

TRANSMISSION OF WARRANT AND CLOSURE OF APRIORITY

Michael McKinsey
Wayne State University

In my 1991 paper, "Anti-Individualism and Privileged Access," I argued that externalism in the philosophy of mind is incompatible with the thesis that we have privileged, nonempirical access to the contents of our own thoughts.¹ One of the most interesting responses to my argument has been that of Martin Davies (1998, 2000, and Chapter _ above) and Crispin Wright (2000 and Chapter _ above), who describe several types of cases to show that warrant for a premise does not always transmit to a known deductive consequence of that premise, and who contend that this fact under-mines my argument for incompatibilism. I will try to show here that the Davies/Wright point about transmission of warrant does not adversely affect my argument.

Before discussing the Davies/Wright point, it will be useful to begin with first, a brief restatement of my argument and second, a general assessment of the dialectical situation regarding the argument and its critics.

1. The argument for incompatibilism.

In my initial argument, I was concerned with a principle of privileged access according to which we necessarily have the capacity to obtain a priori knowledge of the *contents* of our thoughts:

Privileged Access to Content (PAC)

Necessarily, for any person x , if x is thinking that p , then x can in principle know a priori that he himself, or she herself, is thinking that p .²

By 'a priori' knowledge, I meant knowledge that can be achieved 'just by thinking', without perceptual observation or empirical investigation, and without having to make any empirical assumptions (1991a, p. 9; see also McKinsey 1987, pp. 2-3). And I argued that **PAC** is inconsistent with the following externalist thesis:

Semantic Externalism (SE)

Many *de dicto*-structured predicates of the form 'is thinking that *p*' express properties that are wide, in the sense that possession of such a property by an agent logically implies the existence of contingent objects of a certain sort that are external to that agent.³

My argument was a simple *reductio* that considered an instance of 'is thinking that *p*' containing the natural kind term 'water'. Suppose Oscar is thinking that water is wet. Then it follows by **PAC** that

- (1) Oscar can know a priori that he is thinking that water is wet.

Now since **SE** is supposed to apply to cognitive predicates containing natural kind terms like 'water', we also have

- (2) The proposition that Oscar is thinking that water is wet logically implies the proposition **E**,

where **E** is some "external proposition" that asserts or implies the existence of objects external to Oscar. For instance, **E** might be the proposition that Oscar, or members of Oscar's community, have had experiences of water; or **E** might be the proposition that water exists. But the conjunction of (1) and (2) is absurd, whatever **E** might be. If (1) is true, then according to (2), Oscar can know a priori a proposition that logically implies some empirical proposition **E**. If so, then Oscar can just *deduce* **E** from something he knows a priori, and so he can know **E** itself a priori. But this consequence is just absurd, since by assumption **E** is some empirical proposition

that no one could possibly know a priori. Hence the principle that we invariably have privileged access to the contents of our thoughts is inconsistent with semantic externalism.

2. Relational cognitive properties.

When I first gave this *reductio* argument for the incompatibility of semantic externalism and privileged access (in my 1991a), I took no stand on which of these principles is true and which is false. But shortly thereafter, I provided the grounds for an argument that **SE** is true and hence, **PAC** is false. I also proposed a restricted principle of privileged access to replace **PAC**, a principle that, unlike **PAC**, is consistent with **SE**. (See McKinsey 1994.)

SE is shown true by the semantic facts about cognitive predicates that contain small-scope proper names and indexical pronouns. Consider the case of Laura, who when George strolls by, raises her eyebrows and whistles. Witnessing this scene, I turn to a friend and say

(3) Laura is thinking that George is cute.

I believe it is intuitively clear that in uttering (3), I would be using the name 'George' simply to refer to George, and I would be saying that Laura is thinking that he is cute. If the name 'George' as I use it in (3) had a descriptive meaning, then perhaps the name could semantically convey something about the way that Laura is thinking of George.⁴ But in fact, most ordinary names including 'George' have no such descriptive meanings, and so in (3) the name can contribute only its referent to what is said by (3).⁵ As a result, (3) ascribes a cognitive property to Laura that is relational with respect to George: it says that Laura has an occurrent thought about George to the effect that he is cute.⁶

It is perhaps even clearer that cognitive predicates containing small-scope indexicals ascribe relational properties. Consider

(4) Laura is thinking that *he* (or: *that man*) is cute.

Given that in (4) 'he' or ('that man') refers to George, (4) like (3) says that Laura is having an occurrent thought about George to the effect that he is cute. Note that both (3) and (4) ascribe relational properties, even though the occurrences of the relevant singular terms ('George', 'he', 'that man') are assumed to fall both grammatically and logically in the scope of 'is thinking that'. Thus both (3) and (4) are *structurally* (or logically) *de dicto*, but due to the semantic character of the small-scope terms, both sentences end up saying something that is *semantically* relational, or *de re*.⁷

Thus many *de dicto*-structured cognitive predicates containing small-scope names and indexicals express relational properties that are logically wide, and hence semantic externalism (SE) is true. Such predicates also provide straightforward counterexamples to PAC. Suppose for instance that (3) is true. Then given PAC, it follows that Laura can know a priori that she's thinking that George is cute. But this consequence is false. Since what Laura supposedly knows a priori is relational with respect to George, she could just deduce that George exists from something she knows a priori, and so she could also know a priori that George exists, which is obviously absurd.

There has been much resistance to my argument for incompatibilism, mainly from externalists who insist, contrary to my argument, that an externalist thesis should *not* assert that wide cognitive properties logically or conceptually imply the existence of external objects of a certain sort. Instead, these compatibilists claim, the relevant relation between a wide cognitive property and the external facts that make it wide is much weaker than logical or conceptual

implication: perhaps the relation is only that of *metaphysical* implication, or even as weak as *counterfactual* implication.⁸

This kind of objection is always made against applications of my argument to cases of cognitive predicates that contain natural kind terms, such as (1) and (2) above. I think that whatever plausibility this objection might have derives from the general lack of consensus about, or clear understanding of, both the semantics of natural kind terms and the semantics of cognitive predicates that contain such terms. In my opinion, the only clear accounts of the semantics of cognitive predicates containing natural kind terms imply that such predicates express logically wide properties that are relational with respect to external objects or substances.⁹ So in my view the semantic facts about cognitive predicates containing natural kind terms are sufficient to support a strong form of externalism (**SE**) that is clearly inconsistent with the form of privileged access provided by **PAC**.

But even if I'm wrong about this, there should still be general agreement that **PAC** is false. For even in the absence of any generally accepted semantic views of singular terms, natural kind terms, and the cognitive attitude verbs, there should still be a consensus that many simple cases of cognitive predicates containing proper names and indexical pronouns, cases like the predicates found in sentences (3) and (4), express relational, logically wide properties. And these simple cases suffice to show that the *strongest* form of semantic externalism (**SE**) is true, from which it follows by my *reductio* argument that **PAC** is false.

So in my opinion, those who seek to hold on to privileged access in the form of **PAC** by endorsing a weak form of externalism for cognitive predicates containing natural kind terms, are engaged in a pointless exercise, like performing CPR on a long-dead horse. For the simple facts

about cognitive predicates containing names and indexicals show that the strongest form of semantic externalism (SE) is in fact true anyway, and hence that PAC is false, whatever the truth may be about natural kind terms and the cognitive predicates that contain them. Unlike PAC, then, a correct principle of privileged access must be restricted so as to imply at most that we can have a priori knowledge of our possession of certain logically *narrow* properties.¹⁰

3. Transmission of warrant.

While most critics of my argument for incompatibilism object to the argument's main premise - that externalist theses should ascribe logical relations between cognitive properties and external facts - Davies (1998) and Wright (2000) in effect accept my premises (at least for the sake of argument) but claim that my argument is invalid.¹¹ They contend that even when a person both knows a priori that a given proposition **P** is true, and knows a priori that a further proposition **Q** logically follows from **P**, it nevertheless does not follow that the person in question is in a position to know a priori that **Q**.

They base this contention on an important point about the transmission of warrant made by Wright (1985). He and Davies both provide several persuasive examples to show that a person's epistemic warrant for believing a given premise **P** will not necessarily transmit to also provide warrant for a given deductive consequence **Q** of **P**, even when the person knows that **P** logically implies **Q**. Consider, for instance, Wright's nice example of the soccer game (Wright 1985, p. 436; Wright 2000, pp. 141-3, Davies 2000, pp. 397-9). Spectator S is apparently witnessing a soccer game in progress. Seeing a player drive the ball into the net and then seeing the referee turn and point to the center circle, S has warrant for

- (5) A soccer goal has just been scored.

Of course, S also knows that (5) logically implies

(6) A game of soccer is in progress.

But S's warrant for believing (5) cannot be the *basis* of S's having warrant for (6). If S had any doubt as to whether a soccer game really is in progress (perhaps a scene is being staged for a movie), this doubt could not be settled by appeal to (5). For S will have warrant for believing that a goal has just been scored only in the context of having warrant in the first place for believing that a genuine game really is in progress. In a situation like this, any argument that S might give for (6) with (5) as premise would just beg the question.

It would seem that in general all deductively valid arguments that beg the question, or that could beg the question in a given epistemic context, provide examples in which warrant can fail to transmit from a set of premises to a known deductive consequence of those premises. For my purposes here, the important fact is that the very cases to which my argument for incompatibilism most clearly applies, cases where a premise ascribing a wide cognitive property logically implies an empirical conclusion, are also cases in which the relevant inference would beg the question. For instance, we've seen that (3) ('Laura is thinking that George is cute') logically implies that George exists. But of course, the *source* of one's warrant for believing that George exists could not be that one has correctly deduced this conclusion from the relational premise in question, since one would not be warranted in believing the relational premise in the first place, unless one were *already* warranted in believing that George exists.

Thus, in the very cases to which my argument for incompatibilism most clearly applies, warrant does not transmit from the cognitive premise to the externalist consequence. But contrary to what both Davies and Wright contend, this does not show that there is anything

wrong with my argument. For the argument does not assume that warrant is always transmitted from the relevant cognitive premises to the externalist conclusions. Rather, the argument assumes a certain principle to the effect that one's capacity for a priori knowledge is closed under logical implication:

Closure of Apriority under Logical Implication (CA)

Necessarily, for any person x and any propositions **P** and **Q**, if x can know a priori that **P**, and **P** logically implies **Q**, then x can know a priori that **Q**.¹²

When my argument for incompatibilism is applied to a particular relational cognitive premise like (3), we assume for *reductio* that (3) is knowable a priori (by Laura). This assumption, as it turns out, is contrary to fact. But if (3) *were* knowable a priori, then **CA** generates the absurd consequence that Laura could know a priori that George exists. And surely, that is the correct result. For if (3) *were* knowable a priori, then it would be knowable without empirical investigation. Hence any assumption on warrant for which knowledge of (3) is based would itself have to be knowable without empirical investigation. For otherwise, knowledge of (3) would, contrary to our assumption for *reductio*, be based in part on empirical investigation after all, since it would be based in part on whatever empirical investigation is required to provide warrant for the relevant empirical assumption. In this particular case, knowledge of (3) clearly depends on the agent's having warrant for the assumption that George exists. Hence one could not have a priori knowledge of (3), unless one also had a priori knowledge that George exists.

Here then is a clear case in which warrant fails to transmit from a premise to a deductive consequence of that premise, even though my closure principle for apriority **CA** yields the right result. And in general we may conclude that failure of warrant to transmit to a given conclusion

is *not* a good reason to suspect that closure of apriority fails in the same case. Failure of warrant-transmission is perfectly consistent with successful closure of apriority.

This is even more obviously shown by the simple case of conjunction. Any argument from a conjunction **P&Q** to either conjunct as conclusion would of course be blatantly question-begging, and so warrant fails to transmit from premise to conclusion. But just as obviously, closure of apriority *holds* in such a case: since one can't know **P&Q** without knowing **P** (and knowing **Q**), one also cannot know **P&Q** a priori without knowing **P** a priori (and knowing **Q** a priori).

So it is difficult to understand why Davies and Wright should just assume, without argument or discussion, that failure of warrant-transmission implies failure of closure for apriority. This is especially puzzling, since both Davies and Wright are careful to emphasize that failure of warrant-transmission does not imply that either knowledge or warrant fail to be closed under known entailment. (See Davies 1998, p. 349 and 2000, pp. 393-4, and Wright 2000, pp. 140-1 and 157; see also Bob Hale's useful discussion of transmission versus closure in Hale 2000.) Let us briefly consider why failure of warrant-transmission does not imply failure of closure of knowledge or warrant. In cases where warrant for a given premise **P** transmits to a known deductive consequence **Q** of **P**, we can use transmission of warrant to *explain* closure of knowledge and warrant: we can say, for instance, that a person knows (or has warrant for) **Q** *because* he or she knows (or has warrant for) **P** and has correctly deduced **Q** from **P**, or *because* he or she knows (or has warrant for) **P** and knows that **P** implies **Q**. When transmission of warrant fails, then of course this sort of explanation of closure cannot be given, but this doesn't imply that closure must fail. Rather, it implies only that if closure succeeds, then the explanation

of why it succeeds must be *different* from any explanation that assumes transmission of warrant. In short, transmissibility of warrant is a sufficient, but not a necessary, condition for closure of knowledge and warrant. (See Wright 2000, pp. 140-1 and Hale 2000, p. 173.)

We've just seen that this same relation also holds between transmission of warrant and closure of apriority. Sometimes, when warrant successfully transmits from a premise **P** to a deductive consequence **Q** of **P**, we can use this fact to explain how a person can know **Q** a priori, namely, by knowing **P** a priori and correctly deducing **Q** from **P**. When warrant fails to transmit from **P** to **Q**, we no longer have *this* way of explaining why **Q** must be knowable a priori if **P** is. Nevertheless, some other explanation may be available: perhaps, in the given case, we can explain why a priori knowability of premise **P** *requires* a priori knowability of deductive consequence **Q**, even though one's warrant for **P** fails to provide warrant for **Q**. This is the sort of explanation that I gave earlier for the case of conjunction, and for the case in which a priori knowability of (3) ('Laura is thinking that George is cute') requires a priori knowability of (3)'s deductive consequence that George exists, even though warrant for (3) does not transmit to this consequence.

In short, for the same reason that failure of warrant-transmission does not refute closure of knowledge and warrant, it also does not refute closure of apriority. Hence, my *reductio* argument is not shown to be unsound by the Davies/Wright point that warrant fails to transmit from premise to conclusion in many of the cases to which the argument applies. For again, my argument assumes only closure of apriority, and this principle is consistent with failure of warrant-transmission.

4. Strong and weak apriority.

The sense of 'a priori' knowledge in terms of which my argument for incompatibilism was first stated is a fairly strong one, since I meant knowledge that could be obtained without perceptual observation or empirical investigation, and without having to make any empirical assumptions. (Again, see my 1991a, p. 9 and 1987, pp. 2-3.) Davies has recently suggested that, while my *reductio* argument goes through with a priori knowledge understood in my strict sense, the same form of argument will not work to show the incompatibility of externalism and a form of privileged access that is understood in terms of a weaker kind of a priori knowledge. In this weaker sense, a priori knowledge is simply knowledge that is not justificatorily based on empirical investigation, though it might (in ways that Davies does not specify) involve the making of empirical assumptions. (See Davies 2000, pp. 406-8, and Chapter _ above. Wright (2000, p. 152) also seems to be relying on a weak sense of the a priori.)

Davies suggests that my argument works when 'a priori' is given the strict sense but fails when 'a priori' is given the weak sense, because strict but not weak a priori warrant successfully transmits from premises to the known deductive consequences of those premises (2000, p. 407). This is wrong. No matter how strong a sense of 'a priori' we choose, there can be question-begging valid arguments whose premises are knowable a priori in that sense, even though warrant does not transmit from premises to conclusion. For instance, even the strongest a priori warrant for a conjunction will fail to transmit to its conjuncts. Transmission of warrant is not the issue. The only question for my argument that is raised by the distinction between the weak and strong a priori is the question of whether or not my closure principle **CA** remains true under these two different interpretations of 'a priori'.

In discussing this question, I will use explications of the weak and strong a priori derived from those proposed by Hartry Field (1996).¹³ I find Field's proposals salutary because they replace the obscure notion of an 'empirical assumption' that I had used with the somewhat clearer notion of empirical defeasibility. Let us say that a person x 's knowledge that p is *weakly a priori* if and only if x 's knowledge that p is not based, even in part, on either perceptual observation or empirical investigation. And, following Field, let us say that x 's knowledge that p is *strongly a priori* if and only if x 's knowledge that p is both weakly a priori and empirically indefeasible. Here, I understand that x 's knowledge that p is empirically indefeasible if and only if the warrant or justification on which x 's knowledge that p is actually based could not possibly be undermined or outweighed by any additional empirical evidence. In Field's terminology, I mean here that the person's warrant could not be defeated 'in the primary way' by any empirical evidence, where a person's warrant for a belief is defeated by additional evidence in the primary way when the person's actual warrant would no longer support the belief, given the additional evidence. (In the discussion to follow, by 'defeat' I will always mean 'primary defeat'.)¹⁴

Is the capacity for strongly a priori knowledge closed under logical implication? Suppose that x has strongly a priori knowledge that **P**, and that **P** logically implies a proposition **Q**. If x 's empirically indefeasible warrant for **P** transmits to **Q**, then obviously x can have strongly a priori knowledge that **Q**, by deducing **Q** from **P**. But suppose that x 's warrant for **P** does *not* transmit to **Q**. Could it then fail to be possible for x to have strongly a priori knowledge that **Q**? If so, then knowledge that **Q** would have to always be either empirically defeasible, or obtainable only by perceptual observation or empirical investigation. But then it certainly seems that in either case, it would follow that **Q** could be shown false by empirical evidence, and since **P** logically implies

Q, it follows that **P** could *also* be shown false by empirical evidence. There certainly is a tension between this consequence and our assumption that *x*'s knowledge that **P** is empirically infeasible. However, there seems to be no inconsistency here: whatever empirical evidence that there could be against **Q** (and hence against **P**) might in fact be insufficient to outweigh *x*'s empirically infeasible warrant for **P**.¹⁵

So while the Davies/Wright point about transmission of warrant does not show that closure of apriority fails to hold generally, the point does make it difficult to *prove* that closure of apriority does hold generally, whether we take apriority in either the strong or the weak sense. This in turn has the effect of *undermining confidence* in the closure principle on which my *reductio* argument depends, whether the argument is understood in terms of strong apriority or in terms of weak apriority. This difficulty can be overcome by proving restricted principles of closure for both strong and weak apriority. As we shall see, these restricted principles will serve most of the purposes of my original argument.

5. Restricted closure principles for strong and weak apriority.

In stating the restricted closure principles in question, I will make crucial use of an important idea introduced by Wright (2000, p. 143) and further refined by Hale (2000, pp. 177-81), namely, the idea of a person's warrant for a given proposition being *dependent* on a given body of collateral information **I**. Following Hale's definition (p. 181), let us say that a person *x*'s warrant *w* to believe that *p* *strictly depends* upon information **I** if and only if: (i) if it were not the case that for each *q* **I**, *x* has a warrant *w* to believe that *q*, then *w* would not warrant *x* in

believing that p , and (ii) x 's combined warrants to believe the propositions in **I** do not, by themselves, warrant x in believing that p .

Now suppose again that x has strongly a priori knowledge that **P** and that **P** logically implies **Q**. As we've seen, when x 's warrant for **P** transmits to **Q**, x can also have strongly a priori knowledge that **Q**. Moreover, in cases where x 's warrant for **P** does *not* transmit to **Q**, this will typically be because x 's warrant for **P** strictly depends on information containing **Q**, so that any inference from **P** to **Q** would be question-begging. It is easy to see that in such a case, it follows that x can have strongly a priori knowledge that **Q**. For first, x 's warrant for **Q** must be empirically indefeasible, since if this warrant were empirically defeasible, then x 's warrant for **P**, which strictly depends on warrant for **Q**, would itself be empirically defeasible, contrary to our assumption that x 's knowledge that **P** is strongly a priori. Second, x 's warrant for **Q** must be based on neither perceptual observation nor empirical observation; otherwise, x 's warrant for **P**, which strictly depends on warrant for **Q**, would itself be based in part on perceptual observation or empirical investigation, again contrary to our assumption that x 's knowledge that **P** is strongly a priori.

These considerations show conclusively that the capacities for both strongly and weakly a priori knowledge are closed under logical implication in cases where either (i) warrant transmits from premise **P** to conclusion **Q**, or (ii) the agent's warrant for **P** strictly depends on information containing **Q**. So we can be confident in the truth of the following restricted closure principles for strong and weak apriority, which for brevity I state together:

Restricted Closure of Apriority under Logical Implication (R-CA)

Necessarily, for any person x and any propositions **P** and **Q**, if (i) **P** logically implies **Q**, and (ii) x can have strongly [weakly] a priori knowledge that **P** in circumstances in which either (iii) x 's warrant for **P** transmits to **Q**, or (iv) x 's warrant for **P** strictly depends on information **I** such that **Q** **I**, then x can have strongly [weakly] a priori knowledge that **Q**.

Certainly, in the kinds of cases to which my *reductio* argument most clearly applies, use of the restricted principle **R-CA** effectively yields the same results as the unrestricted principle **CA**.

Whenever a given cognitive premise such as (3) ('Laura is thinking that George is cute') is relational with respect to a given object and thus logically implies the existence of that object, a person's having warrant for believing the premise will always strictly depend on the person's having warrant for believing the conclusion. Hence, in such cases, **R-CA** applies to show that if a person can have strongly (or weakly) a priori knowledge of the premise, then that person can have strongly (or weakly) a priori knowledge of the conclusion.

I think it is very likely that any interesting externalist semantic thesis will imply that cognitive ascriptions are logically related to external consequences in such a way that the restricted closure principle **R-CA** will apply to that logical relation. The typical case is one in which a cognitive predi-cate ascribes the having of a thought with a logically wide content, such as a singular proposition. In such cases, the very existence of the content ascribed logically presupposes the truth of some external proposition, and so any warrant one might have for

believing that a given thought has such a content will strictly depend on having warrant for the presupposed external proposition, and thus **R-CA** will apply.

6. Wright's conjecture.

Falling back to the restricted closure principle **R-CA** means that we have to accept at least the logical possibility that some form of semantic externalism might escape my *reductio* argument and perhaps be consistent with privileged access in the form of **PAC**. The situation would look like this: on the externalist thesis in question, (i) a cognitive premise **P** would logically imply an empirical externalist consequence **E**; (ii) warrant for **P** would not transmit to **E**; and (iii) warrant for **P** would also not strictly depend on any information containing **E**. In such a situation, **R-CA** could not be applied to show that given strongly (or weakly) a priori knowledge that **P**, one could (absurdly) have strongly (or weakly) a priori knowledge that **E**. How seriously should we take such a possibility?

Wright (2000) uses Fred Dretske's (1970) famous discussion of scepticism as the basis for a conjecture that externalist theses about certain cognitive premises precisely fit the logically possible model just described. In Dretske's zebra-case, a man, Jones say, is visiting the zoo and sees an animal in the pen before him that looks exactly like a zebra. Jones therefore believes

(7) The animal in the pen is a zebra.

We can assume that Jones's perceptual warrant for (7) is sufficient for knowledge and that (7) is true, so that Jones in fact knows (7). But (7) logically implies

(8) The animal in the pen is not a mule cleverly disguised as a zebra.

As Wright points out (p. 154), Jones's warrant for (7) clearly does not transmit to (8), since this warrant for (7) - which just consists of the way the beast looks - could not provide Jones with a new reason for believing (8) for the first time. Moreover, it seems that Jones's warrant for (7) also does not strictly depend on Jones's having warrant for (8). After all, Jones has done nothing special to make sure that (8) is true. Wright suggests that while Jones lacks warrant for (8), he is nevertheless *entitled* to discount the remote possibility that (8) denies, and so he is entitled, without having warrant, to assume that (8) is true (p. 155).

If Wright's description of Dretske's case is correct, then (7) and (8) provide an example of propositions **P** and **Q** such that **P** logically implies **Q**, but warrant for **P** neither transmits to **Q** nor strictly depends on information containing **Q**. So **R-CA** cannot be applied to a case like this.¹⁶ But Wright claims that the situation regarding cognitive premises to which externalist theses apply is closely analogous to the zebra-case. Consider

(9) Oscar is thinking that water is wet.

And assume some externalist thesis according to which (9) logically implies, say

(10) Oscar has had such-and-such encounters with water.¹⁷

According to Wright, Oscar could have (weakly) a priori knowledge of (9) purely by reflection, even though (9) entails the empirical proposition (10), which Oscar could *not* know a priori in any sense. Wright's rationale for this claim rests crucially on two further claims: first, that Oscar's warrant for (9) would not transmit to (10), and second, that Oscar's warrant for (9) would also not strictly depend on warrant for (10), even though (9) logically implies (10).

While the first of these claims is surely correct, the second is much more tenuous. In the zebra-case, it's plausible to say that Jones's having warrant for believing (7) does not depend on

his having warrant for believing (8), simply because it's also plausible (though not obviously true) to say that Jones would in fact *not* have warrant for (8): since Jones has done nothing special to make sure that he's *not* seeing a cleverly disguised mule, perhaps we should say that Jones has no warrant, but merely entitlement, for assuming that (8) is true. But of course we can't say anything like this about (10). If Oscar is like the rest of us, he has seen, drunk, and washed in water several times a day on every day of his life, and so he would have as much empirical warrant for (10) as it is possible for a person to have. Thus we obviously cannot argue that Oscar's having warrant for (9) does not depend on his having warrant for (10) by appeal to the (obviously false) premise that while Oscar has warrant for (9), he has *no* warrant, only entitlement, for believing (10).

Moreover, it does not really seem *possible* for someone to be entitled to assume that a proposition like (10) is true, in the absence of any empirical warrant for that proposition. Surely, if Oscar had never had perceptual experiences of water, had never even acquired any indirect evidence for water's existence, and had thus never acquired any empirical warrant for believing that he'd had encounters with water, then he would also never have acquired any sort of *entitlement* to assume that he'd had such encounters. It in fact seems precisely as implausible to suppose that one could have unwarranted entitlement to assume that one has had watery encounters, as it is to suppose that one could have a priori knowledge by pure reflection that one has had such encounters.

Why then does Wright at least implicitly endorse the implausible claim that one could have unwarranted entitlement to assume such a proposition as that one has had such-and-such encounters with water? It seems to be due to an analogy that Wright draws between cases of self-

knowledge and cases of perception. When one looks at a wall and sees that it is red, one acquires warrant for believing that the wall is red. In such a case, according to Wright, one is entitled to assume, without having any particular warrant for doing so, that the circumstances are normal, and hence suitable for acquiring knowledge by perception. In other words, Wright is suggesting, one is normally entitled without warrant to assume that one is not the victim of an illusion of the sort that would result from, say, a white wall cleverly illuminated to look red. Similarly, Wright suggests, in cases of self-knowledge, one is normally entitled without warrant to assume that one is not the victim of an 'illusion of content' of the sort that one would have in seeming to think that water is wet, having in fact never experienced water. (Wright 2000, pp. 152-3.)

But how exactly is it supposed to follow from this suggestion that one generally has unwarranted entitlement to assume the externalist presuppositions of one's thoughts?

Presumably, Wright's underlying reasoning must run somewhat as follows:

- (11) When it seems to Oscar that he's thinking that water is wet, Oscar is entitled without warrant to assume that he's not then having an illusion of content.
[Wright's Premise]
- (12) If it seems to Oscar that he's thinking that water is wet, and Oscar is not then having an illusion of content, then Oscar has had such-and-such encounters with water. [Externalist Premise]

Therefore,

- (13) If it seems to Oscar that he's thinking that water is wet, then Oscar is entitled without warrant to assume that he's had such-and-such encounters with water.

Now we've already seen that (13) would be pretty obviously false: while it may well seem to a person that he or she is thinking that water is wet, it's *never* true that a person has unwarranted entitlement to assume that he or she has had such-and-such encounters with water. I would suggest that the inference from (11) and (12) to (13) is simply invalid, because it implicitly depends on a false principle to the effect that unwarranted entitlement is closed under implication (though it is not clear which form of implication is being invoked). That *something* is certainly wrong with this inference is shown by the following analogous inference about perception:

(14) When it seems to Oscar that he's seeing a red wall, Oscar is entitled without warrant to assume that he's not then having a perceptual illusion.

(15) If it seems to Oscar that he's seeing a red wall, and Oscar is not then having a perceptual illusion, then Oscar is seeing a red wall.

Therefore,

(16) If it seems to Oscar that he's seeing a red wall, then Oscar is entitled without warrant to assume that he's seeing a red wall.

While the premises (14) and (15) are both plausible, the conclusion (16) would normally be false. For after all, its seeming to one that one is seeing a red wall is precisely the way that one typically *acquires* warrant for believing that one is seeing a red wall. In such a case of course, (16)'s antecedent would be true and its consequent false.¹⁸

So it is both implausible and unjustified to suggest, as Wright does, that when a person has (weakly) a priori knowledge that **P**, and **P** logically implies some very strong empirical proposition **E**, the person need have no warrant, only entitlement, for assuming **E**. But Wright is

surely correct when he points out (p. 155) that the person in question would have to be entitled to assume the relevant empirical consequence. In the case of Oscar, for instance, if Oscar has only weak, and hence empirically defeasible, a priori knowledge that (9) is true, and (9) logically implies (10), then Oscar must be entitled to assume that (10) is true, since (10)'s falsehood would empirically defeat (in the primary way) Oscar's a priori warrant for (9). However, as we've just seen, Oscar is entitled to assume that (10) is true only if Oscar has warrant for (10). Hence, Oscar's having warrant for (9) would strictly depend on his having warrant for (10). Thus my restricted closure principle **R-CA** applies in this case to show that, given the externalist thesis that (9) logically implies (10), then since Oscar cannot have weakly a priori knowledge that (10) is true, he also cannot have privileged (weakly) a priori knowledge that (9) is true.

So again it seems that forms of externalism will typically be ones to which the restricted closure principle **R-CA** applies, yielding inconsistency with privileged access. For most forms of externalism imply that certain cognitive premises have very strong empirical consequences like (10). In such cases, warrant for the cognitive premise will strictly depend on information containing the empirical conclusion. In Wright's terminology, such externalisms imply that warrant for the cognitive premises in question is always *information-dependent*, where some of this information is empirical and external.¹⁹ But as I have shown, no such cognitive proposition is one of which we can have either weakly or strongly a priori knowledge, and so there can be no privileged access to the truth of such a proposition, in any sense.

Given Wright's discussion, it seems that the only forms of externalism that might escape my *reductio* argument are those which imply at most that some cognitive premises have weak empirical consequences to the effect that various wildly improbable sceptical hypotheses are

false. Perhaps in such cases an agent's (weakly) a priori warrant for the cognitive premise would depend only on the agent's being entitled, perhaps without warrant, to assume that the relevant sceptical hypotheses are false. However, no forms of semantic externalism that I know of are this weak, and it is certainly difficult to see what the motivation for such a view might be.

7. Conclusion.

The Davies/Wright point about transmission of warrant cannot be used to refute my *reductio* argument for incompatibilism, since my argument relies only on closure of apriority, and failure of warrant-transmission is consistent with successful closure of apriority. However, the fact that warrant for a premise can fail to transmit to a known deductive consequence of that premise forces us to reexamine our intuitions about closure of apriority. When transmission of warrant fails for a given premise and conclusion, how could closure of apriority succeed? The answer lies in Wright's conception of a warrant's being information-dependent: when warrant for the premise strictly depends on information containing the conclusion, not only does transmission of warrant *fail*, but closure of both strong and weak apriority *succeeds*. This allows us to have complete confidence in a restricted principle of closure for apriority (**R-CA**), which says that apriority is closed under logical implication precisely when either warrant transmits from premise to conclusion or warrant for the premise strictly depends on warrant for the conclusion. As I've argued here, this restricted principle serves most of the purposes of my original argument, which relied on the unrestricted principle **CA**.

I mentioned earlier that adopting a restricted closure principle for apriority opens up the logical possibility that some specific form of semantic externalism might be found to which the

restricted principle does not apply, so that incompatibility of this form of externalism with privileged access in the form of **PAC** could not be shown. But before diehard compatibilists use this fact as an excuse to rush out and search for the relevant form of externalism, let me reiterate a point that I made earlier: such a search would be a pointless exercise. For again, the semantic facts show that *de dicto*-structured cognitive predicates containing names and indexicals express properties that are relational with respect to external objects, and the undeniable closure principle **R-CA** shows conclusively (by *reductio*) that no one could possibly have either strongly or weakly a priori privileged access to their possession of such a relational cognitive property. Hence, we already know that **PAC** is false, and so there is simply no point in looking for forms of externalism that are consistent with it.²⁰

NOTES

1. The paper (1991a) in which the argument appeared had earlier been presented to the Pacific Division of the American Philosophical Association (Los Angeles, March 1990). From 1978 on, I've regularly presented the argument in classes and seminars at Wayne State. I also gave the argument in a question from the floor at the April, 1985 Oberlin Colloquium after Tyler Burge's presentation of his (1988a) there, and I gave it again in a question from the floor after Burge's presentation of his (1988b) as the Nelson lecture at the University of Michigan, February, 1986.

2. Essentially the same principle as this was more clearly expressed, and endorsed, by McLaughlin and Tye (1998, p. 286):

Privileged access thesis: It is conceptually necessary that if we are able to exercise our normal capacity to have beliefs about our occurrent thoughts, then if we are able to occurrently think that *p*, we are able to know that we are thinking that *p* without our knowledge being justificatorily based on empirical investigation of our environment.

I agree with the salutary qualification that to have the relevant a priori knowledge we must be "able to exercise our normal capacity to have beliefs about our occurrent thoughts." We can take this to be covered by the qualifier 'in principle' in **PAC**.

3. See McKinsey 1991a, p. 15 (principle (Ba)). Here and below I mean 'logically implies' in a broad sense that includes conceptual implication. See McKinsey 1991a, p. 14 and 1991b, p. 152.

4. In McKinsey 1999 I argue that some names do in fact have descriptive meanings, and that as a result, uses of such names in cognitive contexts, unlike the use of 'George' in (3), are non-relational.

5. I believe this in part on the basis of Kripke's (1972) evidence. See also McKinsey 1984, 1994, and 1999.

6. This is in fact the standard externalist view of how cognitive ascriptions containing small-scope names and indexicals should be understood. See for example McDowell 1977, Evans 1982, McKay 1981, Salmon 1986, and Soames 1987.

7. It used to be common to say that demonstratives like 'he' or 'that man' must have *largest scope* relative to any cognitive operator. If so, then a sentence like (4) could not be structurally *de dicto*. But Nathan Salmon (1986, p. 4) effectively showed that this suggestion is a nonstarter. See also McKinsey 1998, pp. 17-18.

8. Brueckner (1992, pp. 113, 114, 116) seems to make *both* of these claims. See also Gallois and O'Leary-Hawthorne 1996, and McLaughlin and Tye 1998. I reply to this sort of objection at length in McKinsey 1994b, 2001, and 2002.

9. In my own work on this topic, I've provided both an explanation of precisely why cognitive predicates containing natural kind terms express logically wide mental properties and an account of what these properties are. I know of no other externalist semantics that does these things. See McKinsey 1987 and 1994a, pp. 321-4.

10. For a statement of what I consider to be the correct principle of privileged access, see McKinsey 1994a, pp. 308-9. See also McKinsey 2001 and 2002.

11. By contrast, McLaughlin and Tye (1998) explicitly say that my argument has a valid form (which they call 'McKinsey's recipe') but object to the argument's main premise. See McKinsey 2001 for my reply to their objection.

12. Most interpretations of my argument do not see it as relying on **CA**, but rather on another very plausible closure principle, namely:

CAK: Necessarily, for any person x and any propositions **P** and **Q**, if x can know a priori that **P** and x can know a priori that if **P** then **Q**, then x can know a priori that **Q**. (See for instance Brown 1995, Gallois and O'Leary-Hawthorne 1996, Boghossian 1997, Davies 1998, McLaughlin and Tye 1998, and Wright 2000.) The serious weakness of versions of my *reductio* that use **CAK** instead of **CA**, is that application of **CAK** in the argument requires the additional, and in my opinion false, assumption that externalist theses like (2) in the text are themselves knowable a priori. For a defense of my use of **CA** rather than **CAK**, and an explanation of why my argument does not have to assume that externalist theses are knowable a priori, see McKinsey 2002. See also my discussion of Boghossian 1997 in McKinsey, forthcoming.

13. Brian McLaughlin (2000 and Chapter _ above) also uses Field's conceptions of the weak and strong a priori in his discussion of issues closely related to my argument.

14. Field (1996, pp. 361-2) distinguishes primary defeat from what he calls 'secondary defeat', where additional evidence does not undermine or outweigh the person's actual warrant, but rather would show that the person does not *have* the warrant in question. For example, suppose that *x* correctly claims to know a priori a theorem of logic to the effect that **P**, by virtue of *x*'s having given a simple proof that **P**. Possible empirical evidence that *x* was on drugs when producing the proof, or that leading logicians say that the proof is faulty, would defeat *x*'s warrant in the secondary way. As Field points out, such examples seem to show that even the most obvious cases of the strongest possible a priori knowledge would be 'empirically defeasible' in the secondary way, and if so, then we should *not* understand the strongly a priori in terms of secondary defeat.

Davies (2000, p. 408) uses an example of Christopher Peacocke's (1999, pp. 244-5) to suggest that there probably is no such thing as self-knowledge that is strictly a priori in my sense. But Peacocke's example appears to be an example of secondary defeat, and if so, then the example is not really relevant to any plausible conception of strongly a priori knowledge. In this connection, see also Burge's discussion (Burge 1993, pp. 463-4) of the role of memory in the acquisition of knowledge by deductive proof.

15. Here I am indebted to McLaughlin, who makes a similar point (McLaughlin 2000, pp. 106-7).

16. Note that, if Wright's description of Dretske's case is correct, it also follows that warrant is not closed under known entailment, since Jones is supposed to have *no* warrant for (8), even though Jones has warrant for (7) and knows that (7) entails (8). Moreover, given that knowledge requires warrant, it also follows that knowledge itself is not closed under known entailment. So if Wright's description of Dretske's case is correct, that case is a counterexample to both closure of knowledge and closure of warrant. This fact is difficult to square with Wright's explicit scepticism (p. 141) as to whether there are any genuine counterexamples to closure. (See also Hale 2000, p. 184.)

17. The sort of externalist consequence that Wright actually considers is not (10) but disjunctions of the form 'Either Oscar or Oscar's speech community has had such-and-such encounters with water' (Wright 2000, p. 156). I use the simpler (10) merely for ease of exposition.

18. In the above discussion, I use the phrase 'is entitled without warrant' to just mean 'both is entitled and has no warrant'. Perhaps a defender of Wright's argument would suggest that the phrase should not be understood so as to imply lack of warrant, but rather should be taken to mean 'is entitled and would still be entitled even without warrant'. Since my case against Wright is unaffected by this change, the reader is free to read the phrase in either sense.

19. In his response to Wright (2000), Alfonso García Suárez (2000) defends my *reductio* argument by contending that externalist theses imply that warrant for the relevant cognitive

premises would be information-dependent. So my reaction to Wright's conjecture is essentially the same as that of Suárez, to whom I am indebted for his useful discussion.

20. For useful conversations about these matters, I am grateful to Martin Davies, Brian McLaughlin, Bruce Russell, and Crispin Wright. I owe a special debt to the papers of Davies (1998, 2000) and Wright (2000). As I hope the text indicates, by thinking about these papers I have learned a great deal about my *reductio* argument.

REFERENCES

- Boghossian, Paul A.: 1997, "What the Externalist Can Know A Priori," in C. MacDonald, B. Smith, and C. Wright (eds.) Knowing Our Own Minds (Oxford: Oxford University Press), 271-284.
- Brown, Jessica: 1995, "The Incompatibility of Anti-Individualism and Privileged Access," Analysis 55, 149-156.
- Brueckner, Anthony: 1992, "What an Anti-Individualist Knows A Priori," Analysis 52, 111-118.
- Burge, Tyler: 1988a, "Cartesian Error and the Objectivity of Perception," in Contents of Thought, R. Grimm and D. Merrill, eds. (Tucson: Arizona University Press), 62-76.
- Burge, Tyler: 1988b, "Individualism and Self-Knowledge," Journal of Philosophy 85, 649-663.
- Burge, Tyler: 1993, "Content Preservation," Philosophical Review 102, pp. 457-88.
- Davies, Martin: 1998, "Externalism, Architecturalism, and Epistemic Warrant," in C. MacDonald, B. Smith, and C. Wright (eds.) Knowing Our Own Minds (Oxford: Oxford University Press), 321-361.
- Davies, Martin: 2000, "Externalism and Armchair Knowledge," in P. Boghossian and C. Peacocke (eds.), New Essays on the A Priori (Oxford: Clarendon Press), 384-414.
- Dretske, Fred: 1970, "Epistemic Operators," Journal of Philosophy 67, 107-122.
- Evans, Gareth: 1982, Varieties of Reference (Oxford: The Oxford University Press).
- Field, Hartry: 1996, "The Apriority of Logic," Proceedings of the Aristotelian Society 96, pp. 359-76.
- Gallois, André, and O'Leary-Hawthorne, John: 1996, "Externalism and Scepticism," Philosophical Studies 81, 1-26.
- Hale, Bob: 2000, "Transmission and Closure," Philosophical Issues 10: Skepticism, E. Sosa and E. Villanueva (eds.). pp. 172-90.
- Kripke, Saul: 1972, "Naming and Necessity," in D. Davidson and G. Harman (eds.), Semantics of Natural Language (Dordrecht: D. Reidel), 253-355.
- McDowell, John: 1977, "On the Sense and Reference of a Proper Name," Mind 86, 159-185.
- McKay, Tom: 1981, "On Proper Names in Belief Ascriptions," Philosophical Studies 39, 287-303.
- McKinsey, Michael: 1978, "Names and Intentionality," The Philosophical Review 87, 171-200.
- McKinsey, Michael: 1984, "Causality and the Paradox of Names," Midwest Studies in Philosophy 9, 491-415.
- McKinsey, Michael: 1987, "Apriorism in the Philosophy of Language," Philosophical Studies 52, 1-32.

- McKinsey, Michael: 1991a, "Anti-Individualism and Privileged Access," Analysis 51, 9-16.
Reprinted in The Twin Earth Chronicles, edited by A. Pessin and S. Goldberg (Armonk, NY: M. E. Sharpe, 1996), and in Externalism and Self-Knowledge, edited by P. Ludlow and N. Martin (Stanford: CSLI Publications, 1998).
- McKinsey, Michael: 1991b, "The Internal Basis of Meaning," Pacific Philosophical Quarterly 72, 143-169.
- McKinsey, Michael: 1994a, "Individuating Beliefs," Philosophical Perspectives 8, Logic and Language, 303-330.
- McKinsey, Michael: 1994b, "Accepting the Consequences of Anti-Individualism," Analysis 54, 124-128.
- McKinsey, Michael: 1998, "The Grammar of Belief," in F. Orilia and W. J. Rapaport (eds.), Thought, Language, and Ontology (Kluwer), 3-24.
- McKinsey, Michael: 1999, "The Semantics of Belief Ascriptions," Noûs 33(November, 1999): 519-557.
- McKinsey, Michael: 2001, "The Semantic Basis of Externalism," in Topics in Contemporary Philosophy: Meaning and Truth, J. Campbell, M. O'Rourke, and D. Shier, eds. (New York: Seven Bridges Press), 34-52.
- McKinsey, Michael: 2002, "Forms of Externalism and Privileged Access," to appear in Philosophical Perspectives 16.
- McKinsey, Michael: forthcoming, "Review of Knowing Our Own Minds, edited by C. Macdonald, B. Smith, and C. Wright," to appear in The Philosophical Quarterly.
- McLaughlin, Brian P., and Tye, Michael: 1998, "Externalism, Twin Earth, and Self-Knowledge," in C. MacDonald, B. Smith, and C. Wright (eds.) Knowing Our Own Minds (Oxford: Oxford University Press), 285-320
- McLaughlin, Brian P.: 2000, "Self-Knowledge, Externalism, and Skepticism," Proceedings of the Aristotelian Society, Supplementary Volume 74, pp. 93-117.
- Peacocke, Christopher: 1999, Being Known (Oxford: Clarendon Press).
- Salmon, Nathan: 1986, Frege's Puzzle (Cambridge, Mass.: The MIT Press).
- Soames, Scott: 1987, "Direct Reference, Propositional Attitudes, and Semantic Content," in N. Salmon and S. Soames (eds.), Propositions and Attitudes (Oxford: Oxford University Press, 1988), 197-239.
- Suárez, Alfonso García: 2000, "On Wright's Diagnosis of McKinsey's Argument," Philosophical Issues 10: Skepticism, E. Sosa and E. Villanueva (eds.). pp. 140-64.
- Wright, Crispin: 1985, "Facts and Certainty," Proceedings of the British Academy 71, 429-472.
- Wright, Crispin: 2000, "Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof," Philosophical Issues 10: Skepticism, E. Sosa and E. Villanueva (eds.). pp. 140-64.