

Anti-Individualism, Content Preservation, and Discursive Justification

SANFORD C. GOLDBERG
University of Kentucky

Abstract

Most explorations of the epistemic implications of Semantic Anti-Individualism (SAI) focus on issues of self-knowledge (first-person authority) and/or external-world skepticism. Less explored has been SAI's implications for the epistemology of reasoning. In this paper I argue that SAI has some nontrivial implications on this score. I bring these out by reflecting on a problem first raised by Boghossian (1992). Whereas Boghossian's main interest was in establishing the incompatibility of SAI and "the *a priori* of logical abilities" (Boghossian 1992: 22), I argue that Boghossian's argument is better interpreted as pointing to SAI's implications for the nature of discursive justification.

1. Introduction to the Basic Dialectic

Semantic anti-individualism (SAI) is the thesis that "the mental natures of many of an individual's mental states and events are dependent for their individuation on the individual's social and physical environments" (Burge 1986: 697). This paper is concerned with one (or a bundle) of the epistemic implications of SAI. However, whereas the preponderance of discussions of SAI's epistemic implications have focused on self-knowledge and external-world skepticism, here I will be interested in SAI's implications regarding the nature of *discursive justification*—the sort of justification that accrues to a belief acquired via inference.

I regard the relatively underexplored topic of SAI's implications regarding discursive justification to be of central importance to the project of identifying SAI's epistemic implications.¹ There are two main reasons for this.

First, this topic serves as a corrective to an impoverished (because overly narrow) focus in the current literature in connection with SAI's epistemic implications. Any superficial review of the literature of the past twenty years on the topic of SAI's epistemic implications will turn up a great number of papers and books that focus exclusively on SAI's implications for one or both of first-person authority and external world skepticism. I regard this focus as overly narrow. For it seems to me that even after it is conceded (as I think it ought to be²) that SAI has no untoward implications for either of these topics, the question remains how *what* one knows, when one knows one's own thoughts, relates to the rest of one's epistemic perspective.³ It is on this score, it seems to me, that a systematic focus on SAI's implications for discursive justification will bring much-needed illumination to this topic.

The second reason for focusing on the relatively underexplored topic of SAI's implications for discursive justification is that, given a popular (though by no means universal) auxiliary assumption, SAI's implications regarding the nature of discursive justification bear directly on the more general topic of SAI's implications for the nature of epistemic justification *simpliciter*. The popular subsidiary assumption is that the epistemic justification of a belief involves the reasoning-based activity by which a subject justifies (or, by searching reflection, could justify) the belief in question.⁴ Given this subsidiary assumption, the justification of a belief is grounded in the reasoning that justifies it. Since this reasoning itself can be represented formally as having the structure of an argument, the reasoning-generated justification will involve a discursive component. The result would be that if SAI has implications for discursive justification, then it has corresponding implications for justification *simpliciter*.

The dialectic of this paper is as follows. I begin (section 2) with an argument Boghossian (1992) has offered for the conclusion that SAI is incompatible with "the *a priority* of logical abilities" (Boghossian 1992: 22). As an argument for this incompatibilist conclusion, however, Boghossian's argument has been met by a response, owed to Schiffer (1992) and Burge (1996b), which notes that reasoning subjects can secure the validity of their deductive inferences, and hence vindicate the *a priority* of their logical abilities, by reasoning with certain univocality intentions. However, while the Schiffer/Burge reply does rebut the incompatibility charge, it comes at a cost. For one thing, a reasoner who reasons in the way suggested by Schiffer and Burge sacrifices the truth of at least one of her premises, and hence the soundness of her argument (section 3). More importantly, given a plausible connection between justification and truth, the further result will be that certain pretheoretically plausible principles of discursive justification will not hold of reasoning of this sort (section 4).

The upshot of these reflections is that, however the anti-individualist responds to the dialectic generated by Boghossian's 1992 charge, SAI has some non-trivial implications regarding discursive justification. I want to make

clear at the outset, however, that I do not advance this conclusion as a reason to reject SAI. On the contrary, I regard the arguments on behalf of anti-individualism (owed above all to Burge) to be compelling, and to remain so even after the present argument is made. I advance my conclusion, rather, to force an honest and systematic assessment of SAI's epistemic implications. My sense is that these implications are much more substantial than has heretofore been appreciated. It is thus that by tracing SAI's implications for discursive justification, I aim to refocus the discussion of SAI's epistemic implications more generally.

2. Boghossian 1992 and the Schiffer/Burge Reply

Boghossian's 1992 argument was designed to show SAI is inconsistent with "the thesis of the *a priority* of logical abilities" (1992: 22). The argument itself is based on a thought experiment devised to bring out the epistemic significance (for reasoning) of SAI's implication that attitude-contents are non-transparent. That SAI has this implication is widely acknowledged, by friends as well as foes of SAI (see Burge 1988a, 1988b, 1996a, 1996b; Boghossian 1989, 1992, and 1994; Owens 1986; Falvey and Owens 1994; Goldberg 1999, 2000, 2003a, and 2003b; and Brown 2000 and 2004). Boghossian takes the failure of content transparency to give rise to the possibility of equivocations of reasoning, which equivocations cannot be detected by any reflection on the part of the reasoning subject.

Boghossian illustrates this possibility with the following thought experiment. 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. . . . 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 esconced 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 twin Domingo, tokens of "Pavarotti" to twin Pavarotti, and so on (Boghossian 1992: 22).

Boghossian then has us imagine a scenario 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" (1992: 22). Thus Boghossian has Peter engage in the following reasoning:

1. Pavarotti once swam in Lake Taupo.
2. The singer I heard yesterday was Pavarotti.
(Therefore)
3. The singer I heard yesterday once swam in Lake Taupo.

The problem as it first strikes us (and as Boghossian described the situation) is this: although each of the premises is true and Peter's reasoning appears (introspectively, to Peter) to be deductively valid, (3) is in fact false (since the two tokens of 'Pavarotti' used in the premises refer to different individuals). Boghossian concluded that, on the assumption of SAI, a subject such as Peter fails to be in a position to determine *a priori* the formal validity (or not) of her own arguments: SAI is inconsistent with "the thesis of the *a priority* of logical abilities" (1992: 22).

In response to Boghossian's argument, both Schiffer (1992: 33-34) and Burge (1996b: 367-68) note that, as a matter of fact, reasoners typically or usually reason in such a way that there will be *anaphoric, reference-preserving* links between the non-indexical word-forms figuring in an extended chain of reasoning—and that as a result their reasoning will not involve introspectively undetectable equivocations.⁵ In that case reasoners whose primary intention was to reason in a way that they could determine *a priori* to be valid would have the option of reasoning in the manner indicated. Then, since the allegation of an incompatibility holding between two doctrines amounts to the claim that there is no possible world in which both doctrines are true, the result would be that, if the Schiffer/Burge reply succeeds in showing that there is a possible world in which both SAI and the doctrine of "the *a priority* of logical abilities" are true,⁶ this reply refutes Boghossian's incompatibility charge. In the remainder of this section I argue that the Schiffer/Burge reply is adequate to refute Boghossian's incompatibility charge.

To argue that the Schiffer/Burge reply is adequate to refute Boghossian's incompatibility charge, it suffices to show that if reasoners reason in the manner indicated, they are in no danger of the sort of undetectable equivocation exhibited by Boghossian's Peter.⁷ This point is not as obvious as it might seem. To reason in the manner indicated is to reason with the primary intention that the non-logical terms one is using (throughout some extended piece of reasoning) be univocal; but it remains to be seen whether one's mere intentions on this score would succeed in rendering one's terms univocal. After all, it may be wondered whether, on the assumption of SAI, one can succeed in rendering one's terms univocal merely by *intending* that they be such. The case of Peter makes clear why this is a worry: Peter's use of 'Pavarotti' at t_1 expresses something other than his use of 'Pavarotti' at t_2 ; so it remains to be seen whether his uses of 'Pavarotti' throughout the reasoning he exhibits at t_3 —reasoning which exploits, or at least purports to exploit, the 'Pavarotti'-beliefs he acquired at t_1 and t_2 —can be rendered univocal merely by his intention that they be such.

I submit that this is indeed the case. Granting that Peter's use of "Pavarotti" at t_1 (in the course of his expressing a thought by uttering (1) at t_1) expresses something different from what is expressed by his use of 'Pavarotti' at t_2 (in the course of his expressing a thought by uttering (2) at t_2), nothing prevents Peter from intending to be using each of his non-indexical terms

univocally at the time of his reasoning ($=t_3$). So unless the assumption of SAI itself introduces something that would have the effect of undermining Peter's univocality intention, the result is that Peter (and, by extension, any reasoning subject) could reason in precisely the manner indicated by Schiffer and Burge, and thereby avoid the sort of equivocation in question. Our question, then, is whether the assumption of SAI would undermine Peter's univocality intention.

To begin, consider what it would be for Peter's univocality intention to be undermined. To put matters formally, Peter's univocality intention is undermined on a given occasion O of extended reasoning if and only if (i) on O Peter uses a single non-indexical expression-type e more than once (call these uses $u_1 \dots u_n$); (ii) for some $i \neq j$ (where $1 \leq i, j \leq n$), he intends uses u_i and u_j of e to be univocal (i.e., to have the same semantic value, to express the same concept⁸); yet (iii) uses u_i and u_j of e are not univocal (i.e., they have different semantic values, express different concepts). SAI undermines Peter's univocality intention on a given occasion, then, iff SAI (together with a complete description of the scenario) entails that on that occasion conditions (i)-(iii) hold. Now one way for SAI to have such an implication would be if SAI were wed to the following *content-fixity* thesis:

CF If at some time t S acquires a belief that she expresses with a sentence β (not involving any indexical elements), then, so long as S 's attempt to recall that very belief produces the disposition in S to express the belief recalled with a token of the very same sentence-type, the token sentence produced at the time of recollection has the very same content as that expressed by the token of β S produced at t .

The idea of CF can be put, intuitively, as follows: once acquired, a belief retains its original content so long as the subject retains the disposition to recall what she would regard as the same belief, using the very same (non-indexical-involving) sentence she used on its original acquisition. Given CF, Peter's univocality intention at t_3 would be undermined by the fact that at t_3 he is purporting to recall what he would regard as the same beliefs he acquired at t_1 and t_2 , and which he expressed (at t_1 and t_2) with (1) and (2), respectively.⁹ But if SAI is not wed to something like CF, then, it seems, SAI does not undermine Peter's univocality intention. For what could possibly undermine his univocality intention at t_3 , other than the facts, first, that he is purporting to recall the very beliefs he acquired at t_1 and t_2 , and that (given SAI) his uses of 'Pavarotti' at t_1 and t_2 were not univocal?

In sum, our question is whether SAI is committed to CF. If so, then SAI undermines Peter's univocality intention; if not, not.

To address this question, I begin by noting that some proponents of SAI (including Burge himself; see his 1993) endorse a doctrine asserting that memory preserves the content of one's original belief. But this does not show that

SAI itself is *committed* to (something like) CF. For one thing, Burge's thesis is a thesis about what memory *can* do, not what it *must* do. (Indeed, as we will see below, Burge's own way of responding to the case of Boghossian's Peter implies that Peter's memory will *fail* to preserve the content of Peter's original belief; rather, as Burge sees it, memorial content preservation works to ensure that the same content is preserved *throughout the interval of the reasoning* Peter is engaged in.) But there is a second reason to doubt that SAI is committed to CF, which is that Burge's own views about memorial content preservation are not required by SAI itself. In fact, in addition to Burge's 1996b discussion of SAI's implications for memory, at least two recent discussions of SAI's implications for memory (Ludlow 1995; Gibbons 1996) suggest that the proponent of SAI should hold that, in at least some cases of world-switching, the contents of one's memory, far from being preserved, actually *shift*. If such positions are coherent (and I see no reason to think they're not), SAI is not committed to CF. But in that case SAI does not undermine reasoners' ability to reason in the manner indicated in the Schiffer/Burge reply to Boghossian. Since such reasoners would be reasoning validly (and hence would not be guilty of any undetectable equivocation), SAI is not inconsistent with "the *a priority* of logical abilities."

3. Semantic and Epistemic Implications of the Schiffer/Burge Reply

I just argued that the Schiffer/Burge reaction to Boghossian succeeds in refuting Boghossian's incompatibility charge. I now want to argue that reasoning in the manner indicated by Schiffer and Burge has some untoward epistemic implications in connection with the issue of discursive justification. To make this case, I begin by making explicit the obvious semantic implication of the Schiffer/Burge reply, and then proceed to argue that this semantic implication has epistemic consequences. In developing these points, we will reach a point of general interest: content-shift proposals, according to which the content of a belief acquired at a given time *t* shifts at some later time owing to the subject's having switched environments, come at a high epistemic cost. This point is of general interest, since it bears against recent proposals regarding the semantics of memory (e.g. in Ludlow 1995 and Gibbons 1996): these proposals have memory do what, on the proposal of Schiffer 1992 and Burge 1998, is achieved by the speaker's univocal intention. So if the argument to follow succeeds in showing that the Schiffer/Burge reply to Boghossian has unacceptably strong implications on the score of discursive justification, it could be adapted to show that the Ludlow/Gibbons anti-individualistic account of the semantics of memory has the same implications.¹⁰

It will be helpful to be able to talk clearly about the dynamics of belief. To do so we will need to specify, regarding a particular belief that an agent has, what epistemic support that belief enjoys (relative to the subject's overall cognitive state) at a given time. I will suppose that whatever support a given

belief enjoys at a given time is provided by *the grounds* on which the belief is held.¹¹ Suppose that at time t S believes that p on grounds g . When and only when g is such that those grounds objectively count towards the truth of p , we will say that S 's belief that p is justified at t . This usage is in keeping with the tradition in epistemology according to which 'justified' is used in the sense of 'having truth-conducive epistemic support'.¹²

Consider now how Boghossian's presentation of his case would be represented in our scheme. At t_1 , Peter acquires a belief he expresses with (1); let us designate the proposition he then believes as p_1 . It can be assumed (given SAI) that

(P1) p_1 is true iff earth-Pavarotti once swam in Lake Taupo.

Peter's grounds at t_1 for believing that p_1 consist of his having seen earth-Pavarotti swimming in Lake Taupo at t_1 . Since these grounds objectively point towards the truth of p_1 , Peter's belief that p_1 is justified at t_1 . Consider next Peter at t_2 , the time at which he acquires the belief he expresses with (2). Let us designate the proposition believed as p_2 . It can be assumed (given SAI) that

(P2) p_2 is true iff Peter heard twin-Pavarotti sing yesterday (= the day before t_2).

Peter's grounds at t_2 for believing that p_2 consist of his having heard twin-Pavarotti sing yesterday. Since these grounds objectively point towards the truth of p_2 , Peter's belief that p_2 is justified at t_2 . Consider finally the belief Peter acquires (at t_3) by reasoning from (1) and (2). Let us designate the proposition in which he comes to believe (on the basis of this reasoning) as p_3 . It can be assumed (given SAI) that

(P3) p_3 is true iff the singer Peter heard sing yesterday (= the day before t_2) once swam in Lake Taupo.

Peter's grounds at t_3 for this belief consist of the two premises of the argument, the grounds supporting each of those premises, and the relation between premises and conclusion (which Peter takes to be one of logical implication but which is not). It would appear that these grounds do *not* objectively point towards the truth of p_3 : since the argument in fact is invalid, the grounds that support the premises do not objectively point towards the truth of the RHS of (P3). Thus it would appear that Peter's belief in p_3 is not justified at t_3 .¹³ This is precisely how we might present matters on behalf of Boghossian (were he to have had an interest in discursive justification rather than "the *a priori* of logical abilities").

The foregoing is a first-blush description of how matters stand if Peter is as Boghossian had described him. How then do matters stand if Peter reasons,

not as Boghossian had described, but rather in the way that, according to Schiffer and Burge, is typical of reasoners engaged in extended chains of reasoning? For one thing, reasoning in this way will affect the *semantics* of the situation. As Schiffer and Burge note, a reasoner engaged in extended reasoning typically does so with the intention to be using each of her tokens of the various non-indexical expression-types univocally, so that there is anaphoric reference-preservation in the extended chain. As noted above, if Peter renders his reasoning valid in this way, he will have done so at the cost of *shifting* the semantic value of one of his uses of 'Pavarotti' at t_3 . That is, he will have succeeded at the cost of making one of the following two inequalities true:

($\neq 1$) The content expressed by Peter's utterance of (1) at $t_1 \neq$ the content expressed by Peter's utterance of (1) at t_3 .

($\neq 2$) The content expressed by Peter's utterance of (2) at $t_2 \neq$ the content expressed by Peter's utterance of (2) at t_3 .

(If talk of content is not welcome here, we can shift to talk of truth-conditions.)¹⁴

The present semantic point regarding shifting contents can be illustrated clearly in connection with the dynamics of belief. (For the sake of expository simplicity I will assume that ($\neq 1$) holds; but a parallel point can be made if one's preferred analysis has it that ($\neq 2$) holds.) Where ' p [Peter, t_1 , (1)]' designates the proposition that Peter expresses with his utterance of (1) at t_1 ($= p_1$), Peter's reasoning in the way characterized by Schiffer and Burge will have the effect that

$$p[\text{Peter}, t_1, (1)] \neq p[\text{Peter}, t_3, (1)]$$

In particular, whereas the truth conditions for $p[\text{Peter}, t_1, (1)]$ ($= p_1$) were given by (P1), the truth conditions corresponding to $p[\text{Peter}, t_3, (1)]$ (call this proposition ' p_1^* ') are as follows:

(P1*) p_1^* is true iff *twin* - Pavarotti once swam in Lake Taupo.

Since Pavarotti and twin-Pavarotti are different individuals, the RHS of (P1) is logically independent of the RHS of (P1*). In fact, in the case described, the RHS of (P1) is true—this accounts for the truth of Peter's belief in (1) at t_1 —whereas the RHS of (P1*) is false. In sum, if Peter reasons with the sort of univocality-ensuring intentions characterized by Schiffer and Burge, he will be reasoning in a way that is *unsound*. This is not a particularly surprising or exciting result.¹⁵ But I want to argue that, in addition to this *semantic* point (about the effect of the univocality intention on the truth of Peter's premises, and hence on the soundness of his argument), the univocality intention has *epistemic* effects as well.

We have just seen that Peter's reasoning with this univocality intention will have the effect that $p_1 \neq p_1^*$. Consider now that at t_1 Peter's belief in p_1 was justified, since at that time it was based on grounds that pointed objectively to the (likely) truth of p_1 . We can stipulate (as part of the thought experiment) that in the interval from t_1 to t_3 (the time of Peter's recollection and reasoning) Peter acquires no new grounds for his belief that p_1 . Then it is plausible that the only grounds Peter has at t_3 for his belief that p_1^* are those preserved in memory, namely, those that were the grounds (at t_1) of his original belief that p_1 . But clearly *those* grounds do not objectively point to the truth of p_1^* . (Peter's having seen earth-Pavarotti in Lake Taupo does not objectively point to the truth of Peter's having seen *twin*-Pavarotti in Lake Taupo, and hence does not point to the truth of the RHS of (P1*.) The result is that at t_3 Peter's belief that p_1^* is *not* justified. Not only will it be the case that Peter's reasoning with the univocality intention (if successful) has the effect Peter's memory fails to preserve the *content* of the belief he expressed at t_1 with his utterance of (1), thereby rendering his argument unsound; *because* of this his reasoning with this intention (if successful) will have the effect that Peter's memory fails to preserve the *justification* of that belief as well.

Admittedly, the foregoing argument might not compel. One who wished to regard Peter's belief that p_1^* (= the shifted content) as justified at t_3 might propose to return to the question of that belief's grounds.¹⁶ Such a view might run as follows. Peter's belief that p_1^* (at t_3) represents *Twin*-Pavarotti as having swum in Taupo; but while the person Peter actually saw swimming in Taupo was Earth-Pavarotti, even so, his present belief that *Twin*-Pavarotti once swam in Taupo is supported by *the appearances* at t_1 . For Peter's experience at t_1 might be described as an experience as of someone *with the appearance of* *Twin*-Pavarotti swimming in Taupo. (Of course, Peter would not have employed this description at t_1 , since at t_1 Peter had not yet made the acquaintance of *Twin*-Pavarotti.) So long as appearances of this sort point objectively to the truth of *Twin*-Pavarotti's having swum in Taupo, we can continue to maintain that Peter's belief, to the effect that *Twin*-Pavarotti once swam in Taupo, is justified at t_3 —even though it is based on Peter's having seen *earth*-Pavarotti swimming in Taupo.

Let the foregoing be granted. The trouble remains that Peter's memory fails to preserve the justification of the original belief. As we might put it: granted that Peter's memory preserves *a* justification for Peter's belief in (1) at t_3 ,¹⁷ even so his memory still fails to preserve the justification Peter's belief in (1) enjoyed at t_1 . This is because the appearances offer a *weaker* sort of justification, than the justification Peter's belief in (1) enjoyed at t_1 .

Where S believes that p on grounds g , we can get a rough index on the strength or degree¹⁸ of the epistemic justification enjoyed by S 's belief, by identifying the worlds inconsistent with p but which g rules out, and then determining the likelihood of S believing that p on similar grounds, in any

of the $\sim p$ -worlds *not* ruled out by g .¹⁹ Consider then the justification that obtains when the ground of a belief is as follows: the fact that p causes S to believe that p , via an intermediary mechanism M that is reliable. Since part of the grounds in this case involve the relevant fact causing S 's belief that p , S 's belief scores very well indeed with respect to the counterpossible worlds ruled out by her grounds. And for good measure we can always add that, since M is reliable, in the nearest worlds in which it produces the belief that p , it is true that p .²⁰ These considerations support the claim that S 's belief that p enjoys a very strong justification. Contrast this with the justification that obtains when the ground of S 's belief that p is an appearance (as in: being appeared to p -ly?). In any world in which the same appearances obtain (i.e., the same appearance-types are tokened), S will form the same belief. But the appearance-types themselves are consistent with various worlds in which $\sim p$. So if any of these apparently-similar-but-actually-non- p worlds are nearby, S will falsely believe that p in those nearby worlds. In that case, the justification enjoyed by S 's appearance-grounded belief is correspondingly weaker than it was in the case in which the relevant fact causes the belief via a reliable mechanism.

With this as preliminary, we can get a clear sense of the loss in strength of the justification of Peter's belief in (1) at t_3 , compared to the strength of justification that was enjoyed by Peter's belief in (1) at t_1 . We might say that at t_1 his belief in (1) enjoys both a strong *externalist* justification (in the sense that it was *the fact itself* that caused his belief, via a reliable belief-producing mechanism) as well as an *appearance-based* justification (in the sense that the perceptual appearances provide some truth-indicative support for his belief in (1)). But by t_3 his belief in (1) has lost its strong externalist justification, precisely because, given the content-shift, the fact that Pavarotti once swam in Taupo (= the fact that elicited Peter's belief in (1) at t_1) is not truth-indicative with respect to what Peter believes in believing in (1) at t_3 (= that *Twin-Pavarotti* once swam in Taupo). This is not to say that at t_3 Peter's belief in (1) is no longer justified. After all, although the 'externalist' support just cited does not provide truth-indicative support for Peter's belief in (1) at t_3 , there remains the support provided the original perceptual appearances, as these appearances are preserved by memory. As a result, whether at t_3 Peter's belief in (1) is justified will turn on whether these appearance-grounds are sufficiently truth-indicative *vis-a-vis* what Peter believes at t_3 . But—and this is the key point—even if we agree that these appearance-grounds *are* sufficiently truth-indicative *vis-a-vis* what Peter believes at t_3 to justify Peter's belief in (1) at t_3 , there can be little doubt that this belief is no longer justified *with the same strength* as that enjoyed by Peter's belief in (1) at t_1 . In this sense, even if memory preserves the justifiedness of Peter's belief in (1), memory fails to preserve *the strength* of the justification for Peter's belief in (1) throughout this interval.

One curiosity is worth noting regarding the failure of Peter's memory to preserve the strength of the justification for Peter's belief in (1). This failure

occurs despite the facts, first, that the grounds on which Peter's belief in (1) was originally acquired are preserved by memory, and so remain intact, and second, that in the interval since the original acquisition of the belief (= t_1), Peter acquires no information that he could reasonably be expected to regard as casting new doubt on his original belief in (1). What is more, at t_3 Peter regards his belief in (1) at t_3 as simply preserving the content of his belief in (1) at t_1 , and although Peter is wrong here he is neither unreasonable nor otherwise epistemically remiss in so regarding his belief in (1) at t_3 .²¹ The result of all of this is that Peter would not be epistemically remiss at t_3 were he to suppose that his belief in (1) is as strongly supported at t_3 as 'it' was at t_1 . Indeed, his confident use of (1) in his reasoning at t_3 suggests that Peter himself does suppose as much. And yet the supposition in question is *false*: the epistemic support for his belief in (1) throughout the interval from t_1 to t_3 is not invariant. In what follows my claims will be that these facts have implications for the discursive justification Peter's belief in (3) enjoys, when this belief is acquired as the result of an explicit inference (at t_3) from his beliefs in (1) and (2); and that these implications, in turn, suggest that SAI has some revisionary implications with respect to the epistemology of discursive justification.

4. SAI and the Epistemology of Discursive Justification

The lesson I seek to derive from the foregoing discussion of the Schiffer/Burge reply to Boghossian can be put as a lesson about the nature of content-preservative memory and its role in the epistemology of reasoning. Boghossian's Peter brings out the fact that there are at least two distinct roles played by content-preserving memory in cases of extended reasoning: such memory can preserve the content of the *original* belief acquired (call this 'originary' content preservation), or else it can underwrite anaphoric reference-preserving links between contents expressed *in the course of an extended piece of reasoning* (call this 'univocality-ensuring' content preservation). Normally, these two roles go hand-in-hand. But it is part of the case of Boghossian's Peter that these two roles are teased apart.²² For his part Boghossian had assumed (implicitly, and without argument) that Peter's memory would be content-preserving in the 'originary' sense; and Schiffer and Burge pointed out (against Boghossian's characterization of Peter) that most reasoners reason with premises that are kept in place by 'univocality-ensuring' content-preserving memory. However, the result from section 3 was that, in cases like that of Boghossian's Peter, reliance on univocality-ensuring content-preserving memory will have the effect that (at the time of reasoning) memory will fail to preserve the strength of the justification, if not the justificatory status itself, of the subject's belief(s) in at least one of the premises of the valid argument. In this final section I argue that this result is relevant to the epistemology of deductive reasoning.

I begin with some brief remarks about the epistemology of deductive reasoning.

It is natural to suppose that, given a logically competent subject *S*, the justification enjoyed by *S*'s belief in the conclusion of what appears to *S*'s searching reflection to be a valid inference has two epistemically internalist features. (In calling these features 'epistemically internalist' I have in mind that they are part of the materials which ground the justification enjoyed by beliefs acquired through inference, and yet are such that their presence or absence from a particular case can be discerned by an agent's searching reflection.) These two epistemically internalist features correspond to the two distinct roles memory plays in deductive reasoning. The first feature is this: it is natural to suppose that, with respect to her belief that *p* acquired via a valid inference at t_r , *S* can tell from the armchair that the following epistemic conditional holds:

- (I1) Unless I [= *S*] am presently [= at t_r] suffering from some relevant cognitive malfunction, my conclusion-belief that *p* enjoys a degree of justification that is at least as great as the degree of justification enjoyed by my belief in the conjunction of the premises at t_r .²³

As I will note below, (I1) reflects features of valid (truth-preserving) reasoning, and hence reflects features of the kind of *univocality-ensuring* content-preserving memory that underwrites valid reasoning. The second epistemically internalist feature of beliefs acquired via deductive reasoning has to do with the *originary* content-preserving role memory plays in reasoning. It is natural to suppose that logically-competent *S* can tell from the armchair that the following epistemic conditional holds of her belief in the conjunction of the premises at t_r :

- (I2) If I [= *S*] am not presently [= at t_r] suffering from a relevant cognitive malfunction and have not acquired any relevant new evidence regarding any of my premise-beliefs since the time I acquired them (= t_{Orig1} , t_{Orig2} , etc.), the degree of justification of my belief in the conjunction of the premises at t_r is exclusively determined by the degrees of justification enjoyed by each of my premise-beliefs, individually, *at the time each was acquired* (= t_{Orig1} , t_{Orig2} , etc.).

The idea behind (I2) is this. In the typical case of reasoning, the premises a subject uses are beliefs she acquired at some earlier time. She uses them because she supposes that (given no new relevant evidence) these beliefs continue to enjoy the justification they did on acquisition. This use reflects the assumption that (given no new relevant evidence) the justification enjoyed by a belief on its acquisition is preserved along with the memorial sustainment of the belief itself. Given limitations in humans' cognitive resources, it is important that this supposition be made: we cannot possibly keep track of each

of our beliefs' original grounds. It is also important that this supposition be warranted: without it, it is unclear how to vindicate the strong intuition that reasoning with beliefs acquired at earlier times often yields conclusion-beliefs that enjoy a high degree of justification, even under conditions in which (at the time of reasoning itself) we cannot recall the grounds for those premise-beliefs themselves. (See Burge 1993 for further discussion.)

Below I will expand a bit on (I1) and (I2). But for now I want to develop the thesis that, if SAI is true, then at the very least there are occasions on which *one of (I1) and (I2) is false of a logically competent reasoner*, where the logically competent reasoner herself cannot determine this by reflection alone. From a certain vantage point, this result should not be surprising. After all, the main result of section 3 was that, given SAI, there are cases (like that of Boghossian's Peter) in which memory's two reasoning-related content-preserving roles come apart, such that there is no way for memory to play both roles at the same time; and (I1) and (I2), respectively, reflect these two distinct content-preserving roles. Of course, if there are cases in which one of (I1) and (I2) is *false* of a logically competent reasoner, then there are cases in which such a reasoner cannot *know* (either from the armchair or in any other way) that both (I1) and (I2) hold of her belief acquired via inference. And if this is so, then discursive justification—in particular, the sort of justification one has for a belief one acquires via what appears to one to be a valid inference—does not have all of the features that, pretheoretically, it is natural to take it to have. Or so I will be arguing in what follows.

I begin with (I1). It is natural to suppose that, with respect to a belief that p acquired via a valid inference, a logically competent subject S can tell from the armchair that the conditional (I1) holds of her belief that p . We can begin here with some comments on the nature of valid inference. Not only is valid reasoning truth-preserving, it is also *justification-preserving*. That is, just as valid reasoning will never take one from a set of true premises to a false conclusion, so too valid reasoning will never take one from a set of justified premises to an unjustified conclusion.²⁴ The basis for the justification-preserving nature of valid deductive reasoning can be brought out as follows. Let G_1, \dots, G_n be the grounds S has for believing each of E_1, \dots, E_n , respectively. And suppose that each of G_1, \dots, G_n provides (truth-conducive) epistemic support for the belief so grounded, so that each of these beliefs is justified. Since each of G_1, \dots, G_n provides (truth-conducive) epistemic support for S 's corresponding beliefs in each of E_1, \dots, E_n , respectively, the result is that, taken as a whole, G_1, \dots, G_n provide (truth-conducive) epistemic support for taking the world to be as one believes to it be, when one believes $E_1 \& \dots \& E_n$. If $\{E_1, \dots, E_n\} \vdash p$ (where ' \vdash ' is the relation of logical implication), then any world in which $E_1 \& \dots \& E_n$ is true is a world in which p is true. Since it is already given that G_1, \dots, G_n provide truth-conducive epistemic support for the conjunctive proposition $E_1 \& \dots \& E_n$, and that any such world is a world in which it is true that p , the result is that G_1, \dots, G_n , taken together,

provide (truth-conducive) epistemic support for S 's belief that p as well—thereby justifying the belief in the validly-inferred-to conclusion. Of course, the foregoing does not ensure that belief in any conclusion of a valid argument will be justified. Rather, it ensures that a belief in the conclusion of a valid argument *each of whose premises are justified* will be justified.

It is important to note that, since the foregoing reasoning (in support of the justification-preserving nature of valid reasoning) turns on facts regarding the truth-preserving nature of valid reasoning, it can be produced and appreciated from the armchair. This important, since it is then natural to suppose that, at least in principle, a subject can know from the armchair that deductively valid reasoning is characterized by these epistemic features. So to the extent that such a subject can know from the armchair that a given piece of her reasoning is valid, she can conclude from the armchair that her conclusion-belief is conditionally justified (= justified if each of the premises are justified), and that the degree of justification enjoyed by her conclusion-belief at the time of her reasoning (= t_r) is at least as great as the degree of justification enjoyed by her belief in the conjunction of the premises at t_r . Of course, there is always the possibility that what she takes to be a valid inference is not; but if she is logically competent (and her reasoning is not very involved), such an outcome would reflect some relevant cognitive malfunction—and this is something she herself can appreciate by reflecting on the possible sources of error in validity judgements. The result is that she can know via reflection that, so long as she is not suffering from some relevant cognitive malfunction, the degree of justification enjoyed by her conclusion-belief at t_r is at least as great as the degree of justification enjoyed by her belief in the conjunction of the premises at t_r . Precisely at (I1) would have it.

Consider then (I2). It is natural to think that a logically competent subject S can tell reflectively that (I2) holds of her belief in the conjunction of her argument's premises. This can be established by establishing two further claims. First, for any of her premise-beliefs that p_i figuring in her inference at t_r , if there has been no new relevant evidence since she originally acquired this belief, then at t_r her belief that p_i continues to enjoy whatever degree of justification it enjoyed on acquisition. Second, this first claim can be arrived at via reflection on the originary content-preserving role of memory in reasoning.

I argue for the first claim by *reductio*. Suppose that S acquires her belief that p_i at t_i on grounds G_i , and that at t_i G_i provides truth-conducive epistemic support to S 's belief that p_i , rendering it justified at t_i . (This latter is inessential to the argument, but makes the point vivid.) And suppose that in the time since t_i there is no new evidence bearing on whether p_i . Now suppose (for our *reductio*) that G_i does not continue to provide the same degree of truth-conducive epistemic support to S 's belief that p_i at t_{present} , as it did at t_i . Now given limitations on memory, for at least a good many beliefs, the grounds that can be explicitly provided for a belief at the time at which the belief is

recalled will not provide the same degree of truth-conducive support as the belief's original grounds provided at the time at which the belief was originally acquired. This will happen, for example, if *S* forms her belief that p_i on excellent grounds, but over time proceeds to forget the grounds on which she acquired her belief that p_i , yet still sustains the belief itself (under conditions in which her memory is properly functioning etc.). In such cases the effect of our *reductio* supposition is clear: the degree of justification such beliefs enjoy *at the time of recollection* will be significantly less than the degree of justification they enjoyed at acquisition.

Now I have no knock-down argument against such a result. But it is worth noting that this result leads to some very unnatural conclusions regarding our reliance on memory in the fabric of our justified beliefs. For in that case, even if one acquires no new relevant evidence, one's beliefs sustained through properly-functioning memory will not, in general, enjoy the same degree of truth-conducive support as they did on attainment, with the further result that many of them will fail to be justified on recollection. And thus it seems that our *reductio* supposition is a recipe for skepticism regarding the prevalence of memorially-sustained justified belief. I take it that such a form of skepticism is implausible on its face. At any rate such a form of skepticism is less plausible than is the claim that (barring new evidence) a belief's original grounds continue to provide the same degree of support for it throughout the period during which the belief is sustained by memory. And this comparative plausibility claim is enough to establish my point, which is that, on the natural view of discursive justification, the grounds on which a belief is acquired *do* continue (barring relevant new evidence) to provide the same truth-conducive epistemic support to a belief throughout the interval during which memory sustains the belief.²⁵ Indeed, the forgoing reasoning involves nothing more than reflection on the role of originary content-preservation in reasoning—thereby confirming our second claim, to the effect that a logically competent subject *S* can tell *by reflection* that (I2) holds of a belief she acquires through deductive inference.

My claim so far has been that, *by the lights of a natural (if never fully articulated) account of discursive justification*, logically competent reasoners can always tell by reflection whether (I1) and (I2) are true of pieces of their own reasoning. The result is that (by the lights of this natural view) such subjects can discern by reflection certain epistemic properties possessed by their belief in the conclusion of a deductive argument. To wit, such a subject can tell by reflection that the degree of justification enjoyed by her conclusion-belief is at least that of the degree of justification enjoyed by her belief in the conjunction of the premises; and that the degree of justification enjoyed by her belief in the conjunction of the premises is determined by the degree of justification enjoyed by her premise-beliefs *at the time of acquisition*, together with the (positive or negative) support provided by any relevant new evidence acquired in the interim. To be sure, not every subject engages in the reasoning

that leads to these conclusions; and we still need to make allowances for the possibility of cognitive malfunctions. But once we do so, we arrive at a picture of discursive justification on which ordinary and univocal-ensuring content-preserving memory function together to enable a logically competent subject to reach a reflective appreciation of these epistemic properties possessed by her belief in the conclusion of an apparently-valid inference.

It is precisely this picture that needs to be revised if SAI is true. This is for the reason anticipated above: on the assumption of SAI, cases are possible in which one of (I1) and (I2) will fail to hold of a subject's belief in the conclusion of an apparently valid inference, yet where the subject herself is not in a position to appreciate this via reflection alone. To see this, we need only review the dialectic that emerges from the case of Boghossian's Peter.

As Boghossian had described the case, Peter's memory had preserved the original contents of his beliefs. In such a case, the degree of justification Peter's belief in the conjunction of the premises, here construed as preserving the original belief-contents acquired, is quite high. We saw this above when we saw that, at the time it was acquired, his belief in (1) enjoyed both a strong externalist justification as well as a justification grounded in the perceptual appearances.²⁶ The same could be said for Peter's belief in (2) at the time it was acquired. Since his memory is working properly, there is no content-shift introduced, and Peter has acquired no new relevant evidence in the interim, it would appear that, at the time of his reasoning ($= t_3$), his belief in (1) and his belief in (2) continue to enjoy a strong externalist justification as well as an appearance-grounded justification. So his belief in *the conjunction* of the two premises enjoys these strong externalist and appearance-grounded justifications: this merely reflects the logical and epistemic properties of conjunction-introduction. And, given that Peter is logically competent, he appreciates the truth- and justification-preserving nature of conjunction-introduction. Unfortunately for him, though, his conclusion-belief in (3) is false; and in any case the facts, first, that he saw Pavarotti swim in Taupo, and second, that he heard twin-Pavarotti sing last night, do not provide his conclusion-belief (that twin-Pavarotti once swam in Taupo) with anything approximating an externalist justification. At best, his conclusion-belief enjoys an appearance-grounded justification. In that case, at the time of his reasoning, his belief in the conclusion of an apparently valid argument does *not* enjoy a justification that is at least as great as the justification enjoyed by his belief in the conjunction of the premises. (I1) is false of him, yet Peter is in no position to appreciate this via reflection alone.

Suppose however that, rather than reasoning in the manner Boghossian describes, Peter were to reason with the univocal-ensuring intention described in the Schiffer/Burge response to Boghossian. Reasoning with such an intention would render Peter's reasoning valid. In that case, (I1) would be true of Peter: his belief in the conclusion of this apparently valid argument enjoys a justification that is at least as great as the justification enjoyed by

his belief in the conjunction of the premises (construed so that they logically imply the conclusion). But it is obvious that the cost of this is that such a result renders (I2) false. In particular, although Peter is logically competent, is suffering from no cognitive malfunction at the time of his reasoning, and has acquired no new relevant evidence in the time since originally acquiring his premise-beliefs, nevertheless at the time of his reasoning his belief in the conjunction of the premises does not enjoy a degree of justification that is exhaustively determined by the degrees of justification that were enjoyed by each of his premise-beliefs, individually, *at the time each was acquired* (t_1 and t_2 , respectively). To see this we need only contrast the degree of justification enjoyed by a belief in the conjunction of the premises *where these are construed as preserving the content of Peter's original beliefs in (1) and (2)*, with the degree of justification enjoyed by a belief in the conjunction of the premises where at least one of the conjuncts *is the result of a content-shift*. Above we saw that the former conjunction-belief would enjoy both a strong externalist as well as an appearance-grounded justification. But the latter conjunction-belief would not enjoy the strong externalist justification: precisely not, since the latter conjunction-belief has a conjunct which is the result of a content-shift, and above we saw that Peter's content-shifted belief is not justified to the same degree as his original belief was. So we see that (I2) is false of Peter, yet he is in no position to appreciate this via reflection alone.

Can this last result be resisted by arguing that in such a case Peter's failure to preserve the originary content is itself a cognitive malfunction? If so, then Peter (construed as reasoning with the univocality intention) is not a counterexample to (I2) after all, since one of the conjuncts in (I2)'s antecedent condition would not then be met. But it seems wrongheaded to think that Peter's failure to preserve the originary content is itself a cognitive malfunction. The case is rather one in which two distinct cognitive aims—valid reasoning,²⁷ originary memorial content preservation—cannot both be met. Consequently, the case is one in which Peter's cognitive system is acting in such a way as to reach one of its goals, in a situation in which realizing both is impossible. But satisfying one goal, in a case involving conflicting aims which preclude the satisfaction of both goals, is not itself a case of cognitive malfunction. Quite the contrary, when the goals are impossible to co-satisfy, satisfying one of them at the expense of the other may be the most rational thing to do.²⁸ So the objection fails.

Might it then be argued that Peter's dealings with twin-Pavarotti constitute *new relevant evidence*—perhaps relevant to his belief in (1)? If so, then one could maintain that (I2) is true of Peter since one of the conjuncts in its antecedent is false. However, such a suggestion brings with it an unhappy epistemic implication that, in effect, is tantamount to the very thesis I am trying to establish. For if Peter is to be counted at t_3 as having acquired relevant new evidence, then his is a case in which a logically competent subject fails to appreciate the epistemic bearing of the belief that constitutes his

evidence on his other beliefs, despite the fact that he has all of the relevant beliefs 'in mind' at the same time, grasps all of the relevant concepts in each, and is suffering from no cognitive malfunction. The result would then be that such a subject would not be in a position to tell by reflection how (if at all) the acquisition of the one belief ought to affect his confidence regarding the truth of the other. Such a result is, if not the same claim, then very close to the thesis for which I am presently arguing, namely, that if SAI is true then subjects will not always be in a position to determine the strength of support given by premises to a conclusion, and so will not always be able to determine the degree of justification enjoyed by their belief in the conclusion of an apparently valid inference. So this objection cannot help one avoid the conclusion of this paper; and I shall regard it as a non-starter.²⁹

I have been arguing that, no matter how one describes the case of Boghossian's Peter, at least one of (I1) and (I2) will fail to hold. I now want to explore why this result is significant; I do so by suggesting what implications it has for the natural picture of discursive justification characterized above. I begin first with a comment on the epistemic import of (I1) and (I2). We can think of these principles as underwriting a claim about the confidence a logically competent subject is entitled to have in the conclusion of an apparently valid argument. It is an epistemic truism that one's confidence in the truth of a belief ought not to outstrip the epistemic good-making properties of that belief. (A familiar version of this truism is that a subject's degree of confidence should not outstrip her evidence.) (I1) and (I2) tell us how this truism might profitably be applied to reasoning-based beliefs. Simply put, in those cases in which (I1) and (I2) hold of a subject's reasoning-based belief, the result is that, if she is in conformity with the truism with respect to each of her premise-beliefs, she will be in conformity with the truism with respect to her conclusion-belief if she has at least as much confidence in the truth of the conclusion, as she has in the truth of the conjunction of the premises.³⁰ In light of this, we might then say that in such cases she is *entitled* to have at least as much confidence in the truth of the conclusion, as she has in the truth of the conjunction of the premises. Now suppose further that a logically competent subject can tell by reflection that (I1) and (I2) hold of her conclusion-belief on a given occasion. Then she is in a position to see via reflection alone that she is entitled to have at least as much confidence in the truth of the conclusion, as she has in the truth of the conjunction of the premises.³¹ This reflects an epistemically internalist dimension to discursive justification.

How do the results obtained above bear on this? Our result need not be taken as bearing against the thesis that, when (I1) and (I2) *do* hold of a subject, she can tell via reflection alone that they do. Whether our result does bear against such a thesis depends on what we take to be the conditions on knowing something by reflection, i.e., on what counterpossibilities the knowing subject must be able to rule out via reflection. But no matter what view we take on such matters, and so no matter what view we take of whether

(with respect to the holding of (I1) and (I2)) the reflective method is *complete*,³² the present result suggests that the reflective method is not *sound*: there will be cases when a logically competent subject thinks via reflection that (I1) and (I2) do hold of her reasoning-based belief that *p* on occasion *O*, arrives at this conclusion under conditions in which she is suffering from no cognitive defect, and yet it is not the case that (I1) and (I2) both hold of her reasoning-based belief that *p* on *O*. So even if a logically competent subject can always tell by reflection that (I1) and (I2) do hold *when they do*, it is not always the case that, if she thinks via reflection that (I1) and (I2) both hold, then (I1) and (I2) both hold. At the very least, this points to a limitation in the epistemically internalist feature noted above: even after it is granted that a logically competent subject can reflectively tell that she is entitled to have at least as much confidence in the truth of her conclusion-belief as she has in the truth of the conjunction of her premise-beliefs, it is not the case that she can reflectively tell when (in particular cases) she is *not* so entitled.

So far I have been granting for the sake of argument that when (I1) and (I2) *do* hold of a subject, she can tell via reflection alone that they do so. Are we then warranted in saying that, in such cases, the subject *knows via reflection* that (I1) and (I2) hold of her? This depends on what counterpossibilities she must be able to rule out via reflection, in order to count as knowing via reflection. If the sort of possibility actualized in the case of Peter is not relevant to our subject, then the fact that the subject cannot rule out such a possibility is irrelevant to whether she knows, and presumably we would be warranted in saying that the subject knows via reflection that (I1) and (I2) hold of her. Of course, where this kind of counterpossibility *is* relevant, then, since she cannot rule it out via reflection alone, she does not count as knowing via reflection that (I1) and (I2) hold of her. So when is the possibility at play in the case of Boghossian's Peter relevant?

Granting that cases of this sort do not occur with great frequency, even so such a possibility might become relevant, e.g., if someone were to introduce such a possibility into a situation, say, by raising it explicitly. One motive for doing so would be the skeptical motive; the result would be a new version of knowledge-skepticism, whose target is (not perceptual knowledge but rather) *discursive* knowledge. But the skeptical motive is not the only motive for introducing the possibility of an introspectively undetectable equivocation in reasoning. Perhaps there are cases in which the hypothesis, that the reasoner has vacillated in the manner of Boghossian's Peter, has become live. Such an hypothesis might become live with respect to the subject's use of *proper names* in an extended chain of reasoning; more generally, such an hypothesis might become live with respect to the subject's use of any expression that has more than one lexical entry in its dictionary definition. Take 'pragmatist' for example. Someone ignorant of the difference between the specialized philosophical sense of 'pragmatist' and the more everyday sense might reason in such a way as to exhibit an introspectively undetectable vacillation with her

use of ‘pragmatist’, with the result that she might come to be in a situation parallel to that of Peter. What is more, most of us are such that, for a great many word-forms w in our lexicon, we are not in a position to determine via reflection whether w has more than one lexical entry in its dictionary definition. But this means that, were the query to come up, we could not rule it out via reflection. In such cases we do not count as knowing via reflection that (I1) and (I2) hold of us. The prevalence of the revisionary implications of our present result may in fact be greater than one would like to suppose.

Still, as significant as this result is, it is important not to overstate the case. What the case of Boghossian’s Peter shows is that discursive justification lacks an epistemically internalist feature that, pretheoretically, it is natural to think it has. We can put the point as follows: even a logically competent subject who is free of cognitive malfunctions cannot always tell by reflection what degree of confidence she is entitled to have with respect to the conclusion of an apparently valid argument. Of course, people are typically ill-placed to determine the level of confidence they are entitled to have, given their evidence for a proposition. But the present case is worse: even if a person were apprized of the level of confidence warranted by the evidence at the time each premise-belief was acquired, knew that there was no new relevant evidence to consider, and knew as well that she was suffering from no cognitive (including memory) malfunction, *still* she could not reflectively determine the level of confidence she was entitled to have in a proposition that she took to be logically implied by these premise-beliefs. The same point can be put more formally. Given SAI, there is no sound and complete reflective method by which the subject can fill in the following, in cases in which the discursive justification is the only justification for her belief that p : given that j_1, \dots, j_n are the memory-preserved degrees of justification enjoyed by premises p_1, \dots, p_n of an apparently valid argument whose conclusion is that p , and given that I am not presently suffering from any cognitive malfunction, the degree of justification of my belief that p is $f(j_1, \dots, j_n)$ (where f is some function defined over ordered n -tuples whose members are (a representation of) degrees of justification).

5. Conclusion

The thesis of this paper is that semantic anti-individualism has some revisionary implications regarding the nature of discursive justification. No matter how one describes Boghossian’s Peter—that is, whether one describes Peter (with Boghossian) in such a way that Peter’s memory is content-preserving in the ordinary sense, or one describes him (in the spirit of the suggestion of Schiffer and Burge) in such a way that Peter’s memory is content-preserving in the univocality-ensuring sense—it is not the case that (I1) and (I2) hold. The result is that, since memory cannot play *both* of its reasoning-supporting content-preserving roles, and since Peter is not in a position via

reflection alone to determine this, Peter is not in a position via reflection alone to appreciate the degree of justification his reasoning-based belief has, as a function of the degrees of justification enjoyed by his premise-beliefs on their acquisition. The further result is that, even if he pegs the confidence he has in the conclusion to the confidence he had on acquiring each of the premise-beliefs, he cannot tell via reflection alone whether he is violating the epistemic truism enjoining him to tailor his degree of confidence in the truth of his belief to, e.g., his evidence. And this is true despite the fact that Peter is a logically competent subject who is suffering from no cognitive defect. This illustrates what I regard as SAI's main revisionary implication for the nature of discursive justification.

This result is not intended to be a criticism of SAI. Rather, it is intended as part of the systematic attempt to identify SAI's epistemic implications. This attempt to date has been overly narrow in its orientation: insofar as most discussions of the epistemic implications of SAI focus only on self-knowledge and external-world skepticism, we have failed to see what are arguably among the most significant of SAI's implications for epistemology. Even after it is allowed that SAI has no untoward implications regarding self-knowledge, the question remains how *what* one knows when one has such self-knowledge relates to the rest of one's knowledge and justified belief, and in particular what roles such knowledge can play as one attempts to keep one's epistemic house in order. On this score SAI's implications have still not been fully appreciated. The objective of the present paper—to assess SAI's implications for the epistemology of reasoning—is meant to be a contribution to that greater project.³³

Notes

¹ The topic of SAI's implications *for reasoning* has been explored not only by Boghossian but also (from the side of defenders of SAI) by Jessica Brown (2004, Chapter 5). However, neither Boghossian nor Brown focuses on the implications of SAI for the sort of *epistemic justification* involved in the acquisition of belief through reasoning. As noted, Boghossian is interested in the *a priority* of logical abilities. And for her part, Brown is interested in whether Boghossian's conclusion (which she accepts) ought to prompt us to surrender SAI itself; focusing on issues of rationality, she argues that it should not.

² See e.g. Burge 1988 as well as various recent defenses of Burge 1988, including (but not limited to) Falvey and Owens 1994, Gibbons 1996, Burge 1996a, Brueckner 1997, McLaughlin and Tye 1998a and 1998b, Goldberg 2003a and 2003b, and Brown 2004.

³ I have attempted to address other aspects of this topic in Goldberg 1999, 2003a, and 2003b.

⁴ I myself do not accept this assumption. But for those who do—and for a very interesting development of the argument-based conception of justification, see Audi 1988—my argument will have the sort of generality indicated.

⁵ Burge allows that there will be cases in which this is not so, but argues that in such cases the flaw in the reasoning will not be a flaw in the reasoner's *logical abilities* but rather a *false presupposition* that she is making (1996b: 367-68). This point strikes me as correct, but as also coming at an epistemic cost not explored by Burge. I address this in section 4.

⁶ This is a world in which all reasoners reason in the manner suggested by the Schiffer/Burge point about anaphoric, reference-preserving links.

⁷ In this and the remainder of my discussion, I will be focusing on reasoning in contexts in which it makes sense to speak of the reasoner's semantic intentions, e.g., to be using her terms univocally throughout the extended reasoning. A reader might infer from this that my comments affect only reasoning that is actually done linguistically (in speech or the written word), but not in thought (where, it might be supposed, talk of semantic intentions makes little sense). In response I have two comments. First, even if this inference were correct (and my argument here concerned only reasoning in language as opposed to thought), the conclusions that I draw would *still* be of interest, since in that case the claim would be that SAI has implications for discursive justification in connection with arguments presented in speech or the written word. (It would then be an interesting implication of this conclusion, that there is an important difference between discursive justification in cases of spoken or written arguments, and that in cases of arguments performed in thought!) But second, I think the inference in question is not acceptable. In any case in which Peter thinks a thought that he would express linguistically with a use of the name 'Pavarotti', something must determine the truth conditions of the thought—in particular, whether those truth conditions regard the antics of earth-Pavarotti or twin-Pavarotti. I submit that what determines this is the analogue (in the case of thought) of what Peter's semantic intentions are in the case of language. See Goldberg 2005 for a development of this point.

⁸ I do not want to get embroiled in a discussion of the nature of semantic values; I am using 'concept' to be the mental analogue of the expression's semantic value. (On some views of semantic value, the concept expressed will *be* the expression's semantic value; I want to remain neutral on this.) My argument turns on what every party should concede: whatever semantic values are (and whatever their relation to concepts expressed), SAI implies that Peter's use of 'Pavarotti' at t_1 (on Earth) has a different semantic value than his use of 'Pavarotti' at t_2 (on Twin Earth).

⁹ This result merely depends on things that any proponent of SAI will grant, namely, that non-indexical expressions with different references express different concepts, and that the token of 'Pavarotti' Peter produced in his utterance of (1) at t_1 has a different reference from that of the token he produced in his utterance of (2) at t_2 .

¹⁰ See Goldberg 2005 for a detailed discussion of the Ludlow/Gibbons account of memory.

¹¹ It might be thought that this begs the question against *coherence theories* of epistemic support. But if the argument to be presented in terms of this *ground-theoretic* conception of epistemic support goes through, a modification of the argument can be adopted as an argument *against* the coherence-theoretic conception of support. As I will present it, the argument's key point is that if the content of 'a' belief shifts via memory, then this will affect the truth-conducive epistemic support that the recalled belief enjoys (relative to the support that 'the' belief enjoyed when it was first acquired). If the coherence theorist denies this, arguing that the recalled belief with the shifted content coheres equally well with the rest of the subject's belief corpus, then *this is so much the worse for coherence theories of epistemic support*. For in that case it will have been shown that coherence-generated support fails to distinguish between content-preserving and non-content-preserving memory processes—and surely *that* is a flaw in a theory of epistemic support (or of epistemic justification)!

¹² I am here using 'justified' in the sense in which it is beliefs, rather than believers, that are said to be justified (in the first instance). If desired, we could formulate a notion of *being justified in believing*; but since my main points can be made in terms of the justification of belief (not believers), I will continue to speak in this fashion—except on those few occasions when clarity will be served by explicitly bringing in the justification of a believer (on a given occasion).

¹³ Below I will consider a proposal according to which Peter's grounds at t_1 for believing that p_1 , and his grounds at t_2 for believing that p_2 , consist of *the appearances* (i.e., its *seeming to him* that he saw someone *just like Pavarotti* swimming in Taupo at t_1 , and its *seeming to him*

that he heard someone *just like twin-Pavarotti* singing last night). On such a view, a case can be made for thinking that, contrary to what I just wrote, Peter's belief in p_3 is justified at t_3 . A slight variation in my argument, however, would succeed in showing that, even so, Peter's belief in p_3 is not justified to a degree equal to that of his belief in the conjunction of the premises. I will develop this point below, where I argue that this constitutes the clearest implication SAI has for the epistemology of reasoning.

¹⁴ Which original belief shifts in content will depend on background assumptions that the proponent of SAI brings to bear in determining the content of the memory-informed belief present at the time of reasoning. Some analyses will treat the content of the memory-informed judgement to be determined by the subject's present linguistic community; this is the view of Ludlow 1995, Gibbons 1996, and Tye 1998. Others, however, will treat the content of the memory-informed judgement to be determined by other factors. Thus Burge 1996b holds that in the case under discussion what fixes the content of the belief used as premises is which premise is presented first (since terms used in later premises are 'calibrated' as it were by reference to earlier terms). If Burge's 1996b views are applied to the present case, ($\neq 2$) holds. However, for my purposes here it does not matter which of the two contents is shifted; the argument could be presented on either account. I should add that there will be those who prefer to analyze the case as one in which the name 'Pavarotti' in Peter's lexicon expresses something like a disjunctive concept, with singular concepts of earth-Pavarotti and twin-Pavarotti as the disjuncts; see e.g. Heal 1998, who suggests an analysis of would-be world-switching cases that might be thought to have this result (though Heal herself does not address Peter's case). But for reasons that I present in Goldberg 2005, I find such a disjunctive analysis independently implausible.

¹⁵ See my Goldberg 1999, where I explicitly take up the issue of SAI's implications for the validity *and soundness* of our arguments.

¹⁶ I thank an anonymous referee for this point.

¹⁷ 'Belief in (1)' is neutral as between the two readings in play. Since the content of the belief originally acquired at t_1 has shifted by t_3 , this locution offers the only way to speak of "the" belief acquired at t_1 and recalled at t_3 .

¹⁸ This talk of 'degrees' or 'strength of justification' can be formalized in many different ways. Perhaps the most obvious way would be in terms of probabilistic epistemology, where strength of justification can be identified with probability on one's evidence.

¹⁹ Other 'externalist' criteria might be substituted for the foregoing; but the point I am trying to make is not affected by which of the various externalist criteria one chooses.

²⁰ This does not rule out the possibility that in the nearest $\sim p$ -world S believes that p . Rather, it makes such a world relatively far away, at least as ordered by M : the reliability of M ensures that, in the nearest worlds in which S acquires the belief that p via M , p is true.

²¹ Here, of course, I am speaking of the *justification for believing*, rather than the justification of a *belief*. I mention this so as to avoid introducing confusion.

²² I should underline that the scope of ordinary content-preservation is jeopardized in the proposals regarding the semantics of memory in Ludlow 1995 and Gibbons 1996. Roughly put (but see Goldberg 2005 for more detailed discussion), the semantics of a memory judgement on the Ludlow-Gibbons line reflect the subject's linguistic environment at the time of recollection, rather than her environment at the time of the original acquisition of the belief. The result is that if one has been embedded in a new environment for a sufficiently long period, one's memories will 'shift' the contents of beliefs whose linguistic expression involve terms whose semantics were affected by the change in environment. These proposals will have precisely the same implications that I am ascribing to the Schiffer/Burge anaphoric-link strategy: they undermine ordinary content-preservative memory in some cases.

²³ Indeed, it is an uncontentious principle of probabilistic epistemology that the probability that an entailed claim is true is *at least as great as* the probability that the conjunction of the premises is true. So if we can identify strength or degree of justification with probability on one's evidence, we would get the result that the strength or degree of justification enjoyed by

an entailed conclusion (relative to subject *S*) is at least as great as the strength or degree of justification enjoyed by *S*'s belief in the conjunction of the premises. (Thanks to Brad Monton for suggesting this way of making the point.)

²⁴ The thesis in question is a closure thesis, to the effect that justification is closed under entailment. As writers such as Crispin Wright and Martin Davies have noted, closure is one thing, justification *transmission* (or what Wright and Davies call 'warrant transmission') another. Thus even if the relevant closure principle is true, it is a *separate* matter whether a given justification-preserving inference *generates* a justification for belief in the conclusion; see Wright 1995, 2003, and Davies 2003. My claim pertains to the *preservation* of justification, rather than its transmission; and, as Wright and Davies note, the former claim is weaker than the latter.

Now there will be some who deny that justified belief is closed under entailment; such people would recoil at the claim that valid deductive inference is a justification-preserving process. However, I submit that the move to deny closure (here, as elsewhere) generates controversy; and this suggests that, at a minimum, the relevant closure principle can be regarded as part of the traditional conception of discursive justification. At any rate those who do not accept closure in this domain can treat the present argument as giving lessons to those proponents of SAI who aim to preserve the relevant closure principle.

²⁵ The claim here is that if we are to retain our natural view about the role of memory in reasoning, then we should reject our *reductio* supposition, in favor of the claim that (barring relevant new evidence) the grounds on which a belief is acquired continue to provide the same truth-conducive epistemic support to a belief throughout the interval during which memory sustains the belief.

²⁶ In speaking of a justification grounded in the appearances, I am being concessive to the Schiffer/Burge reply to Boghossian. For suppose that there is no such justification: the perceptual appearances alone (e.g., without mentioning the fact observed) don't justify Peter's belief at t_1 that Pavarotti swam in Taupo. But if the perceptual appearances don't justify Peter's *true* belief at t_1 , then surely the perceptual appearances don't justify Peter's *false* (because content-shifted) belief at t_3 ! And in that case, the Schiffer/Burge reply to Boghossian is left having to regard Peter's (content-shifted) belief in (1) at t_3 as lacking justification altogether.

²⁷ Perhaps this first cognitive aim ought to be put as follows: reasoning in a way that *it would be rational to suppose* is valid.

²⁸ Although it might be more rational to avoid trying to satisfy either (say, by not engaging in reasoning at all in such cases). Unfortunately, given what is available to him via reflection, Peter is not in a position to appreciate this. But he can hardly be considered to be suffering from a relevant cognitive malfunction for this; and so he can hardly be considered to be suffering from a relevant cognitive malfunction for reasoning in the way that leads to the difficulty.

²⁹ See also Goldberg 1999 and 2005, where I discuss further the position on which this objection is based.

³⁰ 'At least' as much: if the reasoning is the only grounds of her belief in the conclusion, then she should have exactly as much.

³¹ Here I assume that she is in conformity with the truism with respect to each of her premise-beliefs.

³² Let us say that the reflective method in question (regarding whether (I1) and (I2) holds of a subject *S*'s belief that *p* on occasion *O*) is *complete* iff the following conditional holds: for any case in which (I1) and (I2) hold of *S*'s belief that *p* on *O*, *S* can tell via reflection that (I1) and (I2) hold of her belief that *p* on *O*.

³³ I would like to thank Jessica Brown, Tony Brueckner, John Collins, Gary Ebbs, Richard Fumerton, Chris Gauker, Mikkel Gerken, Peter Graham, David Henderson, Michael Lynch, Andrew Melnyk, Brad Monton, Jim Pryor, and David Sosa, for helpful discussions of these matters; a referee from this journal, for comments on an earlier draft; and audiences at the Society for Skeptical Studies and the Southern Society for Philosophy and Psychology, where I have presented early drafts of this paper.

References

- Audi, R. 1986: "Belief, Reason, and Inference." Reprinted in Audi 1993.
- Audi, R. 1988: "Justification, Truth, and Reliability." Reprinted in Audi 1993.
- Audi, R. 1993: *The Structure of Justification*. (Cambridge: Cambridge University Press.)
- Boghossian, P. 1989: "Content and Self-Knowledge." *Philosophical Topics* 17.
- Boghossian, P. 1992: "Externalism and Inference," *Philosophical Issues* 2.
- Boghossian, P. 1994: "The Transparency of Mental Content." In *Philosophical Perspectives* 8: *Language and Logic*, ed. J. Tomberlin (Ridgeview: Atascadero).
- Bonjour, L. 1998: *In Defense of Pure Reason*. (Cambridge: Cambridge University Press.)
- Brown, J. 2000: "Critical Reasoning, Understanding and Self-Knowledge," *Philosophy and Phenomenological Research* 61:3, pp. 659–76.
- Brown, J. 2004: *Anti-Individualism and Knowledge*. (Cambridge, MA: MIT Press.)
- Bruecker, A. 1997: "Externalism and Memory." *Pacific Philosophical Quarterly* 78.
- Burge, T. 1986: "Intellectual Norms and the Foundations of Mind." *Journal of Philosophy* 83: 697–720.
- Burge, T. 1988a: "Individualism and Self-Knowledge." *Journal of Philosophy* 85.
- Burge, T. 1988b: "Wherein is Language Social?" In *Reflections on Chomsky* (Oxford: Basil Blackwell).
- Burge, T. 1993: "Content Preservation." *Philosophical Review* 102:4, pp. 457–88.
- Burge, T. 1996a: "Our Entitlement to Self-Knowledge." In Martin and Ludlow, eds., *Externalism and Self-Knowledge* (Palo Alto: CSLI Press).
- Burge, T. 1996b: "Memory and Self-Knowledge." In Martin and Ludlow, eds., *Externalism and Self-Knowledge* (Palo Alto: CSLI Press).
- Burge, T. 1998a: "Computer Proof, *Apriori* Knowledge, and Other Minds." *Philosophical Perspectives* 12: *Language, Mind, and Ontology*. J. Tomberlin, ed. (Cambridge: Blackwell.)
- Burge, T. 1998b: "Reason and the First-Person." In Wright, C., Smith, B., and Macdonald, C., eds. *Knowing Our Own Minds* (Oxford: Oxford University Press, 1998).
- Davies, M. 2003: "The Problem of Armchair Knowledge." In Nuccetelli, ed., *New Essays on Semantic Externalism and Self-Knowledge* (Cambridge, MA: MIT Press), pp. 23–56.
- Falvey, K., and Owens, J. 1994: "Externalism, Self-Knowledge, and Skepticism." *Philosophical Review* 103: 107–137.
- Gibbons, J. 1996: "Externalism and Knowledge of Content." *Philosophical Review* 105:3.
- Goldberg, S. 1999: "The Relevance of Discriminatory Knowledge of Content." *Pacific Philosophical Quarterly* 80:2, pp. 136–56.
- Goldberg, S. 2000: "Externalism and Self-Knowledge of Content: A New Incompatibilist Strategy." *Philosophical Studies* 100:1, 51–78.
- Goldberg, S. 2003a: "Anti-Individualism, Conceptual Omniscience, and Skepticism." *Philosophical Studies* 116:1, 53–78.
- Goldberg, S. 2003b: "What do you know when you know your own thoughts?" In Nuccetelli, S. ed., *New Essays on Semantic Externalism and Self-Knowledge* (Cambridge, MA: MIT Press, 2003), pp. 241–56.
- Goldberg, S. 2005: "(Non-standard) Lessons from World-Switching Cases." *Philosophia* 32:1, 95–131.
- Heal, J. 1998: "Externalism and Memory II." *Proceedings of the Aristotelian Society* 72 (Supplemental Volume), pp. 95–109.
- Ludlow, P. 1995: "Social Externalism, Self-Knowledge, and Memory." *Analysis* 55:3.
- McLaughlin, B., and Tye, M. 1998a: "Is Content Externalism Compatible with Privileged Access?" *Philosophical Review* 107:3, pp. 349–80.
- McLaughlin, B., and Tye, M. 1998b: "Externalism, Twin Earth, and Self-Knowledge." In C. Wright, B. Smith, and C. MacDonald, eds. *Knowing One's Own Mind* (Oxford: Oxford University Press), pp. 285–320.

- Owens, J. 1986: "Synonymy and the Non-Individualistic Model of the Mental." *Synthese* 66: 361–82.
- Schiffer, S. 1992: "Boghossian on Externalism and Inference." *Philosophical Issues* 2.
- Tye, M. 1998: "Externalism and Memory I." *Proceedings of the Aristotelian Society* 72 (Supplemental Volume), pp. 77–94.
- Wright, C. 1995: "Facts and Certainty." *Proceedings of the Aristotelian Society* 71: 429–72.
- Wright, C. 2003: "Some Reflections on the Acquisition of Warrant by Inference." In Nuccetelli, ed., *New Essays on Semantic Externalism and Self-Knowledge* (Cambridge, MA: MIT Press), pp. 57–78.