

SANFORD C. GOLDBERG

ANTI-INDIVIDUALISM, CONCEPTUAL OMNISCIENCE,
AND SKEPTICISM

ABSTRACT. Given anti-individualism, a subject might have *a priori* (non-empirical) knowledge that she herself is thinking that *p*, have complete and exhaustive explicational knowledge of all of the concepts composing the content that *p*, and yet still need empirical information (e.g. regarding her embedding conditions and history) prior to being in a position to apply her exhaustive conceptual knowledge in a knowledgeable way to the thought that *p*. This result should be *welcomed* by anti-individualists: it squares with everything that compatibilist-minded anti-individualists have said regarding e.g. the compatibility of anti-individualism and basic self-knowledge; and more importantly it contains the crux of a response to McKinsey-style arguments against anti-individualism.

1.

Let anti-individualism (henceforth ‘AI’) be the thesis that at least some propositional attitudes depend for their individuation on (social or physical) facts that do not supervene on the internal states of the subject.¹ Let ‘first-person authority’ (henceforth ‘FPA’) be the thesis that, for all subjects *S* and occurrent thoughts that *p*, *S* has – or is in a position to have – special, non-empirical (or ‘*a priori*’) knowledge of her thought that *p*. The literature regarding the compatibility of these two theses is vast.² Though compatibilism seems to be in the ascendancy, pockets of resistance remain. There are two main arguments that are used to defend incompatibilist conclusions: Boghossian-style arguments from world-switching (see Boghossian, 1989; Goldberg, 1999) and McKinsey-style *reductio* arguments from AI’s anti-skeptical implications (see McKinsey, 1991; Brown, 1995; Boghossian, 1998). In this paper I want to argue that, strange as it may seem, the anti-individualist can actually exploit the intuitions elicited by arguments from world-switching, in order to rebut McKinsey-style *reductio* arguments against AI.³



The intuitions in question can be brought out by a familiar switching story, with one twist: the switched subject is imagined to be conceptually omniscient. The point of this stipulation is to make clear what a subject can be said to know, and what she cannot, in virtue of satisfying FPA (and thereby having special, non-empirical knowledge to the effect that she herself is thinking that p). More specifically, by imagining that the subject of the switching is conceptually omniscient, we make clear that, if there is something that (for lack of a better way to put it) she fails to know regarding her thought, that lack is not something that derives from any incomplete grasp of the concepts involved or from false beliefs about the application conditions of those concepts.⁴ Accordingly, given such a lack, we will need an alternative explanation for that lack. This paves room for the claim that the lack in question is to be traced to her ignorance of certain *empirical* facts. Since the cases we will examine are world-switching cases, the natural suggestion will be that the empirical facts in question are facts regarding the subject's switching regimen. In short, the 'intuitive' result I will invite the reader to accept on the basis of the switching case concerns the role, given AI, that empirical (*a posteriori*) considerations play in a subject's relation to her own thoughts.⁵

The switching case itself is as follows. Jane is conceptually omniscient with respect to English: given any non-logical and non-indexical expression E of English, Jane can correctly and exhaustively explicate the concept expressed by E . (Here I will assume that a full understanding of e.g. 'water' will include knowledge of the chemical composition of water; and, more generally, that a full understanding of any natural kind term will include knowledge of the scientific nature of the property designated by the term.) But Jane is like the rest of us when it comes to her knowledge of the empirical world. There are empirical facts of which she is ignorant; though many of her beliefs regarding empirical matters are true, some are false; and her methods for acquiring beliefs regarding empirical matters are like yours and mine (which is to say that sometimes she makes observations herself, sometimes she takes things on the word of others, and so on). One fine summer's day, Jane is enjoying herself at the beach, and at time t thinks to herself a thought that she expresses with the sentence, 'The water must be

warm today.’ Call this particular thought, *T*. Just as she thinks *T*, however, Jane is informed that she has been the victim of world-switching. She is told that, though she was born on Earth, some time in her past (possibly long ago, possibly not) she began a switching regimen that took her from Earth to Twin Earth, on at least one but possibly many more occasions. She is not told how long she has spent on either planet, or how often she has been switched, or her present world-residency; she remains in ignorance on all of these matters. Under these conditions, can Jane knowledgeably explicate the concepts that figure in *T*? I want to argue that she cannot, and that her knowledge failure on this score should be welcomed by proponents of AI (for eliciting intuitions that are helpful in the battle against McKinsey-style arguments).

2.

I want to begin with my claim that, in the world-switching story just presented, Jane is not in a position to give knowledgeable explications of the concepts that figure in *T*.

We can begin by separating two distinct issues.

First, there is the issue regarding whether she has at her disposal the explicatory resources to explicate the concepts figuring in *T*. Suppose that ‘water’ as used by Jane in her expression of *T* expresses WATER (the concept expressed by the English word-form ‘water’). Then Jane, who is conceptually omniscient regarding English, does have such resources. Now suppose instead that ‘water’ as used by Jane in her expression of *T* expresses TWATER. In such a case would Jane have at her disposal the explicatory resources to explicate the ‘water’-concept figuring in *T*? Not if she is only omniscient regarding English. But if she is omniscient regarding both English and Twin English, then she would have such resources. Since it will be helpful to imagine Jane to have all such resources – for in that way we can most clearly discern the empirical considerations relevant to her relation to her own thoughts – we will imagine Jane to be omniscient regarding both languages: given a thought expressed by a sentence (whether in English or Twin English), Jane has at her disposal all of the explicatory resources to explicate the non-indexical concepts figuring in the thought expressed.

But there is a second issue, regarding whether Jane can *apply* her explicatory resources in such a way as to provide a warranted explication of the concepts figuring in *T*. Here it seems that Jane needs to supplement her complete conceptual knowledge with some empirical knowledge: namely, knowledge regarding the details of her world-switching regimen. For unless she knows whether ‘water’, as it figured in her *T*-expressing utterance at *t*, expresses WATER or TWATER, she does not know whether to explicate the ‘water’-concept by giving the explication associated with WATER, or that associated with TWATER. Then, even granting that she has exhaustive knowledge of both entries, this knowledge by itself does not determine which to apply to the ‘water’-concept figuring in *T*.

In what follows, I will be speaking of a conceptually omniscient subject who nonetheless is not in a position to offer a knowledgeable explication of her ‘water’-concept on a given occasion. The type of situation I have in mind is that in which a subject has exhaustive explicational knowledge of all of the concepts figuring in her repertoire, yet (for lack of some empirical information) is unable to apply this knowledge to a given concept figuring in one of her thoughts on a given occasion *O*, so as to provide a knowledgeable explication of that concept. Admittedly, if we are to treat a subject *S* as exemplifying this type of situation, then several conditions must hold. To begin, there are requirements on the determinacy of *S*’s thought: *S* must be such that on *O* she is thinking a determinate thought, i.e., a thought composed of determinate concepts – and which thought *S* is thinking on *O* must itself be a determinate matter. Moreover, there is what I will call a *referential requirement*: *S* must be using one of her expressions to refer to a type of stuff in such a way that she cannot determine, merely by reflecting on her semantic intentions, which concepts figure in her thought. (Why I call this the referential requirement will emerge below.) Finally, her inability to provide a knowledgeable explication must not undermine either the determinacy condition or the referential condition. But it is clear that, on the assumption of AI, these conditions can be satisfied. Indeed, by developing the case of Jane a bit, we will see that these conditions are satisfied in Jane’s case.

Take the determinacy condition first. It is clear that certain ways of spelling out Jane’s switching history may well undermine the

claim that she is thinking a determinate thought. Suppose she gets switched back and forth between Earth and Twin Earth, spends a roughly equal time on each, forms a roughly equal number of indexical thoughts regarding the watery stuff on each, and then thinks a general thought that she would express with the sentence ‘Water is wet’, without having any intention that would warrant our thinking that she has either water or twater ‘in mind’. Then we might want to say that her use of ‘water’ on this occasion expresses neither WATER nor TWATER but rather WATER-OR-TWATER; or perhaps we will say it expresses WATERY-LIKE LIQUID; or perhaps merely WET, TASTELESS, DRINKABLE LIQUID; or, if she also has the intention to be expressing a natural kind property, then we may not know what to say. Then again, looking at all of the possibilities, we may opt for the conclusion that her thought is not a determinate one at all. Even so, other ways of spelling out Jane’s switching history leave the determinacy of the thought perfectly intact. Suppose Jane was switched only once to Twin Earth, and that after five minutes on Twin Earth she is whisked back to Earth (where she remains for the rest of her life). Suppose further that in her five minutes on Twin Earth she did not form any judgements regarding, and more generally did not direct her attention to, twater. And suppose finally that the interval between the time she returned to Earth and t (the time at which she entertained T), is forty years, during all of which time she remained on Earth. Then it seems quite clear that T is a water-thought, and that this is determinately the case.

Next, take the referential condition. It is easiest to see that this is satisfied in Jane’s case by first considering a case in which it would not be satisfied. Suppose that Jane’s use of ‘water’ at time t was informed by her intention to express the property of the wet, tasteless, drinkable stuff on Earth *whatever planet she happened to be on at t* . Alternatively, suppose that Jane’s use of ‘water’ at t was informed by her intention to express the property of H₂O.⁶ In both of these cases, we should interpret her utterance at t of ‘The water must be warm today’ as expressing a proposition involving WATER. Here, the point would be that she *would* know *a priori* which concept her use of ‘water’ expressed – what she wouldn’t know (*a priori* or otherwise) is whether the stuff she referred to as

'the water' is, as she takes it to be, water (= H₂O). In neither of these cases would the referential requirement be satisfied: her semantic intention to be using 'water' at *t* so as to express the property of the wet, tasteless, drinkable stuff on Earth (whatever planet she happened to be on at *t*) – or, in the other imagined case, her semantic intention to be using 'water' at *t* so as to express the property of H₂O – suffices to establish that this use of 'water' expressed WATER; and so by introspection she could determine that she was thinking a water-thought. But now imagine that she was using 'water' at *t* with the intention of expressing the property of the wet, tasteless, drinkable stuff *around here* (= the spot where she was located at *t*). That is, her semantic intentions regarding her use of 'water' at *t* include a certain referential intention. Then it would seem that, unless she knows on what planet she resides at *t*, she does not know, and mere reflection on her semantic intention will not determine, whether her use of 'water' at *t* expresses WATER or TWATER. In that case the referential condition is satisfied.⁷

I now want to make clear that Jane's inability to provide a knowledgeable explication of her 'water'-concept at *t* need not undermine either the claim that she satisfies the determinacy condition or the claim that she satisfies the referential condition. It will be helpful to distinguish between the objective situation, and the situation as it strikes Jane. Objectively: Jane was switched from Earth to Twin Earth only once, and remained on Twin Earth only for five minutes, during which interval she did not entertain any indexical thoughts regarding twater. Immediately after this brief episode on Twin Earth she returned to Earth where she spent the rest of her life. Forty years went by prior to the scene at time *t* of her uttering (in the presence of water) 'The water must be warm'. Subjectively: prior to be apprized (right after *t*) of her world-switching history, Jane never supposed that she had ever been on a planet other than Earth. Consequently, she never stopped to wonder whether any of her uses of 'water' ever expressed anything other than WATER. To be sure, in a semantics class she once read of Twin Earth, and so (being ever the diligent student) immediately learned how to explicate 'water' as used by Twin Earthians. But she never suspected that she herself spent any time on that planet. When she is apprized of the bare fact of her switching, i.e. without learning of the length of stay on

either planet etc., she is no longer certain which ‘water’-concept she expressed at t . She intended, and knows that she intended, to use ‘water’ so as to express the property of the wet, tasteless, drinkable stuff on the planet on which she then (at t) resided; only she no longer knows on which planet she then resided, or for how long. No amount of introspection will enable to determine her world-residency at t , and so no amount of introspection will enable her to determine whether her use of ‘water’ at t expressed WATER or TWATER. Even though her conceptual knowledge regarding both WATER and TWATER is correct and exhaustive, she is not in a position to knowledgeably apply her conceptual knowledge to the ‘water’-thought she expressed at t . Such is the sorry state she is now in, having been apprized of the bare fact of her world-switching.

3.

The argument so far has established a certain possibility claim: on the assumption of AI, it is possible for a subject to be completely omniscient with respect to the semantics of (say) both English and Twin English, and thus be in a position to offer knowledgeable explications of the concepts expressed by any (non-logical and non-indexical) English or Twin English expression, while at the same time failing to be able to apply this conceptual knowledge to a given thought in such a way as to yield a knowledgeable explication of the concepts figuring in that thought. I will designate this situation as one in which the conceptually omniscient subject fails to be in a position to provide a *knowledgeable explication* of the concepts involved in her thought. As I indicated above, the source of this type of knowledge failure (if I may call it that⁸) lies not in any conceptual lack on the subject’s part; rather, it lies in the fact that the subject lacks some empirical knowledge which she needs to have in order to apply her flawless and complete conceptual knowledge to the case at hand. In this section I want to begin the process of interpreting this result.

My first interpretative claim is this. I maintain (with Burge, 1988; Falvey and Owens, 1994; Brueckner, 2000, and others) that, whatever Jane fails to know, she has (and cannot fail to have) a special *a priori* knowledge of T . The point here is a familiar one,

so I will be brief. On any occasion on which Jane is to be ascribed a determinate thought, we should also be willing to ascribe to Jane an *a priori* knowledge of her thought, merely on the grounds that she can self-ascribe that thought in such a way that the resulting judgement is (to borrow Burge's 1988 phrase) *self-verifying*. The general point is this: For all subjects *S* and propositions *p*, if *S* self-ascribes the thought that *p* using a sentence of the form 'I am currently thinking that *p*', then *S* is thinking that *p*: one thinks the thought in the very act of self-ascribing it (see Burge, 1988, p. 654). For this reason, self-ascriptions of this form will be self-verifying; and so we can see this as the basis for a certain kind of *a priori* self-knowledge of one's thought – a kind of self-knowledge that one cannot fail to have so long as one has the capacity to form self-ascriptive judgements of this form. Upshot: our result in section 2 should not be interpreted so as to cast doubt on what Burge would call Jane's *basic self-knowledge* regarding *T*.⁹

In light of this, it is worth considering how the knowledge failure to which I am pointing compares to other kinds of knowledge failure discussed in the literature on AI and FPA.

Let us begin first with two distinct types of knowledge failure that have been illustrated by Burge. Burge (1986) presented a case in which a thinker fails to be able to provide a correct explication of his 'sofa'-concept, for having formed a false empirical theory regarding sofas, which theory gave rise to false beliefs regarding the application conditions of 'sofa'. The sort of failure of knowledge at play in our result is different: Jane's beliefs regarding the application conditions of 'water' are all correct, based as they are on her complete and correct explicational knowledge regarding WATER. Another knowledge failure identified by Burge (e.g., in his 1989) is a case in which a subject fails to be in a position to provide an explication of the concept figuring in a given thought, for having an incomplete or partial understanding of the concept in question. Again, the sort of failure of knowledge at play in our result differs from this case in one obvious way, since the kind of explicatory failure at play in Jane's case is one that does *not* derive from partial understanding of the concepts involved.

More interesting is the comparison between Jane's knowledge failure and what Falvey and Owens would call a failure of discrimi-

natory knowledge regarding one's thought. Falvey and Owens (1994) were interested in establishing that a thinker's failure to discriminate between the thought she actually entertained from a thought she would have been entertaining had she been in some counterfactual situation, does not undermine her basic self-knowledge of her actual thought.¹⁰ I agree with this point entirely; this is why I claim that Jane's failure of knowledge does not – stronger, *cannot* – undermine Jane's basic self-knowledge of her thought. What is more, it would seem that Falvey and Owens' category of 'discriminatory knowledge of one's thought' is usefully applied to the present case: it is because Jane does not know whether the 'water'-concept in *T* is WATER or TWATER that she cannot provide a knowledgeable explication. If we interpret Jane's failure as a failure of discriminatory knowledge of her thought, then we reach an interesting result: discriminatory knowledge failures of this sort can arise for reasons other than incomplete grasp or false beliefs about a concept's application conditions.

In sum: however it relates to Falvey and Owens' failures of discriminatory knowledge, the kind of knowledge failure which I am designating as Jane's failure of 'knowledgeable explication' of the concepts figuring in *T* is not to be thought of as Jane's failing to have basic self-knowledge regarding *T*. More broadly, since the Burgean considerations supporting the ascription to Jane of basic self-knowledge of *T* are in place given *any* subject and *any* occurrent thought, we should not think that the claim, that (given AI) subjects may fail to be in a position to provide knowledgeable explications of the concepts figuring in their thoughts, does anything to undermine compatibilism regarding AI and basic self-knowledge. On the contrary, failures of knowledgeable explications serve the purpose of pointing out just how insubstantial basic self-knowledge actually is. (I will return to this in section 5.)

4.

So far I have argued that the kind of knowledge failure at play in cases like Jane's is not to be understood as reflecting the sorts of knowledge failure Burge discussed in his 1986 and his 1989, though it may well be assimilable into the (Falvey and Owens, 1994)

category of failures of discriminatory knowledge of one's thought. In this and the following section I want to argue that, however we classify it, the kind of knowledge failure at play in Jane's case does have an interesting property: it appears to offer the proponent of AI with a novel way to rebut McKinsey-style arguments against AI.

McKinsey-style arguments against AI are arguments that purport to establish that AI has unacceptably strong anti-skeptical implications. The crux of such arguments is as follows: given some thinker *S* who is thinking a thought involving a concept which (according to AI) depends for its individuation on conditions that do not supervene on *S*'s internal states, this fact, together with AI and FPA, entail that *S* can have *a priori* knowledge of some fact *f*, where *f* would appear to intuition to be only knowable *a posteriori*. Proponents of McKinsey-style arguments differ over the relevant *f*,¹¹ and some responses to McKinsey-style arguments purport to show that all of the facts that are so derivable are indeed knowable *a priori* (a case of the proverbial "one person's *modus tollens* is another's *modus ponens*").¹² In this section I will be concerned with one of the most recent versions of McKinsey-style arguments, found in Boghossian (1998).¹³ (In the next I will consider how to generalize the case to be presented here.) I am focusing on Boghossian's version, first because it attempts to rectify problems with some of the earlier versions; but second, because it contains what to my mind is an interesting flaw, regarding how the proponent of AI is to characterize concepts. At any rate, on Boghossian's version, the fact *f* that is alleged to be *a priori* derivable from AI and FPA (together with a subject thinking a water-thought) is the fact that water exists.¹⁴ I want to dispute that such a fact is *a priori* derivable from such a basis. To do so, I propose to show how, even assuming AI and FPA, the route to *a priori* knowledge of the existence of water is blocked even to those who have basic self-knowledge of what in fact is a water-thought. My argument here will draw on my earlier claim, to the effect that conceptually-omniscient Jane is not in a position to offer a knowledgeable explication of the 'water'-concept figuring in *T*.

To begin, let us designate as an 'anti-individualistic concept' any concept *C* that is such that, for any speaker *S* who might possess *C*, the individuation of *C* depends on facts that do not supervene on *S*'s

internal states. This characterization is meant to cover such things as natural kind concepts (see Putnam, 1976), as well as concepts for artifacts such as sofas (see e.g. Burge, 1979, 1986). The claim in dispute can now be expressed as follows:

- (A) Where p is a thought-content in which some anti-individualistic concept C figures, and where the possession of C by subject S presupposes or depends upon the existence of object or property E , if S knows *a priori* that she herself is thinking that p , then S is in a position to determine *a priori* that she is thinking a thought involving a concept whose availability depends on the existence of E .

It is clear that McKinsey-style arguments in general, and Boghossian's in particular, depend on (A). What I want to suggest is that the case of conceptually-omniscient Jane appears to provide the anti-individualist with the basis from which to argue that (A) is false, and from this to conclude that something must be wrong with Boghossian's argument.

We can begin with Boghossian's own words. He imagines that, faced with any would-be McKinsey-style argument, the proponent of AI will respond by arguing that a speaker's knowledge that a given word expresses a particular anti-individualistic concept C ¹⁵ "rest[s] on empirical information". To such a suggestion he responds as follows:

Compatibilists are very fond of saying that it does [so depend]; however, it is rare to find their reasons explicitly spelled out. Where exactly do empirical elements intrude into the Twin Earth experiment? (Boghossian, 1998, p. 277)

Of course, for Boghossian this question is rhetorical: he thinks that empirical elements do not so intrude. In particular he thinks, and holds that the proponent of semantic AI is committed to thinking, that a subject has *a priori* access to whether a particular anti-individualistic concept C figures in her thought.

Though I think that the issue in question, regarding a subject's *a priori* access to whether a particular anti-individualistic concept C figures in her thought, is more vexed than Boghossian appears to appreciate (more on which below), it will be helpful to begin with the argument he offers for this contention. The contention

that subjects do have such *a priori* access is supposed to emerge from a line of argument that begins by presenting a difficulty to the proponent of AI. The difficulty is centered around how to describe a case involving a subject on an imagined 'Dry Earth':

... Imagine a planet just like ours on which, although it very much seems to its inhabitants that there is a clear, tasteless, 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 collective mirage. In point of actual fact, the lakes, rivers, and taps on this particular Twin Earth run bone-dry. (Boghossian, 1998, p. 279)

He then asks how a proponent of AI should describe this case: "What concept, if any, should [an anti-individualist] say would be expressed by tokens of 'water' on this Dry Earth?" (p. 280). The difficulty arises when we ask the proponent of AI whether the Dry Earth 'water'-concept is atomic or compound. It is pretheoretically natural to hold that the 'water'-concept on Dry Earth is compound – that it expresses something like THE WET, TRANSPARENT, DRINKABLE LIQUID THAT COMES OUT OF TAPS AND FLOWS IN STREAMS. But this presents a problem for the proponent of AI, since it would appear that the proponent of AI must treat the 'water'-concept on Dry Earth as atomic. After all, the proponent of AI wants to hold that 'water' "expresses an atomic concept under conditions where it has a non-empty extension" (Boghossian, 1998, p. 280); so (Boghossian reasons) we would expect the proponent of AI to hold that 'water' expresses an atomic concept under any condition on which its 'internal syntax' is the same. To deny this would be *ad hoc*, since to do so is to make "facts about compositionality come out a posteriori" (Boghossian, 1998, p. 281). But then, since 'water' as used by Dry Earthians does have the same 'internal syntax' as 'water' used by Earthians, the proponent of AI would thus appear to be forced to endorse the hypothesis that 'water' on Dry Earth expresses an atomic concept. The problem, of course, is that, since 'water' has no referent on Dry Earth, the proponent of AI is without resources to individuate the atomic concept in question. Thus it would appear that, given AI, if a term such as 'water' expresses a determinate concept at all, it must have a non-empty extension; and from here we are off and running to Boghossian's conclusion:

. . . the compatibilist is in a position to conclude – via purely *a priori* reasoning – that if a term expresses a concept in the first place, 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. (Boghossian, 1998, p. 283)

Such is Boghossian’s McKinsey-style argument against AI.

There are two mutually-reinforcing considerations that can be brought to bear against Boghossian’s McKinsey-style argument.

The first is this. Boghossian alleges that “privileged access” assures that

- (*) Subjects who use a term on a given occasion “will be able to tell *a priori* whether or not a given term does express a concept, and indeed, if it does, which one”.

I submit that there are actually two distinct notions of *being in a position to tell which concept figures in one’s own thought*, that Boghossian’s argument trades on having the stronger of the two notions figure in (*), and that the argument from conceptual omniscience (in sections 2 and 3 above) shows that (*) is false on that reading. Let us formulate the weaker notion as follows: a speaker is in a position to ‘weakly tell which’ concept is expressed by her use of ‘water’ on occasion *O* when, were she asked which concept that use expresses, she could provide a trivially correct answer merely by re-using ‘water’ with the intention to express what she had expressed with ‘water’ on *O*. It is rather easy to be in a position to weakly tell which concept figures in your thought: so long as nothing in the interval between *O* and the time of the question prevents you from re-using the same word-form so as to express the same concept, you can be said to be in such a position.¹⁶ However, this is a very weak sense of ‘telling which concept’ figures in one’s own thought. For it is consistent with being in a position to weakly tell which concept, . . . that one is not able to apply one’s conceptual omniscience to the concept so as to produce a knowledgeable explication of the concept in question; and so it is consistent with being in this position that one not know whether to apply the explicational

entry for WATER, or that for TWATER, to the ‘water’-concept in question. From this it is clear that, even if Jane is in a position to weakly tell which ‘water’-concept figures in *T*, she is not in a position to determine *a priori* that water exists. Upshot: Boghossian’s argument succeeds only if (*) is rendered with a reading of ‘tell which concept’ stronger than that of ‘weakly tell which concept’.

To this end, let us say that a speaker is in a position to ‘strongly tell which’ concept is expressed by her use of ‘water’ on an occasion *O* when, were she conceptually omniscient regarding the relevant language,¹⁷ then, if asked to explicate her use of ‘water’ on *O*, she is *ipso facto* in a position to produce a knowledgeable explication of that concept. Now, were Jane to be in a position to strongly tell which ‘water’-concept figured in *T*, then arguably she would be in a position to know *a priori* that water exists. I say ‘arguably’, but let us grant the point to Boghossian; even so, Boghossian’s argument faces trouble. For we can gloss the conclusion of the argument from conceptual omniscience as the claim that Jane fails to be in a position to strongly tell which ‘water’-concept figures in *T*. After all, if she were in a position to strongly tell which ‘water’-concept figured in *T*, then, contrary to what I argued above is the case, she would be in a position to apply her conceptual omniscience so as to provide a knowledgeable explication of that concept. Thus it would seem that the argument from conceptual omniscience shows that (*) is false on the stronger reading of ‘tell which concept’.

I have just argued that Boghossian’s argument depends on the stronger reading of (*), but that so construed (*) is undermined by the argument from conceptual omniscience. But there is another consideration that can be brought to bear against Boghossian’s argument. Boghossian alleges that (on the assumption of AI) “if a term expresses a concept in the first place, it must have a non-empty extension.” But his argument for this allegation makes an assumption that no proponent of AI should be willing to grant. The assumption in question concerns what we might call a *non-trivial specification requirement* on attempts to characterize the Dry Earthian’s ‘water’. I have argued elsewhere that this requirement plays a key role in other would-be arguments against AI, and that the anti-individualist can and should reject such a requirement (see Goldberg, 2002). Here I will restrict myself to showing how

Boghossian's argument trades on this requirement, and how the anti-individualist can respond to Boghossian's argument from Dry Earth by rejecting it.

To begin, Boghossian's argument from Dry Earth succeeds only if Boghossian is entitled to require the anti-individualist to provide a non-trivial semantic characterization of the Dry Earthian's use of 'water'. Let us say that a semantic characterization of the concept expressed by a given subject's use of word-form *W* is trivial when either *W*, or a name of the concept *W* expresses, figures irreducibly (whether used or mentioned) in the characterization itself. As examples of a trivial characterization, consider the following two English sentences:

- (T1) The concept expressed by the English word-form 'dog' = DOG.
- (T2) The concept expressed by the English word-form 'dog' is true of all and only dogs.

(T1) is trivial because it characterizes the concept by simply naming it ('DOG'), whereas (T2) is trivial because it uses the very word that expresses that concept in characterizing that concept's extension. Now I want to suggest that Boghossian's argument depends on assuming that only non-trivial characterizations are admissible, since without being forced to answer to such a requirement, the anti-individualist can characterize the Dry Earthian's use of 'water' as follows. Dry Earthians take their word-form 'water' to apply to what they take to be a clear, tasteless, colorless liquid flowing in their rivers and taps; but there is no such thing. For this reason their word-form 'water' expresses the concept EWATER. EWATER is the (atomic) concept that is just like WATER, TWATER, etc., insofar as it corresponds to the mental state produced in speakers when they are in the presence of what they take to be a clear, tasteless, colorless liquid flowing in their rivers and taps; but EWATER is unlike WATER, TWATER, etc., since, whereas WATER correctly applies to all and only H₂O samples, and TWATER correctly applies to all and only XYZ samples, and, . . . EWATER correctly applies to nothing.

Such a characterization is a trivial characterization of the Dry Earthian's 'water'-concept: like (T1) above, the proposed characterization of the Dry Earthian's 'water'-concept character-

izes the concept in question in such a way that the concept-name ('EWATER') figures irreducibly in the characterization. At the same time, however, this characterization has a number of virtues. First, it succeeds in accommodating the point which Boghossian thinks the anti-individualist must accommodate – namely, that the Dry Earthian word-form 'water' expresses an atomic concept. Second, it makes sense of the Dry Earthian's perspective: it characterizes the conditions that prompt Dry Earthians to enter into mental states involving the relevant concept, *without* implying that when Dry Earthians do enter into such mental states they are employing *what they themselves know to be* an empty concept. Further, the proposed characterization preserves the idea that Dry Earthians have basic self-knowledge of their EWATER-thoughts (they can self-ascribe such thoughts etc.). But perhaps the most important virtue is that the proposed characterization makes clear that Boghossian's argument from Dry Earth is susceptible to a reply from conceptual omniscience. For, given a conceptually omniscient subject who thinks a 'water'-thought in such a way as to satisfy the referential condition, she is not in a position to determine *a priori* whether she is thinking a WATER-thought (as opposed to a TWATER-thought or a . . . or an EWATER-thought) unless she knows a bit about her embedding history. At any rate, given her conceptual omniscience, her inability to provide a knowledgeable explication of the 'water'-concept she expressed is not to be chalked up to any incomplete grasp or false beliefs regarding the application conditions of her concepts. Of course, if she is not in a position to determine *a priori* whether she is thinking a WATER-thought (as opposed to a TWATER-thought or a . . . or an EWATER-thought), then she is not in a position to determine *a priori* that water exists.¹⁸ And so I conclude that, absent a reason to insist on the non-triviality requirement, Boghossian's argument from Dry Earth can be met.

I have presented two objections to Boghossian's McKinsey-style argument against AI. However, a defender of Boghossian might respond as follows. Even if one or both of the preceding objections succeeds, it is a limited success, since both objections address only those cases in which the subject uses 'water' so as to satisfy the referential requirement (section 2). What of those cases in which the subject uses 'water' in such a way as to fail to satisfy the referential

requirement? In such cases, the subject *is* in a position to strongly tell which ‘water’-concept figures in her thought; so if Boghossian’s argument can be run on such a case, the objections above will be simply irrelevant. The trouble (I submit) is that, when a subject uses ‘water’ in such a way that the referential requirement is not satisfied, the subject is not in a position to learn all that much (if anything) about the world merely on the basis of reflection on her ‘water’-thought. This is precisely because her semantic intention regarding ‘water’ is such that the ‘water’-concept expressed would be available to be expressed in a wide range of relevantly different worlds (watery worlds, twatery worlds, dry worlds, etc.). Such cases are not of much help to those who want to pursue McKinsey-style arguments, where the desired conclusion is that the subject can know *a priori* that e.g. water exists. And so I conclude that no help is forthcoming by appeal to cases in which the referential condition is not met. With this point out of the way we can conclude, tentatively, that the McKinsey-style argument of Boghossian (1998) is flawed.

5.

It is now time to determine the generality of the argument I have offered, and to respond to what I see as the main objection that a proponent of a McKinsey-style argument would have to the case I have presented so far.

To begin, I take it for granted that any language rich enough to be a natural language is *ipso facto* rich enough to enable the formulation of non-standard hypotheses regarding the nature of one’s surrounding environment (including one’s linguistic community). For each non-standard hypothesis we imagine, there is a corresponding switching scenario on which, though a subject started out on Earth, she was subsequently switched to another planet whose relevant features are given by that non-standard hypothesis.

A subject who (like Jane) *has* been switched is a subject regarding whom there are at least two possible interpretations of her use of the term in question.¹⁹ (By ‘possible’ I mean ‘epistemically possible to the subject’.) For this reason, once such a subject is *apprized* of her switching, she is not in a position to strongly tell which concept, of the two or more possible concepts that might be

expressed by the term, is actually expressed by the term on that occasion of use. It does not follow, of course, that she does not have basic self-knowledge of the thought in which that concept figures. Rather, the point is that her basic self-knowledge of the thought in which that concept figures does not by itself provide her with sufficient materials to strongly tell which concept figures in the thought.

What should be said of a subject who has not actually been switched? Suppose that we have a subject who, though never having actually been switched, is *wondering whether*, or *entertaining the possibility that*, she has been switched. Such a subject is in a situation analogous to that of the actually-switched subject. This is because the mere entertaining of the thought that one oneself has been switched changes one's *cognitive* environment: to entertain the possibility of having been switched is to entertain the possibility that one of your terms – for concreteness, say 'water' – has an interpretation other than the one that it would have on the planet you (in your non-skeptical moments) have standardly taken yourself to reside on. And so, once a subject is entertaining the possibility of having been switched, she is not in a position to strongly tell, merely on the basis of her *a priori* basic self-knowledge of her 'water'-thought, which 'water'-concept figures in that thought. For that she will need knowledge regarding her embedding conditions – knowledge that is acquirable only in an *a posteriori* way.

But still one can wonder about non-switched subjects who are *not* presently wondering, and perhaps who have *never* wondered, about the possibility of being switched.²⁰ To address how such subjects' cognitive and epistemic perspective ought to be characterized, we need to appreciate the point of last paragraph's appeal to non-switched subjects who are entertaining skeptical hypotheses. The point of this appeal was not to establish something peculiar to such subjects. Rather, it was to establish a claim regarding the nature of the sort of grasp *all* subjects can be said to have of their own thoughts. The claim is that, on the assumption of AI, the grasp that a subject has of her own occurrent thoughts (and which is present in her judgements of the form 'I am presently thinking that *p*') is not sufficient for the task to which McKinsey-style arguments assign to it. In particular, this grasp is not sufficient to provide an *a priori* warrant for the subject's reliance on any inference whose legitimacy

depends on the subject's ability to discriminate the concepts that figure in the thought that p , from other concepts that are (or would be) subjectively indistinguishable to the subject herself.²¹ In the following section I will make clear how McKinsey-style arguments do depend on such inferences. My present point is simply that the appeal to subjects who are actually entertaining world-switching possibilities was meant to establish a general point about the nature of a subject's grasp – *any* subject's grasp – of her own occurrent thoughts. It matters not at all whether the subject is actively entertaining world-switching possibilities.

6.

It is time to make clear the basis for my central contention, to the effect that the grasp we have of our own occurrent thoughts cannot play the role that McKinsey-style arguments would assign to it. This contention is best presented as emerging from a certain dialectical context. Presented with a McKinsey-style argument, the proponent of AI responds by citing the possibility, which she claims is not ruled out by her *a priori* basic self-knowledge of her thought, that she resides on a world in which the relevant term has no reference. But to this the proponent of the McKinsey-style argument has a ready reply: for any anti-individualistic concept C figuring in a thought that p , where ' E ' names the external object or property on whose existence the individuation of C depends, there is a *necessary, conceptual and/or metaphysical* relation between the following two claims:

- (1) S is thinking that p
- (2) E exists.

The McKinsian reply continues by noting that FPA ensures that, whenever you yourself instantiate (1), you have (or are in a position to have) *a priori* knowledge to the effect that you are thinking that p . The result is that, whenever you yourself instantiate (1), you are in a position to know (2) *a priori*. And so it would seem that any proponent of AI who hopes to challenge McKinsey-style arguments is in a bind.

Now I submit that it is at precisely this point in the dialectic that proponents of AI can appeal to world-switching stories involving

a conceptually omniscient subject. For the upshot of the story of Jane, I submit, is that, even if Jane is in a position to know *a priori* that she herself is thinking that *p* (where ‘that *p*’ designates the propositional content of her thought), she is not in a position to know *a priori* that *E* exists, for not being in a position to know *a priori* that the thought that *p* involves *C* (as opposed to some other subjectively-indistinguishable concept). It is worthwhile revisiting how switching stories are supposed to show this. Jane’s story, like switching stories generally, involves a certain intuition-eliciting comparison. We imagine Jane, now with one past history (years on earth, with only one uneventful five-minute stay on Twin Earth), now with another history (born on Earth, switched to Twin Earth in young infancy where she has remained for the rest of her life). On the basis of this, we are invited to conclude that, were she to have formed a ‘water’-thought intending ‘water’ to apply to the watery liquid on the planet on which she happens to reside, then, absent the relevant empirical information, she would not be in a position to strongly tell which ‘water’-concept figured in her ‘water’-thought. This conclusion reflects what would appear to be a strong intuition: even if Jane were conceptually omniscient regarding both WATER and TWATER, if asked to explicate the ‘water’-concept figuring in this ‘water’-thought *she would be at a loss* as to which entry – that for WATER, or that for TWATER – to use. The intuitive explanation for this is that she is in need of certain empirical information – information regarding her embedding condition and history – prior to being in a position to tell which entry is the one to be used. The take-home point, of course, is that even though there is a necessary, conceptual and/or metaphysical connection between the relevant instances of (1) and (2), Jane is not in a position to know *a priori* what the relevant instance of (2) is, precisely because she is not in a position to strongly tell which ‘water’-concept figures in her thought.²² Stronger still: reflection on a switching story involving the case of Dry Earth suggests that Jane is not even in a position to tell *a priori* whether there even is a claim of form (2) that is (conceptually or metaphysically) entailed by the relevant instance of form (1)! What she *can* know *a priori* are a series of conditional claims, regarding various concepts *C* and external conditions *E*, of the general form:

if my thought that p involves C , then E ;²³

But for any claim of this form, she is not in a position to know *a priori* that the antecedent is satisfied – and this, even though she has *a priori* knowledge that she herself is thinking that p !

In short, switching stories elicit intuitions which make vivid what is wrong with the inference, from a claim asserting one's *a priori* knowledge that one oneself is thinking that water is wet, to a claim asserting one's *a priori* knowledge that water exists. Once we appreciate this point, we are in a position to drop the appeal to switching cases and to ask more straightforwardly whether one is ever in a position to determine *a priori* that e.g. water exists, merely in virtue of *a priori* considerations (including one's *a priori* basic self-knowledge of what in fact is a water-thought). The proponent of the McKinsey-style argument is in the unenviable position of having to answer this in the affirmative. We see just how unenviable the situation is when we postulate a conceptually-omniscient subject. For, if the argument so far can be trusted, not even possession of all of the relevant conceptual knowledge will enable the subject to reach the conclusion that, according to McKinsey-style arguments, she is supposed to be able to reach. The appeal to switching stories thus has the effect of eliciting intuitions that enable us to appreciate just how insubstantial basic self-knowledge actually is.

In closing this section, it is worth bringing out what I take to be a central irony. While switching stories were originally introduced by people wishing to *raise trouble* for AI (on the score of self-knowledge), it turns out that proponents of AI should *embrace* switching stories, since such stories produce intuitions that are useful in the battle against McKinsey-style arguments against AI.

7.

In this paper I have made heavy use of switching thought experiments involving a subject imagined to be conceptually omniscient. My main ambition in doing so has been two-fold. First, I claim to have characterized a kind of (for lack of a better way to say it) explicatory knowledge failure that arises on the assumption of AI, which knowledge failure does not have its source in any incomplete grasp or false beliefs regarding the application condition of a given

concept. Second, I argued that the existence of such a knowledge failure should be welcomed by proponents of AI: it squares with everything that compatibilist-minded anti-individualists have said e.g. regarding the compatibility of AI and basic self-knowledge; and it provides a vindication of those proponents of AI who would resist McKinsey-style counter-arguments with the suggestion that (in the context of non-standard hypotheses) empirical knowledge is required to determine which concept one expresses with one's terms.²⁴

NOTES

¹ In this paper I will be focusing on Burge's anti-individualism, sometimes called Social Externalism (see e.g. Bilgrami (1992) for a characterization of the distinction between Social and Non-Social Externalism). This is important, since the arguments that I will be considering bear on Burge's anti-individualist doctrine, but would not bear against the (non-social) version of externalism e.g. in Bilgrami (1992). This is because the argument on offer presupposes something that is accepted by Burge and his followers, but which is denied by Bilgrami's externalism – namely, that there could be two content-distinct thoughts that are nevertheless *subjectively indiscriminable*. The present paper is meant to argue, *first*, that (given Burge's anti-individualism) the doctrine of content-distinct-yet-subjectively-indiscriminable thoughts can be motivated without recourse to (i) incomplete grasp of one's concepts, (ii) agnosticism regarding the application conditions of one's concepts, or (iii) false empirical theorizing; but, *second*, that the implications of this doctrine (in contexts regarding self-knowledge and skepticism) are not as dire as some critics of Burge's anti-individualism suppose. It is worth pointing out that, while the present paper aims to defend Burge's anti-individualism, the first of these claims should be of interest even to *opponents* of Burge: in effect I am arguing that (given Burge's anti-individualism) the phenomenon of content-distinct-yet-subjectively-indiscriminable thoughts cuts more deeply (for not depending on (i)–(iii)) than Burge or his followers appear to have recognized. (For Burge's own views about this phenomenon, see his 1979, 1982, 1986, 1988 and 1989.)

² Some of it is collected in Ludlow and Martin (1998) and Nuccetelli (2003).

³ See also Goldberg (2003b) for an alternative reaction to McKinsey-style arguments. And for a view regarding how the anti-individualist should react to the argument from world-switching, see Goldberg (2003a).

⁴ This is worth making explicit, since to date anti-individualists have identified and discussed only those explicational failures whose source is an incomplete grasp or false beliefs about the application conditions of concepts; see Burge, 1979, 1986, 1989.

⁵ In this way I take myself to be answering a challenge posed by Boghossian (1998). Describing the attitude of externalists towards ‘Twin Earth-eligible’ words, or words on which a Twin Earth experiment can be run, Boghossian writes that “although externalists are very fond of saying that . . . our knowledge that a given word is Twin Earth-eligible rest[s] on empirical information, . . . it is rare to find their reasons explicitly spelled out. Where exactly do empirical elements intrude into the Twin Earth experiment?” (p. 277). The present paper purports to answer this latter question.

⁶ I thank an anonymous referee for pointing out the need to make this latter case explicit.

⁷ As I will argue below, this last point makes it clear that a speaker can use a word such as ‘water’ in such a way as to satisfy the referential condition, and so be in a position on which we can run the present argument, even if she knows that water is H₂O. This is a novelty of the present world-switching argument over previous incarnations, which had always assumed that the world-switched subject does not fully understand the relevant concept. (I thank an anonymous referee for indicating the need for this comment.)

⁸ Nothing hangs on whether this is called a knowledge failure (as opposed to an inability to apply one’s conceptual mastery to a given case at hand); but I find it helpful to talk this way.

⁹ Roughly, to have basic self-knowledge of one’s thought that *p* is to know that one oneself is thinking that *p*. See Burge (1988, p. 649) for a more detailed discussion of what types of state are such that one can have basic self-knowledge of one’s being in that state.

¹⁰ They put the point in terms of the subject’s knowledge of *the content* of the thought, but this difference does not matter to me here.

¹¹ For three such arguments, see McKinsey, 1991; Brown, 1995; Boghossian, 1998; for replies, see Brueckner, 1992 (a response to McKinsey, 1991); Miller, 1997 (a response to Brown, 1995); and McLaughlin and Tye, 1998 (a response to Boghossian, 1998).

¹² See e.g. Warfield (1994, 1998), where it is urged that we accept the anti-skeptical implications of AI.

¹³ For a distinct sort of reply to the argument of Boghossian (1998), see Goldberg (2003b).

¹⁴ Other versions put the point in terms of a disjunction – either water exists and I have been in contact with it, or else I have been in contact with a linguistic community (see e.g. Brown, 1995.) However, Boghossian 1998 is prepared to put this subtlety to the side (Boghossian, 1989, p. 276). In any case, I think that the present argument can be expanded so as to apply cases in which the would-be derived claim is disjunctive – one need merely imagine that the switching story involves switching to a world at which it *appears* that there are other speakers, but in fact there are no such speakers at all. (Characterizing the ‘water’-concept expressed at such a world will have to take this non-existence of others into account.) Considerations of space prevent me from developing this point here.

¹⁵ Boghossian's way of putting this is to speak of the word as 'Twin-Earth eligible'.

¹⁶ Some have suggested that in world-switching cases, once the subject is apprized of her earlier world-switching, she is no longer in a position to express, hence to form thoughts involving, the concept expressed prior to her having been so apprized; see Gibbons, 1996, p. 310. Though I find this dubious in principle (see Goldberg, 2000, 2003a), this matter is irrelevant to my present argument. I am trying to show that, given AI, a speaker in a switching story may (in a soon-to-be-articulated sense) fail to know which concept figures in her thought. If Gibbons' analysis is accepted, then this result falls out trivially: if you are not in a position to express that very concept, then you cannot form any judgement to the effect that that very concept figured in your thought. My present point is that, even if Gibbons is *wrong* and the thinker *retains* her ability to deploy the concept in question, she still fails to know which concept she expressed in the sense relevant to Boghossian's argument.

¹⁷ The relevant language is the language whose lexicon contains the word-form type of which this use of 'water' is a token.

¹⁸ Indeed, she is not in a position to determine *a priori* something much weaker, i.e., that an external world exists. For if (for all the subject herself knows) her 'water'-thought at *t* was an EWATER-thought, then (for all she knows) her 'water'-thought at *t* did not depend for its individuation on any 'external' fact *f*. This point is important, if only because it closes off another possible reaction on the part of a friend of the McKinsey-style argument. The possible reaction in question would be to treat the forgoing argument as having shown that the subject is in a position to acquire *a priori* knowledge of a *disjunctive* proposition, to the effect that either water (= H₂O) or twater (= XYZ) exists. Against such a reaction, the forgoing suggests that among the disjuncts one would have to include the condition corresponding to EWATER, which is to say, the condition of there being nothing that answers to the concept in question. This result depends merely on noting the situation in which our conceptually-omniscient world-switched subject finds herself; it does not depend in any way on the non-triviality requirement presently at issue. I will return to this issue in section 6. (I thank an anonymous reviewer for pointing out the need for this comment.)

¹⁹ I say 'at least' two. Take the case of 'water' involving Earth and Twin Earth. Depending on the details of the story – i.e., the length of time spent by the subject on either planet, the subject's semantic intentions regarding the relevant use of 'water', etc. – the relevant 'water'-concept might be: WATER, TWATER, or WATER-OR-TWATER. Note too that if the subject's intentions do not satisfy the referential condition from section 2 above, the 'water'-concept could be WET, TASTELESS, DRINKABLE LIQUID COMING OUT OF THE TAP. I disregard such cases for the reason given at the end of section 4.

²⁰ I thank an anonymous reviewer for indicating the need to address this kind of case.

²¹ I want to re-emphasize the point that I am *not* insisting on discriminability as a condition on basic self-knowledge. (Quite to the contrary, see section 3.) Rather,

I am insisting on discriminability as a condition on the *a priori* knowability of the legitimacy of certain inferences. The basis for this contention is laid out in section 6.

²² Compare Goldberg (2003b).

²³ In order to keep this point general, we can say that for $C = \text{EWATER}$, ' E ' designates some necessary proposition (e.g. ' $1 + 1 = 2$ ').

²⁴ I want to thank an anonymous reviewer of this journal for very helpful suggestions.

REFERENCES

- Bilgrami (1992): *Belief and Meaning*, Oxford: Basil Blackwell.
- Boghossian (1998): 'What the Externalist Can Know *A Priori*', in Wright, Smith and Macdonald (eds.), *Knowing Our Own Minds*, Oxford: Oxford University Press.
- Brown (1995): 'The Incompatibility of Anti-Individualism and Privileged Access', *Analysis* 55.
- Brueckner (1992): 'What an Anti-Individualist Knows *A Priori*', *Analysis* 52.
- Brueckner (2000): 'Ambiguity and Knowledge of Content', *Analysis* 60, 257–260.
- Burge (1979): 'Individualism and the Mental', Reprinted in D. Rosenthal (ed.), *The Nature of Mind*, Oxford: OUP, 1991.
- Burge (1982): 'Other Bodies', in A. Woodfield (ed.), *Thought and Object: Essays on Intentionality* (pp. 97–120), Oxford: Clarendon Press.
- Burge (1986): 'Intellectual Norms and the Foundations of Mind', *Journal of Philosophy* 83.
- Burge (1988): 'Individualism and Self-Knowledge', *Journal of Philosophy* 85.
- Burge (1989): 'Wherein is Language Social?', in *Reflections on Chomsky*, Oxford: Basil Blackwell, 1989.
- Falvey and Owens (1994): 'Externalism, Self-Knowledge, and Skepticism', *Philosophical Review* 103(1).
- Gibbons (1996): 'Externalism and Knowledge of Content', *Philosophical Review* 105(3).
- Goldberg (1999): 'Self-Ascription, Self-Knowledge, and the Memory Argument', *Analysis* 57(3), 211–219.
- Goldberg (2000): 'Externalism and Self-Knowledge of Content: A New Incompatibilist Strategy', *Philosophical Studies* 100(1), 51–78.
- Goldberg (2002): 'Do Anti-Individualistic Construals of the Attitudes Capture the Agent's Conceptions?', *Noûs* 36(4), 597–621.
- Goldberg (2003a): 'What Do You Know When You Know Your Own Thoughts?', in Nuccetelli (ed.), *New Essays on Semantic Externalism and Self-Knowledge*, Cambridge: MIT Press.
- Goldberg (2003b): 'On Our Alleged *A Priori* Knowledge that Water Exists', *Analysis* 63(1), 38–41.

- Ludlow and Martin (1998): *Externalism and Self-Knowledge*, Palo Alto: CSLI Press.
- McKinsey (1991): 'Anti-Individualism and Privileged Access', *Analysis* 51.
- McLaughlin and Tye (1998): 'Externalism, Twin Earth, and Self-Knowledge', in Wright, Smith and Macdonald (eds.), *Knowing Our Own Minds*, Oxford: Oxford University Press.
- Miller (1997): 'Externalist Self-Knowledge and the Scope of the *A Priori*', *Analysis* 57.
- Nuccetelli (2003): *New Essays on Semantic Externalism and Self-Knowledge*, Cambridge: MIT Press.
- Warfield (1994): 'Knowing the World and Knowing Our Minds', *Philosophy and Phenomenological Research* 55.
- Warfield (1998): 'A *A Priori* Knowledge of the World: Knowing the World by Knowing Our Minds', *Philosophical Studies*.

Department of Philosophy
University of Kentucky
1427 Patterson Office Tower
Lexington, KY 40506-0027
USA