

FORMS OF EXTERNALISM AND PRIVILEGED ACCESS

Michael McKinsey
Wayne State University

In my 1991 paper “Anti-Individualism and Privileged Access” I argued that externalism, or anti-individualism, in the philosophy of mind is incompatible with the traditional Cartesian principle that we each have privileged access to our own thoughts and other mental acts and states.¹ The argument has since received a gratifying amount of discussion, both pro and con. My aim here is to defend and clarify my argument in light of this discussion, and to clarify the overall conceptual situation regarding the logical relationships that hold between various forms of externalism and privileged access.

1. *The argument for incompatibilism.*

To have privileged access to a fact about one’s own mental states, I will assume, is to have the capacity to come to know that fact a priori, where to know something a priori is to know it “just by thinking”, and not on the basis of perceptual observation or empirical investigation. A priori knowledge of this sort, I take it, is neither based on empirical investigation nor requires any empirical assumptions. Although I think it is plausible to suppose that each of us has some form of privileged access to many different types of our cognitive acts and states, including our beliefs, intentions and desires, I will for simplicity confine my attention to principles on which we have privileged access to our occurrent thoughts. In my initial discussion, I considered the compatibility of externalism with the principle that we have privileged access to the *contents* of our thoughts:

Privileged Access to Content (PAC)

It is necessarily true that if a person *x* is thinking that *p*, then *x* can in principle know a priori that he himself, or she herself is thinking that *p*.²

And I took externalism to be the thesis that many *de dicto*-structured cognitive predicates of the form ‘is thinking that *p*’ express “wide” psychological properties. Since this is a semantic thesis, I will call it ‘semantic externalism’:

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 or conceptually implies the existence of contingent objects external to that agent.³

My argument that (PAC) and (SE) are inconsistent was straightforward and simple. I considered instances of ‘is thinking that p’ that are standardly assumed to express wide properties, namely, instances in which the imbedded sentence p contains a natural kind term such as ‘water’. Thus, suppose it is true that Oscar is thinking that water is wet. Then by (PAC), it is also true that

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

Now given that ‘is thinking that water is wet’ expresses a wide property, it will also be true that

- (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 contingent objects of some sort that are external relative to Oscar. However, the conjunction of (1) and (2) is clearly absurd. If Oscar can know a priori that he is thinking that water is wet, and the proposition that he’s thinking that water is wet logically implies E, then Oscar could in principle just correctly *deduce* E from something that he knows a priori, and hence he could come to know E itself a priori. But this consequence is just absurd, whatever proposition E is. For E is by assumption some proposition that asserts or implies the existence of contingent objects of some sort external to Oscar, and Oscar obviously cannot know such propositions a priori. Hence, if the property of thinking that water is wet is wide, then contrary to (PAC), no one can know a priori that he or she is thinking that water is wet. And in general, the principle that we have privileged access to the contents of our occurrent thoughts, is inconsistent with semantic externalism.⁴

Some critics of my original argument against compatibilism have objected to my use of the variable ‘E’ to stand in for whatever empirical, “external” proposition it is whose entailment makes the property of thinking that water is wet a wide property. Thus Brueckner (1992) complains that I am “rather vague” on the question of what E is (p. 112), and demands that my argument be applied to some particular E that is relevant according to Burge’s (1982) view. But this worry is misplaced. I used a variable in my argument so that it would be *general* in its application. I was not interested in showing that some particular form of semantic externalism like Burge’s is incompatible with (PAC), but rather that *any* (interesting) form of semantic externalism is incompatible with (PAC).⁵

In the case of cognitive predicates containing natural kind terms like ‘water’ there are in fact a great variety of actual and possible views as to why such predicates express wide properties. On some views, perhaps, one’s thinking that water is wet might require that one bear some causal relation to water; other views might require, not that one actually live in a world containing water, but only that one have had contact with other speakers who have the concept of water; other views might demand that a certain disjunction of these things be true; and so on. The fact is that it doesn’t at all matter to my argument which empirical, “external” proposition we choose as our instance of ‘E’, since it will be quite clear, for any such choice, that the proposition cannot possibly be known a priori, and this is all that my argument requires.

2. Logical versus metaphysical implication.

Another objection that has been made to my argument concerns its assumption that if a predicate like ‘thinks that water is wet’ expresses a wide cognitive property, then possession of that property must *logically*, or *conceptually*, imply the truth of some external proposition. It has been suggested that the relevant relation between a wide property and the external proposition that makes it wide is much weaker than logical or conceptual implication: perhaps the relation is merely that of *metaphysical* implication, or perhaps it is even as weak as mere *counterfactual* implication.⁶

In my original paper, the justification that I gave for taking the relevant relation to be logical or conceptual is that if we take the relation to be weaker, then the resulting version of externalism becomes trivial and uninteresting. Rather than repeat my argument for this claim, an argument that I now think was rather infelicitously expressed, I would like to make the point in a slightly different, but perhaps clearer way.⁷ Suppose that we define the notion of a “wide” psychological property in terms of metaphysical entailment, for example as follows:

- (3) The psychological property of being S is *wide* if and only if: necessarily, for any person x, x’s being S metaphysically entails that there exists some contingent object y such that y is not identical to x, to any part of x, or to any of x’s mental states, acts, or experiences.⁸

If we assume that a property is *narrow* if and only if it is not wide, then on the definition (3), a given property S* would be narrow only if it satisfies the following condition:

- (4) It is metaphysically possible for there to be a person x who has S*, even though no contingent object exists other than x, x’s parts, and x’s mental states, acts, and experiences.

A psychological property that satisfies this condition would be one that a person could have while in effect existing alone in the universe. But are there

really any such properties? Consider *cognitive* psychological properties, for instance. It may well be that only beings that have achieved a certain level of complexity are capable of cognition, and perhaps any being having such a level of complexity would in turn have to have a *nature* which requires that being to be related to other contingent objects external to itself. Certainly, as I pointed out in the original paper, this is true of every *human* being, since (as Kripke, 1972 first argued) every human's existence metaphysically depends on the existence of that human's biological parents. Thus no human being could (metaphysically) have any psychological property while existing alone in the universe, for the simple reason that no human being could exist alone in the universe, period. But then, for all we know, the same sort of thing might necessarily be true of *any* being that has the degree of complexity required to be capable of cognition. If so, then *all* cognitive properties would fail to satisfy condition (4), and hence all cognitive properties would be "wide" according to the definition (3).

I take it that no defender of internalism would want to be saddled with the contentious and unjustified assumption that there are psychological properties that satisfy condition (4). Thus the defender of externalism who adopts a sense of 'wide psychological property' that is understood in terms of metaphysical dependency, has adopted the uninteresting tactic of defining his opponents out of existence. Would anyone dare disagree with the externalist who asserts that many *de dicto*-structured cognitive predicates, like 'is thinking that water is wet', express metaphysically wide properties? For all we know, what the externalist has asserted is true even of cognitive predicates that express the most intuitively *narrow* of cognitive properties, involving the most intuitively *narrow* of contents.

Consider your favorite paradigm of a narrow cognitive property. I would choose, for example, such properties as *wondering whether God exists*, *thinking oneself to be a thinking thing*, and *believing that $2+2=4$* . Since these properties too might well be metaphysically wide, what the "metaphysical" externalist has asserted about the predicate 'is thinking that water is wet' says nothing about this predicate that, for all we know, is not also true of cognitive predicates that express the most intuitively *narrow* of properties. Thus the externalist has failed to say anything interesting or controversial with which an internalist is likely to disagree. To actually engage with the traditional Cartesian view that sees intuitively narrow properties such as those I've mentioned as fundamental, the externalist will have to adopt a sense of 'wide psychological property' that is understood in terms of logical, or conceptual implication.⁹

3. Relational cognitive predicates and the falsity of (PAC).

In "Anti-Individualism and Privileged Access" I argued merely for the conclusion that semantic externalism (SE) is inconsistent with the principle (PAC) that we have privileged access to the contents of our thoughts. I took no stand on which of these principles is true and which is false. But the semantic facts

about cognitive predicates containing small-scope names, indexicals, and natural kind terms show that (SE) is true and therefore, that (PAC) is false.

Consider the following example. Dave is a new graduate student in the department who hasn't yet met Larry, the department Chairman. One day before the Fall term has begun, Dave sees a middle-aged man in old clothes busily sweeping and dusting in the seminar room. Turning to another student, Dave says "That janitor is certainly hardworking." Overhearing this, I turn to a colleague and say

- (5) Dave thinks that Larry is a janitor.¹⁰

It seems intuitively correct to say that I would be using the name 'Larry' in (5) simply to refer to Larry and to say that Dave believes that he is a janitor. If the name 'Larry' as I use it had some sort of descriptive meaning in English, then perhaps it could be used to semantically convey something about the way Dave thinks of Larry. But in fact the name has no such meaning.

Although I think that there are ordinary proper names that have descriptive meanings in natural languages like English, I also believe, on the basis of the kind of evidence that Kripke (1972) produced in his famous Gödel-Schmidt case and other examples like it, that such names are quite rare and that most ordinary proper names have no descriptive meanings in the public language.¹¹ Thus it is clear, I think, that my use of the name 'Larry' in (5) would have no descriptive meaning of any sort. And so this use can serve only to refer to Larry and serves only to convey that Dave has a thought that is *about* Larry. This holds even when we assume that the relevant occurrence of the name 'Larry' in (5) is both grammatically and logically in the scope of 'thinks that'. Given this assumption, (5) is *structurally* (or logically) *de dicto*. But because of the semantic character of the small-scope name 'Larry', the structurally *de dicto* sentence (5) turns out to be semantically *de re*.

A similar point applies to cognitive contexts containing small-scope indexicals, such as

- (6) Dave thinks that *he* (or: *that man*) is a janitor.
 (7) Dave thinks that *you* are a janitor.
 (8) Dave thinks that *I* am a janitor.

Assuming that in each of (6)–(8) the speaker's utterance of the relevant indexical refers to Larry, then uses of (5)–(8) would all make the same *de re* assertion. They would all ascribe to Larry the same property, namely, the property that an object *x* has just in case Dave thinks that *x* is a janitor. There are various ways of accounting for this fact about such ascriptions as (5)–(8). One popular idea is that each of (5)–(8) says that Dave bears a certain mental relation, which is expressed by the verb 'thinks', to the singular proposition that Larry is a janitor. Another strategy, which I favor, simply takes each of these structurally *de dicto* ascriptions to be semantically equivalent to its structurally *de*

re counterpart, so that each says that Dave has a thought about Larry to the effect that he is a janitor.¹²

For my purposes here, it doesn't matter which of these two approaches we adopt, since the important fact is that on either approach, the sentences (5)-(8) are all *relational* with respect to the referent of 'Larry'. On the first approach, (5)-(8) ascribe to Dave and Larry the relation that an object *x* bears to an object *y* just in case *x* bears the relation of thinking to the singular proposition that *y* is a janitor. On the second approach, the relation ascribed is that which an object *x* bears to an object *y* just in case *x* has a thought about *y* to the effect that *y* is a janitor. On either approach, *de dicto*-structured cognitive predicates containing small-scope names and indexicals typically express properties that are relational with respect to the referents of those terms, just as the predicate 'thinks that Larry is a janitor' expresses a property that is relational with respect to Larry. Cognitive properties like this are wide because their possession by a person *logically*, not just metaphysically, implies the existence of the contingent object with respect to which the property is relational. Thus, the proposition that Dave thinks that Larry is a janitor logically implies that Larry exists.

Cognitive predicates containing proper names and indexicals provide the clearest basis for semantic externalism, and these predicates express wide psychological properties precisely in the sense that their possession by a person logically implies the existence of objects external to the person. A similar result is harder to establish in the case of cognitive predicates containing natural kind terms, but I believe that I have made a strong case for such a view in my own work, where I have provided both an explanation of precisely why such predicates express wide mental properties and an account of what these properties are. (See McKinsey, 1987 and 1994, pp. 321–324.) It is a consequence of my account that cognitive predicates containing natural kind terms also express properties that are relational with respect to external, contingent objects.

The fact that many cognitive predicates of the form 'is thinking that *p*' express logically relational properties shows that (SE) is true and therefore, that (PAC) is false. In particular, as we've seen, a predicate such as 'is thinking that Larry is a janitor' expresses a property that is relational with respect to the referent of the imbedded name 'Larry', and so it provides a straightforward counterexample to the claim made by (PAC) that necessarily, any person who is thinking that *p* can know a priori that he or she is thinking that *p*. In my example, it is true that Dave is thinking that Larry is a janitor; but of course it is false that Dave can know a priori that he's thinking that Larry is a janitor, since if he could, he could know a priori that Larry exists, which is absurd. And in general, one does not have privileged access to one's possession of logically wide mental properties.

4. *The externalism of McLaughlin and Tye.*

McLaughlin and Tye (1998, p. 292) have recently claimed that no one has ever held that we have privileged access to whether our thoughts are about

external objects, or to whether our thoughts have singular propositions about such objects as their contents. Perhaps this is so. But nevertheless, many philosophers, including Burge (1988), and McLaughlin and Tye themselves (1998, p. 289), have quite explicitly endorsed (PAC), and (PAC), given the semantic facts that I've mentioned about many cognitive predicates of the form 'is thinking that p', implies the absurd consequence that persons sometimes do have privileged access to such logically wide facts about their own thoughts.

To avoid commitment to such an absurd consequence, externalists like Burge, McLaughlin, and Tye who endorse (PAC) must of course endorse some form of externalism other than (SE). But such externalists, in being committed to (PAC), also have to assert that (SE) is *false*, even though (SE) again is both the only form of semantic externalism that is either interesting or controversial, and the only form of externalism for which any clear semantic evidence exists. It must be admitted that philosophers like this who explicitly commit themselves to the denial of (SE) by endorsing (PAC), and who also call themselves 'externalists', are a strange breed of "externalist" indeed.

In particular, externalists who endorse (PAC) thereby commit themselves to the following negative semantic thesis:

- (9) No *de dicto*-structured predicate of the form 'is thinking that p' ever expresses a property that is relational with respect to any ordinary contingent object.

There have no doubt been many philosophers of a Fregean persuasion who have consistently endorsed both (PAC) and (9). What is perhaps surprising is to find contemporary defenders of externalism, like McLaughlin and Tye, who are also willing to endorse (9), in spite of the enormous wealth of evidence against it that occurs in the literature over the past twenty years or so, some of which I mentioned above.¹³

McLaughlin and Tye do not really say anything in defense of (9), nor do they discuss any of the evidence that has been mustered against it. They do, however, consider cognitive predicates containing proper names and they state a view on which such predicates are not relational. They consider cases like

- (10) Jones is thinking that Cicero is an orator,

and say that a sentence like this would ascribe a thought whose content is the singular proposition that Cicero is an orator (provided Cicero really exists). But they also insist that, while no one could know a priori that his or her thought has such a singular proposition as its content, anyone who has such a thought would be able to know a priori that he or she is thinking that Cicero is an orator (1998, pp. 291–292, and 299). McLaughlin and Tye thus commit themselves to the view that the predicate 'is thinking that Cicero is an orator' does not express a property that is relational with respect to Cicero. They give no grounds for this view; apparently they hold it simply because they endorse (PAC).

But if the predicate ‘is thinking that Cicero is an orator’ does not express a relational property, then the name ‘Cicero’ does not contribute its referent to the proposition expressed by (10). What, then, does it contribute? Apparently, McLaughlin and Tye think that (10) would not describe Jones’s thought as being about the man Cicero, but rather would describe the thought in terms of some “mode of presentation” of Cicero (p. 294). But as I argued above in the similar example of sentence (5) (‘Dave thinks that Larry is a janitor’), this is surely a false view about the semantics of (10). For the name ‘Cicero’ could not convey anything about a mode of presentation under which Jones thinks of Cicero unless this name expressed such a mode of presentation in English, and the name could not do this unless it had some kind of descriptive meaning in English. But again, examples like Kripke’s Gödel-Schmidt case show that most ordinary names, including ‘Cicero’, have *no* such descriptive meanings.¹⁴

But we don’t need to rely on cognitive predicates containing proper names. For to refute the general negative claim (9) as well as (PAC), all we need is *one* case of a cognitive predicate of the form ‘is thinking that p’ which expresses a relational property. And even if McLaughlin and Tye were right about cognitive predicates containing names, there are many other kinds of cognitive predicates that even more obviously express relational properties. Consider, for instance, uses of cognitive predicates that contain small-scope demonstratives, such as

(11) Dave is thinking that *that man* is a janitor,

where the use of ‘that man’ refers to Larry, as before. Surely, there can be no doubt that such a use of (11) would ascribe a cognitive property that is relational with respect to Larry. A speaker of (11) would be using the demonstrative ‘that man’ solely to refer to Larry, and certainly would not be trying to semantically convey by use of this term any mode of presentation or way that Dave might have of thinking of Larry. Thus the semantic features of sentences like (11) suffice to show that both (9) and (PAC) are false.¹⁵

5. The epistemic status of externalist theses.

An important question that has been raised about my incompatibilist argument, a question that is closely related to the issue of the logical status of the connection between wide properties and external objects, concerns the *epistemological* status of the externalist’s theses regarding such connections. Both defenders of my argument (such as Brown, 1995 and Boghossian, 1997) and critics of the argument (such as Gallois and O’Leary-Hawthorne, 1996 and McLaughlin and Tye, 1998) have interpreted the argument as requiring the assumption that the externalist theses in question must be knowable a priori.¹⁶ The argument’s defenders then try to defend, while its critics attack, the apriority of the relevant externalist theses. In fact, however, my argument requires no particular assumptions to be made about the apriority of the externalist’s

dependency theses. The argument requires that the relevant externalist theses be claims *about* the existence of logical or conceptual relations; but it does not require that these claims themselves should be knowable a priori. And in fact, in my view, the externalist claims in question at least typically are *not* knowable a priori.

Consider the wideness of the property of thinking that water is wet, for instance. My argument for incompatibilism derives an absurd consequence from the conjunction of (1) and (2):

- (1) Oscar can know a priori that he is thinking that water is wet.
- (2) The proposition that Oscar is thinking that water is wet logically implies the proposition E.

(Where again, E is some a posteriori “external” proposition.) The absurd consequence of (1) and (2) is that Oscar can know a priori that E is true. But in deriving this consequence, my argument does not assume that the externalist thesis (2) is knowable a priori. Nor does the argument assume that the conditional corresponding to (2) is knowable a priori; that is, the argument also does not assume that

- (12) Oscar can know a priori that if he is thinking that water is wet, then E.

Rather, the argument assumes only that the 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.

Given (CA), (1) and (2) immediately imply the absurdity that Oscar can know a priori that E. No further assumptions such as (12) are necessary for the derivation.

But the most common interpretations of my argument don’t see it as relying on (CA). Rather, the externalist thesis (2) is taken to imply (12), to which a different closure principle is then applied to yield the absurd consequence that Oscar can know E a priori.¹⁷ This different closure principle might be called

Closure under A Priori Knowable Implication (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.

Now I agree that (CAK) is a perfectly fine closure principle for apriority. And certainly, it seems that (CAK) ought to be relevant to my argument. But the

difficulty is that to use (CAK) to derive an absurdity from (1) and (2), we must somehow get (12) as a line of the inference, and this is not an easy thing to do.

It might seem plausible to suppose that since (2) is a claim about the holding of a logical relation between two propositions, and since such claims are typically knowable a priori, (2) should itself be knowable a priori, and hence (12) would be true as well, if (2) is.¹⁸ But this line of reasoning is defective. For externalist claims like (2) constitute an important class of exceptions to the assumption that logical relations are knowable a priori. Very roughly, the reason is that to know that an implication like (2) holds, one would have to know that the property of thinking that water is wet is relational with respect to some type of external object. But knowledge of this latter sort requires empirical assumptions about the existence of the relevant external objects, and these would be assumptions that one cannot know a priori.¹⁹

The clearest examples of this kind of fact are provided by such *de re* properties as that of thinking that Larry is a janitor. Since this property is relational with respect to Larry, we have

- (13) The proposition that Dave is thinking that Larry is a janitor logically implies the proposition that Larry exists.

Now (13) is *not* knowable a priori. For even though (13) is a meta-proposition about the logical implication of one proposition by another, it is also a proposition that is singular with respect to Larry, and it ascribes a complex logical property to Larry. This is the property that any object *x* has if and only if the proposition that Dave is thinking that *x* is a janitor logically implies the proposition that *x* exists. So the meta-proposition expressed by (13) is itself a singular proposition that does not exist unless Larry does, and hence its truth cannot be known a priori. Similarly, the conditional proposition that if Dave is thinking that Larry is a janitor then Larry exists, is also singular with respect to Larry, and so it too cannot be known a priori. So it is false that

- (14) Dave can know a priori that if Dave is thinking that Larry is a janitor, then Larry exists.

Since (13) is true and (14) is false, (13) does not imply (14). For the same reason, I would say that the externalist assumption (2) of my argument does not imply (12). But again, unless (12) can be derived as a line of the inference, the plausible closure principle (CAK) cannot be applied to derive the absurd consequence that Oscar can know *E* a priori. This is one reason why I did not appeal to (CAK) in my original argument, and instead used the more directly applicable closure principle (CA).

Now I think that the closure principle (CA) is clearly correct, so that my original argument is fine as it stands. Nevertheless, I have come to believe, as so many interpreters of my argument have assumed, that (CAK) probably can be used to a similar effect. (CAK) can be applied by first deriving (12), not

from the externalist thesis (2) alone, but from (2) plus the *compatibilist's* assumption (1). For a plausible case can be made that (12) does indeed follow from the conjunction of (1) and (2) by way of still another closure principle, whose significance has so far been overlooked. We might call this principle

Partial Closure under Logical Implication (PCL)

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 if *P* then *Q*.

The role played in my original argument by the principle (CA) can be played without loss by the two other closure principles (CAK) and (PCL), which together entail (CA). By appeal to (PCL) we can derive (12) from (1) and (2). Then by appeal to (CAK), we can derive from (1) and (12) the absurd consequence that Oscar can know *E* a priori. This new argument for incompatibilism, notice, also reduces the conjunction of (1) and (2) to absurdity without having to assume any falsehoods to the effect that (2) is itself knowable a priori, or that (12) is true, or that (2) implies (12). Rather the argument simply assumes (1) and (2) for *reductio*, and then derives an absurdity from these assumptions using only the closure principles (PCL) and (CAK) as premises.

It remains to explain why I believe that (PCL) is plausible. We've seen that a meta-proposition of the form 'P logically implies Q' can be true, even though the corresponding conditional 'If P then Q' cannot be known a priori. This happens precisely when 'P logically implies Q' is true for *externalist reasons*. Nevertheless, even in such cases, if the additional assumption is made that a given person could know a priori that *P*, then I believe it would follow (contrary to fact) that it *could* be known a priori that if *P* then *Q*.

Consider again the case of

- (13) The proposition that Dave is thinking that Larry is a janitor logically implies the proposition that Larry exists.

The reason why (13) can't be known a priori is that the very existence of the proposition expressed by (13) is itself contingent upon Larry's existence. And for the same reason, the conditional corresponding to (13) also cannot be known a priori:

- (15) If Dave is thinking that Larry is a janitor, then Larry exists.

But now suppose (contrary to fact) that

- (16) Dave can know a priori that he is thinking that Larry is a janitor.

If (16) were true, then Dave *could* know a priori that the proposition that he is thinking that Larry is a janitor, being true, must exist. But then, by logic alone,

Dave can see that the proposition in question logically implies the proposition that Larry exists, and so Dave can know a priori that the conditional (15) is true. In this case, notice that the closure principle (PCL) intuitively holds, *even though the relevant instance of the variable is a singular proposition*.

To put the point another way, when a proposition P logically implies a proposition Q, there is only one obstacle to one's knowing a priori that if P then Q, and that is one's inability to know a priori that P's existence-presuppositions, if any, are satisfied. (We can safely ignore Q's existence-presuppositions, since they are all shared by P, if P logically implies Q.) But if we go on to assume that one can know a priori that P, then it follows that one can know a priori that all of P's existence-presuppositions, whatever they are, are satisfied, and so nothing stands in the way of knowing a priori that if P then Q. For this reason, (PCL) holds even in extreme cases where the instance of P is a singular proposition, and so it is plausible to think that (PCL) holds in general.

So my *reductio* argument can be stated in two ways: the original way, in which the closure principle (CA) is the only premise required to deduce the relevant absurdity; and a second way, in which the role of (CA) is played by the similarly plausible closure principles (CAK) and (PCL). But I should emphasize that neither version of the *reductio* can be applied unless the externalist's dependency theses are assumed to be claims about the existence of logical relations between cognitive properties and external objects. Thus my "trivialization" argument against the "metaphysical" externalist is required before either *reductio* argument can get off the ground.²⁰

And finally, it is also important to emphasize that no form of incompatibilist argument is going to work if the argument itself assumes that the externalist's theses are knowable a priori. For the theses in question are inevitably based in part on common sense assumptions about the existence of such external things as Cicero and water. But again, neither form of *reductio* argument that I've proposed needs to assume that the externalist's theses are knowable a priori. The arguments only need to assume that the externalist's theses are claims about the existence of logical relations between cognitive properties and external objects.

6. Transmission of warrant and closure of apriority.

Martin Davies (1998) has raised an interesting possible line of response to my argument for incompatibilism. Using considerations introduced by Fred Dretske (1970) and Crispin Wright (1985), Davies points out that a person's epistemic warrant for believing a given premise P will not necessarily transmit to a given deductive consequence Q of P, even when the person knows that P logically implies Q. In particular, there are cases where P logically implies Q, but one's *already* having warrant for Q is a crucial precondition of one's having warrant for P. In such cases, one's warrant for Q cannot derive from one's warrant for P, even if one knows that P logically implies Q. Davies describes many such cases, but for my purposes here, the most important type of case is

that in which an ascription of a wide cognitive property logically implies, without providing warrant for, some external proposition. For instance, we've seen that the proposition expressed by (5) ('Dave thinks that Larry is a janitor') logically implies that Larry exists. But of course, the *source* of one's warrant for believing that Larry exists could not be that one has correctly deduced this conclusion from the relational premise in question, for the simple reason that one would not be warranted in believing the relational premise in the first place, unless one were *already* warranted in believing that Larry 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 does this fact, as Davies suggests, show that my argument is unsound? I don't think so. For the argument does not assume that warrant is always transmitted from a premise to the known deductive consequences of that premise. Rather, again, the argument assumes a certain principle of closure for apriority, which I repeat:

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*.

When my argument for incompatibilism is applied to a particular relational cognitive premise like (5) ('Dave thinks that Larry is a janitor'), we have to assume (as part of our assumption for *reductio*) that the premise (5) can be known a priori (by Dave). In such cases, of course, the relevant assumption is *contrary to fact*, as the argument for incompatibilism itself shows, given that (5) really does logically imply the relevant externalist conclusion (in this case, that Larry exists). When applied to our assumption for *reductio*, (CA) generates the absurd consequence that Dave can know a priori that Larry exists. But this is surely correct: if (5) were knowable a priori, then it would be knowable without empirical investigation, and not on the basis of any empirical assumptions. Hence, (5)'s deductive consequences would surely also be knowable a priori, even when warrant for these consequences is presupposed by the warrant for (5). For by our (contrary to fact) assumption for *reductio*, any assumptions on which warrant for (5) depends would not be empirical.

Hence the principle of closure for apriority on which my argument depends is perfectly consistent with the facts about transmission of warrant that Davies emphasizes. The general reason why closure of apriority is consistent with failure of warrant-transmission is that there is more than one type of explanation as to why, given that a person can have a priori knowledge of a certain premise *P*, it follows that the person can also have a priori knowledge of a given logical consequence *Q* of *P*. When warrant successfully transmits from *P* to *Q*, we can use this fact to explain how a given person could know *Q* a priori, namely, by knowing *P* a priori and then correctly deducing *Q* from *P*. When warrant *fails* to transmit from *P* to *Q*, we of course no longer have *this* way of explaining why *Q* must be knowable a priori if *P* is. But this does not show that

closure of apriority must fail in such cases, since there may be another way of explaining why closure of apriority succeeds, an explanation that does not have to assume transmission of warrant. In short, transmission of warrant is a sufficient, but not a necessary, condition for closure of apriority.

The typical reason why a person's warrant for a given premise P does not transmit to a deductive consequence Q of P, is that the agent's having warrant for P *depends on* the agent's already having warrant for Q.²¹ But in all such cases, like the case of (5) above, closure of apriority continues to hold, despite failure of warrant-transmission. For if the agent knows a priori that P, then the agent's knowledge that P is not based on empirical investigation. But then it follows that the agent can know a priori that Q. For otherwise, the agent would have to know Q by empirical investigation, and then since by assumption the agent's having warrant for P *depends on* the agent's having warrant for Q, the agent's knowledge that P would itself depend on empirical investigation, contrary to our assumption that the agent knows a priori that P.

When transmission of warrant succeeds, closure of apriority obviously holds. And when transmission of warrant fails, this is typically because the agent's warrant for a given premise depends on the agent's warrant for the relevant deductive consequence. But in all such cases, as we've just seen, closure of apriority *also* holds. Thus, the fact that warrant sometimes fails to transmit from a given premise to a known deductive consequence of that premise, does not provide any good reason to doubt that closure of apriority holds in general.²²

7. Externalist principles of individuation.

We have seen that some *de dicto*-structured predicates of the form 'is thinking that p' express logically wide properties, and that we in general have no privileged access to our possession of such properties. Since the principle (PAC) implies otherwise, it is an incorrect formulation of the idea that we in general have a privileged way of knowing about our thoughts and their contents. I have suggested elsewhere (McKinsey, 1994) that the correct formulation of this idea restricts the properties of a thought to which one has privileged access to those fundamental semantic properties that make the thought the thought that it is. These are the properties that *individuate* the thought, where a thought that a person x has in a given possible world w is individuated by a property P just in case any person y would have the very same thought in any other possible world w' if and only if in w' y also has a thought that has P. (See McKinsey, 1994, p. 305.)

Then our revised principle of privileged access, stated in terms of individuating properties, is the following:

Privileged Access to Individuating Properties (PAI)

It is necessarily true that if a person's thought is individuated by a given property Φ , then that person can in principle come to know a priori that he or she has a thought that has the property Φ .

We've seen that in general, one never has privileged access to one's possession of logically wide properties. Thus (PAI) implies an important metaphysical principle concerning the nature of the properties that individuate our thoughts, for it implies that all such properties must be logically *narrow*. I will call this principle

Metaphysical Internalism (MI)

It is necessarily true that if a person's thought is individuated by a given property Φ , then Φ is logically narrow.

I endorse both (PAI) and (MI).

I've also said that I endorse semantic externalism, the thesis (SE) that many *de dicto*-structured predicates of the form 'is thinking that p' express logically wide properties. This thesis is of course consistent both with (MI), which says that our thoughts are *individuated* by logically narrow properties, and with (PAI), which says that we have privileged access to the properties that individuate our thoughts. The reason is simply that, being solely a *semantic* thesis, semantic externalism says only that many cognitive predicates express wide properties; it is silent on the metaphysical question of whether or not these properties ever *individuate* our thoughts.

Those who do as I do and restrict their externalism to the semantics of cognitive predicates are thus free to endorse the idea that we have privileged access to the fundamental semantic properties of our thoughts. But most defenders of externalism, it seems to me, have in addition wanted to endorse externalism as a *metaphysical* view about the *nature* of thought and other cognitive attitudes. We can, I suggest, most plausibly take these externalists to be claiming that certain thoughts are individuated, in the sense I've defined, by their wide contents, or by the wide property of having such a content. (See McKinsey, 1994.) We may call this view

Metaphysical Externalism (ME)

In some cases, a person is thinking that p, the content that p is logically wide, and the person's thought is individuated by the property of being a thought that has the content that p.²³

(By a "logically wide" content, I mean an abstract semantic entity, like a singular proposition, whose very existence logically implies the existence of contingent, typically concrete, objects.)

Now it is important to see that (ME) is not directly supported by the evidence that supports semantic externalism. This evidence shows only that in many predicates of the form 'is thinking that p', namely, predicates containing names, indexicals, and natural kind terms, the imbedded sentence p expresses a wide content, specifically a proposition that is singular with respect to the referent of the relevant term. Since surely, many ascriptions containing such predicates are true, it is no doubt also true that the thoughts ascribed have the wide contents

expressed by the relevant imbedded sentences. But it does not follow that any of these wide contents ever serve to *individuate* the thoughts in question, as (ME) asserts. To derive (ME), we must assume the following additional principle about thoughts and their contents:

Individuation by Content (IC)

It is necessarily true that if a person is thinking that p, then this thought is individuated by the property of being a thought that has the content that p.

The semantic facts that support semantic externalism, together with (IC), imply metaphysical externalism (ME). Since (IC) is a plausible principle, metaphysical externalism is a well-motivated view. But now we have a problem. For of course metaphysical externalism (ME) is inconsistent with the principle (PAI) that we have privileged access to the properties that individuate our thoughts. This form of incompatibilism is different from the one we began with between (PAC) and (SE), but it has the same logical basis. The original argument for incompatibilism shows that we can have no privileged access to our possession of logically wide properties. This fact plus (PAI) implies that our thoughts are necessarily never individuated by logically wide properties, contrary to what (ME) says.

Thus we have a choice. We can adhere to (PAI) and to the idea that we have a privileged way of knowing about the properties that individuate our thoughts; in this case we must reject the plausible principle (IC), and say that when the content that p is logically wide, one's thought that p is not individuated by the content that p. Or we can adhere to (IC) and the idea that our thoughts are necessarily individuated by their contents; but if we do this, we must end by rejecting privileged access in all its forms. This is the choice to which we are forced by the undeniable semantic facts that support direct reference theories of names, indexicals, and natural kind terms. For several years, I have been advocating the strategy of adhering to (PAI) and rejecting (IC), and I believe that I have provided enough compelling semantic evidence to justify this strategy. (See McKinsey, 1986, 1991b, 1994, and 1999.) But my main point here is just that my original argument for incompatibilism shows that we *must* make this choice. We cannot consistently assert both that metaphysical externalism is true and that we have privileged access to the properties that individuate our thoughts.

8. Burge's conceptual externalism.

Burge (1979 and 1982) has endorsed a form of externalism that may go beyond mere semantic externalism, and yet which is weaker than, or at least different from, metaphysical externalism as I've described it. At one point (1982, p.97), Burge seems to deny that thoughts would be individuated by *de re*, relational, cognitive properties.²⁴ So perhaps he would agree with my view that thoughts are not individuated by their wide contents, at least in the cases in

which the wide contents are singular propositions involving ordinary contingent objects. Even so, Burge clearly wishes to hold that there are *de dicto*-structured cognitive predicates that express properties that are both individuating and wide, in some important sense. If these properties are not logically relational, then perhaps Burge holds that they are both wide and at the same time purely conceptual. If so, then he endorses a view that we could call ‘conceptual externalism’:

Conceptual Externalism (CE)

In some cases, a person’s thought that *p* is individuated by the property of being a thought that *p*, where this property is purely conceptual and not relational with respect to any contingent object, and yet this property is wide in some important sense.

Let us briefly consider the question of whether Burge’s conceptual externalism is consistent with the principle (PAI) that we have privileged access to the properties that individuate our thoughts.

The main source of support that Burge gives for his view is his well known thought experiment concerning arthritis (Burge, 1979).²⁵ Burge marshals intuitions to support the conclusion that a person with an inadequate grasp of the concept of arthritis could nevertheless have thoughts and beliefs that involve this concept. But he then argues that such a person’s having such thoughts and beliefs would depend on his belonging to a community whose language contains a word that means *arthritis*: had the word ‘arthritis’ had a different meaning in the agent’s language, Burge argues, the agent’s thoughts and beliefs would not have involved the concept of arthritis at all, even if the agent were internally just the same. Burge concludes that the contents of the agent’s thoughts and beliefs are dependent upon his social and linguistic environment. (Burge, 1979, pp. 77–79.)

It is important to see that Burge’s example does not show that cognitive predicates containing the word ‘arthritis’ express properties that are either logically or metaphysically wide. The reason is simple. Suppose that a person *X* has the property of thinking that arthritis is painful. Then if Burge is right, *X*’s having this property will depend on his belonging to a certain sort of linguistic community, *given that X’s grasp of the concept of arthritis is inadequate*. Thus the example cannot show that having the property of thinking that arthritis is painful either logically or metaphysically entails the existence of external objects, such as other speakers. For at most it can show only that having the *conjunction* of this property with the property of having an inadequate grasp of the concept of arthritis, entails the existence of other speakers. (See McKinsey, 1993, pp. 326–327.)

Since Burge’s example fails to show that any cognitive predicate of the form ‘is thinking that *p*’ expresses a property that is either logically or metaphysically wide, the example also fails to support a form of conceptual externalism wherein a cognitive property’s “wideness” is understood in either of

these two senses. However, if Burge's intuitions about his example are correct, then, as Burge points out, the example does support an interesting conclusion, namely that some cognitive properties, involving purely conceptual contents, do not *supervene* upon a person's intrinsic physical properties.²⁶ Let us call a mental property that thus fails to supervene upon a person's intrinsic physical properties an "S-wide" property. Then Burge's conceptual externalism might just be the view that in some cases, the property of being a thought that *p* is not only purely conceptual and individuating, but also S-wide.²⁷

The difficulty with this interpretation is that there are many different types of property that are S-wide, and hence many different possible explanations as to why a given property, such as that of thinking that arthritis is painful, might be S-wide. Thus it becomes very difficult to see, in any clear way, how a form of externalism that invokes the concept of S-wideness, is supposed to be logically related to any traditional principles in the philosophy of mind, such as that of privileged access or that of metaphysical internalism.

Consider *de se* thoughts, for instance. When we say that Descartes is thinking that he himself is a thinking thing, we are ascribing to Descartes the property *P* that an object *x* has just in case *x* is thinking about Descartes from a first person perspective that he is a thinking thing. Now of course only Descartes can have property *P*. In particular, Descartes' molecular duplicate on Twin Earth cannot have *P*. Hence *P* is S-wide. Yet surely, *P* should count as logically *narrow*: Descartes' having this property would logically imply the existence of no object external to himself.²⁸ Moreover, of course, anyone who has this property (that is, anyone who is identical with Descartes) would have privileged access to their having it.

Or consider the *disjunction* of any logically narrow mental property *N* with a mental property *W* that is both logically wide and S-wide. Then the property *N or W* is logically narrow, since having it does not logically imply the existence of external objects (one could have *N or W* by having *N*). But this same property is *S-wide*: since *W* is S-wide, *x* and *y* could be molecular duplicates such that *x* has *W* and *y* does not; suppose that neither *x* nor *y* has *N*; then *x* has the property *N or W* but *y* does not. Note also that this logically narrow, S-wide property is one that does not satisfy a principle of privileged access: persons who have *N or W* because they are *W* but not *N*, would not be able to know a priori that they have *N or W*.

So S-wide properties are a mixed bag. Some are logically wide properties that fail to satisfy a principle of privileged access. But others are logically narrow properties that do satisfy a principle of privileged access, while still others are logically narrow properties that fail to satisfy such a principle.²⁹ This means that a principle of conceptual externalism that uses the logically inchoate notion of S-wideness is strictly consistent with both the principle of privileged access (PAI) and metaphysical internalism (MI). That's just because some S-wide properties are logically narrow and some S-wide properties satisfy a principle of privileged access. On the other hand, since this form of conceptual externalism is consistent with both (PAI) and (MI) and fails to imply (ME), it also

turns out that this view has no interesting logical implications regarding any traditionally held principles in the philosophy of mind, contrary to what its defenders might have intended.

Still, one suspects that a clear semantic explanation, consistent with Burge's intuitions, of precisely *why* a predicate like 'is thinking that arthritis is painful' expresses an S-wide property, might in fact have significant implications for traditional principles. But neither Burge nor any of his many defenders has even hinted at what such an explanation might look like. I have myself suggested that the explanation might well be that the properties expressed by the cognitive predicates in question are *disjunctive* in form (see McKinsey, 1993, pp. 330–332). After all, on Burge's intuitions there are two radically different conditions each of which is sufficient for satisfaction of any (purely conceptual) cognitive predicate: one could satisfy it by completely understanding the concepts involved, or one can lack this complete understanding but make up for it by membership in a community whose language contains words that express the concepts. So Burge's intuitions suggest the hypothesis that the properties expressed by the relevant predicates are S-wide because they are disjunctions of logically narrow and wide properties. Since as we saw earlier, such disjunctive properties are logically narrow but violate privileged access, Burge's claim that such properties are individuating would be consistent with metaphysical internalism (MI) but inconsistent with the principle (PAI) that we necessarily have privileged access to individuating properties.

In this paper, I've defended and explained my original argument for the incompatibility of semantic externalism (SE) and the principle (PAC) that we necessarily have privileged access to the contents of our thoughts. In particular, I've defended the argument's use of conceptual implication in its formulation of externalism, explained why the argument does not need to assume that externalist theses are knowable a priori, and defended the closure principle (CA) about apriority that the argument requires. On the basis of my argument for incompatibilism, I also argued further that since semantic externalism is true, (PAC) is false. I then proposed a restricted principle (PAI) according to which we have privileged access only to the properties that individuate our thoughts. The same considerations I appealed to in my first argument for incompatibilism then show that (PAI) logically implies that metaphysical internalism (MI) is true and that metaphysical externalism (ME) is false. Finally, I considered Burge's conceptual externalism and argued that the evidence Burge has produced for his view supports only a vague principle whose logical relations to such principles as (PAI) and (MI) are obscure. I suggested that if Burge's evidence were clearly explained, the result might well be a view that is consistent with (MI) but inconsistent with (PAI).³⁰

Notes

1. A word about the history of the argument might be in order. The paper (1991a) in which the argument appeared was first presented to the Pacific Division of the Amer-

- ican Philosophical Association (Los Angeles, March, 1990). I used the phenomenon of privileged access to argue against causal theories of cognitive predicates containing proper names in my doctoral dissertation (Indiana University, 1976) and later in my (1978). However, I quickly realized that my original argument for the incompatibility of causal theories and privileged access depended upon an unacceptably strong closure principle, and in a 1978 seminar on the philosophy of language at Wayne State University, I revised the argument to its present form, so that it depends on what I still take to be an acceptably weak closure principle for a priori knowledge (see section 5 below). After that, I regularly taught the argument in classes and seminars at Wayne State on the philosophy of language and mind, and I published an argument of the same structure, but with a different application, in my (1984), p. 501. I also gave the argument in a question from the floor at the April, 1985 Oberlin Colloquium after Tyler Burge's presentation of his paper (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. This principle is quite similar to the thesis of privileged access discussed and endorsed by McLaughlin and Tye (1998, p. 286).
 3. This characterization of "wideness" is vague, though the vagueness should do no harm in the present context. It is actually surprisingly difficult to achieve clear explications of the concepts of "wide" and "narrow" properties. (A property is narrow iff it is not wide.) For an idea of what the difficulties are, the reader should consult my discussion in "The Internal Basis of Meaning" (1991b, pp. 148–155). The definition (D2) on p. 155 of that paper provides, I believe, an adequate clarification of the concept of narrowness and by negation, of the concept of wideness as well. For some improvement over (SE), see the definition (3) below, and my discussion of (3) in note 8 below.
 4. I should note that applications of this form of *reductio* argument to show that specific externalist theses about specific cognitive predicates of the form 'is thinking that p' are inconsistent with (PAC), do not by themselves show that the general thesis (SE) is inconsistent with (PAC). This is because the consequences of conjoining such specific externalist theses with (PAC) are typically stronger, and hence more obviously absurd, than the weaker, perhaps less obviously absurd consequences that result from conjoining (PAC) and (SE). For instance, if we suppose that in the case of 'is thinking that water is wet', the external proposition E is the proposition that water exists, then the conjunction of (1) and (2) implies that Oscar can know a priori that water exists, which is perhaps more obviously absurd than the weaker consequence that one can know a priori that there exist contingent objects external to oneself, which is all that follows from conjoining (PAC) and (SE). I myself find this weaker consequence to also be obviously absurd, and so I take my *reductio* argument to show the inconsistency of (PAC) and (SE). But others may not agree, and may find this weaker consequence to be less than absurd. Even so, it makes little practical difference, since even if my *reductio* argument does not show the inconsistency of (PAC) and (SE), it clearly does show at least that all the specific externalist theses that provide the sole *evidence* for (SE) are inconsistent with (PAC).
 5. For this reason, McLaughlin and Tye (1998, p. 314) are off target when they remark that, "as has been correctly pointed out by Brueckner (1992)," I offered "no textual evidence" for my "charge" that Burge's endorsement of (PAC) is inconsistent with his endorsement of externalism. First, my argument that (PAC) and (SE) are incon-

sistent requires no “textual evidence” from Burge. Second, I never charged Burge with any inconsistency in my (1991a). Rather, I emphasized (on the basis of textual evidence) that Burge endorses a form of externalism other than (SE) and I pointed out that this form of externalism is *consistent* with (PAC) (1991a, pp.12–13). I then argued that Burge’s form of externalism, which unlike (SE) is understood in terms of metaphysical rather than logical implication, results in a trivial, uninteresting view (see section 2 below). Thus, my “charge” against Burge was that he avoids inconsistency with (PAC) only by endorsing a trivial, uninteresting form of semantic externalism. (See McKinsey 1994a, which is a reply to Brueckner 1992.)

6. Brueckner (1992) appears to make *both* of these suggestions, the first on p. 116, and the second on pp. 113 and 114.
7. The following discussion, with some changes, is based on the more detailed discussion in my paper “The Internal Basis of Meaning” (1991b, pp.148–155), as was the briefer argument in my (1991a).
8. My formulation of this definition was influenced by Kirk Ludwig’s statement of externalism (1993, p. 251), though Ludwig, correctly in my view, explicitly states externalism in terms of conceptual or ‘broadly logical’ necessity. But with ‘metaphysically entails’ replaced by ‘logically implies’ in (3), the resulting definition (call it (3*)) is not technically adequate, since its *definiens* provides a sufficient but not a necessary condition for wideness. Consider *de re* mental properties such as the property *thinking of Larry that he’s a professor*. Surely, *de re* properties like this should count as “wide” mental properties, if any properties do. But such properties don’t satisfy the *definiens* of (3*). For *Larry* himself could have the property of thinking of Larry that he’s a professor, and *Larry’s* having this property of course does not logically imply that there exists some contingent object *y* such that *y* is not identical to *Larry*, *Larry’s* parts, or any of *Larry’s* mental states, acts, and experiences. To correct this defect in (3*)—and in Ludwig’s statement of externalism—the following should be added as a disjunct to (3*)’s *definiens*: ‘or there exists an object *y* such that necessarily, for any person *x*, *x*’s being *S* logically implies that *y* exists, and it’s possible for a person *z* to have *S*, even though *y* is not identical to *z* or to any of *z*’s mental states, acts, and experiences’. For the purposes of the present paper, the reader can take the resulting definition of wideness as the one I’d recommend over (3), though my absolute preference would be the negation of the sense of ‘narrow’ given in definition (D2), p. 155 of McKinsey, (1991b). See note 3 above.
9. It has been suggested to me by several philosophers (including Bruce Russell and Tom Stoneham) that a defender of metaphysical externalism could avoid my trivialization argument by somehow restricting the relevant entailments of a “wide” property *S* to those for which only the property *S* is responsible, as opposed to those unwanted entailments which follow from the existence of a person who has *S*. Such a restriction would of course solve the problem. But notice first that my proposal to understand a property’s wideness in terms of logical or conceptual implication has precisely this desired effect, and so my proposal does in fact solve the problem I’ve raised.

Moreover, in my opinion, the problem can only be solved in the manner I’ve proposed. I’ve been unable to find any way of defining wideness in terms of metaphysical implication that both restricts the relevant entailments in the desired way and also avoids trivialization problems of the sort I’ve raised for definition (3). Suppose for instance that we revise (3) by adding a restriction to the effect that the relevant external condition (that the agent *x* is not alone in the universe) is not (meta-

physically) entailed by x 's existence alone. Then it follows from this revised definition that a property S is wide only if, for any person x that has S , it's metaphysically possible for x to exist alone in the universe! So this definition has the undesirable result that there would probably be no *wide* properties. Other more complex attempts that I've tried have similar problems. So again, I suggest that to avoid these problems, we should simply abandon metaphysical implication, and understand the wideness or narrowness of a cognitive property in terms of logical, or conceptual, implication. (I give another, more directly theoretical argument against "metaphysical" externalism in McKinsey, 2001.)

10. I first heard the story of the professor and the janitor many years ago from John Tienson.
11. I argue for this thesis in several places. See McKinsey (1984), (1994), and (1999).
12. For examples of the first view, see for instance McKay (1981), Salmon (1986), and Soames (1987). I argue in my (1999) paper "The Semantics of Belief Ascriptions" that these two ways of doing justice to the intuitive facts about (5)-(8) are not in fact equivalent, and I argue against the first approach and in favor of the second.
13. I have in mind particularly the work of defenders of the "Russellian" theory of belief mentioned in note 12 above, including McKay (1981), Salmon (1986), and Soames (1987). Also relevant is the work of Mark Richard (1983) and (1990). See also McKinsey (1994), (1998), and (1999).
14. I have argued elsewhere that, although they are quite rare, there are names that have descriptive meanings in the public language. Plausible examples might include 'Jack the Ripper', 'Hesperus', 'Phosphorus', 'King Arthur', 'Homer', 'God', and other names whose referents, if they exist at all, are epistemologically remote from all of us, so that we all must refer to them on the basis of the same narrow set of descriptive assumptions. I have also argued that when used in cognitive contexts, these descriptive names could contribute their descriptive meanings in such a way that the cognitive predicate ends up characterizing the agent's way of thinking about the referent (if any), even though the imbedded sentence expresses a singular proposition (when the name has a referent). (See McKinsey, 1986, 1994, and 1999.) Intuitively, a sentence like this, such as 'Jones thinks that God is omniscient' can be true whether or not the imbedded name has a semantic referent, and hence independently of whether or not the imbedded sentence even succeeds in expressing a proposition. Hence, the cognitive predicate contained in such a sentence does not express a property that is relational with respect to the name's referent, even when the sentence ascribes a thought whose content is a singular proposition. So my account of cognitive contexts containing *descriptive* names is quite similar to McLaughlin and Tye's account of contexts containing names generally. Of course, since descriptive names are quite rare, their account has no general application, and as I argue in the text, it is falsified by cases of cognitive predicates containing most ordinary names, which have no descriptive meanings.

Moreover, I seriously disagree with one important feature of their account, even as it applies to cognitive contexts containing descriptive names. They claim that if the name in question has a referent, so that the imbedded sentence does express a singular proposition, then the cognitive ascription is true in any possible world w only if the thought ascribed has that singular proposition as its content in w (1998, p. 299). I think that this claim is obviously false. If the cognitive ascription can be true even if the imbedded name has no referent, as they seem to be saying, then obviously the ascription can be true in a possible world in which the relevant prop-

osition fails to exist and thus is *not* a content of the thought. Contrary to McLaughlin and Tye, I take such examples to imply the important result that a thought can have a singular proposition as its wide content, even though a person could have exactly the same thought in a world in which that proposition fails to exist. In the terminology I introduce below in section 7, such a thought would have a wide content, but it would not be *individuated* by this content. (For a detailed discussion of this point, see McKinsey, 1994). Note that cognitive predicates containing descriptive names, on my account, express completely *narrow* cognitive properties, even when the imbedded sentence expresses a singular proposition. I give a more complete and detailed criticism of McLaughlin and Tye's view in McKinsey (2001).

15. Someone might object to this argument by claiming that a context containing a demonstrative like 'is thinking that that man is a janitor' would not really be of the relevant form 'is thinking that p', since a demonstrative like 'that man' would have to have *largest scope* relative to any cognitive operator. Thus, so the objection goes, the sentence 'that man is a janitor' would not lie entirely in the scope of 'is thinking that' in (11) and so (11) cannot be a counterexample to (PAC). It is important to see that this point is mistaken. It confuses a *semantic* feature of certain kinds of singular terms, namely, that they can only be used referentially, with a structural, or logical, feature of the contexts that contain them. To see the mistake, consider how we would rewrite (11) so as to give 'that man' explicitly largest scope:

(11a) That man is such that Dave is thinking that he is a janitor.

As Nathan Salmon (1986, p. 4) has pointed out, we succeed in giving the term 'that man' large scope in such a sentence only by tying it to the anaphoric pronoun 'he' which in turn must be assumed to have *logically small scope*. But notice that the pronoun 'he', even though it has small scope, is still used in a purely referential way in (11a), and as a result the predicate 'is thinking that he is a janitor' still expresses a property that is relational with respect to a certain object (in this case, Larry). In the end, the idea that purely referential occurrences of terms must always lie outside the scope of cognitive operators simply makes no sense. (See McKinsey, 1994, pp. 309–310, and 1998, pp. 17–18.)

16. As far as I know, the first to interpret my argument this way was Martin Davies, in the 1994 ancestor of his (1998), originally presented to the Eastern Division of the American Philosophical Association (Boston, December, 1994).
17. Again, see Brown (1995), Boghossian (1997), Gallois/O'Leary-Hawthorne (1996), McLaughlin/Tye (1998). See also Davies (1998) and Wright (2000).
18. See McLaughlin and Tye (1998), p. 290, where this assumption is explicitly made.
19. See my discussion of Boghossian (1997) in McKinsey (2002).
20. Again, I give another, more direct argument against the use of metaphysical implication to state the externalist's theses in McKinsey (forthcoming).
21. For further discussion of the dependency relation in question, see Wright (2000), Hale (2000), and McKinsey (forthcoming).
22. For a much fuller discussion of this issue, see McKinsey (forthcoming). Davies (2000, p. 407) has granted that my argument goes through, given the "strict sense" of 'apriori knowledge' in terms of which my argument was originally stated. But both he and Wright (2000) contend that due to failure of warrant-transmission, my argument cannot be used to show the incompatibility of externalism and a form of privileged access that is understood in terms of a weaker sense of 'a priori knowledge'. I respond to this contention in McKinsey (forthcoming).

23. Perhaps some might prefer a weaker form of metaphysical externalism on which thoughts are individuated by properties that are merely metaphysically wide. But again, since probably *all* cognitive properties are metaphysically wide, including all cognitive properties that are logically narrow, this view would be consistent with metaphysical internalism, and would not disagree with what any internalist would want to assert.
24. Martin Davies pointed out this feature of Burge's view to me. See also Davies (1998), p. 332.
25. Burge apparently believes that cognitive predicates containing natural kind terms such as 'water', 'aluminum', and 'tiger' also support his view. (See Burge, 1982). I agree that it is *prima facie* plausible to suppose that cognitive predicates containing natural kind terms could express properties that are both purely conceptual and yet wide. But to substantiate such a supposition requires a clear and justified account of what these properties are, and Burge provides no such account. He provides no semantic view of natural kind terms and no view of the meaning and logical form of cognitive predicates that contain such terms. He provides only some plausible constraints on such a view (in Burge, 1982). However, I have defended an account of cognitive predicates containing natural kind terms that meets Burge's constraints, and as I mentioned earlier, it is a consequence of my view that cognitive predicates containing natural kind terms express properties that are relational with respect to external, contingent objects (see McKinsey, 1987, 1994, and 1999). So if I am right, the correct semantic view of the cognitive predicates in question does not in fact provide any support for Burge's conceptual externalism.
26. I have argued that Burge's intuitions are not correct, so that his example does not really support even this relatively weak conclusion. See McKinsey, 1993, pp. 332–335.
27. Brueckner (1995) seems to be suggesting some such view as this on Burge's behalf.
28. The definition of logical wideness that I suggested in note 8 above entails that P is logically narrow. See also my discussion of the narrowness of *de se* properties in McKinsey (1991b), pp. 146–157.
29. Moreover, some *S-narrow* properties are logically *wide*. Consider, for instance, the property of having a thought that is caused by an event in one's own brain. Surely, this is a property that molecular duplicates would share in common, and so it is not *S-wide*. But it is logically wide, since having the property logically implies the existence of certain physical objects (i.e., brains).
30. Earlier versions of this paper were presented to a graduate seminar on Externalism and Self-Knowledge at Oxford University (February/March, 1999), a conference on Anti-Individualism, Self-Knowledge and Scepticism held at the University of Bristol (February, 1999), to a conference on Internalism/Externalism held at the University of Stockholm (April, 1999), and to Moscow State University (April, 1999). I am grateful to the participants on those occasions for useful comments and discussions, especially to Anita Avramides, Jessica Brown, Greg Currie, Martin Davies, André Gallois, Brian McLaughlin, Joseph Raz, Bruce Russell, Barry Smith, Tom Stoneham, Åsa Wikforss, and Crispin Wright. I owe a special debt to Jessica Brown and Martin Davies, both for their own valuable work on this topic and for many useful discussions of these issues. Work on this paper was supported by a sabbatical leave during 1998–99; I am grateful to Wayne State University for this source of support. I am also grateful to the Sub-faculty of Philosophy of Oxford University, to Corpus Christi College, and to the staff of the Philosophy Centre at

10 Merton Street, for providing me with a superb environment in which to work during 1998–99.

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.
- Brueckner, Anthony: 1995, "The Characteristic Thesis of Anti-Individualism," *Analysis* 55, 146–148.
- Burge, Tyler: 1979, "Individualism and the Mental," *Midwest Studies in Philosophy* 4, 73–121.
- Burge, Tyler: 1982, "Other Bodies," in A. Woodfield (ed.), *Thought and Object: Essays in Intentionality* (Oxford: The Clarendon Press), 97–120.
- 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.
- Davies, Martin: 1993, "Aims and Claims of Externalist Arguments," *Philosophical Issues* 4, Naturalism and Normativity, 227–249.
- 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.
- 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.
- Ludwig, Kirk: 1993, "Externalism, Naturalism and Method," *Philosophical Issues* 4, Naturalism and Normativity, 250–264.
- 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: 1986, "Mental Anaphora," *Synthese* 66, 159–175.
- 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: 1993, "Curing Folk Psychology of 'Arthritis'," *Philosophical Studies* 70, 49–63.
- McKinsey, Michael: 1994, "Individuating Beliefs," *Philosophical Perspectives* 8, Logic and Language, 303–330.
- McKinsey, Michael: 1994a, "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: forthcoming, "Transmission of Warrant and Closure of Apriority," to appear in *New Essays on Semantic Externalism and Self-Knowledge*, S. Nuccetelli, ed. (Cambridge, Mass.: The MIT Press).
- McKinsey, Michael: 2002, "On Knowing Our Own Minds," *The Philosophical Quarterly* 52, 107–116.
- 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.
- Putnam, Hilary: 1975, "The Meaning of 'Meaning'," in K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science* 7 (Minneapolis: The University of Minnesota Press), 131–193.
- Richard, Mark: 1983, "Direct Reference and Ascriptions of Belief," in N. Salmon and S. Soames (eds.), *Propositions and Attitudes* (Oxford: Oxford University Press, 1988), 169–196.
- Richard, Mark: 1990, *Propositional Attitudes* (Cambridge: Cambridge University 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.
- 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.