurate rendering of Moore’s position as regards scepticism, only that it is the standard reading that is applied to his work in this respect (see, for example, Wittgenstein (1969); Wright (1985; 1997); and Davies (1998)).

<sup>9</sup> Of course, if you do think that there is such an *a priori* defence available then you need not be worried about the sceptical implications of this argument at all. Nevertheless, since Moore clearly did not think that a rebuttal formulated along these lines was forthcoming, I shall likewise disregard this avenue of response.

<sup>10</sup> It is important to note that Putnam’s (1981) own version of this argument—which rests upon his claim that his version of semantic externalism illustrates that the statement, “I am a brain-in-a-vat\*”, is necessarily false—is significantly different to that sketched here. (Indeed, although Wright objects to the conclusion of the Extended Putnam-McKinsey Argument, he does accept the basic anti-sceptical argument that Putnam proposes, arguing that there are relevant differences between the basic Putnam argument and the McKinsey-style reformulation (see Wright 2000, 161)). For the key texts in which the Putnam argument is discussed, see Brueckner (1992a); Wright (1994; 1997; 2000a); Forbes (1995); and Warfield (1998). Moreover, it should also be emphasised that Putnam’s primary interest in putting forward this argument was not to disarm a particular form of radical scepticism so much as to undermine a certain philosophical view which he terms “metaphysical realism”. For discussion of Putnam’s argument with this in mind, see Hale & Wright (1999).

<sup>11</sup> In a recent article, Williamson (2000) has argued that in fact our evidence *would* change were a sceptical hypothesis to become true. Although I am sympathetic to this suggestion, since neither Wright nor Davies considers this line of anti-sceptical attack I shall likewise keep to the standard line.

<sup>12</sup> Actually, there is a subtle sleight of hand going on here that tricks the reader into thinking that the Moorean and extended Putnam McKinsey arguments are *exactly* analogous to the standard McKinsey case. The trick is to construe the possibility that one is having thoughts about XYZ rather than H<sub>2</sub>O as a *radical sceptical* possibility, as well it might be, and then directly infer that, like the other arguments which concern radical sceptical possibilities that have been considered, the basic McKinsey argument must also be question-begging. For example, straight after his account of the McKinsey paradoxes, Wright fleshes-out his proposal by considering how it works with various sceptical scenarios, seemingly oblivious to the fact that this is where the disanalogy between the McKinsey argument and the Moorean/Extended Putnam-McKinsey arguments lies—in the fact that the conclusion of the former, since it does not concern a radical sceptical hypothesis, is not necessarily presupposed in the premises of the argument.

<sup>13</sup> This is not to deny, of course, that one could tell a radical sceptical story involving the concept XYZ. Crucially, however, if this were so then it would be an entitlement to believe the denial of the sceptical scenario in question that was being presupposed, not an entitlement to believe that one does not have the concept XYZ, and thus the sort of considerations that were brought to bear against this proposal above would apply.

<sup>14</sup> For further critical discussion—albeit along different lines to that pursued here—of the most recent expression of the Wright thesis in this respect (Wright 2000a; cf. Wright 2000b), see Hale (2000); Sainsbury (2000); and Suárez (2000).

<sup>15</sup> My own view in this regard is that a solution to this paradox can be given by being more exacting about how one’s favoured version of semantic externalism should be formulated and how it affects our traditional construal of privileged access. Kallestrup (2000) offers a development of a view of this sort by arguing that there is an equivocation in the semantic externalist’s position. Essentially, he shows that, insofar as (MK1) is known *a priori*, then (MK2) is not known *a priori*, and that insofar as (MK2) is known *a priori*, then (MK1) is not known *a priori*. In neither case, then, is one able to generate *a priori* knowledge of the empirical claim encapsulated in (MK3).

<sup>16</sup> Indeed, Wright (1991) himself offers an internalist account of warrant (and thus knowledge) which has such an external component. What makes the account internalist is that it does not allow that merely meeting such an external condition could suffice for knowledge.

<sup>17</sup> For more discussion on the externalism/internalism distinction, see Bonjour (1980); Goldman (1980); Alston (1986); Fumerton (1988); and Plantinga (1993).

<sup>18</sup> Indeed, even Goldman (e.g. Goldman 1988) allows an internalist conception of justification to play an important role in his epistemology. For more on this point, see Brandom (1998; cf. Brandom 1994; 1995).

<sup>19</sup> Though he may, of course, have *pragmatic* grounds.

<sup>20</sup> Or at least, *pending some sort of non-standard account of internalist knowledge*, the subject cannot plausibly be thought to be in the possession of internalist knowledge of this proposition. It is difficult to see, however, just how such an account could be made feasible. On the one hand, it is hard to comprehend how one could argue that the presuppositional beliefs in question were foundational in some sense because such foundational beliefs tend to have a direct *empirical* warrant and on the internalist account empirical warrants are out of the question when it comes to belief in these presuppositional propositions. On the other, it is difficult to fathom how one could make the suggestion that one's knowledge of these propositions is *a priori* convincing. Accordingly, since my claim is only that an internalist construal of this distinction between Transmission and Closure is far from plausible, I shall not explore these suggested avenues of response further. I am grateful to an anonymous referee from *Synthese* for pressing me on this point.

<sup>21</sup> At the end of his paper, Davies is actually open to this suggestion but he regards this option as *additional* to the Transmission move, not a replacement of it.

<sup>22</sup> Though not, of course, a knock-down argument for this conclusion, all of the most prominent defenders of the non-Closure thesis—and, in particular, Dretske (1970; 1971) and Nozick (1981)—have been externalists.

<sup>23</sup> There is a third choice available which is to offer a contextualist thesis along the lines of that proposed by DeRose (1995) and Lewis (1996). Space prevents me, however, from giving a considered account of this option. For a recent discussion of the relationship between arguments for non-Closure and contextualist treatments of knowledge, see the exchange between Heller (1999) and Pritchard (2000a).

<sup>24</sup> For recent examples of papers which argue that one could use an externalist account of knowledge to undercut scepticism by explaining how we can know the denials of sceptical hypotheses after all, see Sosa (1991; 1994; 1999) and Hill (1996). Of course, a common motif of the externalist position has been its advocacy of a conception of knowledge which explicitly *disavows* the possibility that such propositions could be known (it is these sorts of externalist accounts which, following Dretske (1970; 1971) and Nozick (1981), have tended to deny the Closure principle). It may well be, however, that an understanding of the failure of Transmission could motivate a retention of Closure in the face of the sceptical challenge, and thus provide support for the thesis offered here that presupposes that such propositions *could* be known. Indeed, this line of argument is further supported by recent work on virtue-theoretic treatments of knowledge, especially Sosa (1991, 1999). He argues for a very mild 'neo-Moorean' modal theory of knowledge that allows for an agent to have knowledge of the denials of radical sceptical hypotheses provided that the actual world is suitably distant, modally speaking, from the nearest sceptical possible world.

<sup>25</sup> I focus on the relationship between Transmission and Closure insofar as it applies to *anti-sceptical* question-begging arguments because these are the types of arguments under consideration. No doubt the account of Transmission that Davies and Wright offers can tell us more about question-begging arguments in general, but such further remarks, whilst undoubtedly important, are not directly relevant to the issue in hand here.

<sup>26</sup> Recall that the externalist need not deny that knowledge can sometimes fulfil the internalist rubric. What the externalist denies is that *all* knowledge must meet this rubric.

<sup>27</sup> This is not to deny that the proponent of a question-begging argument could have internalist knowledge of the conclusion on independent grounds (though this could not, of course, be the case as regards the anti-sceptical arguments we are considering here), only that the premises themselves would provide no support for this sort of knowledge.

<sup>28</sup> It is important to be clear that the claim is *not* that Moore is able to legitimately *reason* his way to this conclusion, such that he is able to cogently argue that he has externalist knowledge of the existence of an external world. Rather, the point is only that insofar as Moore does indeed have knowledge of the premises of his argument, then he will have knowledge of the existence of an external world, albeit knowledge of a purely externalist variety.

<sup>29</sup> Recall that we have accepted, for the sake of argument, the truth of the theses of privileged access and semantic externalism as represented above.

<sup>30</sup> Thanks to an anonymous referee from *Synthese* for comments on an earlier version of this paper, and to Patrick Greenough, Jesper Kallestrup, Patrice Philie and, as ever, Crispin Wright for discussion on this and related topics.

## REFERENCES

- Alston, William P.: 1986, 'Internalism and Externalism in Epistemology', *Philosophical Topics* 14, 181-204.
- Boghossian, Paul: 1989, 'Content and Self-Knowledge', *Philosophical Topics* 17, 5-26.
- Boghossian, Paul: 1997, 'What an Externalist can Know *A Priori*', *Proceedings of the Aristotelian Society* 97, 161-75.
- Bonjour, Laurence: 1980, 'Externalist Theories of Justification', *Midwest Studies in Philosophy* 5, (eds.) P. A. French, T. E. Uehling & H. K. Wettstein, 27-52, University of Minnesota Press, Minneapolis.
- Brandom, Robert: 1994, *Making It Explicit*, Harvard University Press, Cambridge, MA.
- Brandom, Robert: 1995, 'Knowledge and the Social Articulation of the Space of Reasons', *Philosophy and Phenomenological Research* 55, 895-908.
- Brandom, Robert: 1998, 'Insights and Blindspots of Reliabilism', *The Monist* 81, 371-92.
- Brown, Jessica: 1995, 'The Incompatibility of Anti-Individualism and Privileged Access', *Analysis* 55, 149-56.
- Brueckner, Anthony: 1992a, 'Semantic Answers to Scepticism' *Pacific Philosophical Quarterly*, 73, 200-19.
- Brueckner, Anthony: 1992b, 'What an Anti-Individualist Knows *A Priori*', *Analysis* 52, 111-8.
- Burge, Tyler: 1979, 'Individualism and the Mental', *Midwest Studies in Philosophy* 4, (eds.) P. French, T. Uehling & H. Wettstein, 73-121, University of Minnesota Press, Minneapolis.
- Burge, Tyler: 1988, 'Individualism and Self-Knowledge', *Journal of Philosophy* 85, 649-63.
- Chisholm, R. M.: 1989, *Theory of Knowledge*, (3rd Edition), Prentice Hall, Englewood Cliffs, New Jersey.
- Davidson, Donald: 1987, 'Knowing One's Own Mind', *Proceedings and Addresses of the American Philosophical Association* 60, 441-58.
- Davies, Martin: 1998, 'Externalism, Architecturalism, and Epistemic Warrant', *Knowing Our Own Minds: Essays on Self-Knowledge*, (eds.) C. Macdonald, B. Smith & C. J. G. Wright, 321-61, Oxford University Press, Oxford.
- DeRose, Keith: 1995, 'Solving the Sceptical Problem', *Philosophical Review* 104, 1-52.
- Dretske, Fred: 1970, 'Epistemic Operators', *Journal of Philosophy* 67, 1007-23.
- Dretske, Fred: 1971, 'Conclusive Reasons', *Australasian Journal of Philosophy* 49, 1-22.
- Evans, Gareth: 1982, *The Varieties of Reference*, Oxford University Press, Oxford.
- Forbes, Graham: 1995, 'Realism and Skepticism: Brains in a Vat Revisited', *Journal of Philosophy* 92, 205-22.
- Fumerton, Richard: 1988, 'The Internalism/Externalism Controversy', *Philosophical Perspectives* 2, 443-59.
- Goldman, Alvin: 1980, 'The Internalist Conception of Justification', *Midwest Studies in Philosophy* 5, (eds.) P. A. French, T. E. Uehling & H. K. Wettstein, University of Minnesota Press, Minneapolis.
- Goldman, Alvin: 1988, 'Strong and Weak Justification', *Philosophical Perspectives* 2, 51-69.
- Goldman, Alvin: 1999, '*A Priori* Warrant', *Philosophical Perspectives* 13, 1-28.
- Hale, Bob: 1999, 'Putnam's Model-Theoretic Argument Against Metaphysical Realism', (with Crispin Wright), *A Companion to the Philosophy of Language*, (eds.) B. Hale & C. J. G. Wright, 427-57, Blackwell, Oxford.
- Hale, Bob: 2000, 'Transmission and Closure: Some Comments on 'Cogency and Question-

- Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof', *Philosophical Issues* 10, 172-90.
- Heller, Mark: 1999, 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy* 77, 196-208.
- Hill, Christopher: 1996, 'Process Reliabilism and Cartesian Scepticism', *Philosophy and Phenomenological Research* 56, 567-81.
- Kallestrup, Jesper: 2000, 'Incompatibilism and Natural Kind Terms', (*unpublished manuscript*).
- Lewis, David: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.
- McKinsey, M: 1991, 'Anti-Individualism and Privileged Access', *Analysis* 51, 9-16.
- Moore, G. E.: 1925, 'A Defence of Common Sense', *Contemporary British Philosophy* (2nd series), (ed.) J. H. Muirhead, Allen and Unwin, London.
- Moore, G. E.: 1939, 'Proof of an External World', *Proceedings of the British Academy* 25.
- Nozick, Robert: 1981, *Philosophical Explanations*, Oxford University Press, Oxford.
- Plantinga, Alvin: 1993, *Warrant: The Current Debate*, Oxford University Press, New York.
- Putnam, Hilary: 1975, 'The Meaning of "Meaning"', reprinted in his *Mind, Language and Reality*, 215-72, Cambridge University Press, Cambridge, England.
- Putnam, Hilary: 1981, *Reason, Truth and History*, Cambridge University Press, Cambridge, England.
- Pritchard, Duncan: 2000a, 'Closure and Context', *Australasian Journal of Philosophy* 78, 275-80.
- Pritchard, Duncan: 2000b, 'Is "God Exists" a "Hinge" Proposition of Religious Belief?', *International Journal for Philosophy of Religion* 47, 129-40.
- Pritchard, Duncan: 2000c, 'Scepticism and Dreaming', *Philosophia* 28, 373-90.
- Pritchard, Duncan: 2000d, 'Wittgenstein, "Hinge" Propositions, & *On Certainty*', *Rationality and Irrationality: Proceedings of the 23<sup>rd</sup> International Wittgenstein Symposium*, (eds.) Berit Brogaard & Barry Smith, 84-90, Austrian L. Wittgenstein Society, Kirchberg, Austria.
- Pritchard, Duncan: 2001, 'Radical Scepticism, Epistemological Externalism, and "Hinge" Propositions', *forthcoming in Wittgenstein-Studien*.
- Pryor, James: 2000, 'Epistemology c. 1988-2000', *The British Journal for the Philosophy of Science* 51, 1-29.
- Sainsbury, Mark: 2000, 'Comments on Crispin Wright's 'Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof'', *Philosophical Issues* 10, 191-200.
- Sosa, Ernest: 1991, 'Intellectual Virtue in Perspective', in his *Knowledge in Perspective: Selected Essays in Epistemology*, 270-93, Cambridge University Press, Cambridge, England.
- Sosa, Ernest: 1994, 'Philosophical Scepticism and Epistemic Circularity', *Proceedings of the Aristotelian Society* (supplementary vol.) 68, 263-90.
- Sosa, Ernest: 1999, 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.
- Suárez, Alfonso García: 2000, 'On Wright's Diagnosis of McKinsey's Argument', *Philosophical Issues* 10, 164-71.
- Warfield, Ted: 1998, 'A Priori Knowledge of the World: Knowing the World by Knowing our Minds', *Philosophical Studies* 60, 76-90.
- Williamson, Timothy: 2000, 'Scepticism and Evidence', *Philosophy and Phenomenological Research* 60, 613-28.
- Wittgenstein, Ludwig: 1969, *On Certainty*, (eds.) G. E. M. Anscombe & G. H. von Wright,

- (tr.) D. Paul & G. E. M. Anscombe, Blackwell, Oxford.
- Wright, Crispin: 1985, 'Facts and Certainty', *Proceedings of the British Academy* 71, 429-72.
- Wright, Crispin: 1991, 'Scepticism and Dreaming: Imploding the Demon', *Mind* 397, 87-115.
- Wright, Crispin: 1994, 'On Putnam's Proof that we are not Brains in a Vat', *Reading Putnam*, (eds.) B. Hale & P. Clark, 216-41, Blackwell, Oxford.
- Wright, Crispin: 1997, 'On the Acquisition of Warrant by Inference', (*unpublished manuscript*).
- Wright, Crispin: 2000a, 'Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof', *Philosophical Issues* 10, 140-63.
- Wright, Crispin: 2000b, 'Replies to Commentators', *Philosophical Issues* 10, 201-19.