



OXFORD JOURNALS  
OXFORD UNIVERSITY PRESS

---

What the Externalist Can Know a Priori

Author(s): Paul A. Boghossian

Source: *Proceedings of the Aristotelian Society*, 1997, New Series, Vol. 97 (1997), pp. 161-175

Published by: Oxford University Press on behalf of The Aristotelian Society

Stable URL: <https://www.jstor.org/stable/4545260>

---

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

*The Aristotelian Society* and *Oxford University Press* are collaborating with JSTOR to digitize, preserve and extend access to *Proceedings of the Aristotelian Society*

# IX\*—WHAT THE EXTERNALIST CAN KNOW A PRIORI<sup>1</sup>

by Paul A. Boghossian

Even after much discussion, it remains controversial whether an externalism about mental content is compatible with a traditional doctrine of privileged self-knowledge. By an externalism about mental content, I mean the view that what concepts our thoughts involve may depend not only on facts that are internal to us, but on facts about our environment. It is worth emphasizing, if only because it is still occasionally misperceived, that this thesis is supposed to apply at the level of sense and not merely at that of reference: what *concepts* we think in terms of—and not just what they happen to pick out—is said by the externalist to depend upon environmental facts. By a traditional doctrine of privileged self-knowledge, I mean the view that we are able to know, without the benefit of empirical investigation, what our thoughts are in our own case. Suppose I entertain a thought that I would express with the sentence ‘Water is wet’. According to the traditional doctrine, I can know without empirical investigation (a) that I am entertaining a thought; (b) that it has a particular conceptual content, and (c) that its content is that water is wet.

Let us call someone who combines an externalist view of mental content with a doctrine of privileged self-knowledge a *compatibilist*. In this paper, I will present a *reductio* of compatibilism; in particular, I intend to argue that, if compatibilism were true, we would be in a position to know certain facts about the world a priori, facts that no one can reasonably believe are knowable a

1. Earlier versions of the argument of this paper were presented to my seminar on ‘Self-Knowledge’ at Princeton in the Spring of 1991, to my seminar on ‘Mental Content’ at the University of Michigan in the Spring of 1992, and to the plenary session of the Conference on Self-Knowledge at the University of St. Andrews in August of 1995. I am grateful to those audiences for helpful comments and reactions. I am especially grateful to Anthony Brueckner and Stephen Schiffer for detailed comments on a previous draft and to John Gibbons and Christopher Peacocke for numerous helpful conversations on the general topic.

\*Meeting of the Aristotelian Society, held in the Senior Common Room, Birkbeck College, London, on Monday, 24th February, 1997 at 8.15 p.m.

priori. Whether this should be taken to cast doubt on externalism or on privileged self-knowledge is not an issue I will attempt to settle in this paper. Anti-compatibilist arguments with this general form have been attempted in the past, but I believe that those earlier efforts have misstated the case that needs to be made.<sup>2</sup> Before we get into the details, however, it will be useful to outline certain semantical preliminaries.

## I

*Semantical Preliminaries.* In the case of a general term—for instance ‘water’—I recognize a three-fold distinction between its extension, its referent, and its meaning. A term’s extension is just the set of actual things to which it correctly applies. In the case of ‘water,’ it is all the bits of water existing anywhere in the universe. Since we know that those bits of water are just aggregates of H<sub>2</sub>O molecules, we may also say that the extension of ‘water’ consists in the set of all aggregates of H<sub>2</sub>O molecules that exist anywhere (including those aggregates that we may never encounter).

By a term’s referent, I mean the property that it denotes. In the case of ‘water’ it will be natural to say that its referent is the property of being water. It is possible to wonder whether it would be equally correct to say that it is the property of being H<sub>2</sub>O. That depends on whether the property of being water may be identified with the property of being H<sub>2</sub>O, an example of an interesting question in the theory of properties, but not one that I need to settle for present purposes. What is important here is to be able to distinguish between a term’s extension and its referent, so that we are able to say that a term may express a property that nothing actually has. I think of a sentence’s *truth condition* as the proposition it expresses; and I think of the proposition it expresses as composed out of the referents denoted by its terms. Thus, the truth condition of the sentence ‘Water is wet’ is the proposition made up out of the property of being water and the property of being wet and which says that anything that has the one has the other.

2. See, for example, Michael McKinsey, ‘Anti-Individualism and Privileged Access’, *Analysis* 51 (1991), pp. 9–16, and the effective response by Anthony Brueckner, ‘What an anti-individualist Knows A priori’, *Analysis* 52 (1992) pp. 111–118. This style of anti-compatibilist argument is to be distinguished from the ‘traveling case’ arguments discussed in my ‘Content and Self-Knowledge’, *Philosophical Topics* 17 (1989), pp. 5–26.

I distinguish between the property that the term ‘water’ denotes and its *meaning*. The terms ‘water’ and ‘H<sub>2</sub>O’ may have the same referent, but they do not have the same meaning. What do I mean by the meaning of a term? I wish to be as neutral about this as possible and not to presuppose any particular view. I will let the reader decide to what extent I have succeeded in my neutrality.

Finally, I identify a word’s meaning with the concept it expresses, and so I take the meaning of the sentence ‘Water is wet’ to give the content of the belief that a literal assertoric use of the sentence would express. I use quotes to name words and underlining to name the concept those words express: thus, water is the concept expressed by ‘water’. Now, for the argument.

## II

*Externalism and Twin Earth.* Abstractly speaking, externalism is easily enough defined. It is simply the view that facts external to a thinker’s skin are relevant to the individuation of (certain of) his mental contents. So stated, externalism does not commit one to any specific form of dependence of mental contents on external facts, just to some form of dependence or other.

However, philosophers who embrace externalism don’t do so because they regard it as a self-evident truth. They embrace it, rather, because their intuitive responses to a certain kind of thought experiment—Putnamian Twin Earth fantasies—appear to leave them little choice.<sup>3</sup> And that sort of thought experiment motivates externalism only by motivating a specific form of dependence of mental contents on external facts. In particular, it underwrites the claim that, in the case of an atomic, natural kind concept C, the substance actually picked out by C enters into the individuation of C. To put the claim another way: the substances with which a person actually interacts help determine what atomic, natural kind concepts, if any, that person has.<sup>4</sup>

3. In this paper, I will be restricting myself to externalist theses that are motivated by Putnamian Twin Earth experiments concerning natural kind concepts. In particular, I want to put aside for present purposes externalist theses that are motivated by the influential Burge-style thought experiments involving deference to the usage of linguistic communities. I believe that an argument parallel to the one given in this paper can be mounted for those sorts of externalism as well, but will not argue for this here.

4. By the schema ‘x individuates y’, I just mean that if the value of ‘x’ had been different, the value of ‘y’ would have been different, too. By itself, this doesn’t tell us anything about what the value of ‘y’ is for any particular value of ‘x’. More on this below.

To see this, let us remind ourselves how the Putnam thought experiment is supposed to work. Whereas Oscar, an ordinary English speaker, lives on Earth, his molecular and functional duplicate, Toscar, lives on Twin Earth, a planet just like Earth except that the liquid that fills its lakes and oceans, while indistinguishable from Earthly water in all ordinary circumstances, is not H<sub>2</sub>O but some other substance with a different chemical composition—call it XYZ. Going by whatever criteria are relevant to such matters, water and twin water are distinct kinds of substance, even though a chemically ignorant person would be unable to tell them apart. Now, widespread intuition appears to have it that, whereas Oscar's tokens of 'water' apply exclusively to H<sub>2</sub>O, Toscar's tokens of 'water' apply exclusively to XYZ. Widespread intuition appears to have it, in other words, that Oscar's and Toscar's 'water' tokens have distinct extensions. If this intuition is sustained, then that implies either that their 'water' concepts are not individuated individualistically or that they are not individuated in terms of their referents. For Oscar and Toscar are molecular and functional duplicates of each other: they are alike in all internal respects (up to intentional description). Yet the referents of their concepts differ. Hence, either those concepts don't determine what they refer to in some context-independent way (they are not individuated in terms of their referents) or they do determine what they refer to and so are not individuated individualistically.

It is worth emphasizing that a Twin Earth experiment by itself does not get you all the way to an externalism about concepts; it only gets you as far as this disjunction. It is possible to respond to the experiment, and to the intuitions it generates, by opting for the individualistic disjunct and abandoning the idea that concepts are individuated in terms of their referents. That is the response favoured by so-called 'narrow content' theorists. To get an argument for concept externalism you need not only Twin Earth intuitions, you also need to insist that any notion of mental content deserving of the name has to be individuated in terms of its truth conditions, has to determine the conditions for its truth or satisfaction in some context-independent way. Given this further assumption, there is then no option but to say that Earthly and Twin Earthly tokens of 'water' express distinct concepts—water in the case of the former, and let us say, twater in the case of the latter.

Let us make explicit, then, the various presuppositions involved in using the TE thought experiment as a basis for concept externalism. First, and least controversially, water and twater have to be thought of as distinct substances, distinct natural kinds; otherwise, it won't be true that Oscar's word 'water' and Toscar's word 'water' have distinct extensions and referents. Second, the word 'water'—whether on Earth or on Twin Earth—must be thought of as aiming to express a natural kind concept; otherwise, the fact that water and twater are distinct natural kinds will not be semantically relevant. Third, Oscar and Toscar have to be thought of as chemically indifferent, as having no views about the chemical composition of the liquid kinds around them; otherwise, they won't end up as functional duplicates of each other in the way that the experiment requires. Fourth, the concepts expressed by the Earthly and Twearthly tokens of 'water' have to be thought of as atomic concepts, not compound concepts that are compositionally built up out of other concepts in well-defined ways. For example, the experiment presupposes that water can't be thought of as capable of being defined as: A tasteless, odourless liquid that flows in the rivers and faucets. For if it were a compositional concept of that sort, its extension would be determined by the extension of its ingredient parts. Hence, a conclusion to the effect that water and twater have different extensions would have to proceed differently than it does in Putnam's original experiment, by showing that one of the *ingredients* of water—the concept expressed by 'liquid', for example—has a different extension from that expressed by its Twin counterpart. Finally, and as I have recently noted, concepts must be thought of as individuated in terms of their referents.

### III

*The Argument.* Now, let us suppose that Oscar—our prototypical Twin Earth subject—is a compatibilist. I claim that Oscar is in a position to argue, purely a priori, as follows:

1. If I have the concept water, then water exists.
2. I have the concept water.

Therefore,

3. Water exists.

Since the conclusion is clearly not knowable a priori, one of the premises in Oscar's evidently valid reasoning had better either be false or not knowable a priori. The question is: Can Oscar, qua compatibilist, safely count on one or the other claim? I shall argue that he cannot, that he is committed to both premises (1) and (2) and to their being knowable a priori. If I am right, then the compatibilist is committed to the manifestly absurd conclusion that we can know a priori that water exists.

Now, the a priori knowability of premise (2) just *is* the view that I have called the doctrine of privileged self-knowledge, so we don't have to spend any time debating its dispensability for compatibilism. The only real question concerns premise (1), to an extended discussion of which I now turn.

#### IV

*Perhaps: Water is not required for Water.* Two possible objections need to be considered. On the one hand, an opponent might wish to reject the first premise out of hand, on the grounds that it isn't necessary, on an externalist view, that water exist for someone to have the concept water. On the other, he might wish to argue that, although it is true that water is required for water on an externalist view, that fact is not knowable a priori. Which, if any, of these two alternative strategies is available to the compatibilist? Let us begin with a discussion of the first.

How might Oscar have acquired the concept water without actually interacting with some water, according to a Twin Earth externalist? He couldn't have acquired it merely by virtue of its internal functional role, for his duplicate shares that functional role and yet is said not to have the concept water. And he couldn't have acquired it by theorizing that the liquid around him is H<sub>2</sub>O, for it is stipulated that Oscar is no chemist and has no specific views about the microstructure of water.

An externalist could claim that Oscar might have acquired water from other speakers who have the concept. This suggestion harbours a number of difficulties which limitations of space prevent me from discussing here.<sup>5</sup> Even if it were ultimately sustained,

5. Part of what I have in mind here is that not all speakers could reason in this way, for some of them must have acquired the concept without any help from others. But it would be a needless distraction to go into this now.

however, its impact on the argument I'm pursuing would be minimal—it would simply force us to slightly complicate the absurd conclusion that I have claimed the compatibilist is in a position to derive a priori. Instead of (3), we would now have the equally unpalatable disjunction:

3'. Either water exists or other speakers who have the concept water exist.<sup>6</sup>

For now, however, I propose to set aside this complication and say, simply, that if Twin Earth externalism is true, then contact with water is required for possession of the concept water.

## V

*Water is required for water, but that fact is not a priori.* The most important challenge to the line of argument I'm pursuing derives not from opposition to the truth of this claim, but from opposition to its alleged apriority. This opposition can be stated in a number of related ways; I shall present the strongest version I can think of.

According to the externalist, we know that water is required for possession of the concept water because we know, roughly, that 'water' is one of those words on which a Twin Earth experiment can be run. But doesn't our knowledge that a given word is Twin Earth-eligible rest on empirical information? Compatibilists are very fond of saying that it does;<sup>7</sup> however, it is rare to find their reasons explicitly spelled out. Where exactly do empirical elements intrude into the TE experiment? Let us look at this in some detail. What conditions does a word have to meet if it is to be TE-eligible?

As we have seen, it has to be a word that expresses an atomic concept. It also has to aim to name a natural kind. Furthermore, the user of the word must be indifferent about the essence of the kind that his word aims to name, he must be chemically indifferent.

6. It is interesting to note that here we are in agreement with Tyler Burge, if not on the apriority of the disjunction, then at least on its truth, as far as externalism is concerned:

What seems incredible is to suppose that [Oscar], in his relative ignorance and indifference about the nature of water, holds beliefs whose contents involve the notion, even though neither water nor communal cohorts exist.

See 'Other Bodies', in A. Woodfield (ed.) *Thought and Object* (Oxford: OUP 1982), p. 116.

7. Tyler Burge has urged this in conversation; for a statement in print, see Brueckner, *op. cit.*

But aren't all these conditions available a priori to the user of the word? More to the point, wouldn't a compatibilist have to hold that they are?

The answer is perfectly straightforward, it seems to me, in the case of the latter two conditions. Whether or not a person has beliefs about the microstructure of the kinds around him, and whether or not he intends one of his words to name one of those kinds, are matters that not only seem intuitively a priori, but that a believer in privileged access would have to hold are a priori. Notice that we are not asking whether the word actually names a natural kind, but only whether its user intends it to do so. And according to the doctrine of privileged access, the contents of one's intentions and beliefs are available to one a priori.

It might be thought, however, that the question about atomicity is somewhat more delicate. For is it so clear that facts about compositionality are a priori? Haven't we, as philosophers, often been in the unhappy position of assuming that a concept was compositional, investing a lot of effort in seeking its definition, only to conclude that it has none, that it must be deemed atomic after all?

It is important not to conflate apriority with ease. A fact may be a priori but very difficult to uncover, as the example of any number of mathematical or logical theorems might illustrate. We need not claim that facts about atomicity are easy, only that they are not empirical. And in fact it is hard to see how they could be otherwise. What sense can we make of the idea that knowledge of whether a concept is internally structured might depend on empirical information about the external world?

So far, then, we have not come across a TE-eligibility criterion that could plausibly be claimed not to be available a priori. We are now about to consider another criterion, however, which, if it really were a criterion, would definitely make TE-eligibility an empirical matter. The criterion is this: In addition to *aiming* to express a natural kind, a word must *actually* name a natural kind, if it is to be Twin Earth-eligible. One cannot run a TE thought experiment on a word that aims, but fails, to name a kind.<sup>8</sup>

8. I am grateful to my colleague John Gibbons for helping me see the need to confront this objection and the general line of argument that it opens up.

In support of this claim someone might offer the following. Putnam's original experiment is carried out on a term—'water'—in full knowledge that it does refer to a kind, namely, H<sub>2</sub>O. That knowledge plays a central role in the experiment. Twin Earth by itself doesn't speak to what we should say about a term that doesn't name a natural kind. So, for all that Twin Earth overtly commits us to, actually naming a natural kind is a condition on TE-eligibility and that is certainly not a condition that is available a priori. True, Twin Earth teaches us that water is required for the word 'water' to express the concept water, such an objector would concede; but we only learn this because we know—empirically—that water is the kind actually named by 'water'. Hence, TE-eligibility is not a priori.

Now, I think that this objection, as stated, isn't correct; buried within it, however, is another objection that is considerably more challenging. The reason this particular objection doesn't succeed is that it is quite clear that we *can* run a TE experiment on a word that doesn't actually name a natural kind. Suppose we had such a word, W, on Earth. Then, to get a successful TE experiment, all you need to do is describe a Twin situation in which, although the users of the word type W are functional and molecular duplicates of their counterparts on Earth, W does name a kind in the Twin situation. Provided intuition still has it that the extension of Earthly tokens of W are different from the extension of the Twin tokens of W—which of course they will be since the extension of the former will be empty and the extension of the latter won't be—the experiment will succeed.

Now, however, the objector would appear to be in a position to pose a more difficult challenge. For if this is in fact right, and we can run TE experiments even on terms that fail to refer, then how do we know a priori that water is required for 'water' to express water? We can't infer that claim merely from the fact that 'water' is TE-eligible, for we have established that even empty terms are TE-eligible. Maybe water is the concept that 'water' expresses when it fails to name a natural kind, when there is no water for it to name. If we can be said to know that water is required for water, we know that only by virtue of our knowledge that 'water' does name a natural kind, namely, water. And that, of course, is something that we could only have come to know empirically.

Hence, our knowledge that water is required for water is not a priori.

Here, finally, we come across the most important challenge to the line of argument I've been pursuing. It will be interesting to uncover the reason why it doesn't ultimately protect compatibilism from the charge of absurdity.

## VI

*The Empty Case.* I want to approach a response to this objection somewhat indirectly, by focusing on the following question: What should a Twin Earth externalist say about the case where a word aiming to name a natural kind fails to do so? Two sorts of scenario might lead to such an outcome. On the one hand, a word like 'water' may fail to name a natural kind because the liquids to which it is competently applied don't form a natural kind, but rather a heterogeneous motley. On the other hand, a term may fail to name a kind because there fails to be anything at all out there—motley or otherwise—to which it could correctly be said to apply. Here I want to concentrate on the second more extreme sort of case because it throws the issues of interest into sharper relief.

So let us imagine a planet just like ours in which, although it very much seems to its inhabitants that there is a clear, tasteless and colourless liquid flowing in their rivers and taps and to which they confidently take themselves to be applying the word 'water', these appearances are systematically false and constitute a sort of pervasive collective mirage. In point of actual fact, the lakes, rivers and taps on this particular Twin Earth run bone dry. All of this may seem very far-fetched, and no doubt it is. However, the scenario described is not substantially different—except in point of pervasiveness—from what has actually turned out to be true in the case of such terms as 'phlogiston' and 'caloric'; and, anyway, the point isn't to describe a genuine possibility. Rather, it is to inquire how a particular semantical theory proposes to treat cases of reference failure and whether it is committed to treating such cases in a particular way. What *concept*, if any, should a Twin Earth externalist say would be expressed by tokens of the word 'water' on this Dry Earth?

Some may think the answer to be obvious. Since externalism is the view that the concept expressed by a word is individuated in

part by the referent of that word, then it follows, does it not, that if the word has no referent that it expresses no concept?

This reasoning would be far too hasty. It confuses the claim that a concept is individuated in terms of its referent, with the claim that the existence of the concept depends on the existence of a referent. To put matters in terms of a familiar technical vocabulary, it confuses externalist individuation with object-dependence. All that Twin Earth externalism is committed to, strictly speaking, is the claim that, if the referent of a given word were different, the concept it would then express would be different, too. And that is consistent with the claim that the word would express a concept in a case where it fails to refer, provided that the concept it would there express is different from any it would express in a case where it does refer. To say it again, externalist individuation, in the sense in which Twin Earth externalism is committed to it, is just the view that, if two words differ in their referents, then they also differ in the concepts they express; strictly speaking, that is consistent with a word's expressing some concept or other even when it fails to have a referent.

But what concept should we say 'water' expresses under the conditions described, in which there fails to be any natural kind for it to refer to? We may consider options under two main headings: compound and atomic.

We could try saying that under the envisioned dry conditions, 'water' expresses a suitable compound concept made up in the familiar way out of other available concepts. Which compound concept? Most plausibly, I suppose, something like: the clear, tasteless, colourless liquid that flows in the taps and the rivers around here and... It won't matter much for the purposes of this argument how precisely this proposal is fleshed out. On any such view, the word 'water' will contribute a complex property to the proposition expressed by whole sentences involving it, one which, as a matter of contingent fact, nothing in that environment possesses.

Intuitively, this seems to me to be a plausible view of the matter. When I think of a group of people just like us, applying the word 'water' confidently to something that appears to them to be a clear, colourless, tasteless liquid in their environment, when in fact there is no such liquid in their environment, I feel tempted by the sort of

error theory of their linguistic behaviour that the present proposal delivers. It seems plausible to me to say that what these people mean by the word 'water' is this clear, colourless, tasteless liquid etc., which, however and unfortunately, is not to be found in their environment.

The problem is that it is very difficult to see how such a view could be available to the Twin Earth externalist. Remember, the TE externalist is committed, for reasons detailed earlier, to holding that 'water' expresses an *atomic* concept under conditions where it has a non-empty extension, whether that extension be H<sub>2</sub>O or XYZ or whatever. That is one of the presuppositions of the Twin Earth experiment. But, then, how can the very same word, with the very same functional role, express an atomic concept under one set of external conditions and a compound, decompositional concept under another set of external conditions? A concept's compositionality is exclusively a function its internal 'syntax' and can't be contingent upon external circumstances in the way that the present proposal would require.

Let me forestall a possible misunderstanding of this point. My argument here is not that, if the compatibilist were to embrace the compound notion, that would undermine his commitment to privileged access. For although it is true that embracing the compound option for 'water' on Dry Earth, while being committed to its atomicity on Earth, would have the effect of making facts about compositionality come out a posteriori, that would not flout any doctrine of privileged access that I have defined.

Nor is my argument here that the compound option is unacceptable because it runs into conflict with the independently plausible claim that facts about compositionality are a priori, although, as I noted above, that is something I believe and would be prepared to defend.

In fact, my argument here is not epistemic at all, but rather metaphysical. The compound option requires the externalist to say that one and the same word, with one and the same functional role, may express an atomic concept under one set of external circumstances and a compound decompositional concept under another set of external circumstances. But it is hard to see how the *compositionality* of a concept could be a function of its external circumstances in this way. Compositionality, as I understand it, can

only be a function of the internal syntax of a concept; it can't supervene on external circumstances in the way that the compound proposal would require. (This is especially clear on a 'language of thought' picture of mental representation, but is independent of it.)

How do things look with the other main class of available options, that according to which the empty tokens of 'water' express an atomic concept? On this branch, too, we need to answer the question: Which atomic concept will that be, according to the TE externalist?

The externalist will know quite a lot about which concepts it cannot be: in particular, he will know that it cannot be identical with any of the concepts that are expressed by non-empty tokens of 'water'. To suppose otherwise would contradict his overriding commitment to individuating a concept in terms of its referent. But can he tell us, in line with his overriding commitment, what concept *is* expressed by the empty tokens of 'water'?

Unfortunately, there would appear to be a compelling argument showing that the externalist will not be able to say what atomic concept is expressed by the non-referring tokens of 'water', because by his own lights there can't be such a concept. Let me explain.

We have seen that one of the assumptions that is needed to transform a TE experiment into an argument for externalism is the assumption that concepts have context-independent conditions of satisfaction, or, in the case of thought contents, context-independent conditions of truth. So let us ask this: What are the satisfaction conditions for 'water' on Dry Earth, to what sorts of liquid does it apply? By assumption, of course, the actual extension of 'water' is empty on Dry Earth, so there is no liquid in its actual environment to which it applies. But the question I am asking is consistent with the word's actual extension being empty, and consistent even with its extension being empty in all worlds. What I want to know is: What proposition—what truth condition—is expressed by sentences of the form, 'Water is wet', for example, as uttered on Dry Earth? What is it that gets said? Never mind if such sentences are ruled false in the actual world, or even in all worlds.

On the line we are currently investigating, the answer has to be that there is no fact of the matter what truth condition is expressed

by sentences involving ‘water’ on Dry Earth, for there is no fact of the matter what property is denoted by those tokens of ‘water’. Since there is no natural kind at the end of the relevant causal chain leading up to uses of ‘water’ on Dry Earth, there is no fact of the matter what the referent of ‘water’ is and so no fact of the matter what proposition is expressed by sentences involving it.

But on an externalist view, this admission is fatal to the claim that there is a concept there in the first place, for an externalism about concepts is fuelled in part by the conviction that thought contents must possess context-invariant conditions of satisfaction or, as appropriate, of truth. If, in a given context, there is no fact of the matter what the referent of a given concept is, then to that extent there is also no fact of the matter what the concept is.

We have looked at two possible tacks that an externalist might take regarding empty tokens of ‘water’, and we have found them both to be irremediably problematic. Letting the empty tokens express a compound concept, while having the virtue of supplying the word with a property to refer to, runs directly into conflict with the externalist’s commitment to the atomicity of ‘water’. Evading this problem by letting the word express an atomic concept, on the other hand, runs into direct conflict with the externalist’s commitment to the idea that concepts must possess determinate, context-independent, conditions of satisfaction.

What then is the externalist to say about the empty case? The answer would appear to be that he has to say just what the proponent of object-dependence said he should say all along—namely, that the empty tokens simply don’t express a determinate concept. That turns out to be the right thing to say not because TE externalism is conceptually equivalent to object-dependence, but because TE externalism, in conjunction with its other commitments, entails object-dependence.

## VII

*The Argument Completed.* If this is right, then the compatibilist is in a position to conclude—via purely a priori reasoning—that if a term expresses a concept in the first place, that it must have a non-empty extension. Moreover, privileged access assures him that he will be able to tell a priori whether or not a given term does express a concept, and indeed, if it does, which one. In particular, our friend

Oscar will be able to tell non-empirically that his term 'water' expresses a concept, and in particular that it expresses the concept water. Putting these two bits of information together, he is in a position to conclude, a priori, that water must have existed at some time. And that, we are all agreed, is not something he ought to be able to do.<sup>9</sup>

*Department of Philosophy*  
*New York University*  
*New York*  
*NY 10003*

9. To generate our problem for the compatibilist we have had to assume that when Oscar reasons as we have described, his a priori warrant for the premises of his argument transmits, across the a priori known entailment, to the entailed conclusion. Recently, some philosophers have taken to questioning whether this principle is correct. Aren't there cases, they have asked, where although A is known a priori, and although A is known a priori to entail B, nevertheless B is not known a priori. See, for example, the interesting paper by Martin Davies, 'Externalism, Architecturalism and Epistemic Warrant', in MacDonald, Smith and Wright, *Knowing Our Own Minds* (Oxford: Oxford University Press, 1997). I have to say that I would be very surprised if there turned out to be any such cases that survived scrutiny. However, defending this claim in full generality is something that deserves separate treatment and will have to be left for another occasion. Here, I will settle for discussing one such case that has been suggested to me (by Stephen Schiffer). Consider the following inference:

If I have toothache, then teeth exist.

I have toothache.

Therefore,

Teeth exist.

I have defined 'a priori knowledge' as 'knowledge that is obtained without empirical investigation'. Relative to this (admittedly vague and informal) characterization, don't the premises of this argument come out a priori? Can't I know that I have toothache without empirical investigation? And, also, that if I have toothache then that I have teeth? However, the conclusion of this argument is clearly not a priori. Therefore, there must be something wrong with the transmission of warrant principle that we have been assuming.

My perhaps predictable reply is that it is not at all clear that the premises of the toothache argument are a priori, relative to the intended notion of 'a priori'. That we are in pain, and even that we are in a particular kind of phenomenologically classifiable pain (a 'toothache' pain)—these matters seem clearly a priori. But there is no intuitive reason to believe, it seems to me, that we can know a priori that we have toothache, if that is supposed to mean, as it evidently does in the objection under consideration, that we have an ache *in a tooth*. Imagine a toothless person insisting that he has toothache; would we have to defer to his alleged a priori access to that fact?