
THE RELEVANCE OF DISCRIMINATORY KNOWLEDGE OF CONTENT¹

BY

SANFORD C. GOLDBERG

Abstract: Those interested in securing the compatibility of anti-individualism and introspective knowledge of content (henceforth 'compatibilists') typically make a distinction between knowledge of content proper (KC) and discriminatory knowledge of content (DKC). Following Falvey and Owens (1994), most compatibilists allow that anti-individualism is not compatible with introspective DKC, but maintain that nonetheless anti-individualism *is* compatible with introspective KC. Though I have raised doubts about the compatibility of anti-individualism and introspective KC elsewhere (Goldberg, 1997 and forthcoming), here my aim is to suggest the philosophical relevance of DKC itself. My thesis is that there are cases in which a thinker's failing to have DKC will affect the justification which she takes herself to have in drawing various inferences in the course of her reasoning, and so will affect that reasoning itself. After presenting illustrative examples and suggesting why anti-individualists themselves ought to acknowledge this point, I suggest that the examples indicate further work for anti-individualists: formulating what it takes to have DKC, and substantiating the view (widely held by anti-individualists) that anti-individualism's implication that we (often) lack such knowledge is not to be taken as an important weakness of anti-individualism itself.

1. Background

Anti-individualism is the thesis that some contents cannot be individuated in terms of properties of the individual considered in isolation from her social and physical environment. The doctrine of authoritative self-knowledge of content is the view that, for any thinker *S* and occurrent thought that *p*, *S* has introspective knowledge of the content of her thought that

p. Many philosophers wonder about the compatibility of these two doctrines. The initial worry arises when we realize that, if anti-individualism is true, then the content of one's thoughts will depend on features of one's surrounding (social and physical) environment about which one may be ignorant. Burge (1988) attempted to undermine these worries by arguing (correctly, in my view) that a thinker's capacity to form correct, justified self-ascriptive judgements about her own occurrent thoughts derives from the self-referential character of those present-tense judgements. If correct, this would mean that no doctrine – and so by extension no anti-individualist doctrine – can undermine a thinker's capacity to form correct, justified self-ascriptive judgements about her own occurrent thoughts. Incompatibilist worries threaten to resurface, however, when we consider whether *correct, justified self-ascription of a content* can be conflated with *knowing the content one is presently self-ascribing*. Suppose (what at least Burge appears ready to acknowledge)² that knowledge of content requires knowing *what* content one is thinking.³ Then incompatibilist worries would appear to be exacerbated by the fact that most anti-individualists concede, as an implication of their formulation of anti-individualism itself, that:

(*) A thinker S may fail to be able to discriminate her actual thought from other (counterfactual) thoughts she would have had if she had acquired her concepts in a different environment.

Now (*) may appear to jeopardize the doctrine of authoritative self-knowledge, and so lend ammunition to the incompatibilist, insofar as *failing to be able to discriminate the thought expressed by a form of words W from thoughts W would have expressed in other contexts* might be taken as *failing to know what content one expressed with W*. This point would show that one can self-ascribe a thought with *W* yet fail to know the content of the thought one self-ascribed with those very words (for failing to know *what* thought *W* expresses).

It is in this context that the distinction between introspective knowledge of content (KC) and introspective discriminatory knowledge of content (DKC) has its point. Anti-individualists who concede (*) as an implication of their version of anti-individualism insist on this distinction as part of a strategy whose point is to deny that (*) opens up the possibility that S fails to have introspective KC. Rather, on the strength of the view that Burge's point about the self-referential nature of present-tense self-ascriptions still stands, such anti-individualists hold that (*) merely entails the possibility of a failure of introspective DKC.

I think that it is fair to say that this is where the debate stands at present. While incompatibilists continue to formulate thought experiments which by their lights present examples in which a failure of introspective DKC

is *ipso facto* a failure of introspective KC,⁴ compatibilists continue to insist that Burge's point about the self-referential nature of present-tense self-ascriptions inoculates the compatibilist position against any incompatibility claims regarding introspective KC.⁵ Here I want to suggest a way in which we might advance this debate. My strategy shall be to avoid making contentious claims regarding the compatibility of anti-individualism with introspective KC, and to focus instead on the insufficiently examined relevance of introspective DKC. Doing so will enable me to formulate and defend a thesis that I think all sides should agree on: no matter what one's views on the issue of compatibilism are, one ought to accept that anti-individualism has some interesting implications, deriving from anti-individualism's implications on the score of introspective DKC, regarding (practical and theoretical) reasoning. While this thesis bears some similarity to the thesis of Boghossian (1992a), I will present and defend it in a way that does not leave it open to the objections leveled against his argument.

2. *The Relevance of Introspective DKC: the example of Sally*

In this section I aim to present a case in which a thinker's lack of DKC affects the justification which she takes herself to have⁶ in carrying out her (practical and theoretical) reasoning, and so affects that reasoning itself.

2.1

Consider the case of Sally. Throughout the early part of her life Sally is reared in a community (we will suppose) whose members use the term 'pragmatic' and its derivatives in a univocal way, as meaning any attitude which incorporates a realistic view about how to achieve one's goals. Thus among Sally's linguistic peers we hear such varied utterances as "I admire X's pragmatic nature; it enables her to get things done," but also "Y's scheming pragmatism makes me unable to trust him." Over time in this linguistic community, Sally herself acquires (from talking with her peers) various beliefs in whose expression the term 'pragmatic' figures. And, while she herself is agnostic on the application conditions of this term, nonetheless it is reasonable to treat her as expressing the *very same concept* that her linguistic peers express with this word.⁷

I say that this attribution is reasonable: we can imagine various reasons which might support it. For one, we can suppose that Sally herself would defer to community norms regarding the proper explication of this term: she holds her own usage (and the partial explications she would offer for this usage) answerable to those norms. We might further suppose (since this is a thought-experiment) that there *are* such norms, that is, that it is

not the case that her community is one in which the term 'pragmatist' is used in very different ways, according to very different norms. For another, we can imagine cases in which Sally would describe herself as 'believing the same thing' as her linguistic peers, with the result that, to the extent that her peers express one common concept with the term 'pragmatist,' so too does Sally. So it is reasonable to treat Sally as expressing the same concept as her linguistic peers, despite her agnosticism regarding the application conditions of the concept.

Now, some time later, as an undergraduate at college, Sally takes a class in the Philosophy Department, during which she hears her professors speak about 'pragmatists' and 'pragmatism.' She hears the opinionated Philosophy faculty say such things as "Pragmatists are the most honest intellectuals," and "Pierce was the most far-seeing of the pragmatists," but also "Pragmatism represents intellectual laziness of the worst sort," and "No good philosophy has ever come out of the pragmatic tradition." Once again, on the basis of hearing her professors Sally comes to acquire beliefs in whose expression the term 'pragmatist' would figure. And once again she herself is agnostic on the application conditions of 'pragmatist,' deferring such matters to her professors.

Now, on the assumptions (first) that the sense of 'pragmatist' as used in the context of philosophy differs from the general-usage sense of 'pragmatist', and (second) that anti-individualism is true, it would appear that Sally actually possesses two distinct concepts which she expresses with one and the same word-form 'pragmatist.' I will not argue for this two-concept conclusion, but rather I will assume it. After all, that a thinker can come to acquire a second concept and associate it with a word-form which also expresses another (to her, indistinguishable)⁸ concept, and do so under conditions in which she herself may be unaware of having acquired a second concept, is a claim that has been identified and endorsed by many people, as the most plausible way to understand world-switching cases from an anti-individualist perspective (see Burge, 1988 and 1998; Boghossian 1989 and 1992a; Ludlow 1995b; Gibbons 1996). So I will be assuming that there are ways of spelling out Sally's story consistent with the claim that Sally has two distinct but subjectively indistinguishable concepts associated with the term 'pragmatist.'⁹ (Below, in section 2.4, I will suggest that my argument does not depend on Sally's having two *subjectively indistinguishable* concepts, but rather depends on a weaker description of Sally's use of 'pragmatist.')¹⁰

Having assumed, then, that the two-concept thesis is true of Sally, I will be turning my attention to those anti-individualists who, having accepted the two-concept claim regarding Sally's use of 'pragmatist,' also hold that none of what Sally fails to know regarding the two concepts associated with the term 'pragmatist' in her idiolect impugns the knowledge that she has of any of the thought- and belief-contents in whose

expression she uses 'pragmatist.' That is, I will be focusing on *compatibilist* anti-individualists who concede that the two-concept claim is true of Sally. I want to suggest that such anti-individualists must acknowledge the relevance of discriminatory knowledge of content to (practical and theoretical) reasoning.

2.2

To see this, we need to spell out Sally's story in some more detail. Suppose that among the beliefs which she picked up from her pre-college days (that is, from her original linguistic environment) is the belief which she originally expressed, and would continue to express, with the sentence

- (1) "No pragmatists have any principles."

Suppose further that nothing that she learned in college about 'pragmatists' and 'pragmatism' either supported or inquired this belief. To make this plausible we will suppose that, because she was not a philosophy major (she only took two philosophy courses over her college career), she never realized that 'pragmatist' differs in extension according (roughly) to whether the term was used by members of the Philosophy faculty or her earlier linguistic peers; and we will also assume that she holds the mistaken view that the philosophers in their use of 'pragmatist' and the common folk in their usage of 'pragmatist' both defer to the same experts regarding the proper explication of the term. Finally, we will suppose that it was not until years after college that she was told (by a friend in the know) that there is a world of difference between 'pragmatist' in the philosophical sense and in the general sense.

The question on which I want to focus is this. How does Sally's new knowledge about her own term 'pragmatist' – knowledge that she herself associates two concepts with this one world-form, and does so without having any explicational knowledge that would enable her to distinguish the one from the other concept – affect her cognitive situation? Incompatibilists have long sought to show that in cases like this the thinker's failure of discriminatory knowledge should be cast as a failure of knowledge of the thought-contents in whose expression the problematic word figures. In the present context such a conclusion is contentious. Here I want to draw what I hope will strike compatibilists as a more motivated conclusion. The conclusion is that, *whether or not* Sally is to count as 'knowing the content' of her 'pragmatist'-thoughts despite her lack of discriminatory and explicational knowledge, what Sally fails to know will potentially affect her views concerning the justification with which she can draw various inferences in the course of her (practical and theoretical) reasoning, and so will potentially affect her reasoning itself.

Here is one such case. We remarked that Sally acquired a belief which she would express with (1). Let us suppose that she remembers this belief in a content-preserving sort of way,¹¹ but does not remember whether she acquired it in her pre-college linguistic community, or in the Philosophy Department.¹² Now suppose that, after having come to believe that the two-concept claim is true of her use of 'pragmatist', she learns that Jones is a pragmatist in the philosophical sense of the word, and, remembering the belief that she expressed with (1),¹³ she wonders whether she should conclude that Jones has no principles. Well, consider that the following two sentences express two beliefs that she (knowingly) counts herself as having:

- (1) "No pragmatists have any principles."
- (2) "Jones is a pragmatist."

Given that these two sentences express *her own* beliefs, and that she 'knows' that she has these beliefs in the sense that she would self-ascribe the beliefs in question using these very words, it would appear that she should know, *in advance of having to do any research into the content of these beliefs*, whether the contents of these two beliefs entailed¹⁴ a third content, namely, the content that

- (3) Jones has no principles.

Yet it seems pretty clear that, given what Sally does not know, she is not in a position to determine whether (3) follows from the belief-contents expressed by (1) and (2) – at least she is not in such a position *prior to doing some research* into the provenance of the belief-content she expresses with (1) (to determine whether the word-form 'pragmatist' as it occurs in (1) expresses the same concept as it expresses in (2)).

2.3

The preceding argument assumes that Sally can have a memory that preserves the content of the belief she had originally expressed with (1), even in situations in which she no longer remembers in which of two contexts she acquired the belief. Call this 'the assumption.' Perhaps someone might think to challenge the assumption, as follows. Call the belief, which Sally expressed when she originally uttered (1) at the time of acquiring the belief she is presently trying to recall, B_{original} ; and call the concept expressed by 'pragmatist' in B_{original} , c_{original} . Since Sally does not know whether B_{original} was acquired in her pre-college linguistic community or during college, Sally does not know whether c_{original} is the philosophical concept or the general concept. But (the objection proceeds) since she

does not know which concept is the concept that figures in B_{original} , the result is that, even if she presently utters (1) claiming to recall B_{original} , she would not count as expressing B_{original} : at most she should be credited with the belief that *whatever she meant when she uttered (1) in the expression of B_{original} is true*. That is, the proposal on the table is that, given what she knows (that the two-concept hypothesis is true of her) and what she fails to know (in which community she acquired B_{original}), we must provide a *metalinguistic construal* of the belief she expresses when she presently utters (1).

It would be unprincipled, however, for any anti-individualist who is a compatibilist to reject the assumption in this way. To see this, suppose Sally claims to recall B_{original} , and does so *prior to* finding out about the difference between the philosophical and the general sense of 'pragmatic.' In such a case, I submit, there is no good reason to deny that the principle of content-preservation applies – that is, no good reason to deny that she recalls the very belief-content of B_{original} (the belief-content she had expressed when she originally uttered (1), in her pre-college days).¹⁵ Now, to be sure, after she finds out about the two senses of 'pragmatic,' she is in a position to wonder which of the two senses she had in mind (which of the two concepts she expressed by uttering 'pragmatic'). But – and this is the crucial part – the *very same point* could be made in present-tense self-ascriptions, where compatibilist anti-individualists have *always* rejected the idea that the mere failure to be able to distinguish one's content from another content undermines one's ability to think that content. On the contrary, in such situations, compatibilists have always appealed to the distinction between KC and DKC to suggest that (not only can one think a content without being able to distinguish it from another, subjectively indistinguishable content, but what is more) one can have *introspective KC* of that content without being able to distinguish that content from other subjectively indistinguishable contents! Let me explain the parallel here at greater length, in order to undermine any lingering temptation by compatibilist anti-individualists to reject the assumption in the metalinguistic manner indicated.

To begin, the compatibilist position has all along insisted that, in the context of *present-tense* self-ascriptions of thought, a thinker's inability to distinguish between the thought-content that p from other thought-contents does not undermine her ability to think that p 'knowingly.' It is not difficult to understand why the compatibilist position has insisted on this. Typical incompatibilist arguments (Boghossian, 1989; Brueckner, 1990; Butler, 1997; Goldberg, 1997; Goldberg (forthcoming)) proceed by establishing that a thinker can think an occurrent thought, and can even self-ascribe that occurrent thought, but fail to be able to discriminate her actual thought-content from other thought-contents – thought-contents she *would* have been thinking had she acquired her concepts in a relevantly-different

environment. Against this, the traditional compatibilist response (Heil 1987; Burge 1988; Falvey and Owens 1994; Gibbons 1996; Brueckner (forthcoming))¹⁶ has always been that this point is *true but irrelevant* to the issue of the thinker's introspective KC. Rather, such compatibilists have suggested, all that the would-be incompatibilist's point shows is that the thinker fails to have introspective DKC. Notice, however, that the compatibilist position here *presupposes* the claim that one's inability to discriminate one's actual thought-content from another distinct thought-content does not undermine one's ability to think that thought-content, let alone undermine one's knowledge (KC) of that thought-content. Call this presupposed claim, the *compatibilist's point*.

Now if the compatibilist's point holds with respect to *present-tense* self-ascriptions, as the compatibilist position has strenuously insisted, then I cannot see why it does not also hold with respect to *past-tense* self-ascriptions, that is, in recollection cases. In recollection cases the compatibilist's point would show that the mere inability of a thinker to distinguish between the thought-content she is attempting to recall from thought-contents she is not attempting to recall, does not undermine the claim that nonetheless the thinker's memory has preserved the very content of the thought she is attempting to recall. Rather, the situation appears to be this: if there are reasons to think that Sally's memory *has* in fact preserved the content of the thought she is attempting to recall (in this case, the content of B_{original}), then Sally's inability to distinguish between the content of B_{original} from other (to her subjectively indistinguishable) contents does not undermine this claim about preservative memory. And we just saw that, had she claimed to recall B_{original} *prior to* finding out about the distinct senses of 'pragmatist,' then she *would* count as recalling B_{original} in a content-preserving way. In short, we see that the compatibilist bent on denying the assumption (that Sally can have a memory that preserves the content of B_{original} , even in situations in which she no longer remembers in which of two contexts she acquired B_{original}) cannot do so by appeal to Sally's discriminatory failure. What is more, since there are reasons to think that Sally's memory has in fact preserved the content of B_{original} (as seen in the case in which she claims to recall this belief prior to finding out about the two senses of 'pragmatist'), the burden of proof is surely on those who would have us reject the assumption.¹⁷

2.4

So far I have merely tried to make the point that, given anti-individualism and a suitably situated (and suitably memory-deprived) thinker, a failure of knowledge regarding the provenance of one's belief can affect the justification which one takes oneself to have (or not to have) regarding the making of certain inferences, and so can affect one's reasoning. I have

suggested that this conclusion can be established on grounds accepted by compatibilist anti-individualists. In this section following this, however, I will consider and reject two ways in which anti-individualists have thought to reply to an argument much like the present argument. But for now I want to consider the *generality* of the conclusion I have derived. My question is this: how often is it the case that a thinker (i) remembers a belief in the content-preserving kind of way, (ii) fails to remember when and under what conditions she acquired the belief, and yet (iii) has reason to think that she associates two distinct yet subjectively indistinguishable concepts with one or more of the words that figure in the expression of the belief in question? If such cases arise very infrequently, then perhaps this alone will make the problem I have identified insignificant, even if it is conceded to be real.

To begin, it should not be contested by anyone that we often do satisfy condition (i). As Burge (1993) makes explicit, memory provides us with a source of belief-contents using which we can construct arguments. Furthermore, for the very reasons that Burge provides there, we should not doubt that there are many cases in which memory *preserves the very content* of the beliefs recalled. Since anti-individualism is perfectly compatible with this point, and since the point strikes me as very plausible on independent grounds, I will simply assume that there are many cases in which we satisfy (i).

Restricting ourselves to those cases where we do satisfy (i), in how many of these cases do we also satisfy (ii), that is, fail to remember the conditions under which we acquired the recalled belief? Again, I think that it is relatively uncontroversial that such cases, if they do not make up the majority of the cases of recollection, will nonetheless make up a significant minority of the cases of recollection. So I will treat this claim as uncontroversial.

The real question, then, concerns the extent to which those thinkers who satisfy both (i) and (ii) simultaneously satisfy (iii). I think that the answer to this question depends in part on how thinly the anti-individualist theory of individuation slices concepts: the more thinly concepts are sliced, the more likely it will be that one has come into contact with two distinct concepts expressed by one and the same word – with the result that one will be correspondingly more likely to encounter some information in support of the possibility that one possesses two distinct yet subjectively indistinguishable concepts associated with that single word-form.

To illustrate this, consider for example the thesis from Ludlow (1995a), to the effect that so-called ‘world-switching’ is much more prevalent than one might suspect. Ludlow began by noting that different linguistic groups often differ over the experts to whom they defer, and experts from different groups often endorse distinct explications of the concepts in question. Suppose that we have a case in which two groups of experts from two distinct communities *do* endorse distinct explications of the concept

expressed by a given word. In such a case, Ludlow argued, any individual who moves back and forth between these communities has in fact undergone a case of 'world-switching.' The result is that world-switching is much more 'prevalent' than many have supposed, with the result that cases of acquiring information attesting to one's own previous world-switching would appear to be correspondingly more numerous.¹⁸

Two final comments on the satisfaction of condition (iii) are in order. First, Sally's having two *subjectively indistinguishable* concepts (in the sense indicated in note 8 above) is not the *only* way in which the two-concept thesis can be true of Sally, in such a way that Sally lacks some kind of (reasoning-relevant) discriminatory knowledge regarding her concepts. For Sally might (a) know that she has two concepts associated with one and the same term 'pragmatist,' (b) be able to provide clear explications of each of these concepts, so as to capture the way in which these concepts differ, and yet (c) fail to know, on a given occasion on which she uses 'pragmatist,' *which* is the concept she has expressed (where this 'failing to know which' is a matter of not knowing which of the two concepts she knows she has, and which she can explicate perfectly well, is the concept she is presently expressing).¹⁹ Indeed, this much is obvious from the points I made above concerning memory: all that is required, in order for Sally to satisfy condition (c), is that Sally forget the provenance of the belief she presently expresses with (1). Then, even if Sally *can* exhaustively explicate the difference between the two concepts she associates with 'pragmatist,' still, because she does not know from whom she picked up the belief she expresses with (1), she will not know which of these two concepts (each of which she can explicate perfectly well) is expressed in the belief in question. In such a case, while her two concepts will not be subjectively indistinguishable to her, still, she does not know which of the two concepts figures in the belief she expresses with (1).²⁰

The first point regarding condition (iii), then, is that subjective indistinguishability is not the only way for Sally to lack some (reasoning-relevant) kind of discriminatory knowledge regarding her concepts. This point is very important, for two reasons. First, I believe that the argument I have been offering can be run on *either* of the kinds of 'failure of discriminatory knowledge,' with the result that those who are suspicious of the very possibility of subjectively indistinguishable concepts, or who are skeptical that Sally's case is aptly so described, cannot dismiss my argument on these grounds: my argument does not depend on Sally's having subjectively indistinguishable concepts, but depends rather on Sally's not knowing which concept is expressed in her utterance of (1).²¹ Second, that there is more than one kind of 'failure of discriminatory knowledge' regarding ones concepts only *increase* the relevance of DKC, since it only increases the sorts of case that can be described in terms of a failure of DKC.

In sum, condition (iii) should be revised so as to acknowledge that we can accept the claim that the two-concept thesis is true of Sally, in such

a way as to raise the possibility of a (reasoning-affecting) failure of discriminatory knowledge, even if Sally does *not* lack the relevant explicational knowledge required to distinguish her two concepts. So perhaps (iii) would be more accurately formulated if we said that a thinker satisfies this condition just in case she has reason to think that she associates two distinct concepts with one or more words that figure in the expression of the belief in question, in such a way that she cannot say in any illuminating way²² which of the concepts is the one she had expressed on this particular occasion. The case of having two subjectively indistinguishable concepts would be *one* way for a thinker to satisfy condition (iii), but as we just saw it is not the *only* way.

The second comment I want to make concerning condition (iii) is this. Sally finds herself in a situation in which she is not certain whether what *appears* to be a valid argument (having as its premises belief-contents she knows herself to have) *is* in fact a valid argument. It bears emphasis that one need not have actually been world-switched in order to be in Sally's position. Condition (iii) is explicit about this: one need only *suspect* (or *have reason to believe*) that one has been world-switched, in order to be in a position relevantly like Sally's. Let me illustrate briefly and then explain the significance of this point.

Here is a variant of Sally's story. Suppose that (contrary to the story above) the Philosophy faculty with whom Sally interacts never actually spoke in her presence about philosophical pragmatism (the imperious chair hated any mention of pragmatism in front of impressionable undergraduates), but that from time to time the Philosophy faculty did use the word 'pragmatic' in its sense of 'practical attitude' – that is, in accord with the sense of the term used by members of Sally's pre-college linguistic community. And suppose (as in the earlier example) that at some time long after college a friend apprizes Sally of the distinction between 'pragmatic' in its general usage and in its philosophical usage. Then, it seems to me, Sally is in the same cognitive situation vis-à-vis the inference from (1) and (2)²³ to (3), as she was in the earlier example, despite the fact that in the present example she has not actually been world-switched (even by the lights of Ludlow's liberal account of the conditions under which one is world-switched). In this manner we see that what is crucial to these cases is not that the thinker has *actually* been world-switched, but rather that she has some reason to *think* that she has been world-switched.

Now, if cases such as Sally's can arise even in the absence of actual world-switching (as this variant of Sally's story suggests), then even those anti-individualists who would reject Ludlow's liberal views about the extent of actual world-switching should acknowledge that anti-individualism has some potentially far-reaching implications regarding reasoning. To repeat, the extent of situations like Sally's will go beyond cases of actual world-switching, to include any case where a person comes to suspect that she

may have been world-switched. That *anti-individualism* itself is responsible for generating these types of situations becomes clear, once we see that what grounds Sally's belief in the possibility of a vacillation in her use of 'pragmatist' is *nothing other than her implicit acknowledgement of anti-individualism itself*: it is *because* she acknowledges (in some sense) that she has acquired the concept(s) she expresses with 'pragmatist' from her linguistic community, that her belief that she has been a member of two distinct linguistic communities is *ipso facto* a reason to believe in the possibility of a vacillation with the term as it occurs in two premises of the argument.²⁴ In short, given anti-individualism, any reason to suspect that one's linguistic community has changed in relevant respects²⁵ (or that one oneself has moved across linguistic communities) is *ipso facto* a reason to suspect a vacillation in one's terms, with the result that a failure to remember the provenance of one's beliefs can potentially affect one's reasoning in precisely the way that Sally's reasoning is affected in the example above.

To be sure, an anti-individualist can respond to the foregoing by rejecting (what Ludlow assumes) the thin manner of slicing concepts, and with it reject the whole line of argument starting from Ludlow's thesis asserting the prevalence of world-switching. Since no such anti-individualist theory has yet to be articulated, however, it remains to be seen whether such a view can be independently motivated. Given the state of such discussions, however, it would seem reasonable at present to conclude that the type of situation to which I have pointed is more prevalent – and arguably *much* more prevalent – than one might suspect.

I should point out as well one alternative reaction by anti-individualists. On this reaction, the anti-individualist would simply *accept* this conclusion, and argue that anti-individualism's implications regarding reasoning, as exemplified by the case of Sally, are defensible on independent grounds.²⁶ Indeed, this is a view that I share. However, as I will indicate below, taking such a line requires doing more philosophical work: providing such independent arguments (for one of which see Owens 1990), and giving more philosophical attention to the notion of discriminatory knowledge of content. If the present argument only establishes the nature and importance of the work that remains, it has done its job.

3. *Boghossian on Anti-individualism and Inference*²⁷

So far I have been arguing that, given anti-individualism, there will be cases where a thinker's lack of DKC will affect her (practical and theoretical) reasoning. The case that I presented for this thesis bears some important similarities to a thesis from Boghossian (1992a). While Boghossian's argument has been criticized by Schiffer (1992) and Burge (1998), the claim that I am making is *weaker* than the central claim

Boghossian tried to make, in ways relevant to Schiffer's and Burge's criticisms. As a result, my argument is not susceptible to their criticisms, but rather shows that these criticisms do not generalize to cover arguments which appeal to Sally-style examples.

To begin, Boghossian (1992a) had tried to argue that anti-individualism "undermines our ability to tell *a priori* whether any particular inference of ours satisfies" one of the valid forms of inference (Boghossian 1992a, p. 22n). To do so, he used various examples which (he claimed) show that anti-individualism entails the possibility of undetected equivocations in reasoning. In one such example Peter, an opera fan who is hiking in New Zealand,

comes across Lake Taupo and is startled to see the famous tenor Luciano Pavarotti floating on its pristine waters ... Understandably enough for an opera buff, this experience of Peter's gives rise to many subsequent memories on his part, and to beliefs based upon them ... Some years go by and Peter moves to twin earth and becomes happily ensconced there. Of course, he maintains his interest in opera and so continues to read and hear about his favorite performers. Eventually, some of the tokens of his mental names come to refer to the twin counterparts of the familiar earthly performers: tokens of "Domingo" refer to Domingo, tokens of "Pavarotti" to twin Pavarotti, and so on (p. 21).

Boghossian then argues that, given such world-switching, there will be cases in which "true premises conspire, through a fallacy of equivocation that Peter is in principle not in a position to notice, to produce a false conclusion" (p. 22). So for example Peter might use the following two sentences to express beliefs he takes himself to have:

- (4) "Pavarotti once swam in Lake Taupo."
- (5) "The singer I heard yesterday is Pavarotti."

and from these conclude that:

- (6) (Therefore) the singer I heard yesterday once swam in Lake Taupo.

The conclusion Boghossian draws from this is that, given anti-individualism and a thinker situated like Peter, the result is that "inferences that look to be, 'from the inside', valid, aren't. And the thesis of the *a priori* of logical abilities is shown, thereby, to be inconsistent with externalist assumptions" (p. 22).

In some ways, my argument from section 2 is similar to Boghossian's argument. Like him, I aim to show how anti-individualism may well affect our reasoning. Also like him, I have tried to argue for this on the basis of certain kinds of 'switching' examples. On the strength of these examples I have claimed that anti-individualism entails the possibility that there

are occasions on which a thinker lacks DKC, where her lacking such knowledge affects her decision whether to draw certain (theoretical and/or practical) inferences.

At the same time, there are some important differences between Boghossian's argument and mine. For one, I do not want to hang anything on the notion of '*a priori*' or on the question of what it means to be determinable '*a priori*.' This is important, since Burge (1993 and 1998, p. 365) criticizes Boghossian's argument on the grounds that it illegitimately conflates '*a priori* determinable' with 'internally detectable'. But my example of Sally, and the conclusion I base on it, are unaffected by Burge's criticism of Boghossian's conception of the *a priori*. My aim has been to establish that a failure of DKC is sometimes relevant to reasoning; whether (given anti-individualism) this relevance then entails anything about the *a priori* status of the relevant reasonings is a *further* question – one on which I am not taking a stand. In short, to the extent that Burge's criticisms of Boghossian's argument turn on Boghossian's conception of the '*a priori*', those criticisms are not applicable here.

This difference between Boghossian's argument and mine on the score of the *a priori* is relevant to a second, related criticism that has been brought to bear against Boghossian's argument. This criticism, presented in Schiffer (1992) and reiterated in Burge (1998), seeks to show that anti-individualism has no untoward consequences regarding undetectable equivocation. After identifying three distinct ways one may interpret statements like (4), Schiffer suggests that (even assuming anti-individualism) there is no problem of equivocation in Peter's reasoning: in the context of producing an argument of the sort above (Schiffer argues), Peter will typically intend to be using his terms univocally, with the result that whatever fixes the reference of 'Pavarotti' as this term occurs in Peter's utterance of (4) in some particular argumentative context, will *also* fix the reference of 'Pavarotti' in Peter's utterance of (5) in that same argumentative context (Schiffer 1992, pp. 33–4).²⁸ On this construal, Peter's argument involving (4)–(6) is indeed valid (but unsound). The important point here is that (if true) such a claim about reference-fixation – call it *Schiffer's thesis* – might be thought to be applicable to the sorts of cases I have described.²⁹

In fact, it is easy to see that the kind of criticism embodied by Schiffer's thesis has only a limited use against the kind of examples I am describing in section 2. But before I suggest why this is so, let me set the stage by emphasizing the difference of emphasis between Boghossian's argument and mine, regarding the perspective from which we construct our examples. Boghossian, unlike me, takes what we might call the perspective of *the logician*: the examples he uses, and the points he raises with them, appear to be motivated by a dominant interest in the logical status of certain inferences. I do not deny having such an interest; on the contrary,

I think that these issues go to the heart of issues of 'knowledge of content.' But I do think that, in the context of trying to establish claims regarding anti-individualism's implications for human reasoning, the emphasis on logic (and with it the *a priori*) is not the best one to take. This is for the simple reason that, if one's primary interest concerns the logical status of one's inferences, then it would seem that an insistence on Schiffer's thesis is exactly the right way to ensure that anti-individualism will not entail equivocations in reasoning. But then the perspective of the logician obscures a point that *can* be made against Schiffer's thesis: the point is not that Schiffer's thesis is false, but rather that it does not generalize to cover the cases in which the dominant interest of the reasoner goes beyond ensuring that she is only making deductively valid inferences. In the above cases, what is of primary interest – both to the thinkers described in the example, and to those of us reflecting on these examples – is not a desire to perform deductively valid inferences *per se*, but rather to use reasoning in the production of sound (practical) belief-formation.³⁰

Having said this, I am now in a position to suggest how Schiffer's remarks against Boghossian's argument fail to speak to such a context, by returning to the example of Sally. For the sake of the argument we are granting Schiffer's thesis – that a thinker who wanted to ensure that she made only deductively valid inferences could do so by intending to use each word-form that occurs more than once in the premises in a univocal manner. Assume that Sally is told this. If Sally is rational, this point should strike her as correct, but irrelevant to her situation. For Sally the goal is not (or at least not only) to have deductively valid arguments, but rather to determine whether she should believe that Jones has no principles. And this she will not know until such time as she knows whether, given the contents of the two premises as they come to her through content-preserving memory, these two premises employ 'pragmatist' in a univocal fashion. If she has reason to think that these two premises have not employed 'pragmatist' in a univocal sense, then she would be foolish to employ the method by which Schiffer and Burge would have a thinker ensure a deductively-valid argument – for such a move would open her up to the risk of endorsing a *false belief* (albeit one acquired in a deductively valid fashion). The conclusion can be generalized: in the context of any case where one reasons as part of one's efforts at true belief-formation (as distinct from an interest merely in ensuring that one makes deductively valid inferences), then Schiffer's thesis is largely irrelevant.

Again, this is not a criticism of Schiffer's thesis, as a claim concerning how to ensure against equivocation in one's reasoning. Rather, my claim is that the method in question does not generalize to cover cases where we need to reason in practical situations. Indeed, Burge himself appears to be aware that there are such situations, as when he writes (1998) that:

Perhaps there are cases where the reasoner does not tie the parts of an argument together [by *stipulating* that her terms are used univocally]. That is, the reasoner's intended reasoning does not close the question of whether the concepts expressed by the same word-sound in an argument that is otherwise syntactically-valid are the same. (Burge, 1998 pp. 367–8)

In fact, having acknowledged this, Burge himself goes on to acknowledge something very close to the point I have been anxious to make regarding the relevance of DKC:

Insofar as the reasoner's intentions in reasoning are not dominant in requiring "anaphorically" that the same concept be used through the reasoning, and insofar as we think that there is a gap between the premises that the reasoner has not made explicit, it would seem obvious that the reasoner tacitly and mistakenly presupposes that the concepts apply to the same objects ... To fully capture the reasoner's cognitive state in a case where the reasoner does presuppose (mistakenly) that the concepts apply to the same objects, one would have to supply for the reasoner the mistaken presupposition ... Again, there is no mistake in reasoning, only a mistake in presupposition (p. 368).

Suppose that this *is* the right way for an anti-individualist to describe the cognitive situation of a reasoner who uses a word to express two distinct concepts but fails to realize this.³¹ This would acknowledge that a failure of DKC is tantamount to a mistaken presupposition, which presupposition is relevant to the reasoning itself. As such Burge is acknowledging that there are cases in which a failure of DKC must be represented in a formalization of a thinker's reasoning.

But I think we must acknowledge more than Burge has acknowledged here. To see this, I begin by pointing out that, in the section of Burge (1998) in which he deals with Boghossian's argument, he is anxious to show that anti-individualism does not have any untoward implications regarding the *a priori* status of reasoning. Quite reasonably, he inherited the example he considers from Boghossian himself. In a sense, then, Burge's discussion has the same limited perspective Boghossian's has. True, Burge is aware that there are ("perhaps") cases in which one will not reason in such a way as to stipulate univocal uses of one's own terms. But this is acknowledged only in passing. The result, I think, is that even if Burge has *rebutted* the charge that anti-individualism entails cases of invalid-yet-undetected equivocation, he has not made explicit what anti-individualism *does* entail regarding our reasoning. This was precisely the point of my example, to show that, given anti-individualism, there will be cases in which a thinker fails to have DKC, which failure of knowledge affects her reasoning. If this is merely the (positive) flip side of the (negative) claim Burge himself made against Boghossian, I think it is definitely worth making, if only to dampen whatever enthusiasm there is amongst anti-individualists for thinking that KC is the only notion of knowledge of content relevant to our reasoning.³² I now turn to elaborating on this point as a way to conclude this paper.

4. *The Significance of the Thesis*

I began by noting that compatibilists typically respond to would-be incompatibilist arguments by distinguishing between introspective KC and introspective DKC. This distinction, I remarked, is meant to be part of a strategy whereby the compatibilist *concedes* that anti-individualism has implications regarding possible failures of introspective DKC, but *denies* that this conceded point shows that anti-individualism is incompatible with authoritative knowledge of content (in the sense of KC). I submit that the thesis of the paper – that anti-individualism's implications regarding the possibility of a thinker's failing to have introspective DKC will potentially affect a thinker's reasoning – shows that anti-individualists who make the KC/DKC distinction should not take themselves to have said all that needs to be said against would-be incompatibilist arguments. There are two reasons for this. First, we want to know why, on the assumption that individualism preserves introspective DKC,³³ the conceded point is not a significant weakness of anti-individualism vis-à-vis its individualist opponents.³⁴ Second, we want to know more about DKC: what it is to have introspective DKC,³⁵ and how DKC relates to KC itself.

Here I will only comment on the second of these questions. If my present argument is correct in its assertion that anti-individualism entails possible failures of DKC, which failures are potentially relevant to a thinker's reasoning, then it would show that there are issues regarding 'knowledge of content' that are not settled by any notion of KC that does not involve DKC. In particular, the present argument would establish that if by 'knowledge of content' we mean to include *whatever knowledge a thinker has of her thoughts, such that she can use these thoughts as premises in arguments or practical syllogisms*, then any notion of KC which does not involve DKC will not be sufficient to the task. This result puts anti-individualists on notice that the well-worn strategy of securing the compatibility of anti-individualism and KC by appeal to the self-referential nature of self-ascriptive judgements (for which see Heil 1988; Burge 1988; Falvey and Owens 1994; and Gibbons 1996), does not yet provide us with a full theory of 'knowledge of content.' My aim in making this point is constructive: I believe that the project of motivating and articulating an anti-individualist notion of DKC is one that should begin to receive some attention.

Department of Philosophy
Grinnell College and the University of Kentucky

NOTES

¹ I thank Akeel Bilgrami, Paul Boghossian, Tony Brueckner, Tyler Burger, Kevin Falvey, York Gunther, Sidney Morgenbesser, Bill Robinson, and Michael Root, for

discussions of these and related matters; Joseph Owens, for having given me some constructive criticisms of an earlier version; those who attended the session of the 1998 Iowa Philosophical Society at which I gave this paper; and Grinnell College, for a generous summer stipend to work on this paper.

² See Burge (1988) p. 662. I should point out, however, that Burge himself is explicit about employing two distinct notions of the 'knowing what thought' locution; see p. 662 and Burge (1998) p. 357. I review and comment on these two notions in Goldberg (unpublished).

³ The suggestion that knowledge of content should be understood as a kind of knowing *what* one thinks was originally stressed to me by Sidney Morgenbesser. He suggested to me that knowing *that* one oneself thinks that *p* should not (or, at least, should not *necessarily*) be conflated with knowing *what* one oneself thinks in thinking that *p*. However, since Morgenbesser's suggestion employs a notion of 'knowing what,' which Burge and his defenders would reject, and since in this paper I aim to formulate an argument which should be acceptable to Burge and his follows, I will not press Morgenbesser's suggestion here. (I develop this suggestion at length in my unpublished manuscript.)

⁴ See for example Butler (1997), Goldberg (1997), and Goldberg (forthcoming).

⁵ Brueckner (forthcoming) levels such a response against Goldberg (1997). (He does so, I should point out, not as a compatibilist, but rather as someone who has some very elaborate views about how one can *and cannot* generate skepticism regarding knowledge of content.)

⁶ This formulation of the point is admittedly awkward; it is meant to be neutral with respect to issues of internalism and externalism in the theory of epistemic justification. Thus I speak of the justification which Sally *takes herself* to have; I do not make any further claims about justification *simpliciter* in her case. It is enough for my purposes that a failure of introspective DKC lead Sally to question whether she can make certain inferences; her doubts here will affect her reasoning, even if from the perspective of some theory of justification she is (unbeknownst to her) ultimately justified in making the inferences she presently questions.

⁷ Sidney Morgenbesser has pointed out to me (in conversation) that putting matters in terms of sameness and difference of *concepts* needlessly complicates the point I wish to make; he has suggested that the point is better put in terms of sameness and difference of associated *extensions*. To those who share his scruples about the lack of clarity surrounding the expressions 'concept' and 'same concept,' I would suggest that they reformulate my present argument by substituting 'extension' and 'same extension' throughout; I do not think that this would seriously alter the nature of the argument itself. But I will continue to use 'concept,' if only because I find it more natural.

⁸ To say that these concepts are subjectively indistinguishable to Sally is to say that Sally cannot distinguish her concepts, i.e., that she is not in a position to provide clear explications which capture the way in which the concepts differ (for more on which see 2.4 below). As I formulate the notion of subjective indistinguishability, to say that two concepts are subjectively indistinguishable to Sally is compatible with saying that Sally can think thoughts with one and not the other concept.

⁹ To say this is not to say that her use of this term on every given occasion expresses a disjunctive concept. On the contrary, what concept(s) she expresses on a given occasion of uttering 'pragmatist' and its derivatives depends in large part on her semantic intentions. I develop this point both in Goldberg 1997 and (in greater detail) in Goldberg (forthcoming); see also Schiffer 1992 and Burge 1998, pp. 352–4.

¹⁰ In particular, I will suggest that my argument depends on (a) the claim that the two-concept thesis is true of her, and (b) a suitable sort of lapse in memory concerning the provenance of the beliefs in whose expression she uses the term 'pragmatist.' The fact that my argument only requires this more restricted description of Sally's case is important. Joseph Owens has pointed out that, so long as we insist that Sally does not have enough explicational knowledge to distinguish in an illuminating way between the two concepts

expressed by 'pragmatist,' it is not at all clear that we *should* count Sally as having acquired the two concepts associated with the respective linguistic communities. But if I am right that my argument can be formulated consistent with the idea that Sally *does* have sufficient explicational knowledge (regarding the two concepts expressed by 'pragmatist') to warrant the ascription of the two concepts to Sally, then Owens' point can be accepted without affecting the status of my argument. See 2.4 for the details.

¹¹ See Burge (1993) for a characterization of content-preserving memory. Roughly, content-preserving memory is memory that enables us to access a past attitude in such a way as to preserve the content of the attitude; that is to say, the content of the memory (what is remembered) *preserves*, and so is identical to, the content of the attitude remembered (what one believed and so on).

¹² Below (section 2.3) I will consider arguments that reject the idea that content-preserving memory can be operative in such cases; I will find such arguments wanting.

¹³ In describing Sally's belief in this awkward way I do not mean to imply that the belief which she has in virtue of her memory is a metalinguistic belief, as, for example, it would be if she believed *that what she meant earlier by "No pragmatists have any principles" is true*. On the contrary, I take it (following Burge, 1993) that her memory gives her *the very belief* (the very belief-*content*) that she had originally expressed when she uttered "No pragmatists have any principles." I will return to this issue in 2.3.

¹⁴ I am assuming that it is *contents themselves* that enter into entailment relations with one another; my argument could be easily modified if one thinks that the items that enter into logical relations are not the contents themselves but rather are *other* items (interpreted logical forms, for example).

¹⁵ Burge (1998) pp. 361–2 would appear to agree with this judgement: he suggests that a thinker can recall an aluminium-thought in a content-preserving way, even in circumstances in which she has been switched without her knowledge to (and has been living for years in) a world in which the term 'aluminium' is used to refer to a different (albeit superficially similar) metal, twaluminium.

¹⁶ Once again, Brueckner himself is not a compatibilist, although he has criticized as ultimately unsuccessful various would-be incompatibilist arguments; one such criticism is found in Brueckner (forthcoming).

¹⁷ I would point out that Burge (1998) would appear to acknowledge the assumption (see especially pp. 360–4). The reason why, despite his acknowledgement, I nonetheless went to the trouble of *arguing* for the assumption, is to convince those anti-individualists who have not seen their way to his (to my mind, wise) acknowledgement.

¹⁸ Further substantiation of the prevalence of world-switching comes once we recognize that the technologies that permit remote communication – the telephone and e-mail in particular – enable us to acquire beliefs from others without being in a position to be able to determine very much about our interlocutors' linguistic community. This only makes the possibility of world-switching both more prevalent (one can do it without leaving one's office!) and less easy to detect at the time at which one acquires the belief in question – with the result that there will be correspondingly more opportunities to acquire subsequent information indicating one's own previous world-switching, and so more situations in which one's own case is akin to Sally's.

¹⁹ See Goldberg (1997) for a defense of the claim that Sally's mere ability to use the term to express one and not the other concept is not sufficient for her counting as 'knowing which' concept she has expressed.

²⁰ I thank Sidney Morgenbesser for emphasizing the differences between the two types of discriminatory failure.

²¹ I thank Joseph Owens for bringing to my attention the need to make clear that my argument does not depend on the attribution to Sally of two distinct *but subjectively indistinguishable* concepts.

²² 'in any illuminating way' might be rendered as 'in any way that indicates to the thinker whether the inference in question is deductively valid.' For an alternative gloss, see Goldberg (1997).

²³ Strictly speaking, this should read, "*what is expressed by* (1) and (2)," since (1) and (2) are *sentences*, not the claims made by those sentences.

²⁴ Saying that anti-individualism is to blame for this case is compatible with saying that there are cases like Sally's where the failure of DKC arises for reasons *independent of* anti-individualism. Whether there are such cases is a matter that ought to receive more attention than it has (although I point to Owens (1990) as the rare paper that does take up this issue). In any case I thank Joseph Owens for the suggestion that, arguably, there *are* cases of a failure of DKC that derive from grounds independent of anti-individualism.

²⁵ Here, 'relevant respects' means respects relevant to the term(s) at issue: obviously, a shift from one scientific community to another will most probably not affect the concept one expresses with 'pragmatism.'

²⁶ I thank both Bill Robinson and Joseph Owens, each of whom (independently) pointed out to me the need to acknowledge this reaction.

²⁷ Boghossian spoke of 'externalism,' not 'anti-individualism,' but this is not a difference that need concern us here. I will continue to use the latter expression, for the sake of expository clarity.

²⁸ That there is an anaphora-like way in which the reference of a term is fixed as this term occurs in subsequent premises of an argument is also made in Burge (1998). There, Burge examines the scope and limits of the parallel between the reference-fixing of terms and the anaphoric nature of "pronomial back-reference," pp. 357–9.

²⁹ Schiffer's thesis has been seconded by Burge himself (1998, pp. 366–7).

³⁰ In fairness to Boghossian, he himself is aware that the practical reasoning perspective I am now hinting is the best place to consider anti-individualism's implications on reasoning; see for example Boghossian's use of examples involving practical syllogisms and his comments about the rationalizing role of the attribution of content (1992a, pp. 26–8). My criticism of Boghossian, then, is that this recognition is spoiled by his emphasis in 1992a on the logical (*a priori*) status of certain inferences. I should add that Burge, too, is aware of the practical reasoning perspective (see Burge 1998 p. 368); I will return to Burge's position below.

³¹ For a disagreement that Burge is entitled to say this, consistent with compatibilism, see Bilgrami (1992).

³² What is more, I take myself to have established this claim in such a way that it should be accepted even by those who endorse Burge's views about the compatibility of anti-individualism and introspective KC. Thus, this paper is offered in the spirit of serving as an anti-individualist elaboration on the expression 'knowledge of content,' rather than as presenting an argument antagonistic to anti-individualists. I am simply saying that notion of knowledge of content employed by most anti-individualists – that is, KC – is not sufficient to do everything that we expect from a notion of knowledge of content.

³³ But see my unpublished manuscript, where I argue that no plausible version of individualism does succeed in preserving DKC.

³⁴ I should point out once again that Owens (1990) contains an argument relevant to this issue. He presents an argument, independent of anti-individualist considerations, for the conclusion that DKC should not be thought to be introspectible. If his argument is sound, then it should come as good news to anti-individualists: Owens' reasoning could be used to argue that, even if (contrary to my unpublished manuscript) there is an individualist doctrine that *does* preserve introspective DKC, this is no virtue of that version of individualism, since (Owens established that) we have independent reasons to think that (there are many cases in which) DKC cannot be had on the basis of introspection alone.

³⁵ It is tempting to take the simplest view, that to have introspective DKC is to be able to discriminate one's thought from *any* content-distinct thought. However, in Goldberg (forthcoming) I argue against such a demanding conception.

REFERENCES

- Bilgrami, (1992). "Can Externalism be Reconciled with Self-Knowledge?" *Philosophical Topics*, 20(1).
- Boghossian, (1989). "Content and Self-Knowledge," in *Philosophical Topics*, 17(1).
- Boghossian, (1992a). "Externalism and Inference," *Philosophical Issues*, 2, pp. 11–28.
- Boghossian, (1992b). "Reply to Schiffer," *Philosophical Issues*, 2, pp. 39–42.
- Brueckner, (1990). "Scepticism about Knowledge of Content," *Mind*, 99, pp. 447–51.
- Brueckner, (1997). "Externalism and Memory," *Pacific Philosophical Quarterly*, 78, pp. 1–12.
- Brueckner, (forthcoming). "Difficulties in Generating Skepticism about Knowledge of Content," *Analysis*.
- Burge, (1988). "Individualism and Self-Knowledge," *Journal of Philosophy*, 85.
- Burge, (1993). "Content Preservation," *The Philosophical Review*, 102(4).
- Burge, (1996). "Our Entitlement to Self-Knowledge," *Proceedings of the Aristotelian Society*, 96, pp. 91–116.
- Burge, (1998). "Memory and Self-Knowledge," in Peter Ludlow and Norah Martin (eds), *Externalism and Self-Knowledge*. Stanford: CSLI Publications, pp. 351–70.
- Butler, (1997). "Externalism, Internalism, and Knowledge of Content," *Philosophy and Phenomenological Research*, 57(4), pp. 773–800.
- Falvey and Owens, (1994). "Externalism, Self-Knowledge, and Skepticism," in *The Philosophical Review*, 103(1).
- Gibbons, (1996). "Externalism and Knowledge of Content," *The Philosophical Review*, 105(3).
- Goldberg, S. C. (1996). Introduction to Section IV, entitled "Self-Knowledge," in Pessin and S. C. Goldberg, (eds) *The Twin Earth Chronicles*. New York: M. E. Sharpe.
- Goldberg, S. C. (1997). "Self-Ascription, Self-Knowledge, and the Memory Argument," *Analysis*, 57(3), pp. 211–19.
- Goldberg, S. C. (forthcoming). "Externalism and Authoritative Knowledge of Content: A New Incompatibilist Strategy," *Philosophical Studies*.
- Goldberg, S. C. (unpublished manuscript). "The Problem of Content-Identification."
- Heil, (1988). "Privileged Access," in *Mind*, 97, pp. 238–51.
- Ludlow, (1995a). "Externalism, Self-Knowledge, and the Prevalence of Slow Switching," *Analysis*, 55(1).
- Ludlow, (1995b). "Social Externalism, Self-Knowledge, and Memory," *Analysis*, 55(3).
- Owens, (1990). "Cognitive Access and Semantic Puzzles," in C. A. Anderson and J. Owens (eds) *Propositional Attitudes: The Role of Content in Logic, Language, and Mind*. Palo Alto: CSLI Press, pp. 147–73.
- Schiffer, (1992). "Boghossian on Externalism and Inference," *Philosophical Issues*, 2, pp. 29–38.